

Scrappy: SeCure Rate Assuring Protocol with Privacy

Kosei Akama
Keio University
akama@keio.jp

Yoshimichi Nakatsuka
ETH Zurich
yoshimichi.nakatsuka@inf.ethz.ch

Masaaki Sato
Tokai University and Keio University
masaaki@tsc.u-tokai.ac.jp

Keisuke Uehara
Keio University
kei@sfc.keio.ac.jp

Abstract—Preventing abusive activities caused by adversaries accessing online services at a rate exceeding that expected by websites has become an ever-increasing problem. CAPTCHAs and SMS authentication are widely used to provide a solution by implementing *rate limiting*, although they are becoming less effective, and some are considered privacy-invasive. In light of this, many studies have proposed better rate-limiting systems that protect the privacy of legitimate users while blocking malicious actors. However, they suffer from one or more shortcomings: (1) assume trust in the underlying hardware and (2) are vulnerable to side-channel attacks.

Motivated by the aforementioned issues, this paper proposes **Scrappy: SeCure Rate Assuring Protocol with Privacy**. **Scrappy** allows clients to generate unforgeable yet unlinkable *rate-assuring proofs*, which provides the server with cryptographic guarantees that the client is not misbehaving. We design **Scrappy** using a combination of DAA and hardware security devices. **Scrappy** is implemented over three types of devices, including one that can immediately be deployed in the real world. Our baseline evaluation shows that the end-to-end latency of **Scrappy** is minimal, taking only 0.32 seconds, and uses only 679 bytes of bandwidth when transferring necessary data. We also conduct an extensive security evaluation, showing that the rate-limiting capability of **Scrappy** is unaffected even if the hardware security device is compromised.

I. INTRODUCTION

The invention of the World Wide Web has facilitated the hosting of many services online. While many services typically anticipate users to access their resources at a moderate rate, malicious users often attempt to exceed these limits. Consequently, online services employ techniques to slow down such users, commonly referred to as *rate-limiting*. Several scenarios highlight the usefulness of rate-limiting users:

- **Online polls and product ratings:** to prevent manipulation of poll outcomes and ratings.
- **Services using third-party APIs (e.g., AI services, SMS text messages):** to prevent overuse and potential high billing.
- **Services with free trials:** to increase the conversion rate for sustainability.

- **Preventing dictionary attacks:** to limit the number of login attempts.
- **Online crawlers:** to limit the collection of data.

Two major techniques commonly adopted to slow down adversaries or limit their actions are SMS authentication and CAPTCHAs. SMS authentication authenticates users via phone numbers over the web while CAPTCHA (Completely Automated Public Turing test to tell Computers and Humans Apart) [14] distinguishes between humans and a machine (i.e., bot) through puzzles that are easy for humans to solve but difficult for bots. Although both methods were not originally designed for rate-limit users, many services employ them for this purpose, as observed by [47] and [51]. However, SMS authentication has privacy issues as all user actions are tied to a user's phone number, and studies have shown that CAPTCHAs degrade user experience [22], [59].

In response to these challenges, there have been proposals for rate-limiting alternatives to SMS authentication and CAPTCHAs with improved security, privacy, and efficiency. Opaak (OPen Anonymous Authentication framework) [47] allows users to prove possession of a unique phone number without revealing it to the server. CACTI (Captcha Avoidance via Client-side TEE Integration) [51] is a system that allows users to generate CAPTCHA-avoiding rate-proofs proving to servers that they are not engaging in abusive actions without compromising their privacy, utilizing client-side Trusted Execution Environments (TEEs). CAP (Cryptographic Attestation of Personhood) [48] and PrivacyPass [16], [30] provide CAPTCHA alternatives by using FIDO authenticators [17] in CAP and anonymous cryptographic tokens in PrivacyPass. However, these systems suffer from one or more of the following issues: (1) Security of the system relies on the security of the underlying hardware and/or (2) Vulnerability to side-channel attacks.

Motivated by the aforementioned challenges, this paper presents **Scrappy**, a novel system overcoming the limitations of prior work while providing rate-limiting capabilities. When contacted by users, the server defines a time window t and requests a *rate-assuring proof* over t . In essence, rate-assuring proofs show within which time window the user sent a request to a certain server. These proofs are *unforgeable* and, in **Scrappy**, are produced using widely available hardware security devices, acting as a source of *uniqueness*. Malicious users attempting to produce fake proofs are easily detectable.

As with prior work, ensuring *user privacy* is one of the important goals of **Scrappy**. This is achieved through

digital signatures that remain the same for the server-defined t but become *unlinkable* once the rate-assuring proof uses a different t or targets different servers. This prevents malicious servers from tracking users while blocking malicious users from accessing the server’s service too frequently.

Another highlight of *Scrappy* is that its rate-limiting feature *does not* rely on the security of the underlying hardware security device. Even if either the essential secret key proving the device’s uniqueness or the secret key used to generate rate-assuring proofs is leaked, servers can still block malicious users.

Scrappy is *agnostic* concerning the underlying hardware security device. We show this by implementing *Scrappy* using three different devices: TPM, hardware security token, and smartphone. The TPM implementation can be deployed straight out of the box, while the other implementations require a small modification to software and specifications. We also list devices other than the above three that can be used in *Scrappy* in Table VII.

To contribute to reproducible research, we have open-sourced our implementation of *Scrappy* using TPMs at [44]. We hope this will allow future research to take advantage of our system and help reduce the friction between users and services due to rate-limiting.

Contribution. The anticipated contributions of this work are:

- 1) The introduction of *rate-assuring proofs*, a security primitive allowing servers to prevent abusive activities from users while respecting their privacy.
- 2) *Scrappy*, a novel rate-limiting protocol that uses rate-assuring proofs and the Direct Anonymous Attestation protocol as building blocks and generic hardware security devices as a source of uniqueness.
- 3) Three proof-of-concept implementations of *Scrappy*, with one capable of being deployed out of the box.
- 4) Contribution towards reproducible research by open-sourcing implementation of *Scrappy* over a TPM.
- 5) A comprehensive analysis of the security of the protocol, notably showing some guarantee of privacy and rate-limiting even if keys are leaked.
- 6) An extensive evaluation of latency, bandwidth, and storage of *Scrappy*, showing that our system is practical for real-world usage.

II. RELATED WORK

This section summarizes the most relevant rate-limiting schemes. A side-by-side comparison of these schemes alongside the proposed method is shown in Table I.

A. Commonly used systems

SMS authentication [63] is a method of identification and authentication using phone numbers. A server first sends a random, variable-length number to a preregistered phone number via SMS. By entering the random number into the server’s website, users prove that they own the phone number. In addition to user authentication, servers use SMS authentication to block users who excessively access their resources. SMS authentication raises privacy concerns because all user actions

are tied to a specific user via their phone number. Moreover, it is trivial for an attacker to obtain multiple phone numbers [64].

CAPTCHA [14] is a mechanism that distinguishes a bot from a human. This is achieved by challenging a user with a task that is difficult for a bot to solve but easy for a human. If the user fails they are determined to be a bot, and their request is dropped. Due to the recent advancements in machine learning, studies have demonstrated successful attacks against CAPTCHAs [34]. In addition, CAPTCHA farms have diminished the efficacy of CAPTCHAs [53]. Despite these issues, [51] observed that CAPTCHAs are currently used as rate-limiters, raising the bar for attackers by forcing them to spend monetary and computational resources. However, CAPTCHAs have been criticized for reducing website usability [22], [59].

B. Techniques using hardware-assisted security

CAP [48] is a privacy-preserving CAPTCHA alternative utilizing the FIDO framework [17]. Users are served a random challenge and asked to use their FIDO-compatible device (i.e., authenticator) to sign that challenge. This makes such signatures unforgeable, as the private key is protected by the authenticator’s secure element (SE) [35].

CACTI [51] is a privacy-preserving rate-limiting system that utilizes client-side trusted execution environments (TEEs) [25]. This is facilitated via rate-proofs that allow the server to understand the number of user actions conducted in a given time window. CACTI maintains a counter of the number of users’ actions and generates the rate-proof by signing the counter, which the server uses to check whether the counter exceeds the threshold. Rate-proofs are difficult to forge since CACTI stores the private key used for the digital signature within the TEE. Moreover, CACTI uses group signature schemes, preventing servers from linking rate-proofs to users and other rate-proofs, thereby protecting user privacy.

The limitation of CAP and CACTI is that the rate-limiting functionality depends on the security of the client-owned device. Protecting the secret key within the SE is crucial for CAP, as it is impossible to revoke the key, which is replicated among many FIDO authenticators [56]. This is similar to CACTI, as an adversary is able to forge rate-proofs once the secret key is extracted from the TEE.

C. Cryptographic techniques

Privacy Pass [30], [16] allows users to obtain anonymous cryptographic tokens each time they participate in a task (e.g., solving a CAPTCHA). They can then “spend” this token each time they encounter such a task until they run out of tokens. Privacy Pass protects its users’ privacy by implementing this token using blind signature schemes. However, Privacy Pass is vulnerable to timing correlation attacks, which use the time difference between the generation and usage of tokens.

Opaak (OPen Anonymous Authentication framework) [47] provides rate-limiting for mobile phone users by employing a cryptographic primitive called periodic k-times anonymous authentication (periodic k-TAA). Combining this cryptographic primitive with SMS authentication, Opaak limits the number of private keys each user can obtain, thus limiting the number of times a user can generate signatures. There are several

TABLE I. COMPARISON OF THE PROPOSED METHOD (HIGHLIGHTED IN BOLD) AND RELATED WORK

	CAPTCHA [14]	SMS Auth	CAP [48]	CACTI [51]	Privacy Pass [30]	Opaak [47]	This paper
Private key storage	-	-	SE	TEE	Undefined	*1	TPM
Resource for uniqueness	-	Phone number	Attestation Secret key*2	Provisioning Key	Undefined	Phone number	Endorsement Key
Resistance to timing correlation attacks	-	-	Strong	Strong	Weak	Strong	Strong
Rate-limiting depends on device security	No	No	Yes	Yes	No	No	No

*1: Files encrypted with a master password, *2: Secret key installed by manufacturer

downsides to Opaak. First, Opaak assumes that a user cannot possess many phone numbers. This assumption is unrealistic due to services such as SMS farm services [64]. Second, the Opaak architecture encrypts the private key using a user-defined password. This puts a high burden on users as they must create a password that is secure enough to withstand brute-force attacks. Finally, the private key gets exposed to untrusted memory. This makes Opaak vulnerable to attacks such as side-channel attacks.

III. BACKGROUND

A. Group signature scheme

Group signature schemes [23] allow signers to prove group membership without revealing their identity by associating multiple private keys with a single public key. Since signatures generated using any private key can be verified by the public key, group signature schemes make it impossible for verifiers to distinguish users based on their signatures. A group signature scheme consists of the following entities:

- **Group Manager (GM)** managing group members (i.e., signers) by provisioning and revoking them;
- **Signer** generating signatures for verifiers; and
- **Verifier** verifying signatures produced by the signer.

B. Direct Anonymous Attestation (DAA)

DAA [19], [46] (depicted in Figure 1) is a special group signature scheme that provides the following additional properties on top of those provided by group signature schemes: (1) Neither a group member nor the GM can link a signer to a signature (**anonymity**), and (2) Given two signatures generated for two different verifiers, it is computationally hard to determine whether they were signed by the same signer (**unlinkability**). DAA also satisfies the **pseudonymity** property, allowing the signer to choose whether to allow multiple signatures directed to the same verifier to be linked with each other. We briefly describe the details of the DAA protocol below using the parameters shown in Table II¹.

Setup. GM generates gpk and gsk .

Join. Signer obtains usk and $cred^2$ by following the procedure below.

- 1) Signer generates a key pair upk and usk .
- 2) Signer sends upk and $EKCert$ to GM.
- 3) GM verifies $EKCert$.

¹The details described here follow the DAA protocol [19] which uses TPM-specific terminology. For the ECDAA protocol, see [46].

² $cred$ is treated similarly to usk and therefore it is not sent to the verifier.

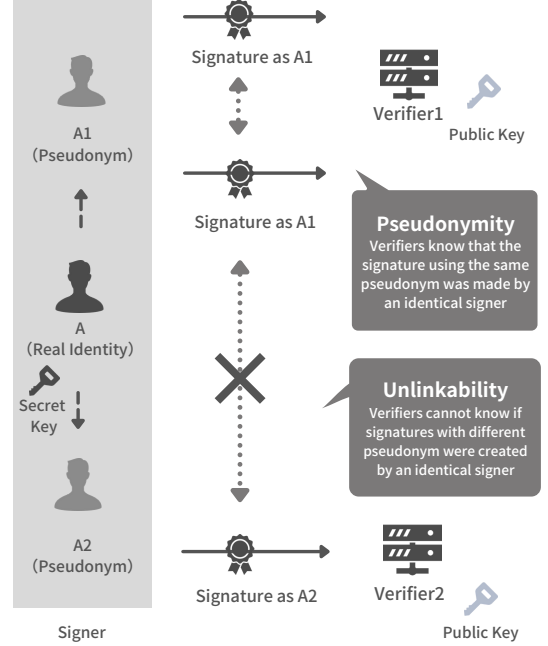


Fig. 1. Pseudonymity and Unlinkability of DAA

TABLE II. DAA PARAMETERS [19]

Notation	Description
gpk	Group public key.
gsk	Private key of the GM.
upk	Signer's public key generated by the device.
usk	Signer's private key generated by the device.
$cred$	Signer's credential generated by GM.
msg	Message to be signed.
bsn	Basename (chosen by Verifier or device).
σ_r	Probabilistic part of the signature, computed over msg using usk .
σ_d	Deterministic part of the signature, computed over bsn using usk .
σ	Signature. $\sigma = (\sigma_r, \sigma_d)$
RL	Revocation list storing revoked private keys usk .
EK	Endorsement Key, a unique private key for each device.
$EKPub$	EK public key, corresponding to EK .
$EKCert$	EK certificate for $EKPub$.

- 4) GM checks that it has not seen $EKCert$ in the past.
- 5) GM generates $cred$ based on upk and gsk .
- 6) GM computes ciphertext $c = Enc(cred, EKPub)$ using $EKPub$ contained in $EKCert$.
- 7) GM sends c to signer.
- 8) Signer obtains $cred = Dec(c, EK)$.

Sign. Signer computes $\sigma = DAA_Sign(msg, bsn, usk, cred)$

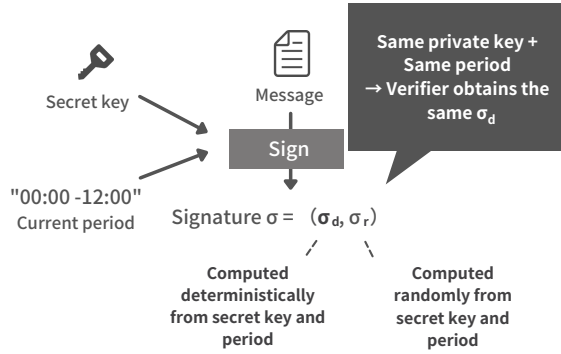


Fig. 2. How periodic k-TAA works.

over msg and bsn , using usk , $cred$ and gpk . σ and msg are then sent to the verifier.

Verify. Using msg , bsn , RL , and gpk , verifier checks the validity of σ by calculating $0|1 = DAA_Verify(\sigma, msg, bsn, RL, gpk)$. Signer is allowed access to the requested resource if the verification returns 1.

Revoke. The signer and GM revoke usk as follows:

- 1) The signer sends usk to the GM.
- 2) GM adds usk to RL .

C. Periodic k-times anonymous authentication scheme (Periodic k-TAA)

Periodic k-TAA [47] is a special group signature scheme with the following features:

- **Rate-limiting:** The signer cannot create more than one signature that is valid for a certain time period.
- **Anonymity:** No one, including the GM, can trace the signer from the signature.

In the periodic k-TAA scheme, depicted in Figure 2, the verifier can determine whether the same secret key has been used to generate a signature more than once by comparing the σ_d of the current signature with the one sent in the past. Any σ_d that is calculated using the same secret key during the same time window will generate the same value.

IV. SYSTEM AND THREAT MODEL

Our proposed system consists of four types of entities: GM, signer, verifier, and hardware security device. We assume that the number of hardware security devices is the same as the signers, and each signer owns exactly one hardware security device. Note that a signer can always purchase multiple such devices and integrate them into their platform, but there exists a clear limit that implies higher costs to the signer. This is discussed further in Section X. Manufacturers of the hardware security device are assumed to provide one attestation key (e.g., EK) per device. Additionally, the GM is assumed to not provide more than one private key to the device when provisioning the device. Finally, we assume that all cryptographic primitives are implemented correctly and that the random oracle assumption holds.

All entities other than the hardware security device are considered untrusted, i.e., they may act maliciously. The goal of a malicious signer is to generate as many signatures as possible in a short period of time in an attempt to maximize access to resources. The goal of a malicious verifier is to violate the signer’s privacy. The goal of a malicious GM is to collude with either the signer or verifier to help achieve their objectives.

The hardware security device acts as a source of *uniqueness* and we assume it has the following properties: (1) Tamper-resistant memory for storing secret data, (2) Is capable of remote attestation, (3) Implements basic cryptographic primitives, including key generation, encryption/decryption, digital signatures, and (4) Supports a group signature scheme that provides privacy and pseudonymity. Most importantly, *Scrappy does not* rely on the security of the hardware security device. Although we assume that the device has tamper-resistant memory, we also consider powerful adversaries that can extract secret data stored in such memory. We show in Section VIII that certain properties of the proposed system are not affected even in the presence of these powerful adversaries.

Finally, we consider DoS attacks out-of-scope and do not consider malware on the signer’s platform.

Our threat model yields the following security requirements for the anticipated system:

- **Rate-limiting:** Signers cannot send requests that exceed the verifiers’ threshold.
- **Unforgeability:** Signers cannot forge or modify signatures.
- **Unlinkability:** Given two signatures generated for two different verifiers or at two different points in time, it is hard to determine whether the signatures were generated by the same signer.

We also consider the following non-security goals:

- **Latency:** The required latency should be minimal.
- **Bandwidth:** The amount of data transferred between the entities should be minimal.
- **Storage:** Data stored in each entity should be minimal.

V. DESIGN & CHALLENGES

A. Conceptual Design

Rate-assuring proofs. In the heart of our design lies the concept of *rate-assuring proofs*. The idea of the proof is as follows: The signer provides the hardware security device with a time window t obtained from “some source” and the verifier’s basenane bsn . The device then generates a signature over the two pieces of data. This proof can be interpreted as the signer assuring the verifier that it has requested access to a specific bsn within a time window t . On the other hand, the verifier keeps a “log of proofs” from various signers and checks whether it has received a similar proof in the past. If so, the verifier assumes that the signer requests resources too frequently and drops the request.

Note that rate-assuring proofs generated by a signer for *different verifiers* (thus resulting in different identities) should generate indistinguishable proofs, meaning that no entity can

know whether the proofs came from the same signer. Similarly, rate-assuring proofs generated by a signer within a *different time window* for the same verifier should also be indistinguishable, thus preventing adversaries from linking the two proofs to the signer.

B. Design Challenges

In realizing the design shown above, we identified the following challenges.

Limitations in existing cryptographic primitive. Deciding what cryptographic approach to use to produce the rate-assuring proof is a challenge. The best fitting cryptographic primitive is periodic k-TAA, which meets the requirements raised in Section IV. However, there are several issues with using periodic k-TAA in tandem with hardware security devices. First, periodic k-TAA requires support for a special RSA modulus. Specifically, it uses $n = (2p + 1)(2q + 1)$ instead of the widely supported $n = p \cdot q$, where p, q are primes. Second, periodic k-TAA needs to compute $r \cdot sk + c \bmod n$, where r is a random number, sk is a secret key, c is a hash, and n is the aforementioned special RSA modulus. This calculation is used to generate Zero Knowledge Proofs, although hardware security devices do not normally support it.

Obtaining t without compromising security and privacy. So far, we have assumed that the hardware security device obtains t from “some source”. One source could be a clock on the hardware security device owned by the signer. However, many hardware security devices do not have a clock. Moreover, the clock must be synchronized frequently to prevent any skew. This process is difficult and, if done incorrectly, may lead to privacy leakage due to differences in t caused by small skews. The hardware security device could also obtain t from an external source. However, since the source may be malicious (e.g., the signer trying to cheat the system), the device must first verify the authenticity of t . This is a drawback, as it requires the hardware security device to be programmable and also puts trust in the device, which does not align with our threat model.

Managing logs without losing functionality. The verifier must maintain a log of past interactions with signers for the proposed system to function correctly. This obviously puts strain on the verifier’s storage over time, as it must interact with multiple signers.

C. Realizing the Design

This section presents a detailed design that addresses the aforementioned design challenges.

Cryptographic primitive. Instead of the periodic k-TAA scheme, we chose to use DAA to realize *Scrappy*. This is because it has been the standard authentication scheme for TPMs, enabling it to perform well even on resource-constrained hardware security devices. Another key factor for choosing DAA is that other hardware security devices also use a similar scheme. For instance, FIDO has made the ECDA specification [46] publicly available, allowing future FIDO devices to use DAA during attestation. Additionally, Intel

SGX [27], a widely available TEE, uses a similar primitive called EPID [20], which could also be used instead of DAA³.

However, DAA does not provide the most important *rate-limiting* property, i.e., it cannot be directly used to implement rate-assuring proofs. To overcome this issue, we modify the input to the *DAA_Sign* function to cryptographically bind t to the pseudonym used to generate the signature. This allows verifiers to link multiple signatures generated during the same t via the pseudonym. Note that once the t changes, the signature also changes and cannot be linked to the previous one, providing unlinkability between two time periods. We describe how this works in detail in Section V-D.

Obtaining t . In the proposed system, the signer obtains t from the verifier. A malicious verifier may provide an invalid t or a maliciously crafted t to track the signer (see Section VIII for more discussion). Therefore, the signer *must* check whether the t is valid. Similarly, a malicious signer may ignore the t provided by the verifier and use a fake one to cheat the system. Therefore, the verifier *must* check whether the t included in the rate-assuring proof is valid.

Managing logs. To reduce the impact on the verifier-side storage, we designed *Scrappy* so that verifiers can delete old entries that are no longer useful. In simple terms, any entries created before the current t can be deleted without hindering the system’s functionality.

In addition to the verifier, the signer maintains a log of past t received from the verifier. This is to allow the signers to check whether they have produced a rate-assuring proof within the same time window and prevent them from generating rate-assuring proofs that the verifier may reject. *Scrappy* is designed so that the signer can also delete old records in its log to reduce storage.

D. The *Scrappy* Protocol

The sequence diagram of the *Scrappy* protocol is shown in Figures 3 and 4. We assume that the communication channels between the signer and verifier are protected via standard means, such as TLS. This means that the verifier holds valid certificates signed by trusted CAs that bind their domain name to a certain public key. In addition to the parameters shown in Table II, the proposed scheme uses three additional parameters:

- t : A verifier-defined time window that encompasses the current time of signing (e.g., 10:00–11:00). The time window is inclusive at the start and cannot overlap with prior or future time windows.
- Log_s : A list of pair (t, bsn) received from the verifier in the past, i.e., $Log_s = \{(t_1, bsn_1), \dots, (t_n, bsn_n)\}$ where n is the number of items in the list.
- Log_v : A list of pair (t, σ_d) received from the signer in the past, i.e., $Log_v = \{(t_1, \sigma_{d1}), \dots, (t_n, \sigma_{dn})\}$ where n is the number of items in the list.

Scrappy consists of the following four phases: (1) initialization, (2) proof generation, (3) proof verification, and (4) cleanup. In this section, we describe each phase in detail.

³Although Intel has officially announced to deprecate SGX from consumer-grade CPUs [55], there is a possibility that the CPUs may still support EPID, allowing them to be used for *Scrappy*.

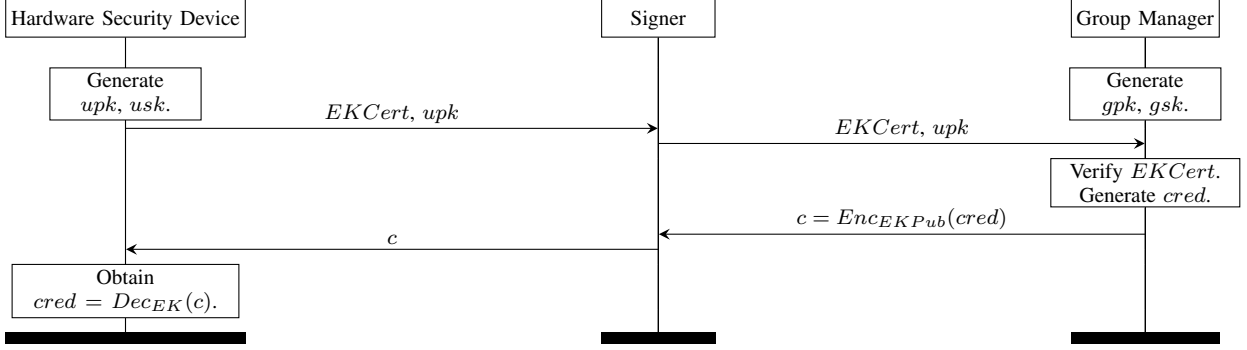


Fig. 3. Initialization Phase of `Scrapy` protocol. $Enc_k(m)$ and $Dec_k(c)$ denotes a cryptographic encryption/decryption operation over m/c using key k , respectively.

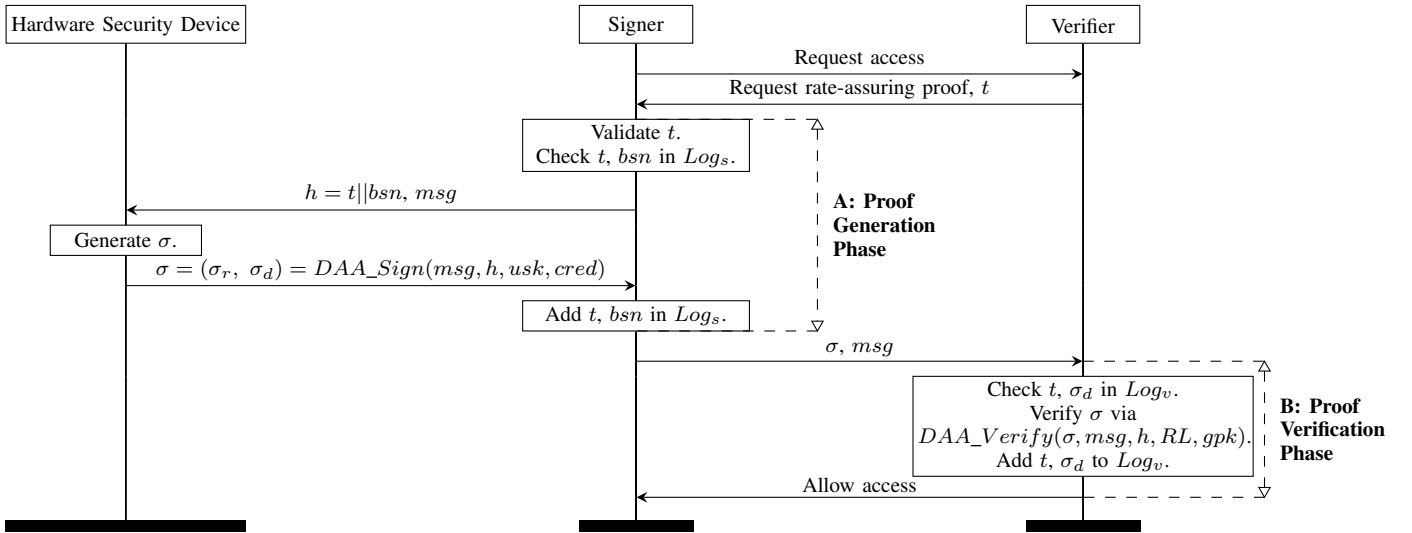


Fig. 4. Proof Generation and Verification Phases of `Scrapy` protocol. $DAA_Sign()$ and $DAA_Verify()$ are the DAA-specific operations described in Section III-B.

Initialization phase. During this phase, the GM⁴, verifier, and signer initialize various variables as follows.

- **GM:** generates gpk and gsk .
- **Verifier:** initializes its log $Log_v = \{\}$.
- **Signer:** generates usk and upk , obtains $cred$ from GM, and initializes its log $Log_s = \{\}$.

Proof generation phase. After the signer requests access to the verifier's resource, the verifier responds by requesting a rate-assuring proof over a t . The proof is generated using t , bsn , usk , and $cred$, following the process shown below.

- 1) Check whether t is valid, e.g., check that it covers the current time. Abort signing if not.
- 2) Obtain bsn from the domain included in the verifier's certificate and check its validity. Abort signing if not.

- 3) Check whether the pair (t, bsn) exists in Log_s . If so, abort signing.
- 4) Compute value $h = bsn || t$.
- 5) Forward msg and h to hardware security device to generate σ ⁵:

$$\sigma = (\sigma_r, \sigma_d) = DAA_Sign(msg, h, usk, cred)$$

- 6) Add new pair (t, bsn) to Log_s .
- 7) Send σ and msg to verifier.

Proof verification phase. Upon receiving σ and msg from the signer, the verifier proceeds with the verification phase by following the below process.

- 1) Check whether t is valid and reject proof if not.
- 2) Check whether the pair (t, σ_d) exists in Log_v . If so, the signer has reached the rate limit, and the request must be rejected.

⁴The GM would be required to maintain a list of certificates of trusted hardware security device vendors, which is used during verification of attestation proofs. It is worth noting that the number of vendors is expected to be small (26 for TPM [41]).

⁵In the actual implementation, msg is represented as an empty string. Nevertheless, we leave msg present here in the design, anticipating that it may be used as an extension in the future.

- 3) Compute value $h = bsn \parallel t$.
- 4) Verify σ via $DAA_Verify(\sigma, msg, h, RL, gpk)$. If verification fails, reject the signer request.
- 5) Add pair (t, σ_d) for the verified signature to Log_v .
- 6) Return success and allow access to the signer.

Cleanup phase. To reduce the storage overhead of the signer and verifier, the two entities periodically “clean up” their logs. Signers delete pairs (t, bsn) for a certain verifier bsn that are older⁶ than the newest t . On the other hand, verifiers delete all entries as soon as the current t expires.

VI. IMPLEMENTATION

We implemented a Proof of Concept (PoC) of the proposed design presented in the previous section. The baseline implementation consists of browser extension, signer application, TPM (see Section VI-E for implementations over additional devices), and verifier server.

All entities use SHA256 [29] for the hash function and BN256 [39] for the pairing curve. Note that BN638 [39] is recommended over BN256 due to security reasons [62], and the only reason we use BN256 is that it is the only pairing curve supported by our TPM.

The verifier is set to allow one access per minute. Therefore, all t is expressed as Unix timestamps, rounded up to the nearest minute. Note that this can be changed to adopt other formats of t .

The implementation of `Scrappy` and the ECDSA library can be found at [44]. Note that this release only includes the TPM implementation and we plan to release the other implementations in the future.

A. Browser extension

The browser extension of the PoC is built for the Chrome browser (version 100.0.4896.60) [37]. Note that many browsers provide similar interfaces, allowing our extension to be easily ported to other browsers. The browser extension consists of 2 scripts: content script and background script.

Content script: responsible for obtaining the data sent from the verifier. To do this, it looks for a special HTML tag in the verifier’s HTML file. The script parses this tag, obtains the contained parameters and the verifier’s origin, and finally sends all the data to the background script. For the PoC, we use the following implementation: `<input period="t" value="" />`, where t is the server-defined time window.

The content script also sends the data received from the background script to the server. We use the same HTML tag shown above for this purpose. Specifically, the parameters of the HTML tag are filled in as follows: `<input period="t" value="σ" />`, where t is the period provided by the verifier, and σ is the signature generated by the hardware security device.

Background script: forwards the data obtained by the content script to the signer application. This script is necessary because the content script is designed so that it cannot directly communicate with external applications. The communication

between the background script and the signer application uses the standard I/O and follows the Chrome Native Messaging Protocol [31].

B. Signer application

The signer application running on the signer is responsible for (1) generating the credential by cooperating with the TPM and GM, (2) producing the rate-assuring proof with the help of the TPM using the data received from the browser extension, (3) storing the period t and signature σ_d in its log, and (4) returning the proof to the browser extension.

All signer application implementations are implemented in Golang [9] due to its high performance and memory safety. In addition, we used the AMCL (MIRACL/core) [58], [7] library for cryptographic operations. To communicate with the TPM, the signer application uses a modified version of the `go-tpm` library [3], which includes support for the `TPM2_Commit` command [40].

We chose to use an SQLite [13] database to implement the logging system. This database has a table with one column named `BASENAME`, which stores bsn and t as `VARCHAR` (56 bytes). To fix the size of the otherwise variable data, the signer application calculates $hash(bsn) \parallel “_” \parallel t$ and stores the result in the database.

Before the signer application generates the rate-assuring proof, it queries this database to check whether a proof for the same period t and origin bsn has been created in the past. After the rate-assuring proof has been generated, the signer application also makes sure to update the t for the bsn .

C. TPM

TPM was our first choice for the hardware security device because it supports DAA by default and does not require any hardware modifications to be used for `Scrappy`. This enables `Scrappy` to be deployed in the real world, allowing both users and services to immediately take advantage of the various functionalities provided by `Scrappy`. Moreover, TPM on client-side devices is becoming more widely adopted, especially since Microsoft officially announced that the Windows 11 OS requires a TPM [50], and we envision that more devices support TPMs (or similar cryptographic hardware) as time progresses. We refer the reader to Section VI-E for implementations using devices other than TPM.

We used the ST33TPHF2ESPI TPM manufactured by ST Microelectronics [49]. The TPM clock speed is estimated to be around several tens to hundreds of MHz. The TPM is housed within a machine owned by the signer and generates rate-assuring proofs upon receiving requests from the signer application. The signer application uses the following TPM functions [38], [40] when generating the credential or the proof: `TPM2_CreatePrimary`, `TPM2_NVReadPublic/TPM2_NVRead`, `TPM2_AcivateCredential`, `TPM2_Create`, and `TPM2_Commit/TPM2_Sign`⁷.

⁷Although FIDO specification adopts `TPM2_Certify`, we use `TPM2_Sign` instead of `TPM2_Certify`, as it supports any message type [11].

⁶“Old” time windows are those no longer encompassing current time.

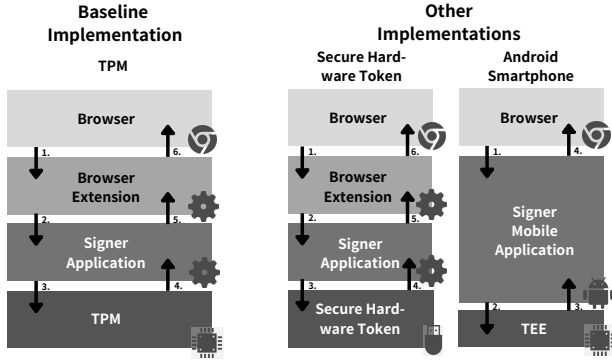


Fig. 5. Implementation of signer-side components and their relationship.

TABLE III. COMPARISON OF THE TYPE OF IMPLEMENTATION

	Baseline	Hardware Token	Smartphone
Device	PC (carrying TPM)	Hardware Token	Android Phone
Private key storage	TPM	Hardware Token	Encrypted File using TEE
Unique Resource	Endorsement Key	Manufacturer-Installed Private Key	Serial Number, IMEI/MEID *1
Method of Proving Unique Resource	TPM Attestation	Proof of Possession of Private Key	Android ID Attestation (ARM TrustZone)
Rate-limiting Depends on Device Security	No	No	Yes

*1 Serial Number of Phone, IMEI (International Mobile Equipment Identifier) and MEID (Mobile Equipment Identifier)

D. Verifier server

The verifier server conducts the following: (1) Sends the data necessary for the rate-assuring proof to the signer by embedding it into the HTML tag, and (2) verifies the proof sent by the signer. Similar to the signer application, the server was written using Golang and uses the `gin` web framework [2]. SQLite stores the logs and AMCL (MIRACL/core) for cryptographic operations. The database has a table with one column named “K” as VARCHAR (45bytes), and stores concatenated values $\sigma_d || \text{“_”} || t$.

E. Implementing *Scrappy* on other devices

TPM was our device of choice when implementing *Scrappy*. However, signers may only have access to computers that do not have TPM or not even have access to computers at all. To accommodate such users, we implemented *Scrappy* on two other devices: a secure hardware token and a smartphone. Although we had to make slight modifications to the specifications and the software stack during the implementation, this shows the versatility of *Scrappy* and its potential for being deployed to many users. Figure 5 provides an overview of the implementation with respect to the relationship of each component regarding the different devices.

1) *Secure hardware token*: We used Solokey Hacker [60] for this implementation. The token includes an STM32L432KC microprocessor with an Arm Cortex-M4 MCU running at 80 MHz, 64 kB of RAM, 256 kB of flash memory, and a true random number generator. Solokey supports the FIDO API [61], and we extended the

Client to Authenticator Protocol (CTAP) API to support the FIDO ECDA [18] functionalities. We chose to use the Xaptum ECDA library [68], as it is compatible with FIDO ECDA version 1.1 [67]. The library uses the Miracl AMCL cryptographic library [66] for the necessary BN256 curve. Our custom Solokey firmware was written in C language.

Solokey exposes a custom ECDA_SIGN API through a modified FIDO CTAP1 (U2FHID) API that takes `message` and `basename` as inputs and outputs an ECDA digital signature over the two pieces of data. To simplify the implementation process, we assume that the secure hardware token has already run the ECDA Join protocol and has obtained a valid `cred`.

Challenges and solutions. The challenge of implementing *Scrappy* on secure hardware tokens is that no tokens support FIDO ECDA in the wild, including Solokey. This is a drawback, since it does not enable users to purchase a token and start using *Scrappy* immediately. However, we presume that some tokens will support ECDA in the future, which we simulated by adding ECDA functionalities to the FIDO software stack on Solokey.

Another challenge is that the FIDO ECDA specification does not allow the use of basenames when producing digital signatures to ensure signer privacy [18] (see Section VIII for more discussion). This required the *Scrappy* implementation on the token to slightly deviate from the official specification, which, again, prevents this implementation from being used in the real world. Nonetheless, the Xaptum ECDA [12] implementation supports the usage of basenames, allowing us to implement the PoC.

Security of device. Although current secure hardware tokens share the same public/private key pair to ensure device privacy [56], we envision tokens used in *Scrappy* to have unique key pairs per device. This is possible because ECDA allows such tokens to achieve the same level of privacy provided to current tokens. Moreover, with this assumption, we can guarantee the rate-limiting functionality of *Scrappy* even if per-device unique key pairs get leaked from tokens, as GMs can reject multiple Join requests from the same public key.

2) *Smartphone*: We used an Android [10] smartphone for this implementation. The signer application was implemented as an Android mobile application (hereafter referred to as “Signer mobile application”). The main Signer mobile application was written in Kotlin [5], while the ECDA library and other cryptographic schemes were written in Golang due to security and efficiency reasons. Since Golang is not supported natively on Android, we used Gomobile [36] to convert these libraries. The Signer mobile application also used Android Room [8] and SQLite to store its log.

Challenges and solutions. The challenge of implementing *Scrappy* on a smartphone is that most mobile browsers do not support browser extensions. To overcome this, we used a custom URL to communicate between the browser and the Signer mobile application. Briefly, the URL allows the browser to open the Signer mobile application and share the necessary information. Similarly, the application jumps back to the browser by following a callback URL, which allows the browser to send the generated proof back to the verifier

server. Note that the user must check the callback URL to prevent generating proofs for unintended basenames.

Another challenge is storing private keys securely. This is an issue for `Scrappy` because Android does not have any API that supports the storage and usage of ECDSA keys within hardware, unlike TPM. We overcame this by encrypting the ECDSA private key using standard methods and decrypting it when needed. This way, the encryption key can be stored securely using the Android Keystore API [1], which uses a TEE for storage. However, it is worth noting that this requires the plaintext ECDSA private key to briefly exist on the smartphone memory. Although this occurs only for a short time, it makes this implementation less secure against side-channel attacks compared to the baseline implementation. We envision that this will no longer be an issue once the Android Keystore API supports ECDSA signing operations within a TEE.

An additional challenge is integrating smartphone identification schemes to `Scrappy`. This is required as the smartphone acts as a source of uniqueness for `Scrappy`, similar to attestation keys in TPM and secure hardware tokens. Although we could potentially use Android ID Attestation⁸ [4], integrating it into the ECDSA Join protocol is no easy task. Since this requires a large implementation effort, we leave this for future work. Therefore, similar to the secure hardware token implementation, we assume that the smartphone has already obtained a valid *cred*.

Security of device. The Android ID attestation documentation does not clearly describe whether the public/private key pair used to generate the proof of the device serial number is unique to the device. However, the fact that the public key is not listed in the identifying information in the documentation leads us to believe that the key pair is shared across multiple devices. Therefore, once an adversary can extract the private key protected by ARM TrustZone via side-channel attacks, they can forge valid proofs of identifiers of phones that do not exist. We, thus, determine that the rate-limiting functionality of `Scrappy` using Android phones depends on the security of the TEE.

VII. PERFORMANCE EVALUATION

This section presents the latency, bandwidth, and storage evaluation results of the proposed system.

For the TPM and secure hardware token evaluation, we used a Lenovo ThinkPad A285 [45] with an AMD Ryzen 5 PRO 2500U CPU (4 cores, 2 GHz) with 16 GB of RAM, representing a standard consumer-grade laptop. For this evaluation, the signer application ran within a Docker container. The containers ran Ubuntu 22.04 and utilized standard Docker APIs to mount necessary hardware devices to the signer container. For the TPM it was `/dev/tpm` and for the secure hardware token it was either `/dev/ttyUSB0` or `/dev/ttyACM0`. For the smartphone evaluation, we used an Android Pixel 6a smartphone with a Google Tensor CPU (2 × 2.8 GHz, 2 × 2.25 GHz, 4 × 1.8 GHz) and 6 GB of RAM, running Android version

⁸Android ID Attestation relies on ARM TrustZone for protecting resources proving device uniqueness. Thus, adversaries may steal this resource via side-channel attacks. Table III reflects this fact.

TABLE IV. END-TO-END LATENCY EVALUATION OF `SCRAPPY` (CRYPTOGRAPHIC OPERATION + DATABASE OPERATION)

Proof generation		Proof verification
Hardware Device	Measurement [ms]	Measurement [ms]
TPM (baseline)	243.16 (236 + 7.16)	84.1 (73.7 + 10.4)
Hardware Token	2771 (2764 + 7.16)	
Smartphone	26.4 (23 + 3.4)	

13. The signer application ran as a smartphone application. All verifier-side experiments ran on an AS-2024US-TRT/3Y machine with two AMD EPYC 7513 CPUs (32 cores, 2.6 GHz) with 512 GB of RAM, representing a standard server-grade machine. Unless otherwise stated, the numbers reported are an average of 20 measurements.

A. Latency Evaluation

1) *Baseline latency evaluation:* We first present the proof generation latency of the three devices as well as the proof verification latency under simulated baseline conditions. For the setup, we do not populate the signer-side log with any entries, simulating a situation where a user has not visited websites with `Scrappy`. On the other hand, the verifier-side log is populated with 100,000 entries to simulate a situation where the verifier has been running for some time. The results are presented in Table IV.

For the proof generation phase, we can see that the smartphone is the fastest, taking 26.4 ms (cryptographic operation: 23 ms, database operation: 3.4 ms), the TPM the second fastest, taking 243.16 ms (236 ms, 7.16 ms), and the hardware token the slowest, taking 2771 ms (2764 ms, 7.16 ms).

It is interesting to observe that the TPM exhibits a speed that is an order of magnitude faster than that of the secure hardware token when generating signatures, despite both devices having processors with similar frequencies. We hypothesize that this is because TPM has optimized hardware for producing DAA signatures, whereas the secure hardware token does not.

The average latency for the proof verification phase is 84.1 ms, with the verification operation taking about 73.7 ms and the database lookup and insert taking 10.4 ms. This is an acceptable latency since the AMCL cryptographic library benchmark test results on the same server-grade machine were in the same order of magnitude (i.e., tens of milliseconds). Additionally, [69] reports similar numbers for the DAA verification process on a CPU with a faster clock frequency. Note that the verification operation currently uses only one CPU core and could benefit from multi-threading as well as standard load-balancing techniques when dealing with a large number of verification requests.

2) *Latency evaluation under extreme conditions:* Next, we evaluated the latency of proof generation and verification phases under certain extreme conditions. For the proof generation phase, the signer-side log was populated with a large number of entries, and the verifier was required to check a revocation list with a large number of entries for the proof verification phase.

The results are presented in Table V. Differences between Table IV and V are highlighted in bold text. We observe that

TABLE V. END-TO-END BASELINE LATENCY EVALUATION OF SCRAPPY UNDER EXTREME CONDITIONS (CRYPTOGRAPHIC OPERATION + DATABASE OPERATION). DIFFERENCES ARE HIGHLIGHTED IN BOLD.

Condition	Proof generation [ms]	Proof verification [ms]
Large signer log	243.23 (236 + 7.23)	84.1 (73.7 + 10.4)
Large revocation list	243.16 (236 + 7.16)	151.4 (141 + 10.4)

while the overhead of having a large signer-side log is minimal, having a large revocation list introduces a large overhead. A thorough analysis of these results is shown below.

Latency of signing with a large number of signer logs. This evaluation simulates where a signer visits a large number of websites that support `Scrappy`. For this experiment, we populated `Logs` with 1,000 entries and measured how long it took to generate a signature. 1,000 entries are significantly larger than the number of distinct web pages⁹ visited by a user per day (163 web pages) [28], representing an upper bound.

The average latency for this measurement was 243.23 ms. The signing operation took 236 ms, while database lookup and insertion took 7.23 ms. Comparing this to the results presented in Section VII-A1, we can see that the database operation latency increased by 0.07 ms, which is negligible. This shows that the signer-side latency does not get affected even if signers visit a large number of websites.

Latency of verifying with a large number of revocations. This evaluation simulates where many signers had their group private keys revoked. The number of entries in the revocation list was set to 50, which is the number of maximum allowed entries in a revocation list in EPID [57].

The average verification latency with 50 entries in the revocation list was 151.4 ms, with the verification operation taking 141 ms and the database lookup operation taking 10.4 ms. This is nearly two times slower than the latency reported in the baseline evaluation, although the overall end-to-end latency is still within the order of hundreds of milliseconds. Even though it is highly unlikely that there are this many entries in the revocation list, this shows that verifiers must take into account such a case when designing their services.

3) *Comparison of latency results:* Overall, we conclude that `Scrappy` introduces minimal latency overhead, even under extreme conditions. Table VI compares the latency of `Scrappy` with the numbers reported in related work. While evaluation environments vary across different papers, we can observe that the latency does not drastically differ (except for Opaak). It is worth noting that our baseline evaluation uses a resource-limited TPM, whereas others use more powerful devices.

B. Bandwidth Evaluation

We measured the amount of data transferred between the signer and verifier when using `Scrappy`. Recall that the verifier sends t to the signer when requesting the rate-assuring proof, and the signer responds by sending σ to the verifier.

The size of t is 19 bytes, which is a 64-bit Unix timestamp encoded as a string. The size of the raw signature σ is 261

⁹Note that “distinct web page” refers to a unique URL and not necessarily a unique domain, i.e., the actual number of distinct domains may be smaller.

TABLE VI. COMPARISON OF LATENCY MEASUREMENTS BETWEEN SCRAPPY (HIGHLIGHTED IN BOLD) AND RELATED WORK.

Work	Proof generation [ms]	Proof verification [ms]
Scrappy (TPM, baseline)	243.16	84.1
Scrappy (Hardware Token)	2771	84.1
Scrappy (Smartphone)	26.4	84.1
CACTI [51]	211.9	27.3
Privacy Pass (N tokens) [30]	341.48 + 180.87 × N	57.8
Opaak [47]	2550 (combined measurement)	

bytes, consisting of five 33-byte elliptic curve points and three 32-bit big integers. Note that we use Base64 when encoding the raw data and gob (Golang encoding format) when bundling the signature parameters. Thus the actual size of σ is 660 bytes. We can observe that this number is small, thus deploying `Scrappy` should only introduce minimal bandwidth overhead.

C. Storage Evaluation

Finally, we measured the amount of storage it takes to store the logs for both the signer and verifier sides.

The amount of storage required to store 1,000 entries in the signer-side log is 94.2 KB and 6.64 MB for the 100,000-entry verifier-side log. This indicates that our PoC stores only a small amount of data for a large number of logs and can be further optimized.

VIII. SECURITY EVALUATION

In this section, we provide a security analysis of `Scrappy` based on the threat model and requirements defined in Section IV.

Signature forgery attacks. A malicious signer may attempt to generate a fake rate-assuring proof to cheat the system. The proposed system prevents this, as it is impossible for signers to receive a valid `cred` from GMs without a genuine hardware security device. Even if the signer skips the above step altogether and generates a fake `cred`, it will still get caught since the `cred` used to generate the rate-assuring proof will not be valid.

Timestamp forgery attacks. A malicious signer may attempt to use a t that the verifier did not provide. For instance, if the valid t is 10:00–11:00, the malicious signer may try to create a rate-assuring proof on 10:00–10:10 in an attempt to get more access to the verifier’s resource. This is preventable by the verifier, since the proof provided by the signer will include a different t than the one provided.

Generating multiple proofs upfront. A malicious signer could generate multiple rate-assuring proofs with different t upfront and send them all at once to the verifier. However, this will not grant the signer with the same amount of access as the number of generated proofs, since the t in the proof will not match the one expected by the verifier and will be rejected.

Obtaining multiple credentials. A malicious signer may attempt to generate multiple `usks` and obtain credentials over them, allowing the signers to generate multiple rate-assuring proofs for a given set of t and `bsn`. This is impossible, as

the GM will know that the same signer is requesting multiple credentials via their attestation key (i.e., EK).

Device Reset attacks. Malicious signers can always reset their hardware security device and re-run the Join protocol claiming that they are a fresh device. The proposed method prevents this by having the GM limit the number of times a device can run the Join protocol. This is possible because the GM can keep track of the number of times it has observed the same EK .

Signer tracking via proofs. A malicious verifier may attempt to track signers using their rate-assuring proofs in the following two ways: (1) linking two proofs from two periods in time; and/or (2) colluding with another verifier and linking proofs that were sent to the two verifiers. These attacks are not possible because using either a different t or bsn will generate a signature that is indistinguishable from other signatures.

Signer tracking via t . A malicious verifier may attempt to track signers using t which they provide during the proof generation phase. For instance, a verifier may provide a signer of interest with t' while providing the rest with t . Sometime in the future, a verifier suspects that the same signer is requesting access, and provides the suspected signer with t' . If the signer cannot respond with the proof, the verifier's suspicion is correct and has successfully identified a signer. This is easily detectable since t' is from the past.

Similarly, a malicious verifier may provide a signer with a t which is valid for an extremely long time. If a verifier suspects that the same signer is requesting access, they provide the signer with t' that overlaps with t . If the signer is in fact the same, they cannot generate proof, allowing the verifier to confirm its suspicion. This is also detectable since the signer observes that the verifier is providing overlapping time windows, which violates the protocol.

In either case, signers should report such malicious activities to third-party auditors so that malicious verifiers are held accountable.

The use of basenames in Scrapy and its privacy implications. We have discussed in Section VI-E1 that the FIDO ECDAA specification prohibits the use of basenames due to privacy concerns. This is because FIDO ECDAA aims to provide full anonymity to its users and the use of basenames may cause signatures to be linked to each other (i.e., pseudonymity, see Section III-B for more details).

This pseudonymity property is necessary for Scrapy as it is what enables verifiers to know that a certain signer is conducting multiple actions within a given time window. Although this may seem that Scrapy is unable to provide full anonymity to its signers, this is not the case. This is because malicious signers are the only entities that produce signatures that can be linked to each other. On the other hand, legitimate signers do *not* send multiple rate-assuring proofs within the same time window. In addition, the use of Log_s during the proof generation phase described in Section V-D allows legitimate signers to keep track of whether they have generated proofs for the same verifier within the same time window. Therefore, as long as signers abide by the Scrapy protocol they can maintain full anonymity.

Violating signer privacy via side-channel attacks. Malicious actors may attempt to steal the usk from signers using side-channel attacks. This is an issue since once an attacker steals the usk , they are able to link rate-assuring proofs to an individual signer, by calculating the pseudonym used by the signer. Although it is not possible to prevent this via cryptographic approaches, the use of secure hardware in Scrapy significantly increases the difficulty of the attack for adversaries (except for the smartphone implementation, see Section VI for more details).

Network attacks. A malicious party on the network may change the values for t and/or bsn in an attempt to mount a DoS attack against an unsuspecting signer. This is preventable since the signer checks whether the t and bsn are valid before generating the proof. Furthermore, an honest signer will detect that there is a network adversary, as their requests will be denied despite legitimately producing the proof.

Compromised Devices. Malicious signers may attempt to extract the EK and/or the usk from hardware security devices owned by other signers through various means, including side-channel and physical attacks. We investigate the consequences of an adversary obtaining the two keys.

Leaked EK . A large number of devices may have their EK extracted while in possession of other parties after manufacturing (e.g., during shipping). This allows a malicious signer to generate valid rate-assuring proofs for a given t and bsn , granting them an abnormal number of accesses. However, this is *detectable* by the victim, as a GM will refuse to issue a $cred$ as it has already done so to the same EK in the past. Moreover, even if a GM issues a valid $cred$ to the victim, the leaked EK does *not* violate the privacy of the victim, since the usk generated by a victim cannot be linked to EK .

Leaked usk . Extracting usk from the hardware security device allows the adversary to not only generate valid rate-assuring proofs but also violate the privacy of the victim. This is because the adversary can link victim-generated proofs to usk . However, extracting usk is extremely challenging for the adversary. First, the adversary is required to have access to the victim's device (e.g., via malware). This is because the victim generates usk only when the device is within their possession. Second, the hardware security device stores usk within its tamper-proof memory. This would be even more challenging if the device is detachable (e.g., secure hardware token). In essence, the tamper-proof assumption of the hardware security device can be seen as a means to protect the *signer privacy* property of Scrapy.

Leaked EK and usk . The unforgeability and unlinkability properties of Scrapy would be violated if both EK and usk are stolen. However, we would like to emphasize that even in this situation, the *rate-limiting* property remains intact. The intuition behind this is that the verifier will accept a proof generated by *either* the adversary *or* the victim, but not both.

Rogue GMs. A malicious GM may attempt to compromise signer privacy by colluding with a malicious verifier. However, this is prevented by the use of DAA, as it is impossible to link EK and usk , even by the GM.

A malicious signer may also attempt to produce multiple rate-assuring proofs within the same period of time for the

same verifier with the help of a malicious GM. This is done by having the malicious GM produce credentials for the multiple *usks* generated within the single hardware security device. However, once caught (e.g., through community reporting), verifiers will stop verifying proofs that use *gpk*s issued from these GMs, essentially losing the trust of the public. Therefore, we assume that a *rational* GM will not engage in such activities and will be heavily penalized if discovered.

GM with only one member. It is easy to imagine that malicious verifiers will be able to track a signer that generates proofs that are verifiable using a *gpk* of a group of which they are the sole member. Since *Scrappy* does not regulate such behavior, it cannot prevent such incidents from happening. Signers must be aware of the GM they are interacting with and should consider joining a group that is reputable and popular. GMs can use this information to stand out from other GMs, for example by publishing the number of current group members.

Summary. Overall, we claim that the proposed system successfully meets the security requirements defined in Section IV. Specifically, the **Unforgeability** requirement is satisfied since it is impossible for the signer to generate rate-assuring proofs without a genuine device and using valid *t* and *bsn*. The **Unlinkability** requirement is also satisfied because of the use of DAA and the rate-assuring proofs not leaking any information regarding the signer’s identity. The **Rate-limiting** requirement is met because rate-assuring proofs that carry invalid timestamps will not be accepted by the verifier, thus preventing any access above the rate set by the verifier. Moreover, any signer attempting to collude with a malicious GM will be detected and punished.

IX. USABILITY ANALYSIS

In this section, we provide an analysis of *Scrappy* from the perspective of signer usability.

Installation. Installing *Scrappy* on the signer platform requires the *Scrappy* browser extension and native application to be downloaded and installed. We implemented *Scrappy* to ensure that the installation process does not impose excessive overhead on the user. However, the process can be further improved by integrating the extension logic and native application functionalities directly into the browser. This way, *Scrappy*’s signer-side logic will work “out of the box”, providing a seamless user experience.

Signer-perceived latency. Minimizing the time a signer needs to spend when using the system is crucial for its usability. As shown in Tables IV and VI, the end-to-end latency of *Scrappy* ranges between 100 and 2800 ms, comparable to other state-of-the-art rate-limiting systems. Therefore we claim that the latency introduced by *Scrappy* is minimal and does not degrade user experience.

Signer involvement. *Scrappy* is designed to be fully transparent to the signer, i.e., signer input is not required when creating rate-assuring proofs. However, there are cases where signer input may be useful.

For instance, a malicious verifier may use *t* to track signers (see Section VIII). Such an attack can be automatically detected via the aid of signer-side software (e.g., *Scrappy* browser extension). However, there may be cases where the

software is not confident enough to either reject or accept a *t*. In such a case, signer input would be required to make the final decision.

Another example is malware on signer-side devices. This allows malicious signers to use other legitimate signers’ devices to generate valid rate-assuring proofs (i.e., cuckoo attack). Preventing such an attack would require signer confirmation via I/O which is hardwired to the hardware security device (e.g., secure I/O in ARM TrustZone).

In short, incorporating signer input improves security. However, requiring frequent input degrades user experience due to UI fatigue. To strike a balance between user experience and security, we envision signer-side software allowing signers to choose from the following options: (1) *Ask every time*: Requires signer input each time it generates a proof (default). (2) *Ask for untrusted*: Requires signer input only for websites that are not trusted by the signer. (3) *Do not ask*: Does not require any signer input.

Usability of supported devices. Supporting many devices is crucial for adoption, as it allows signers to use their preferred device. Here, we assess the usability of the three devices used in this paper. Signers that want absolutely no involvement during the proof generation phase are recommended to use devices with TPM. Signers comfortable with minimal user involvement may opt to use smartphones. Secure hardware tokens require high signer involvement as they need to be plugged in every time.

Summary. Overall, we claim that *Scrappy* is designed and implemented with usability in mind. The installation process is minimal, with potential optimization, and signer-perceived latency is comparable to other related work. In addition, signer involvement in *Scrappy* is optional. While important when mitigating some attacks, we have demonstrated that the level of involvement can be adjusted according to the signer’s preference. Furthermore, signers have a wide selection of devices to choose from when participating in the *Scrappy* system.

X. DISCUSSION

A. Fallback options

Signers who cannot participate in the *Scrappy* protocol due to lack of access to applicable devices or do not want to participate in the protocol will fall back to the current rate-limiting services provided by the verifier. Note that this does not change the current state of security and privacy provided to the verifiers and signers. Moreover, as we have shown in Section VI, *Scrappy* can be implemented and deployed on various types of devices. We are certain that users will have at least one device compatible with *Scrappy*.

B. Deployment considerations and incentives

1) *CDN and 3rd party integration*: Although *Scrappy* was designed in mind to reduce deployment effort by using well-known components (e.g., DAA, SQLite, Golang), we expect that verifier operators may not be familiar with the underlying technology stack used in *Scrappy*. To accommodate such needs, we envision Content Delivery Networks (CDNs)

and other 3rd parties operating the `Scrappy` verifier on behalf of its customers.

CDNs host the content of their customers in servers distributed across the globe, reducing delivery latency to users. Recently, CDNs started to offer security services to their customers, especially protection against abusive actions. The three most popular CDNs – Cloudflare [24], Akamai [15], and Fastly [32] – all offer rate-limiting services, using various methods ranging from basic threshold-based limiting to CAPTCHAs. `Scrappy` can be easily integrated into the CDN gateways, providing stronger rate-limiting capabilities for customers who are interested.

2) *Website operator incentives:* Websites have several incentives to integrate `Scrappy` into their services. First of all, `Scrappy` provides strong guarantees of rate-limiting through proven cryptographic schemes. Secondly, `Scrappy` requires minimal latency when generating and verifying rate-assuring proofs, limiting the impact on user experience. Thirdly, the unlinkability property of `Scrappy` enables users to participate in a private manner.

Website operators can leverage `Scrappy`'s privacy aspect and promote their commitment to safeguarding users' privacy. This can have a positive impact on their service, as it may attract additional privacy-conscious users.

3) *GM operator incentives:* The role of the GM in `Scrappy` is to check the genuineness of the hardware security device and issue a credential if confirmed. Since this only occurs once every lifetime of a hardware security device (unless it gets reset, see Section VIII for more information about this), the operation of the GM is relatively lightweight. Various organizations could run a GM, each with different incentives, such as online identity providers providing federated login services and non-profit organizations (e.g., Let's Encrypt [6]).

C. Supporting multiple GMs

So far, we have not considered more than one GM in the `Scrappy` ecosystem, but given the incentives shown above, many GMs may want to benefit from `Scrappy`. `Scrappy` does not put a limit on the number of GMs it supports. However, this allows a malicious signer to obtain multiple credentials for a pair of *upk* and *EKCert*. This is an issue since malicious signers would be able to generate multiple rate-assuring proofs that are unlinkable to each other.

Depending on the services they operate, verifiers may want to choose one of the following solutions: (1) For verifiers that cannot allow more than one access to a certain resource per signer at a given point in time (e.g., online ticket sellers), they may want to notify the signers that they will only accept signatures from certain single GM to prevent such issue; or (2) For verifiers that have some tolerance, they may want to opt to support as many GMs as possible to attract signers.

It is easy to imagine that verifying the proof one by one against a list of GM public keys is not efficient. Therefore, once support for multiple GMs in the ecosystem is introduced, we propose to make the following modifications to the proof generation phase: (1) When submitting the rate-assuring proof to the verifier, the signer must also include the identifier of the GM (e.g., the domain name). (2) The verifier must maintain a

one-to-one mapping of GM identifiers and their group public keys, which allows efficient lookup and verification of the proof. (3) When the verification fails, verifiers can either choose to cross-check the proof against all public keys within the table or deny access to the service.

D. Different TPM implementations

There are four types of TPMs: discrete, integrated [26], firmware-based [54], and virtual/software-based [52]. Discrete TPMs are dedicated cryptoprocessors separate from the CPU and are normally plugged into the motherboard. Integrated TPMs are CPUs that provide functionalities and security guarantees similar to those of TPMs without a dedicated processor. Firmware TPMs (fTPMs) use TEEs to store necessary secret keys and to protect the integrity of important data. Virtual/Software-based TPMs are a full software implementation of TPM functionalities and do not utilize hardware to store secrets nor carry valid proof of uniqueness.

Discrete and integrated TPMs satisfy the four requirements discussed in Section IV, while firmware and virtual/software-based TPMs do not satisfy some requirements. In particular, fTPMs cannot protect secret keys such as *EK* from side-channel attacks due to using TEEs (e.g., [43]). However, fTPMs can still be used in `Scrappy`, as the underlying TEE acts as a source of uniqueness. On the other hand, virtual/software-based TPMs do not provide proof of uniqueness and, thus, do not have valid *EK*. Therefore, they would not be able to be used in `Scrappy`.

E. Balancing deployability and rate-limiting functionality

Supporting different types of hardware security devices is crucial for deployability so that services would be more willing to employ `Scrappy`. As we have shown in Section VI, `Scrappy` can be deployed on a wide range of devices. However, there is a clear trade-off between increasing deployability and providing rate-limiting functionality. This is because as more devices support `Scrappy`, the easier it is for an attacker to operate farms that employ a large number of such devices. In this section, we take a closer look into device farms and their effect on `Scrappy`.

Similar to CAPTCHA farms, a device farm is an operation that uses a large number of devices to generate numerous valid rate-assuring proofs, attempting to attack the rate-limiting capability provided by `Scrappy`. Although `Scrappy` does not (and cannot) prevent such farms from being operated, it certainly can discourage them from forming. This is because (1) each device can only produce one proof for the same *t* and *bsn*, and (2) the cost of purchasing such a large amount of devices would be significantly higher than other alternatives. For example, some CAPTCHA farms charge as little as 0.5 to 1 USD for solving 1,000 text-based CAPTCHAs and 3 USD for 1,000 reCAPTCHAs [53]. In comparison, TPM chips usually cost around 20 USD each, requiring the adversary to spend 20,000 USD as an initial investment to provide the same number of rate-proofs.

It is important to note that these farms are not a unique issue to `Scrappy`, as other related work also cannot defend against such farms (e.g., CACTI farms [51]). What is unique to `Scrappy` is, however, the fact that supporting a wide variety

of devices for deployability allows adversaries to operate different kinds of farms. For instance, secure hardware tokens may be easier to farm compared to TPM chips since they use standard USB interfaces rather than manufacturer-specific ones. Moreover, USB hubs are extremely easy to obtain and make such tokens easier to connect to a host computer on a large scale.

To deal with device farms, we envision verifiers to accept rate-proofs according to the GM public key provided by the signer, similar to the discussion in Section X-F. For example, for actions requiring stricter rate-limiting policies (e.g., account creation, purchasing high-value items), verifiers may choose to accept proofs from devices that are more expensive and difficult to farm (e.g., firmware TPMs that cost several hundred USD per chip).

F. Supporting additional devices

Not all users may have access to the three devices we have shown in this paper. In this section, we discuss additional devices¹⁰ that may potentially be used in *Scrappy*. Table VII compares the list of devices.

SIM cards. SMS authentication provides proof of possession of a certain phone number. Similar to Opaak [47], SIM cards can be used as a source of uniqueness in *Scrappy*.

Intel SGX CPUs. Intel CPUs from 6th generation to 10th generation support Intel SGX [27], which is a form of a TEE. Such CPUs have secret keys that are used to prove the genuineness of the platform as well as the code running inside the TEE. Using Intel SGX in *Scrappy* is fairly straightforward: prove to the GM during the Join protocol that it has not registered in the past and produce rate-assuring proofs within the TEE.

PUFs. Physical Unclonable Function (PUF) [33], [42] is a one-way function that is unique per device. PUFs utilize the subtle differences in the physical characteristics of the components of the device to provide such uniqueness. Attestation of a PUF uses a standard challenge-response protocol that is similar to HMAC. Using PUFs in *Scrappy* is challenging because PUF keys are symmetric. Therefore, we envision that either the vendor of the PUF device will operate a GM or provide third-party GMs with proof of verification of responses.

Assuming that we have deployed *Scrappy* that accepts all the different devices we have listed above, several points must be taken into consideration. First, *Scrappy* does not (and cannot) prevent GMs from accepting a source of uniqueness. As some resources are easier to obtain than others, verifiers are advised to decide whether to accept rate-assuring proofs based on the signer’s group affiliation. For example, verifiers that can tolerate high-rate access (e.g., API access, website viewing) may choose to allow signers with easily obtainable devices. Others may choose GMs that only provide *cred* to hard-to-obtain devices (e.g., integrated TPMs running on Intel Xeon CPUs). Second, the security and privacy guarantees of devices differ. Some devices require the device to be trusted to provide

¹⁰Devices that cannot provide proof of uniqueness *cannot* be used in *Scrappy*, including virtual TPM [52], [65] and virtual secure hardware tokens [21]. However, such devices could be used if they can be cryptographically bound to a source of uniqueness shown in this section.

TABLE VII. COMPARISON OF ADDITIONAL SOURCES OF UNIQUENESS AND THEIR SECURITY

	SIM Card	Intel CPU [27]	SGX	PUF [33] [42]
Unique Resource	Phone Number	Provisioning Secret		PUF Response
Method of Proving Unique Resources	SMS Authentication	Remote attestation		Challenge and Response Authentication
Rate-limiting Depends on Device Security	No	No		No

necessary security and privacy guarantees (see Section VIII for more discussion). Moreover, vulnerabilities of the underlying hardware may also cause further disruption to the guarantees (e.g., TEE side-channel attacks, attacks on PUF).

However, it is worth emphasizing that the security of the devices *do not* impact the rate-limiting functionality of *Scrappy*. The reason behind this is similar to that discussed in Section VIII: Even if the source of uniqueness gets stolen from the device, the GM will only provide a valid *cred* to only one device, be that the legitimate signer’s or the attacker’s.

G. Device interoperability

Supporting many types of hardware security devices may also raise interoperability concerns. This is not an issue for *Scrappy* as long as the device and the signer owning the device adhere to the following restrictions: (1) The device satisfies the four properties shown in Section IV, and (2) The signer does not deviate from the protocol in Section V-D.

XI. LIMITATIONS

Features other than rate limiting: *Scrappy* is a system that provides rate limiting. This implies that it cannot fully replace existing systems such as CAPTCHAs and SMS authentication. For instance, websites that rely on CAPTCHAs to provide human-bot distinction or SMS authentication to provide second-factor authentication cannot utilize *Scrappy*. However, we believe certain use cases exist where websites can benefit from *Scrappy*; for example, servers that can tolerate a certain amount of bot activity.

Requires supported devices: *Scrappy* requires signers to own at least one of the devices shown in the paper. In theory, *Scrappy* can incorporate any device that satisfies the requirements listed in Section IV, and there are a wide variety of devices to choose from, as we have shown in Table VII. However, in reality, the GM is the entity that decides whether to support a particular device. It is highly unlikely that a GM will support every single type of device since it will increase the maintenance cost. Therefore, we presume that GMs will support only a handful of devices, restricting the choice of devices available for use in *Scrappy*.

Actions are time-window-based: Signers who have taken action within a time window are required to wait until the next time window. This is by design, as it is a clear violation of rate-limiting if a signer can produce multiple valid rate-assuring proofs within the same t for the same origin. However, waiting for the next time window introduces user friction and constitutes another limitation of *Scrappy*.

XII. CONCLUSION & FUTURE WORK

This paper proposes `Scrappy`, a novel rate-limiting protocol that uses unforgeable yet privacy-preserving rate-assuring proofs allowing verifiers to have strong assurance that signers are not acting abusively. In contrast to previous approaches, the rate-limiting capability of our method does not rely on the security of the underlying hardware security device. Through a comprehensive security evaluation of `Scrappy`, we show that the proposed system can defend against adversaries targeted in our threat model. We also provide an extensive performance evaluation of our proof-of-concept implementation, showing its practicality under both realistic and extreme conditions. We identify the following as potential directions for future work:

- Extend `Scrappy` to support EPID and other hardware security devices,
- Explore methods to protect secret keys from side-channel attacks,
- Optimize the system to further reduce latency and storage overhead, and
- Investigate the deployment of `Scrappy` in real-world scenarios.

ACKNOWLEDGMENT

We thank the anonymous reviewers and the shepherd for providing us with valuable feedback on prior versions of this paper. The first author was supported in part by the Cybozu Lab Youth program. Authors from Keio University received support from the Keio SFC Internet Research Laboratory. The second author was supported in part by The Nakajima Foundation.

REFERENCES

- [1] Android Keystore system — Android Developers. [Online]. Available: <https://developer.android.com/training/articles/keystore?hl=en>
- [2] Gin Web Framework. [Online]. Available: <https://gin-gonic.com/ja/>
- [3] google/go-tpm. [Online]. Available: <https://github.com/google/go-tpm>
- [4] Key and ID Attestation — Android Open Source Project. [Online]. Available: <https://source.android.com/docs/security/features/keystore/attestation>
- [5] Learn the Kotlin programming language — Android Developers. [Online]. Available: <https://developer.android.com/kotlin/learn?hl=en>
- [6] Let's Encrypt. [Online]. Available: <https://letsencrypt.org/>
- [7] miracl/core: MIRACL Core. [Online]. Available: <https://github.com/miracl/core>
- [8] Save data in a local database using Room — Android Developers. [Online]. Available: <https://developer.android.com/training/data-storage/room?hl=en>
- [9] The Go Programming Language. [Online]. Available: <https://go.dev/>
- [10] What is Android — Android. [Online]. Available: https://www.android.com/intl/en_us/what-is-android/
- [11] xaptum/ecdaa: A C implementation of elliptic-curve-based Direct Anonymous Attestation (DAA) signatures. Created to support the Xaptum Edge Network Fabric, an IoT Network Solution. [Online]. Available: <https://github.com/xaptum/ecdaa/blob/master/doc/IMPLEMENTATION.md>
- [12] xaptum/ecdaa: A C implementation of elliptic-curve-based Direct Anonymous Attestation (DAA) signatures. Created to support the Xaptum Edge Network Fabric, an IoT Network Solution. [Online]. Available: <https://github.com/xaptum/ecdaa>
- [13] (October, 2021) About SQLite. [Online]. Available: <https://www.sqlite.org/about.html>
- [14] L. v. Ahn, M. Blum, N. J. Hopper, and J. Langford, "CAPTCHA: Using hard AI problems for security," in *International conference on the theory and applications of cryptographic techniques*. Springer, 2003, pp. 294–311.
- [15] Akamai. (2022, 03) Rate limiting. [Online]. Available: <https://techdocs.akamai.com/adaptive-media-delivery/reference/rate-limiting>
- [16] C. A. W. Alex Davidson, Jana Iyengar. (2022, October) draft-ietf-privacypass-architecture-08 - The Privacy Pass Architecture. Internet Engineering Task Force. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-privacypass-architecture/>
- [17] F. Alliance. FIDO Alliance Specifications Overview - FIDO Alliance. [Online]. Available: <https://fidoalliance.org/specifications/>
- [18] —. FIDO ECDAAs Algorithm Overview. [Online]. Available: <https://fidoalliance.org/specs/fido-uaf-v1.1-id-20170202/fido-ecdaa-algorithm-v1.1-id-20170202.html#overview>
- [19] E. Brickell, J. Camenisch, and L. Chen, "Direct anonymous attestation," in *Proceedings of the 11th ACM conference on Computer and communications security*, 2004, pp. 132–145.
- [20] E. Brickell and J. Li, "Enhanced Privacy Id: A Direct Anonymous Attestation Scheme with Enhanced Revocation Capabilities," in *Proceedings of the 2007 ACM Workshop on Privacy in Electronic Society*, ser. WPES '07. New York, NY, USA: Association for Computing Machinery, 2007, p. 21–30. [Online]. Available: <https://doi.org/10.1145/1314333.1314337>
- [21] Bulwark. Virtual FIDO. [Online]. Available: <https://github.com/bulwarkid/virtual-fido>
- [22] E. Bursztein, S. Bethard, C. Fabry, J. C. Mitchell, and D. Jurafsky, "How Good Are Humans at Solving CAPTCHAs? A Large Scale Evaluation," in *2010 IEEE Symposium on Security and Privacy*, 2010, pp. 399–413.
- [23] D. Chaum and E. v. Heyst, "Group signatures," in *Workshop on the Theory and Application of Cryptographic Techniques*. Springer, 1991, pp. 257–265.
- [24] Cloudflare. (2022, 11) Configuring Cloudflare Rate Limiting. [Online]. Available: <https://support.cloudflare.com/hc/en-us/articles/115001635128-Configuring-Cloudflare-Rate-Limiting>
- [25] C. C. Consortium, "Confidential Computing: Hardware-Based Trusted Execution for Applications and Data," *A Publication of The Confidential Computing Consortium*, January 2021.
- [26] I. Corporation. (2014) Strengthening Security with Intel (R) Platform Trust Technology. [Online]. Available: <https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/enterprise-security-platform-trust-technology-white-paper.pdf>
- [27] V. Costan and S. Devadas, "Intel SGX explained," *Cryptology ePrint Archive*, 2016.
- [28] K. Crichton, N. Christin, and L. F. Cranor, "How Do Home Computer Users Browse the Web?" *ACM Transactions on the Web (TWEB)*, vol. 16, no. 1, pp. 1–27, 2021.
- [29] Q. H. Dang *et al.*, "Secure hash standard," 2015.
- [30] A. Davidson, I. Goldberg, N. Sullivan, G. Tankersley, and F. Valsorda, "Privacy Pass: Bypassing Internet Challenges Anonymously." *Proc. Priv. Enhancing Technol.*, vol. 2018, no. 3, pp. 164–180, 2018.
- [31] G. Developers. (2018, May) Native Messaging - Chrome Developers. [Online]. Available: <https://developer.chrome.com/docs/apps/nativeMessaging/>
- [32] Fastly. Rate limiting. [Online]. Available: <https://developer.fastly.com/learning/concepts/rate-limiting/>
- [33] Y. Gao, S. F. Al-Sarawi, and D. Abbott, "Physical unclonable functions," *Nature Electronics*, vol. 3, no. 2, pp. 81–91, 2020.
- [34] D. George, W. Lehrach, K. Kinsky, M. Lázaro-Gredilla, C. Laan, B. Marthi, X. Lou, Z. Meng, Y. Liu, H. Wang *et al.*, "A generative vision model that trains with high data efficiency and breaks text-based CAPTCHAs," *Science*, vol. 358, no. 6368, p. eaag2612, 2017.
- [35] I. GlobalPlatform. (2018, May) Introduction to Secure Elements. [Online]. Available: <https://globalplatform.org/wp-content/uploads/2018/05/Introduction-to-Secure-Element-15May2018.pdf>
- [36] Go. gomobile. [Online]. Available: <https://pkg.go.dev/golang.org/x/mobile/cmd/gomobile>
- [37] Google. Google Chrome - Download the Fast, Secure Browser

- from Google. [Online]. Available: https://www.google.com/intl/us_en/chrome/
- [38] T. C. Group. (2019, November) Trusted Platform Module Library Part 1: Architecture. [Online]. Available: https://trustedcomputinggroup.org/wp-content/uploads/TCG_TPM2_r1p59_Part1_Architecture_pub.pdf
- [39] ——. (2020, June) TCG Algorithm Registry. [Online]. Available: https://trustedcomputinggroup.org/wp-content/uploads/TCG-_Algorithm_Registry_r1p32_pub.pdf
- [40] ——. “Trusted Platform Module Library Part 3: Commands,” https://trustedcomputinggroup.org/wp-content/uploads/TCG_TPM2_r1p59_Part3_Commands_pub.pdf, June 2020.
- [41] ——. (2023, August) Vendor ID Registry. [Online]. Available: <https://trustedcomputinggroup.org/resource/vendor-id-registry/>
- [42] C. Herder, M.-D. Yu, F. Koushanfar, and S. Devadas, “Physical Unclonable Functions and Applications: A Tutorial,” *Proceedings of the IEEE*, vol. 102, no. 8, pp. 1126–1141, 2014.
- [43] H. N. Jacob, C. Werling, R. Bühren, and J. Seifert, “faultTPM: Exposing AMD fTPMs’ Deepest Secrets,” in *8th IEEE European Symposium on Security and Privacy, EuroS&P 2023, Delft, Netherlands, July 3-7, 2023*. IEEE, 2023, pp. 1128–1142. [Online]. Available: <https://doi.org/10.1109/EuroSP57164.2023.00069>
- [44] A. Kosei. akakou/scrappy: SeCure Rate Assuaring Protocol with PrivacY. [Online]. Available: <https://github.com/akakou/scrappy>
- [45] Lenovo. (2020, July) ThinkPad_A285_Spec.PDF. [Online]. Available: https://psref.lenovo.com/syspool/Sys/PDF/ThinkPad/ThinkPad_A285/ThinkPad_A285_Spec.PDF
- [46] R. Lindemann. (2017, February) FIDO ECDA Algorithm. FIDO Alliance. [Online]. Available: <https://fidoalliance.org/specs/fido-uaf-v1.1-ps-20170202/fido-ecdaa-algorithm-v1.1-ps-20170202.pdf>
- [47] G. Maganis, E. Shi, H. Chen, and D. Song, “Opaak: using mobile phones to limit anonymous identities online,” in *Proceedings of the 10th international conference on Mobile systems, applications, and services*, 2012, pp. 295–308.
- [48] T. Meunier. (March, 2021) Humanity wastes about 500 years per day on CAPTCHAs. It’s time to end this madness. Cloudflare Inc. [Online]. Available: <https://blog.cloudflare.com/introducing-cryptographic-attestation-of-personhood/>
- [49] S. Microelectronics, “ST33TPHF2ESPI Product Brief,” 2019.
- [50] Microsoft. Windows 11 requirements. [Online]. Available: <https://learn.microsoft.com/en-us/windows/whats-new/windows-11-requirements#hardware-requirements>
- [51] Y. Nakatsuka, E. Ozturk, A. Paverd, and G. Tsudik, “{CACTI}: Captcha Avoidance via Client-side {TEE} Integration,” in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 2561–2578.
- [52] R. Perez, R. Sailer, L. van Doorn *et al.*, “vTPM: virtualizing the trusted platform module,” in *Proc. 15th Conf. on USENIX Security Symposium*, 2006, pp. 305–320.
- [53] Prowebscraper. (2017, December) Top 10 Captcha Solving Services Compared. [Online]. Available: <https://prowebscraper.com/blog/top-10-captcha-solving-services-compared/>
- [54] H. Raj, S. Saroiu, A. Wolman, R. Aigner, J. Cox, P. England, C. Fenner, K. Kinshumann, J. Loeser, D. Mattoon, M. Nystrom, D. Robinson, R. Spiger, S. Thom, and D. Wooten, “fTPM: A Software-Only Implementation of a TPM Chip,” in *25th USENIX Security Symposium (USENIX Security 16)*. Austin, TX: USENIX Association, aug 2016, pp. 841–856.
- [55] A. Rao. (2022, January) Rising to the Challenge – Data Security with Intel Confidential Computing. [Online]. Available: <https://community.intel.com/t5/Blogs/Products-and-Solutions/Security/Rising-to-the-Challenge-Data-Security-with-Intel-Confidential/post/1353141>
- [56] E. T. Rolf Lindemann. (2017, February) FIDO UAF Protocol Specification. [Online]. Available: <https://fidoalliance.org/specs/fido-uaf-v1.1-ps-20170202/fido-uaf-protocol-v1.1-ps-20170202.html>
- [57] X. Ruan, “Privacy at the next level: Intel’s enhanced privacy identification (epid) technology,” in *Platform Embedded Security Technology Revealed*. Springer, 2014, pp. 117–141.
- [58] M. Scott, “The Apache Milagro Crypto Library (Version.”
- [59] A. Searles, Y. Nakatsuka, E. Ozturk, A. Paverd, G. Tsudik, and A. Enkoji, “An Empirical Study & Evaluation of Modern CAPTCHAs,” in *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*, J. A. Calandrino and C. Troncoso, Eds. USENIX Association, 2023. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity23/presentation/searles>
- [60] SoloKeys. SoloKeys. [Online]. Available: <https://solokeys.com/>
- [61] ——. SoloKeys Solo1. [Online]. Available: <https://github.com/solokeys/solo1>
- [62] P. I. E. Staffs. (2018, August) Security Concerns Surrounding WebAuthn: Don’t Implement ECDA (Yet) - Paragon Initiative Enterprises Blog. [Online]. Available: <https://paragonie.com/blog/2018/08/security-concerns-surrounding-webauthn-don-t-implement-ecdaa-yet>
- [63] J. Sumrak. (2021, December) SMS Verification: What It Is & How It Works — Twilio. Twilio inc. [Online]. Available: <https://www.twilio.com/blog/what-is-sms-verification>
- [64] TextFree. Free Texting and Calling — TextFree. [Online]. Available: <https://textfree.us/>
- [65] VMWare. (2021, 08) Virtual Trusted Platform Module Overview. [Online]. Available: https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.vm_admin.doc/GUID-6F811A7A-D58B-47B4-84B4-73391D55C268.html
- [66] Xaptum. amcl. [Online]. Available: <https://github.com/xaptum/amcl>
- [67] ——. Details of the ECDA Algorithm Used. [Online]. Available: <https://github.com/xaptum/ecdaa/blob/master/doc/IMPLEMENTATION.md>
- [68] ——. ecdaa. [Online]. Available: <https://github.com/xaptum/ecdaa>
- [69] L. Xi, D. Feng, Y. Qin, F. Wei, J. Shao, and B. Yang, “Direct Anonymous Attestation in practice: Implementation and efficient revocation,” in *2014 Twelfth Annual International Conference on Privacy, Security and Trust, Toronto, ON, Canada, July 23-24, 2014*, A. Miri, U. Hengartner, N. Huang, A. Jøsang, and J. García-Alfaro, Eds. IEEE Computer Society, 2014, pp. 67–74. [Online]. Available: <https://doi.org/10.1109/PST.2014.6890925>