

# SMART Credentials in the Multi-queue of Slackness (or Secure Management of Anonymous Reputation Traits without Global Halting)

Jack P. K. Ma

Department of Information Engineering  
The Chinese University of Hong Kong  
Shatin, N.T., Hong Kong

Sherman S. M. Chow

Department of Information Engineering  
The Chinese University of Hong Kong  
Shatin, N.T., Hong Kong

**Abstract**—Anonymous credentials encourage online communication without fear of surveillance, but may invite misbehavior like hate speech. Previous updatable anonymous credentials keep a chronological queue of authenticated sessions and a global pointer to the last chunk of judged sessions. This design allows efficient authentication for proving over only a subset of sessions. However, complications in subjective evaluation often introduce hard-to-judge sessions, which halt all users since sessions that come after the global pointer cannot be redeemed, eventually exceeding the queue size that limits the creation of bad sessions. Such a global-halting loophole may also make judgments overly harsh and hasty.

We propose SMART (slack management of anonymous reputation traits), maintaining multiple queues so the server could issue interim judgments many times before finalization. Such slackness removes the binary judgment of old methods and mitigates the global-halting issue. Prior schemes only allow score upgrades (WPES '14) or require proving against a global session list since the last checkpoint (S&P '22). Our authentication time is linear in the number of queues or sessions in a designated queue for immediate revocation.

## 1. Introduction

Securing anonymous access to internet services such as social media and collaborative platforms fosters participation via privacy protection, among other benefits. For example, users may risk censorship and legal/social persecution for posting “sensitive” content online. Unfortunately, anonymous access can also facilitate harmful behavior by allowing malicious users to evade detection and carry out evil deeds such as spreading misinformation or cyberbullying. Service providers (SPs) like Wikipedia [2] restrict what users connecting via Tor [23] can do since its anonymous connection makes it hard to prevent vandals.

Some systems address the tension between anonymity and accountability with trusted third parties (TTPs). TTP-based revocation (e.g., group signatures [15], [21], [38], traceable signatures [17], [28] or “revocable anonymous credentials” [1], [14]) often assumes that the misbehavior has been deanonymized by external means. A TTP could use the identifying information to recover some tracing tokens for revocation. Some systems define misbehavior

with objectively evaluable criteria, e.g., double-voting [19] or exceeding the limit of  $k$ -times authentication [9]. It takes no external mechanism to judge. Such misbehavior can often be detected quickly.

Many systems allow diverse forms of user interactions. Subjective assessments are needed to identify misbehavior. The SP may resort to “external” evaluation, say, voting among the clients, user reports, judgment of moderators, or a professional review by domain experts such as lawyers. This process often takes time, especially when there are controversies or complexities like translation, e.g., “Zhemaos” hoaxes on simplified Chinese Wikipedia.

Completely blocking users may not always be desirable. It is important to incentivize positive contributions by maintaining reputation scores. It also helps alleviate sybil attacks. Users who have yet to build up a reputation might be deemed less trustworthy than those who have. Similar to judging bad deeds, (finalizing the) subjective evaluation can be time-consuming. Systems with slow/sporadic reviews, e.g., Wikipedia, might employ a “time-dependent model” (cf., blockchain consensus) – if a contribution has not been flagged inappropriate for, say, six months, it is deemed solid, and a score will be assigned.

### 1.1. Two Major Paradigms

**Judging since Genesis.** Blacklistable anonymous credentials (BLAC) [33] feature TTP-free subjective judgment, enforced by zero-knowledge proof (ZKP) against a global blacklist  $\mathcal{L}$ . A user generates a ticket  $\text{tag} = \text{PRF}_x(\text{sID})$ , a pseudorandom function (PRF) output of its credential secret  $x$  over an assigned session identifier  $\text{sID}$ . If a session is judged to be bad, the SP puts  $\text{sID}$  and its tag into  $\mathcal{L}$ . BLAC is readily extensible to BLACR [8] with reputation via proving membership (claiming positive scores) on top of non-membership proof. Its authentication costs  $O(|\mathcal{L}|)$ .

SNARKBlock [30] improves the non-revoked proof in BLAC with zero-knowledge succinct non-interactive arguments [27] (zkSNARK) for  $O(\log |\mathcal{L}|)$ -time verification. Since the user proving complexity remains linear, hidden common input aggregate proofs [30] is proposed, which allows for the aggregation of existing proof for different chunks of the blacklist. However, if we update  $\mathcal{L}$ , say, using it for reputation [8] and allowing updatable scores, it invalidates the corresponding (offline) precomputation. As an amortized approach, infrequent users need to catch up with linear proof for all sessions that have happened

*Sherman Chow (corresponding) is supported by General Research Fund (CUHK 14210621 and 14209918) and is grateful for the Early Career Award on “Accountable Privacy in Online Communication” (439713) from RGC, Hong Kong. We thank reviewers for their helpful comments.*

since they walked away. New users cannot benefit either unless they reveal their joining time to skip older chunks. Their authentication is set apart from other users and degrades their anonymity. BLACR-express [8] faces similar shortcomings. To mitigate the scalability bottleneck due to the ever-growing global list, we study alternative designs.

**A Fixed-Size Queue Design.** PEREA [34] proposes to store only the  $K$  most recent session identifiers as *tickets* in a credential and use  $K$  (non-)membership proofs for claiming positive scores or showing they are guilt-free. Reputation can only be transient. PERM [7] keeps tickets in a user credential<sup>1</sup>. So, (infrequent/new) users only prove w.r.t. their own fixed-size credentials but not the ever-growing list of global sessions. Unfortunately, this design comes with a dealbreaker. Adversaries are motivated to authenticate rapidly to *wash out* problematic tickets from the fixed-sized queue. A possible remedy is to tightly couple another *anonymous* rate-limiting mechanism with the credential, which is non-trivial.<sup>2</sup> Rate-limiting also unnecessarily curbs the enthusiasm of honest users [37].

## 1.2. The Unspoken Global Halting Issue

To support persistent reputation scores besides blocklisting, users must absorb judgments into their credentials upon each interaction, even if judgment comes later. Later designs [7], [35] speed up authentication by skipping sessions that have been absorbed or proven irrelevant.

PERM [7] extends PEREA with a global *judgment pointer* pointing to the last one in a *continuous* chronological sequence of judged sessions, enabling an easy proof of judged/unjudged status of tickets in a credential instead of proving w.r.t.  $\mathcal{L}$ . Meanwhile, the last unjudged ticket in a credential cannot be too far away from such a pointer. Without this rule as a part of authentication semantics, malicious users can inject many “bad” sessions before being caught, leading to a *denial-of-service (DoS)* vulnerability of overloading the session evaluation operation.

FARB [35] explicitly makes the sequential order assumption on session identifiers (higher IDs for later sessions). Every ticket after the global pointer must be unjudged. A single range proof can replace PERM’s linear disjunctive proofs (judged/unjudged) of each ticket.

Cryptographic enforcement of such authentication semantics using the global pointer is incompatible with the time-consuming evaluation based on subjective criteria. A *single* controversial session that takes long to finalize stalls the moving of *global judgment pointer* and halts *all users*, not just the one who originated the session, no matter how many later sessions have been judged. The technical reason for such *global halting* is that no users can redeem (and remove) any sessions beyond the global judgment pointer (pointing to the *earliest* unjudged session) since their status cannot be inferred without direct explicit proof.

Once the global pointer moves a bit slower than the authentication rate (and hence the birth rate of sessions),

1. Non-membership proof via an accumulator takes constant verification costs. However, creating a witness is linear and requires recomputation when the accumulated values change. PERM replaces PEREA’s  $K$  such proofs w.r.t. the public accumulator storing all blocked tickets to  $K$  disjunctive proofs, allowing faster authentication with  $K > 10$  [34].

2. The rate-limiting of  $k$ -time anonymous authentication [9] is periodic. A burst of  $2k$  authentications can be run across the epoch boundary.

tickets accumulate in the credential and reach the queue size limit for DoS resilience. Such a global-halting vulnerability also affects the operational aspect of judgments, forcing the subjective evaluation operates like a single-thread process on a first-come-first-served basis.

**Notable Exceptions.** PE(AR)<sup>2</sup> proposed by Yu *et al.* [37] is an interesting attempt to achieve the best of both worlds, namely, maintaining  $K$  pending sessions in the credential but without mandating sequential judging. It uses a blacklist to revoke [33] instead of assigning a very negative score [35]. Apart from a faster non-membership proof for a set of elements, it supports “soft” rate-limiting by storing a variable number of unjudged sessions and proving that it is less than some given threshold, albeit at a degraded level of anonymity PE(AR)<sup>2</sup> does not support negative/downgradable scores. ARBRA [36] supports negative scores, but the user authentication cost is again linear in  $|\mathcal{L}|$  plus the number of the user’s unredeemed tickets. We defer our concurrent work [20] to concluding remarks.

## 1.3. Revocation with Mercy/Elasticity or Urgency

Anonymous blocklisting/reputation systems should allow “revocation with mercy.” When accused, anonymous blocklisting suspends users instead of permanently revoking access. If cleared, users can be removed from blocklists. It is exactly the existence of “difficult-to-judge” sessions that motivates the need for anonymous blocklisting.

The sequential requirement of PERM/FARB mandates the SP to judge sequentially and only once. The one-time judgment forces SPs to catch misbehaving users in time, or otherwise, they may escape from punishment forever. These system constraints severely limit the applicability. Consider contact tracing; a user can be “forgiven” if the place s/he checked in is indeed safe (and vice versa). Given the system constraints, overly cautious and harsh judgments for preemptive measures seem logical.

Enabling multi-judging, specifically downgrading, can be challenging since users have no incentive to voluntarily redeem worse ratings. Naïve attempt probably involves inefficient ZKP enumerating all sessions ever associated with a credential, not to mention the wash-away problem.

PERM/FARB also does not readily allow immediate revocation. The SP must rapidly *finalize* the judgment of all tickets prior to the one needing immediate attention.

## 1.4. Motivating the Design Choices of SMART

We design SMART for “secure management of anonymous reputation traits” with a few distinctive features.

**Sublinear Authentication.** SMART mandates users to redeem the (temporary) score of the earliest session across all queues (independent of their number of transient judgments). The authentication cost scales linearly with the number of queues, or sublinear in the number of user sessions  $NK$  when each queue has  $K$  of them, as unjudged tickets of a queue can be shown via range proofs.

**Multi-queue of Slackness.** A credential maintains multiple queues, each for sessions with the same number of transient judgments made but with different transient (*i.e.*, not finalized) scores. Making one more transient judgment

on a session moves it to the next queue. When the score is finalized, the user redeems it out of the credential.

Supporting multiple queues brings us many benefits. Scores of different sessions can keep updating (up to  $(N - 1)$  times) and possibly at different paces. They could also be finalized at different times. The SP now has more flexibility when finalizing the session scores. Our design thus fits with different operational characteristics.

Maintaining multiple queues seems an intuitive and arguably “trivial” solution to avoid one single queue from getting stuck due to global halting. Note that the global halting problem happens at the head of a queue but not the end. Enlarging the buffer alone does not help much.

**Custom Rate-Limiting.** More importantly, we argue that the apparent quantitative difference from a single queue actually contributes to qualitative differences in two ways.

FARB has built-in rate-limiting from range proof. In single-queue FARB, rate-limiting must be enforced; otherwise, it is vulnerable to DoS attacks of rapidly inducing many bad sessions before the blocklisting becomes effective, *i.e.*, when the earliest such session is put into the blocklist and the authenticating user is forced to redeem it.

Note that these newly spawned sessions only appear at the end of a queue. In our multi-queue design, they all appear in  $\mathcal{Q}_0$ . In other words, SMART only needs to rate-limit  $\mathcal{Q}_0$ . Note also that the SP knows at which queues after  $\mathcal{Q}_0$  those difficult-to-judge sessions are. The SP can set a different rate limit accordingly or even remove the limit altogether. This ensures the correctness of SMART’s cryptographic enforcement. Namely, transient judgments move the tickets to the next queue, which will never be full, in contrast to the (only) base queue  $\mathcal{Q}_0$  (in FARB).

**Immediate Revocation.** FARB features constant-time authentication due to its clever use of range proof instead of proving about each ticket in the single queue  $\mathcal{Q}_0$ . With multiple queues, SMART can assign different “responsibilities” to them. Of particular interest is an emergency queue for immediate revocation. Specifically, any severely serious session will be sent there. This queue does not need to be long. Authentication now runs a special linear scan of the emergency queue. Here, we can do reputation-based blocklisting via assigning a very negative score, or a direct blocklisting that merely a single ticket there will block the credential from further authentication. Our easiest implementation is to make this queue the last one.

Rendering this idea with FARB runs into a dilemma. This queue cannot be long for efficiency with the forced linear scan (which spoils the constant-time authentication time of FARB). This queue cannot be short either for being the only storage for unjudged sessions. In short, prior schemes with sequential judging require the SP to judge once-and-for-all. An immediate revocation of a new session can only be performed by removing all previously unjudged sessions. The SP either blocks or forgives all of them. SMART allows transient judgments to keep the affected sessions and extend to support microscopic queue management, *e.g.*, designated queues to handle these bans.

**Mitigating Global-Halting.** Sequential judging still exists to a certain degree since we move a sequential chunk of tickets to the next queue. However, the backlog of non-finalized sessions is split across different queues.

There does exist a pathological case that our system will reach the “limit.” Still, the possibility is dramatically trimmed down along three dimensions – individual transient judgment, enlarging the buffer, and customizable queue-specific rate-limiting. The SP can make informed choices in “slacking off” along these dimensions. In short, the SP can better manage the operations for subjective evaluation with all these additional features of SMART.

**Revocations with Mercy and Elasticity.** Transient scoring naturally supports revocation with mercy/elasticity.

**Versatile Applicabilities.** Consider contact tracing; a user checked in a place can have an initial (neutral) score indicating the risk level. For visits later identified as posing a high risk, the health department can downgrade the score (*e.g.*, to the minimum). After the negative results of some antigen tests, it can be upgraded. With more negative results from regular tests, the risk of a visit will eventually vanish. The number of changes before finalization is limited, fitting the time-dependent model.

Consider content reviewing (*e.g.*, videos or photos); an uploaded post may go through an initial review and is put aside. The SP may wish to temporarily reward/punish the uploader by issuing transient judgments. The multi-queue design allows a multi-layered (periodic) review process, *e.g.*, the post can be judged monthly, and finalizing and blocking can be done when needed. This also fits some (anonymous) point reward systems that take a few days to review each point collection/transaction.

**Cryptographic Enforcement.** Ensuring the efficiency of authentication requires careful formulations of how users are identified and how their sessions are scored. SMART uses multiple queues to ease the tensions between operational characteristics of judgment and its corresponding cryptographic enforcement. It boils down to moving sessions across credential queues (rather than simply adding them to a single queue). This aligns with how judgments are updated across the SP-side global queues. We extensively use zero-knowledge proof of knowledge (PoK) of credential signatures (signatures with efficient protocols) and homomorphic commitments for credential updates. To prevent attackers from reusing signatures and rolling back credentials, we rely on nonces and data-structure invariants, namely, monotonic increasing session identifiers and the head/tail pointers of the queues.

## 2. Preliminaries

Let  $\lambda$  be the security parameter and  $\mathbb{G}$  be a cyclic group of order  $p$ , where  $p$  is a  $\lambda$ -bit prime. PPT means probabilistic polynomial time (in  $\lambda$ ). The set of integers  $\{0, \dots, N - 1\}$  is denoted by  $[N]$ . If  $X$  is a finite set,  $x \leftarrow X$  denotes sampling an element  $x$  uniformly at random from the set  $X$ . An empty output is denoted by  $\perp$ .

Let  $\mathcal{R} : \{0, 1\}^* \times \{0, 1\}^* \rightarrow \{0, 1\}$  be a binary relation which takes a statement  $x$  and witness  $w$ . Let  $\mathcal{L}_{\mathcal{R}} = \{x \mid \exists w \text{ s.t. } \mathcal{R}(x, w) = 1\}$  be a language associated with the relation  $\mathcal{R}$ , which can also include a common reference string  $\text{crs}$  as a part of the statement  $x$ .

We use  $P(\mathcal{U}(in_{\mathcal{U}}), \mathcal{S}(in_{\mathcal{S}})) \rightarrow (out_{\mathcal{U}}, out_{\mathcal{S}})$  to denote parties  $\mathcal{U}$  and  $\mathcal{S}$  interact via the protocol  $P$  taking  $in_i$  and outputting  $out_i$  for  $i \in \{\mathcal{U}, \mathcal{S}\}$ , respectively.

## 2.1. Homomorphic (Vector) Commitments

In a commitment scheme  $\Pi^{\text{com}}$ , the sender can commit to a message  $m$  via a commitment  $C_m = \text{Com}(m; r)$  for some randomness  $r$ . It can later be opened by providing  $(m, r)$  (possibly in ZKP). For an ordered tuples of message  $M = (m_1, \dots, m_n)$ , we denote the set of committed messages by  $C_M = \{\text{Com}(m_i; r_i)\}_{i=1}^n$ . Note that this set of commitments does not preserve the message-index relation  $(m_i, i)$  in the vector, *i.e.*, it is not position binding.

**Definition 1.** A commitment scheme is defined as follows:  $\text{Setup}(1^\lambda) \rightarrow \text{pp}$ : It takes as input the security parameter  $1^\lambda$ , and outputs a public parameter  $\text{pp}$ .

$\text{Com}_{\text{pp}}(m; r) \rightarrow (C, o)$ : Given the public parameter  $\text{pp}$  (as an implicit input), the message  $m$ , and randomness  $r$ , it outputs a commitment  $C$  and opening  $o = (m, r)$ .

$\text{Vf}(\text{pp}, C, m, o) \rightarrow b$ : The verification algorithm outputs 1 if  $o$  is a valid opening for the commitment  $C$ , 0 otherwise.

It should satisfy the hiding and binding properties.

**(Perfect) Hiding.** Given  $\text{pp}$ , any (computationally unbounded) adversary who chooses  $(m_0, m_1)$  cannot distinguish  $b$  when given  $c = \text{Com}(m_b; r_b)$  for random  $b$ .

**Binding.** No PPT adversary can produce valid openings  $o_0, o_1$  that open a commitment  $c$  to distinct  $m_0, m_1$ .

In a vector commitment scheme, one can commit to a vector of message  $\vec{m} = (m_1, \dots, m_n)$ . The commitment should be hiding and (position)-binding.

The (generalized) Pedersen commitment [29] allows  $m$  to be a vector of field elements. Let  $\{g_i\}_{i=1}^{n+1}$  be independent generators of  $\mathbb{G}$ . A commitment of  $\vec{m} \in (\mathbb{Z}_p)^n$ , denoted by  $C_{\vec{m}} = \text{Com}(\{m_i\}_{i=1}^n; r)$ , is  $g^{r_{n+1}} \prod_{i=1}^n g_i^{m_i}$ , where  $r$  is randomly picked from  $\mathbb{Z}_p^*$ . It also features homomorphism, *i.e.*,  $\text{Com}(\vec{a}) \cdot \text{Com}(\vec{b}) = \text{Com}(\vec{a} + \vec{b})$ .

One can prove that the set of committed messages  $\{C_{m_i} = \text{Com}(m_i)\}_{i=1}^n$  commits to the entries of a generalized Pedersen commitment  $C_{\vec{m}}$  on vector  $\vec{m}$ . We use Pedersen commitments for our credential attributes.

## 2.2. Zero-Knowledge Proof or Argument

A prover can convince a verifier via a zero-knowledge proof of knowledge (ZKPoK) protocol that a statement  $x$  is true without revealing  $w$  and any other information. A (non-interactive) proof system  $\Pi$  is defined as follows.

**Definition 2** (Proof System). A proof system for a language  $\mathcal{L}$  consists of three algorithms  $\Pi = (\text{Setup}, \text{P}, \text{V})$ :  $\text{Setup}(1^\lambda) \rightarrow \text{crs}$ : It outputs the common reference string.

$\text{P}(\text{crs}, x, w) \rightarrow \pi$ : It outputs a proof for the statement  $x$  and witness  $w$ , such that  $(x, w) \in R$  given  $\text{crs}$ .

$\text{V}(\text{crs}, x, \pi) \rightarrow b$ : It outputs 1 if  $\pi$  is a valid proof for  $x$ .

A proof system must be complete and sound.

**Completeness.**  $\forall x \in \mathcal{L}$  and  $w$  where  $(x, w) \in R$ , the prover can always output  $\pi$  such that the verifier accepts.

**Soundness.**  $\forall x \notin \mathcal{L}$ , no prover can output  $\pi$  to make a verifier accepts.

A zero-knowledge proof (argument) of knowledge is a proof system (denoted by ZKP) that satisfies knowledge

soundness (against computationally bounded provers) and zero-knowledgeness (defined in Appendix A.1).

**Knowledge Soundness.** No cheating provers can convince an honest verifier that for some false statement  $x \notin \mathcal{L}$ . It requires the existence of an extractor  $\text{Ext}$  that can extract the witness  $w$  from any proof  $\pi$  passing the verification.

**Zero Knowledge.** For all statements  $x \in \mathcal{L}$ , there exists a simulator  $\text{Sim}$  that can simulate a proof indistinguishable from the real proof, *i.e.*, the proof does not leak additional information about the witness except the statement is true.

**ZKP Notation.** We use the standard notation introduced by Camenisch and Stadler [15]. For example,  $\text{ZKP}\{x : y = g^x\}$  denotes ZKPoK for proving  $x$  such that  $y = g^x$ . The values on the left of the colon are not known by any verifier. Symbols on the right are public values.

**2.2.1. Non-Interactive Proof and Disjunctive Proof.**  $\Sigma$ -protocols [31] are three-move ZKP, which can be made non-interactive with the Fiat-Shamir transform.

Disjunctive ZKP of two or more statements can be realized by the secret-sharing technique [5], [22]. For  $R_1 \vee R_2$ , the prover can prove  $R_1$  with the witness under the challenge  $e - e_1$ , where  $e$  is picked by the verifier and  $e_1$  is self-picked randomness, then simulate the ZKP for  $R_2$ .

**2.2.2. zkSNARK for Range & Circuits.** A succinct non-interactive argument of knowledge (SNARK) features succinct proofs sublinear in the statement size. The arithmetic circuit over secret (witness) inputs and public inputs and outputs can be translated into rank-1 constraint systems (R1CS) and proved via (zk)SNARK. Notable zkSNARK examples include Groth16 [27] and Bulletproofs [11].

Bulletproofs support succinct range proof using inner product arguments. It is special honest-verifier zero-knowledge for the relation  $\mathcal{R} : ((\text{pp}, C_v, A, B), (v, r) : C_v = \text{Com}(v; r) \wedge A \leq v < B)$  where  $A = 0, B = 2^n$  and  $\text{Com}$  is a Pedersen commitment ( $g^v h^r$ ) with  $\text{pp} = (g, h)$ . It takes a logarithmic number of rounds in the bit-length  $n$  and can be made non-interactive with a logarithmic proof size via Fiat-Shamir transform. With only a  $\log(m)$  overhead,  $m$  range proof can be aggregated.

**2.2.3. Commit-and-Prove (CP) ZKP and zkSNARK.** LegoSNARK [16], using a trusted setup, enables composing different zkSNARK for modular designs. Let  $\Pi^{\text{com}} = (\text{Setup}, \text{Com}, \text{Vf})$  be a commitment scheme. A CP-SNARK  $(\text{KGen}, \text{P}, \text{V})$  is a zkSNARK for the relation:

$$\mathcal{R}_{\text{CP}} : ((x, (c_1, \dots, c_\ell)), ((m_j, o_j)_{j \in [\ell]}, w) : \bigwedge_{j \in [\ell]} \Pi^{\text{com}}. \text{Vf}(\text{pp}, c_j, m_j, o_j) = 1 \wedge R(x, (m_j)_{j \in [\ell]}, w) = 1)$$

where  $\text{pp} \leftarrow \Pi^{\text{com}}. \text{Setup}(1^\lambda)$ ,  $c_j = \Pi^{\text{com}}. \text{Com}(m_j; o_j)$  and  $R$  is a relation over statement  $x$  on committed inputs  $(m_j)$  and witness  $w$ .

## 2.3. Credential Signatures with Protocols

**Definition 3.** A credential signature scheme  $\text{Cred} = (\text{KGen}, \text{Sig}, \text{Vf})$  for message vectors, associated with a vector commitment scheme  $\Pi^{\text{com}} = (\text{Setup}, \text{Com}, \text{Vf}')$  with public parameter  $\text{pp}$ , is defined by the algorithms/protocols below, all taking  $\text{pp}$  as implicit input.

$\text{KGen}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$ : The probabilistic key generation algorithm takes the security parameter  $\lambda$  and outputs a public key  $\text{pk}$  and a secret key  $\text{sk}$ .

$\text{Sig}(\mathcal{U}(\vec{m}; \rho), \mathcal{S}(\text{aux}, \text{sk})) \rightarrow (\sigma, \text{aux}')$ : In the signature-issuing protocol, the user  $\mathcal{U}$  hides the message  $\vec{m}$  with randomness  $\rho$  using  $\Pi^{\text{com}}$ . The signer  $\mathcal{S}$  inputs  $\text{sk}$  and  $\text{aux}$ . The user receives a signature  $\sigma$ . The signer receives  $\text{aux}'$ .

Both  $\text{aux}$  and  $\text{aux}'$  represent auxiliary information, one for input and one for output. For example,  $\text{aux}$  can be commonly known messages or ZKPs involving previously issued signatures, and  $\text{aux}'$  can be commitments of the messages  $\text{Com}(\vec{m})$ ,  $\{\text{Com}(m_i)\}$  and proofs of opening.

$\text{Vf}(\text{pk}, \vec{m}, \sigma) \rightarrow b$ : The deterministic verification algorithm takes as input the public key  $\text{pk}$ , the message vector  $\vec{m}$ , and a signature  $\sigma$ . It outputs a result bit  $b$ . For brevity,  $(\vec{m}, \sigma)$  might be collapsed into a single input.

It also supports efficient (non-interactive) proof for the knowledge of a signature over a vector commitment of  $\vec{m}$  or over a subset of committed messages, e.g.,  $\text{ZKP}\{(\vec{m}, \rho, \sigma) : \text{Vf}(\text{pk}, \vec{m}, \sigma) = 1 \wedge C_{\vec{m}} = \text{Com}(\vec{m}; \rho)\}$ .

It should be existentially unforgeable against adaptive chosen-message attacks. No PPT adversary can output a valid signature on a message not in the queried message set, given that it can run the signing protocol with the signer for queried messages polynomially many times. The detailed definition is presented in Appendix A.2.

We use the BBS+ signature of Au, Susilo, Mu, and Chow [9]. Camenisch, Drijvers and Lehmann [13] proposed a PoK for signature on partially disclosed messages without target-group operations in the (Type-III) pairing setting, by also publishing signer secrets in the exponent.

In SMART, the user commits the attribute vector  $\text{attr}$  (and its update  $\text{attr}'$ ) and sends it to the SP. For flexibility (e.g., using the most efficient ZKP for a specific relation), the user can also send a set of individual commitments on  $C_{v_i}$  for  $v_i \in \text{attr}$  and proves in ZK that each  $C_{v_i}$  corresponds to the  $i$ -th attribute committed in  $C_{\text{attr}}$ .

### 3. Anonymous Reputation

Anonymous credentials with reputation revoke users by assigning a low score to their previous sessions, barring them from fulfilling authentication policies in the future. A secure design should mandate that users redeem all redeemable sessions before authentication. Redemption can be done one by one or in a batch. It can be a standalone protocol or integrated with the authentication.

Our system allows unblocking by supporting upgrading the score of non-finalized sessions. Meanwhile, the score can also be downgraded. Not all systems offer all these features. For example, PEREA/PERM does not support score downgrading or separate credential updates from authentication; each authentication removes one ticket from the credential, which could even be unjudged!

#### 3.1. Syntax

**Definition 4.** An anonymous credential system with reputation consists of the PPT algorithms/protocols below.

$\text{Setup}(1^\lambda) \rightarrow (\text{pp}, \text{sk}, \text{st})$  is the SP setup algorithm. It inputs a security parameter  $1^\lambda$  and outputs the public parameter  $\text{pp}$ , a secret key  $\text{sk}$ , and public state  $\text{st}$ .

In all protocols below, both the SP and users take the public state  $\text{st}$  as implicit input. If it fails to run, both the SP and users get  $\perp$ ; otherwise, the SP updates  $\text{st}$  into  $\text{st}'$ . Policies  $\text{AP}$ ,  $\text{UP}$ , and update function  $f$  also take  $(\text{pp}, \text{st})$ .

$\text{Reg}(\mathcal{U}(\text{pp}), \mathcal{S}(\text{pp}, \text{sk})) \rightarrow (\text{cred}, \text{st}')$ : The SP  $\mathcal{S}$  uses its secret key  $\text{sk}$  to register the user  $\mathcal{U}$  by creating a credential  $\text{cred}$  with initial attributes defined by  $\text{pp}$ .

$\text{Auth}(\mathcal{U}(\text{pp}, \text{cred}, \text{AP}, f), \mathcal{S}(\text{pp}, \text{sk}, \text{AP}, f)) \rightarrow (\text{cred}', \text{st}')$ : A user  $\mathcal{U}$  with a credential  $\text{cred}$  authenticates to the SP  $\mathcal{S}$ . If it passes authentication policy  $\text{AP}$ , i.e.,  $\text{AP}(\text{cred}) = 1$ ,  $\mathcal{U}$  obtains an updated credential  $\text{cred}'$  on updated attributes  $\text{attr}' \leftarrow f(\text{cred})$  as generated by  $\mathcal{S}$  using its secret key  $\text{sk}$ . The created authenticated session is captured by  $\text{st}'$ .

$\text{StMgn}(\text{pp}, \text{sk}, \text{st}) \rightarrow \text{st}'$ : The SP runs the state management algorithm to render the subjective judgment (not explicitly listed as an input) into a redeemable state. It uses the secret key  $\text{sk}$  to update current state  $\text{st}$  into  $\text{st}'$ .

$\text{Redeem}(\mathcal{U}(\text{pp}, \text{cred}, \text{UP}, f), \mathcal{S}(\text{pp}, \text{sk}, \text{UP}, f)) \rightarrow (\text{cred}', \text{st}')$ : Any user  $\mathcal{U}$  runs this protocol to update  $\text{cred}$  to redeem what is pending to redeem, i.e.,  $\text{UP}(\text{cred}) = 1$ , and barring  $\text{cred}$  from authentication, i.e.,  $\text{AP}(\text{cred}) = 0$ . The update does not generate new authenticated sessions.

Throughout this work, the attributes  $\text{attr}$  are the main credential data (e.g., the data structure storing tickets), while  $\text{cred}$  contains a credential signature on  $\text{attr}$  (and  $\text{attr}$  itself) and possibly other (certified) auxiliary data. In particular, update function  $f$  in SMART also returns “enqueuing data” for the SP to issue enqueueing signatures as a part of the user credential.

#### 3.2. Security Requirements

An anonymous credential system with reputation should provide the following security properties:

**Completeness.** An honest user can authenticate (redeem) with  $\text{cred}$  if it satisfies the access (update) policy, i.e.,  $\text{AP}(\text{cred}) = 1$  ( $\text{UP}(\text{cred}) = 1$ ), when talking to an honest SP. Upon completion,  $\text{cred}$  updates according to  $f(\text{cred})$ .

**Anonymity.** A corrupted SP cannot distinguish any two users who request to authenticate or redeem with credential  $\text{cred}_0$  and  $\text{cred}_1$ , respectively, where  $\text{AP}(\text{cred}_0) = \text{AP}(\text{cred}_1)$  or  $\text{UP}(\text{cred}_0) = \text{UP}(\text{cred}_1)$ , i.e., they hide among the anonymity set indexed by the policy  $\text{AP}$  or  $\text{UP}$ .

**Soundness.** Users with an invalid  $\text{cred}$  cannot authenticate or redeem with SP, i.e.,  $\text{Auth}(\text{Redeem})$  returns  $\perp$  for  $\text{AP}(\text{cred}) = 0$  ( $\text{UP}(\text{cred}) = 0$ ).

Security games are deferred to the full version.

It remains to specify soundness regarding  $\text{AP}$ ,  $\text{UP}$ , and credential update function  $f$  specifically used by SMART.

(i) **Redeemable.** Authentication updates a credential to include a new, unique session. Redeeming updates the credential to incorporate the finalized score of the session.

In more detail, every successful authentication  $\text{Auth}$  creates a unique session identifier (ID), which is added into  $\text{attr}'$  of  $\text{cred}'$  as updated by  $f$ . A ticket of the same ID with an unjudged status will be added into  $\text{st}'$  too, which allows later SP judgment via  $\text{StMgn}$ . User with  $\text{cred}$  can run  $\text{Redeem}$ , where  $\text{UP}(\text{cred}) = 1$ , to accumulate the judgment from  $\text{st}'$  into  $\text{cred}'$  via  $f$ . The uniqueness of tickets prevents users from “sharing” a session.

(ii) **Non-escaping.** The authentication policy AP disallows further authentication of credential cred, *i.e.*,  $AP(cred) = 0$ , until cred is free from any unredeemed but judged session as specified by st. Meanwhile, Redeem only removes finalized tickets from cred according to UP and incorporates its score into cred' via  $f$ , where the SP specifies the score in st' via StMgn.

(iii) **Transient Judging.** A ticket  $t$  can be temporarily judged with some score via StMgn until it is finalized. In more detail, when cred that originated  $t$  is used in Auth, its evaluation of AP will take into account the score of  $t$  as reflected by st on top of the current score of cred. Note that Redeem does not remove such non-finalized tickets.

## 4. Technical Overview of SMART

### 4.1. Data Structures for Credential & Judgments

SMART extends the single implicit queue design of FARB [35] to multiple queues while carefully ensuring that no user can escape punishment even when they are anonymous. In our “baseline” scheme below for illustration, we consider users specify the queue to redeem. Still, the identifier of the ticket being redeemed remains private. Table 1 lists some major notations of the whole paper.

**4.1.1. Our Baseline Design from FARB.** A main credential is  $(x, \nu, s, \text{head}, \text{tail}, \sigma)$ , with  $s$  as its current score and  $\sigma$  is an SP's signature signing everything else. To prevent reuse, each authentication/redemption reveals the nonce  $\nu$ . The user then chooses a fresh value for the SP to blindly issue an updated credential.

After each authentication, the credential will be equipped with a new ticket, meaning that the user also gets *enqueueing signatures*  $\{\text{enQ}_x^{(i)} = (x, t_i, i, \sigma_x^i)\}_{i=1}^{|\mathcal{Q}|}$ . The ticket number  $t_i$  comes from a global counter  $tc$  that increments by 1 for each authentication. The secret  $x$  links to enqueueing signatures in an implicit ticket queue  $\mathcal{Q} = (\{\text{enQ}_x^{(\text{head}+1)}, \dots, \text{enQ}_x^{(\text{tail})}\}, \text{head}, \text{tail})$  with  $k = (\text{tail} - \text{head})$  tuples. Due to the *ticket order invariant*, *i.e.*, the increasing order of  $t_i$ 's, the user only needs to prove the first ticket  $t_{\text{head}+1}$  is unjudged via proving  $jp < t_{\text{head}+1}$ , where  $jp$  is the *global judgment pointer* pointing to the last judged authentication. FARB assumes a continuous order of judgment, *i.e.*, all tickets smaller than  $jp$  are judged, and all bigger than  $jp$  are unjudged.

A judgment is  $(t, s_t, \sigma')$  with an SP signature  $\sigma'$  signing on a ticket  $t$  with score  $s_t$ . Redeeming a judgment on  $t_{\text{head}+1}$  with  $\text{enQ}_x^{(\text{head}+1)}$  performs dequeue on  $\mathcal{Q}$ , which updates  $\mathcal{Q}$  to  $(\{\text{enQ}_x^{(\text{head}+2)}, \dots, \text{enQ}_x^{(\text{tail})}\}, \text{head}+1, \text{tail})$ , and  $s$  in cred absorbs its score  $s_{t_{\text{head}+1}}$ . Authentication increments tail into  $\text{tail}' = \text{tail} + 1$ , and the SP blindly issues an enqueueing signature on  $(x, t_{\text{tail}'}, \text{tail}')$  to enqueue  $t'_{\text{tail}'}$ . The queue of signatures becomes  $\mathcal{Q}' = (\{\text{enQ}_x^{(\text{head}+1)}, \dots, \text{enQ}_x^{(\text{tail})}\} \cup \{\text{enQ}_x^{(\text{tail}')}\}, \text{head}, \text{tail}')$ .

Figure 1 illustrates the global-halting issue of FARB and its failure to support immediate revocation. Suppose each of the two users authenticated five times; session  $t = 2$  is hard to judge, blocking the advancement of the global pointer  $jp$ . Alice is then halted if the queue size limit is 5. Also, for the SP to immediately revoke the user who spawned  $t = 8$ , it needs to rush through  $t \in [2..7]$ .




Figure 1. FARB's global halting and delayed revocation: i) Suppose only  $jp = 1$  is judged and  $t = 2$  is “hard-to-judge”: Alice cannot authenticate if the queue size limit has been reached, even though she did not originate session  $t = 2$ . ii) Immediately revoking ticket 8 requires the SP to rush in finalizing judgments on  $[2..7]$  to advance  $jp$ .

**4.1.2. SMART's Three Types of Signatures.** We extensively use signatures on commitments to maintain the data structure in credentials. This part focuses on the data they sign, and so do many parts in Section 5. We sometimes use an abused notation that omits the signature for brevity.

- i) Main credential stores  $N$  pair of pointers, *i.e.*,  $cred = (x, \nu, s, \{\mathcal{Q}_j.\text{head}, \mathcal{Q}_j.\text{tail}\}_{j \in [0, N-1]})$ , *e.g.*,  $\sigma$  in Figure 4.
- ii) Enqueueing signatures on  $\{\text{enQ}_x^{(i,j)} = (x, t_i, i, j)\}$  associate ticket  $t_i$  to the  $i$ -th position in the  $j$ -th (implicit) queue, *e.g.*, the columns at  $\mathcal{Q}_0, \mathcal{Q}_1$  in Figures 4-6.
- iii) Judgment is a signature on the state (elaborated below) of a ticket possibly with a score, *e.g.*, list  $\mathcal{J}_0$  in Figures 5-6. Unlike the above two, the SP maintains this in public. We also call lists  $\mathcal{J}_j, \forall j \in [0, N-1]$ , by *score lists*.

### 4.1.3. Multiple Queues for Transient Judgment.

$$\mathcal{Q}_j = (\{\text{enQ}_x^{(\text{head}+1,j)}, \dots, \text{enQ}_x^{(\text{tail},j)}, \text{head}, \text{tail}\}), j \in [N]$$

are queues implicitly maintained by  $(\mathcal{Q}_j.\text{head}, \mathcal{Q}_j.\text{tail})$  in the credential.  $t \in \mathcal{Q}_j$  refers to ticket  $t$  that is encoded in some  $\text{enQ}_x^{(i,j)} \in \mathcal{Q}_j$ , with  $i \in [\mathcal{Q}_j.\text{head} + 1, \mathcal{Q}_j.\text{tail}]$ .

Each judgment is a tuple  $(t, j, b_{\text{next}}, s_{(t,j)})$  for ticket  $t$  residing at  $\mathcal{Q}_j$  with score  $s_{(t,j)}$  (or simply  $s_t$ ). It can be:

- i) not judged (the initial status, or judged  $j = 0$  times),
  - ii) skipped ( $b_{\text{next}} = 1$ , as an interim decision), or
  - iii) finalized ( $b_{\text{next}} = 0$ , to redeem out of a credential).
- A signature (by the SP) on the judgment tuple changes the state of ticket  $t$ , moving it from the  $j$ -th queue to the  $(j+1)$ -th queue (if  $b_{\text{next}} = 1$ ) of a credential or out of the credential (if  $b_{\text{next}} = 0$ ) when the user redeems it.

For system parameter  $N$  (the total number of queues), a session can be “skipped” at most  $(N-1)$  times, each corresponds to an interim decision or *transient judgment*. Correspondingly, SMART maintains  $N$  judgment pointers  $jp_j, \forall j \in [0, N-1]$ , which is simply the largest ticket (ID) in the  $j$ -th score list  $\mathcal{J}_j$ , not like the single  $jp$  in FARB.

**4.1.4. Intra-Queue and Inter-Queue Invariants.** Sequential judging *within each queue* is still needed to maintain the *intra-queue invariant*. A judgment on  $t$  can only be done after a (transient) judgment on each of the previous tickets in the *same* queue has been made.

For simplicity, we enforce an *inter-queue invariant*

$$t \in \mathcal{Q}_i \implies t > t' \forall t' \in \mathcal{Q}_{i+1} (\forall i \in [N]).$$

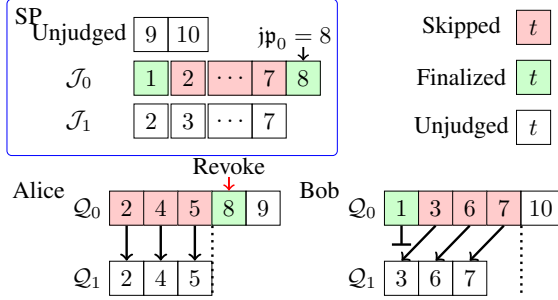


Figure 2. The SP decides to make a transient judgment for each  $t \in [2..7]$ , which moves them up from  $Q_j$  to  $Q_{j+1}$ , as denoted by the arrows (or  $\dashv$  for a finalized judgment). Tickets  $t \in [2..7]$  are then pending judgment on  $J_1$ . (The dotted lines depict  $Q_j.\text{head} \geq Q_{j+1}.\text{tail}$ .)

It means only tickets with smaller identifiers can go to the next (upper) queue. Since they will be larger than those in the upper queue, they will be appended to the end to maintain the intra-queue invariant.

## 4.2. Multiple Queues in Action

Authentication (enqueueing  $t$  to  $Q_0$ ) works like FARB but across multi-queues, which requires  $t_{Q_j.\text{head}+1} > \text{jp}_j$ ,  $\forall j \in [N]$ . All newly spawned tickets are enqueued to  $Q_0$ .

Redeeming in SMART depends on the judgment type: i) redeeming a finalized ticket at  $Q_j$ , which removes it from the  $j$ -th queue and hence the credential;

ii) redeeming a skipped ticket at  $Q_j$ , which dequeues it from  $Q_j$  and enqueues to  $Q_{j+1}$ , when  $j < N - 1$ .

We illustrate how the tickets are moved in credentials, corresponding to the judgments made by the SP as reflected by the multiple global score lists  $J_0, \dots, J_{N-1}$ . Figure 2 focuses on the credential of Alice and Bob, respectively holding  $(2, 4, 5, 8, 9)$  and  $(1, 3, 6, 7, 10)$  in  $Q_0$ .

First, note that when  $\text{jp}_0 = 1$ , Alice can authenticate as long as  $Q_0$  is not yet full since  $t_{Q_0.\text{head}+1} = 2 > \text{jp}_0$ .

Suppose the SP skips all tickets  $t \in [2..7]$ . This can be responding to any of the plausible scenarios below.

- i) Those sessions now receive their first transient judgment, *e.g.*, in the time-dependent model, one month has passed, and no negative report has been filed against them.
- ii) Those sessions are difficult to judge, but the SP notices that the global list  $J_0$  has been accumulating. To avoid halting all users like in FARB, the SP moves them to  $J_1$ .
- iii) Similar to the plausible scenario above, the SP further notices that session  $t = 8$  is “especially” bad and would like to revoke the user behind as soon as possible.

The SP makes the first transient judgments on  $t \in [2..7]$ , which moves the global judgment pointer of  $J_0$  from  $\text{jp}_0 = 1$  to  $\text{jp}_0 = 8$ . Figure 2 depicts the (projected) movements to  $Q_1$  of the affected tickets. Alice is now forced to redeem  $t \in [2..7]$ . This is to make  $t_{Q_0.\text{head}+1} > \text{jp}_0$  or Alice cannot authenticate anymore. (The conditions will be made more precise in Section 4.3.) Likewise, Bob redeems  $t \in \{1, 3, 6, 7\}$ , moving them from  $Q_0$  to  $Q_1$ , which makes  $t_{Q_0.\text{head}+1} = 10 > \text{jp}_0$ . Subjective evaluation can now have flexibility, either judging (the first one in)  $\{2, \dots, 7\} \in J_1$  or the unjudged ones in  $\{9, 10\}$ .




Figure 3. The set of skipped tickets  $\{2, 3, 5\}$  continues to be judged one more time, which moves those in user credentials to the next queue  $Q_2$ . If  $N = 3$ , their next judgment must be final for being in the last queue.





Figure 4. Alice authenticates when  $\text{tc} = 4$ . The new ticket  $t = 5$  is enqueued to  $Q_0$  at index 3. She obtains  $\text{cred}'$  with  $Q_0.\text{tail}' = 3$  and a new enqueueing signature (in green).

Figure 3 illustrates subsequent actions after some sessions have been “skipped” to the next global score list  $J_1$ , and the corresponding changes to Alice’s session queues. Suppose the SP further skips 2, 3, 5 and finalizes 4 in  $J_1$ . Alice’s last queue will store tickets 2, 5 after redemptions.

## 4.3. Illustrations Involving Proof of Signatures

We further illustrate the involvement of signatures in our cryptographic protocols in a simpler example. Suppose Alice has authenticated at  $\text{tc} = 2, 4$ , a global ticket counter maintained by the SP. Tickets  $\{1, 3\}$  do not belong to her.

### 4.3.1. Authentication. AP checks for each queue:

- i)  $Q_j$  is empty ( $Q_j.\text{head} = Q_j.\text{tail}$ ) or
- ii) the first ticket at  $Q_j$  is unjudged ( $\text{jp}_j < t_{Q_j.\text{head}+1}$ ) and  $Q_j$  is not full ( $Q_j.\text{tail} - Q_j.\text{head} < K_j$ ).

In Figure 4, without any judged tickets, Alice proves the credential and the enqueueing signatures  $\text{enQ}_x^{Q_j}$  for the first ticket (if it exists) in each queue  $j$  for authentication. The SP then increments  $\text{tc}$  (to 5) and binds a new ticket  $t' = \text{tc} = 5$  to the  $(Q_0.\text{tail} + 1)$ -th position (*i.e.*, 3) of  $Q_0$  of her credential via issuing an enqueueing signature on  $(x, 5, 3, 0)$ , where the commitment of  $Q_0.\text{tail}$  is sent by Alice during the proof of AP. Her  $\text{cred}$  is also updated.

**4.3.2. SP Judging.** SP sequentially judges unjudged sessions to appear in  $J_j$ , *e.g.*,  $t \in [1..5]$  for  $J_0$  in Figure 5. Suppose the SP wants to finalize ticket 1 with score  $s_1$  and revoke 3 with a very negative score  $s_3$  as a non-final judgment. Any unjudged tickets  $< 3$  need to be moved to the next queue or finalized. In this case, it is only ticket 2, with a transient score, say,  $s_2 = 0$ .

TABLE 1. NOTATIONS

Notation	Description
$N, K_i$	# of queues, the maximum size of the $i$ -th queue
$t$	Session/Ticket (identifier)
$x, \nu, s$	Credential secret, nonce, and score
$\text{enQ}_x^{(i,j)}$	Enqueuing signature of $t_i$ at $i$ in the $j$ -th queue $\mathcal{Q}_j$ ,
$(x, t_i, i, j, \sigma)$	( $\text{enQ}_x^{(i,j)}$ denotes $t$ is at the head ( $i = \mathcal{Q}_j.\text{head} + 1$ ))
$\mathcal{Q}_j$	( $\{\text{enQ}_x^{(\text{head}+1,j)}, \dots, \text{enQ}_x^{(\text{tail},j)}\}, \text{head}, \text{tail}$ )
$\mathcal{J}_j[t]$	A judged ticket ( $t, j, b_{\text{next}}, s_t, \sigma'$ ) in score list $\mathcal{J}_j$
$s_{(t,0)}^*$	Score for a finalized judgment on ticket ( $t, 0$ )
$\mathbb{T}_0^S/\mathbb{T}_j^S/\mathbb{T}^F$	Global $0$ -/ $j$ -time judged/finalized ticket list
$\mathcal{V}$	The set of used credential nonces
$\text{tc}$	A counter for the new session/ticket identifier
$\text{jp}_j$	Judgment pointer for $\mathcal{J}_j$

The SP gives judgment signatures on:

- i)  $(1, 0, b_{\text{next}} = 0, s_1)$ ,
- ii)  $(2, 0, b_{\text{next}} = 1, s_2)$ , and
- iii)  $(3, 0, b_{\text{next}} = 1, s_3)$ , respectively.

They are put into the public score list  $\mathcal{J}_0$ .

The SP also updates the global judgment pointer  $\text{jp}_0$  to 3, meaning that the third and its prior sessions (in  $\mathcal{J}_0$ ) have been judged and must be redeemed (from  $\mathcal{Q}_0$  of any credential). This is how SP ensures a user will redeem any judged sessions or the credential cannot pass AP.

**4.3.3. Redeeming.** Alice checks her credential against the score list to see if it has judged tickets ( $\text{jp}_j \geq t_{\mathcal{Q}_j.\text{head}+1}$ ). If so, she fetches those (by downloading everything) and runs the Redeem protocol on  $\mathcal{Q}_j$  with signatures below:

- i)  $\text{cred} = (x, \nu, s, \{\mathcal{Q}_j.\text{head}, \mathcal{Q}_j.\text{tail}\}_{j=0}^{N-1})$ ,
- ii)  $\text{enQ}_x^{\mathcal{Q}_j} = (x, t, \mathcal{Q}_j.\text{head} + 1, j)$ , and
- iii) a judgment  $(t, j, b_{\text{next}}, s_t)$ .

Upon completion,  $\text{cred}$  will be updated with:

- i) a fresh nonce  $\nu'$  and a new score  $(s + s_t)$ ,
- ii)  $\mathcal{Q}_j.\text{head}++$  (dequeue), and
- iii)  $\mathcal{Q}_{j+1}.\text{tail} := \mathcal{Q}_{j+1}.\text{tail} + b_{\text{next}}$  (enqueue if  $b_{\text{next}} = 1$ ).

An enqueuing signature  $(x, t, \mathcal{Q}_{j+1}.\text{tail} + b_{\text{next}}, j + b_{\text{next}})$  will be blindly issued<sup>3</sup> upon redemption on  $(t, j)$  too.

In Figure 5, Alice's  $\mathcal{Q}_0$  has a transient judgment  $(2, 0, b_{\text{next}} = 1, s_2)$ . She thus redeems ticket 2 in  $\mathcal{Q}_0$ , which dequeues 2 from  $\mathcal{Q}_0$  and enqueues it to  $\mathcal{Q}_1$ , resulting in  $\mathcal{Q}_0 = (4, 5)$ ;  $\mathcal{Q}_1 = (2)$ .  $\mathcal{Q}_0.\text{head}$  advances which voids the old  $\text{enQ}_x^{\mathcal{Q}_0} = (x, 2, 1, 0)$  since  $\mathcal{Q}_0.\text{head} + 1 > 1$ . For  $\mathcal{Q}_1$ , an enqueuing signature on  $(x, 2, 1, 1)$  is issued;  $\mathcal{Q}_1.\text{tail}$  is also incremented since  $b_{\text{next}} = 1$ .

**4.3.4. Finalizing a Transient Judgment.** The SP now needs to judge 4, 5, and 2 that will be posted on  $\mathcal{J}_0, \mathcal{J}_1$ . Suppose the SP finalizes ticket 2 (Figure 6), it posts:

- i) a final judgment  $(2, 1, b_{\text{next}} = 0, s_{(2,1)})$  on  $\mathcal{J}_1$  and
- ii) [optional] a final judgment<sup>4</sup>  $(2, 0, 0, s_{(2,0)}^*)$  on  $\mathcal{J}_0$ .

For example, if an initial score  $s_{(2,0)} = s_2 = 3$  has been given and a finalized score of  $s_{(2,0)}^* = 10$  is to be given; the SP finalizes it by setting  $s_{(2,1)} = 10 - 3 = 7$ . If Alice

3. This makes redeeming a final/transient judgment indistinguishable, except for the last  $\mathcal{J}_{N-1}$  having only final judgments.

4.  $s_{(t,j)}^*/s_{(t,j)}$  distinguish scores for finalized/skipped judgments in the same score list, e.g., ticket 2 in  $\mathcal{J}_0$  in Figure 6. It reduces the number of redeeming requests when Alice waits until  $(2, 0)$  is finalized.

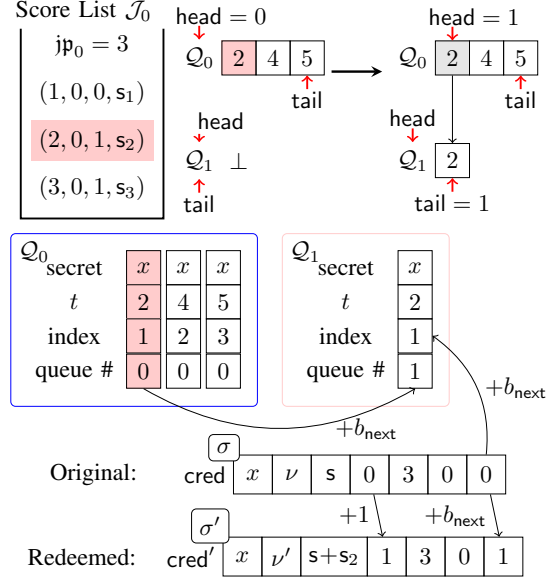


Figure 5. Alice redeems a judgment (red) on  $t = 2$  (a transient one with score  $s_2$ ) via PoK of i) judgment signature on  $(2, 0, b_{\text{next}} = 1, s_2)$  and ii) enqueuing signature  $(x, 2, 1, 0)$  (stored in the  $\mathcal{Q}_0$  box).  $\mathcal{Q}_0.\text{head}$  advances to remove ticket 2 at  $\mathcal{Q}_0$  and  $\mathcal{Q}_1.\text{tail}$  advances by  $b_{\text{next}} = 1$  for accommodating the moved ticket.  $s$  becomes  $s + s_2$ .

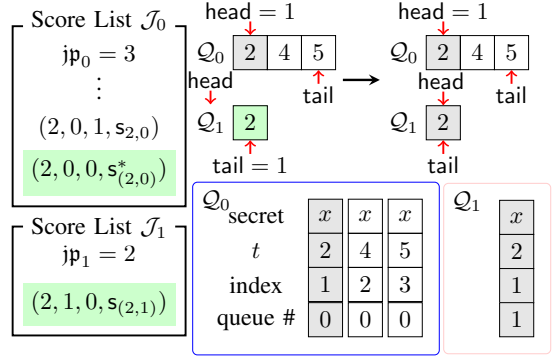


Figure 6. Alice redeems a finalized ticket 2 at  $\mathcal{Q}_1$  with  $(2, 1, 0, s_{(2,1)})$ .  $\mathcal{Q}_1.\text{tail}$  advances by  $b_{\text{next}} = 0$ , i.e., unchanged.  $\mathcal{Q}_1.\text{head}$  advances to  $\mathcal{Q}_1.\text{tail}$ , so  $\mathcal{Q}_1$  is empty.

absorbs  $s_{(2,0)}, s_{(2,1)}$  via two updates in Figures 5 and 6, they sum up to 10. The SP also updates  $\text{jp}_1 = 2$  to indicate all tickets  $t \in \mathcal{J}_1, t \leq 2$  are judged (ticket 1 is finalized at  $\mathcal{J}_0$  and thus will not appear in subsequent score lists).

**4.3.5. Redeeming a Finalized Judgment.** With new  $\text{jp}_1$ , Alice needs to redeem judgment  $(2, 1, 0, s_{(2,1)})$  with  $\text{enQ}_x^{\mathcal{Q}_1} = (x, 2, 1, 1)$  since her  $\mathcal{Q}_1$  contains a judged ticket 2. (Alternatively, if she waits until ticket 2 is finalized, cf. Figure 5, she uses  $(2, 0, 0, s_{(2,0)}^*)$  to redeem and no ticket will be enqueued to  $\mathcal{Q}_1$ .) After redeeming,  $\mathcal{Q}_0 = (4, 5)$ ;  $\mathcal{Q}_1 = (\perp)$ , with  $\mathcal{Q}_1.\text{head} = \mathcal{Q}_1.\text{tail} = 1$  in  $\text{cred}'$ . Thus, Alice can prove tickets in  $\mathcal{Q}_0$  are unjudged ( $t_{\mathcal{Q}_0.\text{head}+1} > \text{jp}_0$ ) and  $\mathcal{Q}_1$  is empty ( $\mathcal{Q}_1.\text{head} = \mathcal{Q}_1.\text{tail}$ ). Figure 6 illustrates the resulting changes.



## 5. SMART Credential Construction

### 5.1. Overview of Actions taken by Users/SP

**5.1.1. User's Actions.** A user registers with the SP to obtain a signature encoding a starting credential. It has all empty queues, so no enqueueing signature is issued.

To authenticate, the user first checks its credential with the score lists for judged tickets. If there are judged tickets, the user redeems them prior to actual authentication.

The user redeems with a valid credential, an enqueueing signature, and a judgment on  $(t, j)$ , where  $t$  is the first ticket at  $Q_j$ . We assume the user redeems one ticket at each time and  $j$  is revealed in the baseline construction.

When  $j < N$ , the redeemed ticket would be moved to the next queue, so the user should receive a (valid) new enqueueing signature on  $(x, t, Q_{j+1}.tail + 1, j + 1)$ .

**5.1.2. SP's Actions.** The SP setups the system parameters, *e.g.*, the number of queues  $N$  and their bounds  $\vec{K}$ . During authentication and redemption, the revealed nonce is stored in  $\mathcal{V}$ , and a new ticket (ID) generated in Auth is added to the public state  $st$ . The SP can issue transient or finalized judgments on tickets via  $StMgn$  to update  $st$  (which consists of score lists and pointers  $\{(\mathcal{J}_j, j\mathfrak{p}_j)\}_{j=0}^{N-1}$ ).

### 5.2. Setup

The SP calls  $Setup(1^\lambda)$  to run the following steps.

1. Determine the following system parameters appropriate for the application scenario:

- $N$ : the number of queues;
- $\vec{K} \in \mathbb{N}^N$  (or  $\{K_j\}_{j=0}^{N-1}$ ): the maximum number of sessions in each of the  $N$  queues in a credential;
- $s^{\min}$ : the minimum score for the authentication policy AP. (SMART could be easily extended to cover  $\ell$  score categories.)

2. Initialize the following sets to be an empty set  $\emptyset$ :

- $\mathcal{V}$ : the set of credential nonces that the SP has seen in user authentication and redeeming;
- $\{\mathcal{J}_i\}_{i \in [N]}$ : the set of judgments that have been skipped/(temporarily-)judged for  $i$  times.
- $\{\mathbb{T}_j^S\}_{j \in [N]}, \mathbb{T}^F$ : the queues of tickets that are judged  $j$  times, and the queue of finalized tickets.

3. Initialize the following values to be 0:

- $\{j\mathfrak{p}_i\}$ : global judgment pointers indicating the last judged tickets of the  $i$ -th score list  $\mathcal{J}_i$ ;
- $tc$ : a ticket counter indicating the last session ID assigned to an authenticated session.

4. Setup the cryptographic systems:

- $crs \leftarrow ZKP.Setup(1^\lambda)$ : the common reference string for the zero-knowledge proof/argument;
- $\{(pk_k, sk_k) \leftarrow Cred.KGen(1^\lambda)\}_{k \in \{c, t, j\}}$ : the key pair(s) of the credential signature scheme for signing tuples of credentials, ticks, and judgments.

5. Output:

- $pp := (crs, \{pk_k\}_{k \in \{c, t, j\}}, N, s^{\min})$ ;
- $st := (tc, \mathbb{T}_0^S, \{\mathbb{T}_j^S\}, \mathbb{T}^F, \mathcal{V}, \{K_i, \mathcal{J}_i, j\mathfrak{p}_i\})$ , where  $j \in [N]$  and  $i \in [N]$ .
- Return  $(pp, (sk_c, sk_t, sk_j), st)$ .

$Reg(\mathcal{U}(pp), \mathcal{S}(pp, sk))$

$\mathcal{U} : x, \nu \leftarrow \mathcal{S}, s := 0$

$\mathbf{attr} := (x, \nu, s, \{Q_i.head, Q_i.tail\}_{i \in [N]})$

$\mathcal{U} \leftrightarrow \mathcal{S} : \mathbf{run} \text{ Cred.Sig}(\mathcal{U}(\mathbf{attr}; \rho), \mathcal{S}(pp, sk_c)) \rightarrow (\sigma, \perp)$

*/ the user-chosen secret and nonce are hidden in a commitment*

$\mathcal{U} : \mathbf{return} (cred := (\mathbf{attr}, \sigma))$

$\mathcal{S} : \mathbf{return} st$

Figure 7. Registration Protocol

$AP(cred)$

**parse**  $pk_c, pk_t, pk_j, \{j\mathfrak{p}_i, K_i\}_{i \in [N]}$  **from**  $pp, st$

**parse**  $(x, \nu, s, \{Q_i.head, Q_i.tail\}_{i \in [N]}, \sigma)$  **from**  $cred$

$\mathbf{attr} := (x, \nu, s, \{Q_i.head, Q_i.tail\}_{i \in [N]})$

**parse**  $\{enQ_x^{Q_i} = (x, t_{Q_i.head+1}, Q_i.head + 1, i, \sigma_x^i)\}_{i \in [N]}$

**if**  $(Cred.Vf(pk_c, \mathbf{attr}, \sigma) = 0)$  **then return** 0

**if**  $(s < s^{\min})$  **then return** 0 */ minimum score*

**for**  $i$  **in**  $[N]$

**if**  $(Q_i.tail = Q_i.head)$  **then continue** */  $Q_i$  is empty*

**if**  $(t_{Q_i.head+1} \leq j\mathfrak{p}_i)$  **then return** 0 */ unjudged first ticket*

**if**  $(Q_i.tail - Q_i.head \geq K_i)$  **then return** 0

**if**  $(Cred.Vf(pk_t, enQ_x^{Q_i}) = 0)$  **then return** 0

**endfor**

**return** 1

Figure 8. Authentication Policy Checking

### 5.3. Registration

Figure 7 presents the registration protocol. The initial attributes of the user credential are set to  $\mathbf{attr} = (x, \nu, s, \{(Q_j.head, Q_j.tail)\}_{j \in [N]})$ , with initial score  $s$  and  $N$  empty queues  $(Q_j.head, Q_j.tail) = (0, 0)$ .

The user engages in the signature issuance protocol of Cred with the SP, proving the commitment of attributes (using some random  $\rho$ ) is well-formed, *i.e.*, everything in  $\mathbf{attr}$  other than  $x$  and  $\nu$  is 0. The secret  $x$  can be jointly sampled by the user and SP, or related to a long-term user public key, *e.g.*,  $g^x$ . The user then gets a credential  $cred$ .

### 5.4. Authentication

To authenticate, the user credential must pass the authentication policy AP in Figure 8 on the following:

- 1) Main credential  $cred$  is valid with a fresh nonce  $\nu$ .
- 2) The scores satisfy<sup>5</sup>  $s^{\min}$  announced by the SP.
- 3) Each queue  $i$  has no unredeemed tickets, *i.e.*,  $Q_i$  is empty (via a simple equality check), **or** the first ticket at  $Q_i$  is unjudged (via the enqueueing signature and a range proof against  $j\mathfrak{p}_i$ ).
- 4) Each queue  $i$  has not reached its full capacity  $K_i$ .

If it passes, the parties run the protocol in Figure 9, *i.e.*,  $Auth(\mathcal{U}(pp, cred, AP, f), \mathcal{S}(pp, sk, AP, f))$ , as follows.

5. We use  $s \geq s^{\min}$  to denote score satisfaction. The policy can be all categories' scores or weighted some of them being above the threshold.

$\text{Auth}(\mathcal{U}(\text{pp}, \text{cred}, \text{AP}, f), \mathcal{S}(\text{pp}, \text{sk}, \text{AP}, f))$
$\mathcal{U} : \text{parse cred} := (\text{attr}, \sigma, \{\text{enQ}_x^{\mathcal{Q}_i}\})$ $\text{parse attr} := (x, \nu, s, \{\mathcal{Q}_i.\text{head}, \mathcal{Q}_i.\text{tail}\}_{i \in [N]})$ <b>if</b> !AP(cred) <b>then return</b> $\perp$ $\pi \leftarrow \text{ZKP.P}(\text{pp}, (\text{AP}, \text{st}), \text{cred})$ $\text{attr}' \leftarrow f(\text{cred})$
$\mathcal{U} \rightarrow \mathcal{S} : (\pi, \nu)$
$\mathcal{S} : \text{if}$ !(ZKP.V(pp, (AP, st), $\pi$ ) $\wedge$ ( $\nu \notin \mathcal{V}$ )) <b>return</b> $\perp$ $\mathcal{V} := \mathcal{V} \cup \{\nu\}, \text{tc} := \text{tc} + 1, \mathbb{T}_0^S := \mathbb{T}_0^S \cup \{\text{tc}\}$
$\mathcal{U} \leftrightarrow \mathcal{S} : \quad / \text{Run in parallel}$ $\text{run Cred.Sig}(\mathcal{U}(\text{attr}'; \rho), \mathcal{S}((\pi, f), \text{sk}_c)) \rightarrow (\sigma, \perp)$ $\text{run Cred.Sig}(\mathcal{U}((x, \text{tc}, \mathcal{Q}_0.\text{tail} + 1, 0); \rho'), \mathcal{S}((\pi, f), \text{sk}_t)) \rightarrow (\sigma_1, \perp)$
$\mathcal{S} : \text{return st}' \quad / \text{Updated } \mathcal{V}, \text{tc}, \mathbb{T}_0^S$
$\mathcal{U} : \text{enQ}'_x := (x, \text{tc}, \mathcal{Q}_0.\text{tail} + 1, 0, \sigma_1)$
$\mathcal{U} : \text{return cred}' := (\text{attr}', \sigma; \text{enQ}'_x)$

Figure 9. Authentication Protocol

$f(\text{cred}) \text{ for Auth}$
$\text{parse attr} := (x, \nu, s, \{\mathcal{Q}_i.\text{head}, \mathcal{Q}_i.\text{tail}\}_{i \in [N]})$ $\nu' \leftarrow \mathbb{Z}_p$ $\text{attr}' := (x, \nu', s, \mathcal{Q}_0.\text{head}, \mathcal{Q}_0.\text{tail} + 1, \dots, \mathcal{Q}_{N-1}.\text{tail})$ $\text{enQ}' := (x, \text{tc}, \mathcal{Q}_0.\text{tail} + 1, 0) \quad / \text{where tc is from st}$ $\text{return cred}' := (\text{attr}'; \text{enQ}')$

Figure 10. Credential Updating in Auth

1. The user generates ZKP  $\pi$  for satisfying AP:

$$\text{ZKP} \left\{ \begin{array}{l} (\text{cred}, \{\text{enQ}_x^{\mathcal{Q}_i}\}_{i=1}^N, \nu') : \\ (C_x = \text{Com}(x)) \wedge (C_s = \text{Com}(s)) \\ \wedge (C_{\nu'} = \text{Com}(\nu')) \\ \bigwedge_{i \in [N]} \left\{ \begin{array}{l} C_{\mathcal{Q}_i.\text{head}} = \text{Com}(\mathcal{Q}_i.\text{head}) \\ \wedge C_{\mathcal{Q}_i.\text{tail}} = \text{Com}(\mathcal{Q}_i.\text{tail}) \\ \wedge C_{t_{\mathcal{Q}_i.\text{head}}} = \text{Com}(t_{\mathcal{Q}_i.\text{head}}) \end{array} \right\} \\ \wedge (\text{Cred.Vf}(\text{pk}_c, \text{cred}) = 1) \wedge (s \geq s^{\min}) \\ \bigwedge_{i \in [N]} \left\{ \begin{array}{l} ((\mathcal{Q}_i.\text{head} = \mathcal{Q}_i.\text{tail}) \vee \\ ((\mathcal{Q}_i.\text{tail} - \mathcal{Q}_i.\text{head} \in [0, K_i - 1]) \\ \wedge (t_{\mathcal{Q}_i.\text{head}+1} > \text{jp}_i) \\ \wedge \text{Cred.Vf}(\text{pk}_t, \text{enQ}_x^{\mathcal{Q}_i}) = 1 \end{array} \right\} \end{array} \right\}.$$

The statement involves the opening of (individually) committed attributes in cred, the new nonce  $\nu'$ , the enqueueing signatures  $\text{enQ}_x^{\mathcal{Q}_i}$ , and the score passes the threshold  $s^{\min}$ . The implicit queue design and the intra-queue order invariant allow checking only the first ticket (knowledge of  $\text{enQ}_x^{\mathcal{Q}_i}$ ) in  $\mathcal{Q}_i$  with  $\text{jp}_i$  or  $\mathcal{Q}_i.\text{head} = \mathcal{Q}_i.\text{tail}$  (for an empty  $\mathcal{Q}_i$ ) via an OR-composition proof. The user also proves the queue size is within the public bound  $K_i$ .

- The user sends the nonce  $\nu$  and above ZKP  $\pi$  to the SP, which proves that its credential satisfies the policy.
- The SP confirms that  $\nu \notin \mathcal{V}$  and the validity of  $\pi$ . Then it updates  $\mathcal{V} := \mathcal{V} \cup \{\nu\}$ , the ticket counter  $\text{tc} := \text{tc} + 1$ , and the unjudged ticket set  $\mathbb{T}_0^S := \mathbb{T}_0^S \cup \{\text{tc}\}$ .
- The user prepares the update as  $f(\text{cred})$  in Figure 10:

$\text{Commitments for Signing } f(\text{cred}), C_1 \leftarrow \text{Com}(1)$
$C_{\text{attr}'} := (C_x, C_{\nu'}, C_s, C_{\mathcal{Q}_0.\text{head}}, C_{\mathcal{Q}_0.\text{tail}}, C_1, \dots, C_{\mathcal{Q}_{N-1}.\text{tail}})$ $C_{\text{enQ}'} := (C_x, C_{\text{tc}}, C_{\mathcal{Q}_0.\text{tail}}, C_1, C_0) \quad / C_0 \leftarrow \text{Com}(0)$

Figure 11. SP Computation of  $f(\text{cred})$  in Auth

- randomly picks a new nonce  $\nu' \leftarrow \mathbb{Z}_p$ ;
- increments  $\mathcal{Q}_0.\text{tail}$  in  $\text{attr}'$  for adding a slot for the new ticket at the end of the 0-th queue;
- prepares the enqueueing data for ticket ID  $\text{tc}$  at the  $i$ -th slot,  $i = \mathcal{Q}_0.\text{tail} + 1$ , of the 0-th queue.

5. The parties reuse commitments from ZKP  $\pi$  (namely,  $(C_x, C_{\nu'}, C_s, \{C_{\mathcal{Q}_i.\text{head}}, C_{\mathcal{Q}_i.\text{tail}}\}_{i \in [N]})$  for the old credential and  $C_{\nu'}$  for the new nonce) to invoke in parallel 2 credential-signature issuance instances over commitments  $C_{\text{attr}'}$  and  $C_{\text{enQ}'}$  for SP's signatures on  $\text{attr}'$  and  $\text{enQ}'$ .  $C_{\text{attr}'}$  and  $C_{\text{enQ}'}$  can be created as in Figure 11.

The user stores the two obtained signatures: the updated credential signature  $\sigma$  on new attributes  $\text{attr}'$  and the enqueueing signature for  $\text{enQ}'_x$ , which enqueues  $t = \text{tc}$ .

6. The SP updates state  $\text{st}$  to  $\text{st}'$  with new  $\text{tc}$ ,  $\mathbb{T}_0^S$ , and  $\mathcal{V}$ .

## 5.5. Session Judging and Queue Management

A successful authentication appends a new ticket to  $\mathbb{T}_0^S$ , the unjudged ticket set. In StMgn, the SP puts judgments into  $\{\mathcal{J}_j\}_{j \in [N]}$  for  $\mathbb{T}_0^S$ ,  $j$ -time judged sets  $\{\mathbb{T}_j^S\}_{j \in [1..N-1]}$ , and the finalized ticket list  $\mathbb{T}^F$ , all in st. (In Figure 2,  $\mathbb{T}_0^S = (9, 10)$ ,  $\mathbb{T}_1^S = \emptyset$ ,  $\mathbb{T}_2^S = (2, \dots, 7)$ .)

With the sequential judgment of each queue, the SP should judge tickets at the heads of  $\{\mathbb{T}_j^S\}$  (e.g., 2, 9 in Figure 2). For a transient (or final) judgment of a ticket  $t$ , the SP posts it to score list  $\mathcal{J}_j$ , moves  $t$  to  $\mathbb{T}_{j+1}^S$  (or  $\mathbb{T}^F$ ), and advances judgment pointer  $\text{jp}_j$  to  $t$ . Specifically,

- Transient judgment moves  $t$  from  $\mathbb{T}_j^S$  to  $\mathbb{T}_{j+1}^S$  and puts a signature on  $(t, j, b_{\text{next}} = 1, s_{(t,j)})$  to  $\mathcal{J}_j$ .
- Final judgment moves  $t$  from  $\mathbb{T}_j^S$  to  $\mathbb{T}^F$  and puts a signature on  $(t, 0, 0, s_{(t,0)})$  to  $\mathcal{J}_0$ , where  $s_{(t,j)}$  is computed using previous transient scores, e.g.,  $s_{(t,j)} = s_{(t,0)}^* - \sum_{i=0}^{j-1} s_{(t,i)}$ .

As an example, a ticket  $t$  that is skipped twice goes through  $\mathbb{T}_0^S$  (unjudged),  $\mathbb{T}_1^S$  and  $\mathbb{T}_2^S$  (2 transient judgments in  $\mathcal{J}_0, \mathcal{J}_1$ ). To finalize  $t$ ,  $(t, 2)$  is added to  $\mathbb{T}^F$  where the SP judges  $(t, 0), (t, 2)$  with  $b_{\text{next}} = 0$  and adjusted scores.

## 5.6. Session Redeeming or Credential Updating

When a valid credential has  $t \in \mathbb{T}_j^S$  in its  $\mathcal{Q}_j$ , the user needs to redeem judged tickets in every queue  $j$  before the next authentication. Redeeming a ticket updates the credential. Like enqueueing a new ticket by Auth, the credential (and other signatures) is updated by obtaining a signature on the committed multi-block message.

The user can redeem  $t \in \mathcal{Q}_j$  by revealing  $j$  or keeping  $j$  private. They have different update functions  $f$ .

**5.6.1. Redeeming a Ticket in a Known Queue.** Redeem in Figure 13 proceeds mostly like the steps in Auth.

- When  $j$  is public, the user checks if the credential satisfies the specific redeeming policy  $\text{UP}^j$  parameterized

by  $j$  in Figure 12, *i.e.*, eligible to redeem  $t_{Q_j.\text{head}+1} \in Q_j$ . If so, the user proves the ZKP below by proving the knowledge of signatures on each of the following:

- i)  $\text{cred} = ((x, \nu, s, \{Q_i.\text{head}, Q_i.\text{tail}\}_{i \in [N]}, \sigma), \sigma)$ ,
- ii)  $\text{enQ}_x^{Q_j} = ((x, t_{Q_j.\text{head}+1}, Q_j.\text{head} + 1, j), \sigma_x^j)$ , and
- iii)  $\mathcal{J}_j[t_{Q_j.\text{head}+1}] = ((t_{Q_j.\text{head}+1}, j, b_{\text{next}}, s_t), \sigma'')$  for:

$$\text{ZKP} \left\{ \begin{array}{l} (\text{cred}, \text{enQ}_x^{Q_j}, \mathcal{J}_j[t], \nu') : \\ C_x = \text{Com}(x) \wedge C_{\nu'} = \text{Com}(\nu') \wedge \\ C_j = \text{Com}(j) \wedge C_s = \text{Com}(s) \wedge \\ \left\{ \begin{array}{l} C_{Q_i.\text{head}} = \text{Com}(Q_i.\text{head}) \wedge \\ C_{Q_i.\text{tail}} = \text{Com}(Q_i.\text{tail}) \end{array} \right\}_{i \in [N]} \wedge \\ C_{s_t} = \text{Com}(s_t) \wedge C_{b_{\text{next}}} = \text{Com}(b_{\text{next}}) \wedge \\ \text{Cred.Vf}(\text{pk}_c, \text{cred}) = 1 \wedge \\ \text{Cred.Vf}(\text{pk}_t, \text{enQ}_x^{Q_j}) = 1 \wedge \\ \text{Cred.Vf}(\text{pk}_j, (t, j, b_{\text{next}}, s_t), \sigma'') = 1 \end{array} \right\}.$$

Secret  $x$  and  $Q_j.\text{head}$  in  $\text{cred}$  link to those in  $\text{enQ}_x^{Q_j}$ , which contains the ticket ID  $t = t_{Q_j.\text{head}+1}$  and queue index  $j$  that judgment  $\mathcal{J}_j[t]$  is for.

2. The user sends the proof  $\pi$  and current credential nonce  $\nu$  to the SP, whom confirms  $\nu \notin \mathcal{V}$  and  $\pi$  is valid. If so, the SP updates the nonce set  $\mathcal{V} \cup \{\nu\}$ .

3. The parties also invoke in parallel 2 credential-signature issuance instances with the individual homomorphic commitments in ZKP  $\pi$  as auxiliary inputs to compute updated  $\text{attr}' \leftarrow f(\text{cred})$  as in Figure 14.

- samples a new nonce  $\nu' \leftarrow \mathbb{Z}_p$ ;
- absorbs  $s_t$  by updating  $s$  as  $s + s_t$ ;
- increments  $Q_j.\text{head}$  which dequeues  $Q_j$ ;
- adds  $b_{\text{next}}$  to  $Q_{j+1}.\text{tail}$  in  $\text{attr}'$  (*i.e.*, adding a slot if  $b_{\text{next}} = 1$ );
- prepares the data to enqueue  $t_{Q_j.\text{head}+1}$  to the  $(i = Q_{j+1}.\text{tail} + b_{\text{next}})$ -th slot of the  $(j + b_{\text{next}})$ -th queue.

All the involved terms, including  $Q_{j+1}.\text{tail}$ , are from the credential ( $b_{\text{next}}, s_t$  are from the judgment signature) and have been proven by ZKP already. The SP computes the commitment of  $\text{attr}'$  as in Figure 15 (which also lists the individual commitments from  $\pi$ ) for blind signature issuance. Here, the SP can pick the  $Q_j$  head and  $Q_{j+1}$  tail from  $C_{\text{attr}}$  and update accordingly by adding  $C_1 \leftarrow \text{Com}(1)$  and  $C_{b_{\text{next}}}$  since  $j$  is revealed. The user receives:

- an updated credential  $\text{cred}'$  on  $\text{attr}' \leftarrow f(\text{cred})$ :  $(x, \nu', s + s_t, \dots, Q_j.\text{head} + 1, \dots, Q_{j+1}.\text{tail} + b_{\text{next}}, \dots, Q_{N-1}.\text{tail})$ ;
- an enqueueing signature on  $(x, t_{Q_j.\text{head}+1}, Q_{j+1}.\text{tail} + b_{\text{next}}, j + b_{\text{next}})$ .

A transient judgment has  $b_{\text{next}} = 1$ ; or 0 if final. A user redeeming a transient judgment obtains a new enqueueing signature on the redeemed ticket binding to  $Q_{j+1}.\text{tail} + 1$  and index  $j + 1$ . (In Figure 5, after redeeming a transient judgment on ticket  $t = 2$  at  $Q_0$ , Alice gets an enqueueing signature on  $(x, 2, 1, 1)$  for  $t = 2$  at  $(Q_1.\text{tail} + 1)$ .)

On the other hand, the enqueueing signature obtained for redeeming a finalized judgment will be “outdated” for any queues since it binds  $t$  to the current queue  $j$  with index  $Q_{j+1}.\text{tail} < Q_j.\text{head} + 1$  (the inter-queue invariant), *e.g.*, the (invalid) enqueueing signature Alice obtains in

```

UPj(cred)
-----
parse (pkc, pkt, pkj), (ipj, Jj) from pp, st
parse ({Qi.head, Qi.tail}i ∈ [N]) from cred
parse (x, tQj.head+1, Qj.head + 1, j, σxj) from cred
enQxQj := (x, tQj.head+1, Qj.head + 1, j, σxj)
parse Jj[tQj.head+1] from Jj / returns ⊥ if not exist
if (Cred.Vf(pkc, cred) = 0) then return 0
if (Cred.Vf(pkt, enQxQj) = 0) then return 0
if (Cred.Vf(pkj, Jj[tQj.head+1]) = 0) then return 0
return 1

```

Figure 12. Redeeming/Update Policy Checking

```

Redeem(U(pp, cred, UPj, f), S(pp, sk, UPj, f))
-----
U : parse cred := (attr, σ, {enQxQj})
   parse attr := (x, ν, s, {Qi.head, Qi.tail}i ∈ [N])
   parse Jj[tQj.head+1] := (t, j, bnext, st, σ'') from Jj
   if !UPj(cred) then return ⊥
   π ← ZKP.P(pp, (UPj, st), cred)
   attr' ← f(cred)
U → S : (π, ν)
S : if !(ZKP.V(pp, (UPj, st), π) ∧ (ν ∉ V)) then
   return ⊥; else V := V ∪ {ν}
U ↔ S : / Run in parallel, using (Com(Qj+1.tail), Com(bnext)) in π
   run Cred.Sig(U(attr'; ρ), S((π, f), skc)) → (σ, ⊥)
   run Cred.Sig(U((x, tQj.head+1, Qj+1.tail + bnext,
   j + bnext); ρ'), S((π, f), skt)) → (σ', ⊥)
S : return st' / Updated V
U : enQxQj := (x, tQj.head+1, Qj+1.tail + bnext, j + bnext, σ')
   return cred' := (attr', σ; enQxQj)

```

Figure 13. Redeeming Protocol

```

f(cred) for Redeem
-----
parse attr := (x, ν, s, {Qi.head, Qi.tail}i ∈ [N])
parse enQxQj = (x, tQj.head+1, Qj.head + 1, j)
parse Jj[tQj.head+1] := (tQj.head+1, j, bnext, st)
ν' ← sp
attr' := (x, ν', s + st, ..., Qj.head + 1,
   ..., Qj+1.tail + bnext, ..., QN-1.tail)
enQ' := (x, tQj.head+1, Qj+1.tail + bnext, j + bnext)
return cred' := (attr', enQ')

```

Figure 14. Credential Updating in Redeem

Figure 6 will be  $(x, 2, 0, 1)$ , where  $i < Q_1.\text{head} + 1 = 1$ , assuming  $Q_2$  is empty. Knowing  $j$  is the last queue, the signature issuance protocol needs not to be run.

4. The user stores the updated credential signature  $\sigma$  on the new attributes  $\text{attr}'$  and the enqueueing signature on  $t_{Q_j.\text{head}+1}$  (if  $b_{\text{next}} = 0$ , the user can discard the signature).
5. The SP updates its state  $\text{st}$  with the new nonce set  $\mathcal{V}$ .

**5.6.2. Hiding the Queue Index.** If  $j$  is also hidden in a commitment, the *user*, instead of the SP, will compute the

Commitments for Signing $f(\text{cred})$	
<b>parse from</b> $\pi : C_{\nu'}$ ,	
$(C_x, C_{\nu'}, C_s, \{C_{Q_i.\text{head}}, C_{Q_i.\text{tail}}\}_{i \in [N]}), / C_{\text{attr}}$	
$(C_x, C_{t_{Q_j.\text{head}+1}}, C_{Q_j.\text{head}+1}, C_j), / C_{\text{enQ}_j}$	
$(C_{t_{Q_j.\text{head}+1}}, C_j, C_{b_{\text{next}}}, C_{s_t}) / C_{\mathcal{J}_j[t_{Q_j.\text{head}+1}]}$	
$C_{\text{attr}'} := (C_x, C_{\nu'}, C_s \cdot C_{s_t}, C_{Q_0.\text{head}}, \dots, C_{Q_j.\text{head}} \cdot C_1,$	
$\dots, C_{Q_{j+1}.\text{tail}} \cdot C_{b_{\text{next}}}, \dots, C_{Q_{N-1}.\text{tail}})$	
$C_{\text{enQ}} := (C_x, C_{t_{Q_j.\text{head}+1}}, C_{Q_{j+1}.\text{tail}} \cdot C_{b_{\text{next}}}, C_j \cdot C_{b_{\text{next}}})$	

Figure 15. SP Computation of  $f(\text{cred})$  in Redeem

TABLE 2. POSSIBLE TICKET JUDGMENTS FOR THE  $j$ -TH QUEUE

Type	$b_{\text{next}}$	Next Queue Index $j'$
Transient	1	$(j+1) \neq (N-1)$
Final	0	$j$
Exigent	1	$N-1$

commitment(s) of  $\text{attr}'$  and  $\text{enQ}$ , and proves to the SP that they are updated according to update function  $f$ . The user prepares the update  $f(\text{cred})$  similar to Section 5.6.1 (with differences highlighted in blue):

- $\nu' \leftarrow \mathbb{Z}_p$ ;  $s_t := s + s_t$ ;
- increments  $Q_j.\text{head}$  to dequeue  $Q_j$ ;
- adds  $b_{\text{next}}$  to  $Q_{j+b_{\text{next}}}.\text{tail}$  in  $\text{attr}'$ ;
- prepares the data to enqueue  $t_{Q_j.\text{head}+1}$  to the  $i'$ -th slot of the  $(j+b_{\text{next}})$ -th queue, where  $i' = b_{\text{next}} \cdot (Q_{j+b_{\text{next}}}.\text{tail} + b_{\text{next}})$ .

Unlike when  $j$  is known, now the SP does not know which commitment to apply additive homomorphism to compute  $f(\text{cred})$ , and hence it relies on the user. Besides,  $j$  could be the last queue where  $Q_{j+1}$  is not well-defined. This explains the formula of  $i'$ . For a final judgment where  $b_{\text{next}} = 0$ ,  $i' = 0$ . “Enqueuing” to the 0-th slot of any slot means nothing, as it is an invalid slot for any credential.

The user proves that  $C_{\text{attr}'}$  updates  $Q_j.\text{head}$  and  $Q_{j+b_{\text{next}}}.\text{tail}$  in Figure 16 with two ZKPs on  $\mathcal{R}_{\text{sel}} = ((C_{\bar{m}}, C_{\bar{m}'}, C_{j'}, C_b), (\bar{m}, \bar{m}', j', b) : \bar{m}'[j'] = \bar{m}[j'] + b)$ , which proves the  $j'$ -th slot of both  $\bar{m}'$  and  $\bar{m}$  differ by  $b$ :

- $(Q_j.\text{head}) \mathcal{R}_{\text{sel}}$  holds for  $\bar{m} = \text{attr}$ ,  $\bar{m}' = \text{attr}'$ ,  $j' = 3 + 2j + 1$ ,  $b = 1$  (public);
- $(Q_{j+b_{\text{next}}}.\text{tail}) \mathcal{R}_{\text{sel}}$  holds for  $\bar{m} = \text{attr}$ ,  $\bar{m}' = \text{attr}'$ ,  $j' = 3 + 2(j+1+b_{\text{next}})$ ,  $b = b_{\text{next}}$ .

For  $C_{\text{enQ}}$ ,  $\mathcal{R}_{\text{sel}'} = ((C_{\bar{m}}, C_{j'}, C_{i'}, C_b), (\bar{m}, j', i', b) : i' = b \cdot (\bar{m}[j']))$  is used to prove  $C_{i'}$  (in  $C_{\text{enQ}}$ ) commits the product of  $b = b_{\text{next}}$  and the  $j'$ -th element of  $\bar{m} = \text{attr}'$ .

If all three ZKPs above are valid, the SP runs  $\text{Cred.Sig}()$  on the commitments  $(C_{\text{attr}'}, C_{\text{enQ}})$ , and the user will obtain an updated credential over  $(\text{attr}', \text{enQ})$ .

## 5.7. Emergency Revocation

In the baseline design, tickets in  $Q_j$  can only be moved to  $Q_{j+1}$  (or out of the credential). When a judgment can easily be deemed to be final, slowly advancing one queue at a time is not needed. Allow moving a ticket from  $Q_j$  directly to the finalized queue  $Q_{N-1}$  also enables emergency revocation. *i.e.*, the SP wants to revoke the user

Proof of $\text{attr}' \leftarrow f(\text{cred})$ when $j$ is hidden	
<b>parse</b> $C_{\text{attr}'}, C_{\nu'}$ <b>from</b> $\pi$	
<b>parse</b> $(C_x, C_{t_{Q_j.\text{head}+1}}, C_{Q_j.\text{head}+1}, C_j)$ <b>from</b> $\pi$ / $\text{enQ}$	
<b>parse</b> $(C_{t_{Q_j.\text{head}+1}}, C_j, C_{b_{\text{next}}}, C_{s_t})$ <b>from</b> $\pi$ / Judgment	
$C_{j'} := C_{3+2j+1}$ / computes homomorphically	
$C_{j''} := C_{3+2(j+1+b_{\text{next}})}$ / index of $Q_{j+b_{\text{next}}}$ 's tail	
$\mathcal{U} : C_{\text{attr}'} := (C_x, C_{\nu'}, C_s \cdot C_{s_t}, C_{Q_0.\text{head}}, \dots,$	
$C_{Q_j.\text{head}} \cdot C_1, \dots, C_{Q_{j+b_{\text{next}}}.\text{tail}} \cdot C_{b_{\text{next}}}, \dots, C_{Q_{N-1}.\text{tail}})$	
$C_{\text{enQ}} := (C_x, C_{t_{Q_j.\text{head}+1}}, C_{i'}, C_j \cdot C_{b_{\text{next}}})$	
$\pi_1 \leftarrow \text{ZKP.P}(\text{pp}, \mathcal{R}_{\text{sel}}, (\text{attr}, \text{attr}', j', 1))$	
$\pi_2 \leftarrow \text{ZKP.P}(\text{pp}, \mathcal{R}_{\text{sel}}, (\text{attr}, \text{attr}', j'', b_{\text{next}}))$	
$\pi_3 \leftarrow \text{ZKP.P}(\text{pp}, \mathcal{R}_{\text{sel}'}, (\text{attr}', j'', i', b_{\text{next}}))$	
$\mathcal{U} \rightarrow \mathcal{S} : (C_{\text{attr}'}, C_{\text{enQ}}, \pi_1, \pi_2, \pi_3)$	

Figure 16. Proving  $f(\text{cred})$  in Redeem for Hidden  $j$

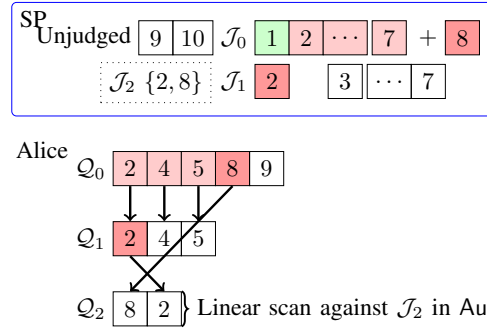


Figure 17. Let the emergency queue be  $Q_2$ . Dark-red boxes (8 in  $Q_0$ , 2 in  $Q_1$ ) are exigently-judged tickets. Redeeming them moves them directly to  $Q_2$ . Authentication will prove each ticket in  $Q_2$  w.r.t.  $\mathcal{J}_2$ .

concerned immediately, without any delay when the concerned ticket is not in the first slot of each queue. We thus designate the last queue as the “emergency” queue. The tickets there can only be redeemed and incorporated into the credential score, instead of moving to other queues.

To realize this technically, we start by enabling an SP to move a ticket to the last queue from previous queues (via exigent judgments); then, we describe the linear scan in authentication for proving each ticket in this queue.

### 5.7.1. Augmented Judgment and Queue Advancement.

During redemption,  $j$ ,  $b_{\text{next}}$  in the judgment  $(t, j, b_{\text{next}}, s_t)$  determines the destination queue  $(j+b_{\text{next}})$ . We extend the judgment from  $(t, j, b_{\text{next}}, s_t)$  to  $(t, j, b_{\text{next}}, j', s_t)$ , which adds the next queue index  $j'$  and is set according to Table 2. The update function  $f$  is the same as  $f$  in Section 5.6.2, except the new enqueueing signature’s queue index  $(j+b_{\text{next}})$  is replaced by  $j'$  (highlighted in blue):

- prepares the data to enqueue  $t_{Q_j.\text{head}+1}$  to the  $(i = Q_{j'}.\text{tail} + b_{\text{next}})$ -th slot of the  $j'$ -th queue.

Indices-hiding can still be achieved as in Section 5.6.2.

### 5.7.2. Emergency “Queue” with a Linear Scan.

Tickets in the emergency queue are likely out of order. For example (*cf.*, Figure 17, where  $N = 3$ ), the SP decides to move ticket 8 and ticket 2 at  $Q_0$  and  $Q_1$  to  $Q_2$ . Thus,

it posts two exigent judgments  $(8, 0, b_{\text{next}} = 1, j' = 2, s_8)$ ,  $(2, 1, 1, 2, s_2)$  on  $\mathcal{J}_0, \mathcal{J}_1$ . Alice must redeem them to authenticate again. Whether she redeems from  $\mathcal{Q}_0$  or  $\mathcal{Q}_1$  first is not enforced, so the intra-queue invariant does not hold for  $\mathcal{Q}_2$ . Nonetheless, it is no longer needed when authentication checks each ticket in  $\mathcal{Q}_2$  against the score list  $\mathcal{J}_2$  instead of a single range proof against  $\text{jp}_2$ .

To authenticate, the user proves for the last queue  $\mathcal{Q}_2$  in AP that  $|\mathcal{Q}_2| \leq K_2$  where  $|\mathcal{Q}_2| = \mathcal{Q}_2.\text{tail}$ , and each ticket  $t_i$  in  $\mathcal{Q}_2$  satisfies some predicate  $P(\mathcal{J}_2, t_i)$ . If  $|\mathcal{Q}_2|$  shall be revealed, the user proves the knowledge of  $|\mathcal{Q}_2|$  enqueueing signatures and  $P(\mathcal{J}_2, t_i) = 1 \forall t_i \in \mathcal{Q}_2$ .

The range  $0 \leq |\mathcal{Q}_2| \leq K_2$ , is where  $|\mathcal{Q}_2|$  can be hidden within. For each  $i$  in  $[1, K_2]$ , we can use OR-composition proof to prove either the possession of (message-)signature (pair)  $\text{enQ}_x^{(i,2)}$  for  $i \leq |\mathcal{Q}_2|$ , or  $\mathcal{Q}_2.\text{tail} < i$  (exceeding the user's  $\mathcal{Q}_2$ ) for  $|\mathcal{Q}_2| < i \leq K_2$ .

Specifically, for known  $(\mathcal{J}_2, \text{pk}_t, C_{\mathcal{Q}_2.\text{head}}, C_{\mathcal{Q}_2.\text{tail}})$ , we prove  $\mathcal{R}_{\text{scan}} = ((\{\text{enQ}_x^{(i,2)}\}_{i=1}^{|\mathcal{Q}_2|}, \mathcal{Q}_2.\text{head}, \mathcal{Q}_2.\text{tail}) : \bigwedge_{i=1}^{K_2} [(\text{Cred.Vf}(\text{pk}_t, \text{enQ}_x^{(i,2)}) \wedge P(\mathcal{J}_2, t_i) = 1) \vee (\mathcal{Q}_2.\text{tail} < i)] \wedge (\mathcal{Q}_2.\text{tail} \leq K_2))$  holds.

Note that the emergency queue size would be small for its special purpose.

### 5.7.3. Blocklist/Reputation Check.

- $P(\mathcal{J}_2, t_i)$  can be:
- (If  $\mathcal{J}_2$  stores blocked tickets :)  $t_i$  is a non-member of the revoked ticket set  $\mathcal{J}_2$ , e.g., via an accumulator [1], i.e., one is revoked for holding any ticket in this queue; or
  - There is an up to date judgment on  $t_i$  in  $\mathcal{J}_2 = \{(t, s_t, \tau)\}$  where  $\tau$  is the current timestamp.

The latter reputation-based approach can still incorporate the transient/finalized judgment using timestamp  $\tau$ . Each exigent judgment can first be judged with a (very negative) initial score  $s_{(t, \tau)}$ , and can be updated to  $s_{(t, \tau^*)}$  afterward; until a special  $\tau^*$  is marked, denoting the finalized judgment. The user can redeem finalized ones out of the emergency queue with a more complex operation, such as using verifiable shuffle [20]. This also requires a loosely-synchronized clock for the timestamps across all judgments or a signature expiry mechanism [20]. One can see this as adopting the idea in our concurrent work [20] for a specific queue instead of the entire credential.

## 6. Instantiation and Performance

### 6.1. Related Zero-Knowledge Proofs

An authenticating user needs to compute ZKP that each ticket queue  $\mathcal{Q}_j$  is either empty (equality check over commitments) or  $t_{\mathcal{Q}_j.\text{head}+1}$  is unjudged. Given a  $\Sigma$ -protocol for the equality check, signature verification, and range proof over committed values, the OR-proof technique [22] could be applied directly. We use the following ZKP for different relations:

- ZKP<sub>equal</sub> $\{(a, b, r, s) : C_a = \text{Com}(a; r) \wedge C_b = \text{Com}(b; s) \wedge a = b\}$
- ZKP<sub>score</sub> $\{(a, r) : C_a = \text{Com}(a; r) \wedge A \leq a < B\}$
- ZKP<sub>diff</sub> $\{(a, b, r, s) : C_a = \text{Com}(a; r) \wedge C_b = \text{Com}(b; s) \wedge a - b \in [0, K - 1]\}$
- ZKP<sub>inner</sub> $\{(\vec{a}, \vec{b}, c, r, s, t) : C_{\vec{a}} = \text{Com}(\vec{a}; r) \wedge C_{\vec{b}} = \text{Com}(\vec{b}; s) \wedge C_c = \text{Com}(c; t) \wedge \langle \vec{a}, \vec{b} \rangle = c\}$

- ZKP for knowing BBS+ signature over individual commitments of  $\vec{m}$  [7], [13].

The range proof can be instantiated using the signature-based set membership proof [12] or Bulletproofs [11], [25] which directly works with Pedersen commitment. One can also utilize the compressed- $\Sigma$  protocol theory [4]–[6], which applies Bulletproofs to  $\Sigma$ -protocols for opening linear forms on a committed vector.

A range proof of  $v \in [0, 2^n - 1]$  (supported natively by Bulletproofs) can be converted to an arbitrary range  $[A, B]$  using two range proofs [12] on subtracted values to  $A, B$ . Assume  $2^{n-1} < B < 2^n$ , we can prove that  $v - B + 2^n \in [0, 2^n - 1] \wedge v - A \in [0, 2^n - 1]$ . Thus, we set a soft upper bound on the maximum value of the tickets and score to be  $2^{64} - 1$ .

For our combined use with BBS+ signature, we work on a pairing-friendly curve for the pairing operations, e.g., BLS12-381. The proof of message equality in two commitments is straightforward. A redeeming user needs to prove the satisfaction of redeeming policy by ZKP of three valid signatures. The details on ZKP of the BBS+ signature can be found in Appendix B.

### 6.2. Using CP-SNARK

We also utilize the commit-and-prove paradigm. Most of the efficiently-implemented zkSNARKs focus on proving general circuits or rank-1 constraint systems. ZKP of  $\text{AP}(\text{cred}) = 1$  on committed attributes of  $\text{enQ}_x^{\mathcal{Q}_i}$  and  $\text{cred}$ , and public parameters  $\text{jp}_i$  and  $K_i$ , can be separated into

$$\bigwedge_{i \in [N]} \left\{ \left( \begin{array}{l} (\mathcal{Q}_i.\text{head} = \mathcal{Q}_i.\text{tail}) \vee \\ ((\mathcal{Q}_i.\text{tail} - \mathcal{Q}_i.\text{head} \in [0, K_i - 1]) \wedge \\ (t_{\mathcal{Q}_i.\text{head}+1} > \text{jp}_i)) \end{array} \right) \right\} \text{ and } \bigwedge_{i \in [N]} ((\mathcal{Q}_i.\text{head} = \mathcal{Q}_i.\text{tail}) \vee \text{Cred.Vf}(\text{pk}, \text{enQ}_x^{\mathcal{Q}_i}) = 1),$$

where the first and second nested relations can be proved by CP-SNARK (e.g., LegoGro16 [16] on circuits written with `circCom` library<sup>6</sup>) and proof of partial knowledge respectively. The committed values in the two relations are linked with the commit-and-prove paradigm.

**Instantiating ZKP for  $\mathcal{R}_{\text{sel}}$  and  $\mathcal{R}_{\text{sel}'}$ .** In Section 5.6.2, ZKP on  $\mathcal{R}_{\text{sel}}$  and  $\mathcal{R}_{\text{sel}'}$  are used:  $\mathcal{R}_{\text{sel}} = ((C_{\vec{m}}, C_{\vec{m}'}, C_{j'}, C_b), (\vec{m}, \vec{m}', j', b) : \vec{m}'[j'] = \vec{m}[j'] + b)$  is used to prove correct update of queue pointers in  $f(\text{cred})$ . Using zkSNARK,  $\vec{m} = \text{attr}$ ,  $\vec{m}' = \text{attr}'$ , the committed index  $j'$  and bit  $b = b_{\text{next}}$  are the inputs to a circuit that checks if  $\mathcal{R}_{\text{sel}}$  is satisfied, where the commitments are used for the proof of UP. One can utilize the one-hot encoding  $\vec{e}$  of  $j'$  (as an implicit input), where the circuit checks:

- $\langle \vec{e}, \vec{1} \rangle = 1 \wedge \langle \vec{e}, \vec{1} - \vec{e} \rangle = 0$  ( $\vec{e}$  is a weight-1 bit vector);
- $\langle \vec{e}, (0, \dots, 3 + 2N) \rangle = j'$ ;
- $\langle \vec{e}, \vec{m}' - \vec{m} \rangle = b$ .

ZKP on  $\mathcal{R}_{\text{sel}'}$  =  $((C_{\vec{m}}, C_{j'}, C_{i'}, C_b), (\vec{m}, j', i', b) : i' = b \cdot (\vec{m}[j']))$  is used to prove the index  $i'$  in the enqueueing data, where  $\vec{m} = \text{attr}'$ . Utilizing the one-hot encoding  $\vec{e}$ , ZKP on  $\mathcal{R}_{\text{sel}'}$  can be instantiated by proving the first and second conditions with two more conditions:  $\langle \vec{e}, \vec{m} \rangle = \vec{m}[j']$  and  $b \cdot (\vec{m}[j']) = i'$ . The above relations are readily supported by zkSNARKs.

6. <https://github.com/iden3/circom>

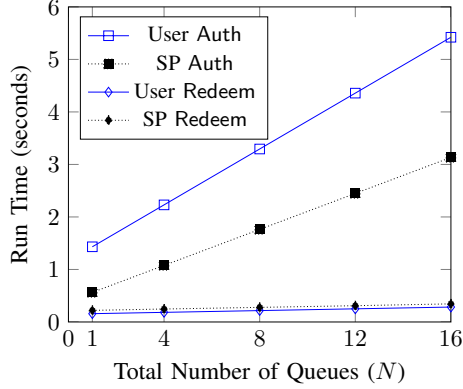


Figure 18. Computation times for Auth, Redeem (using Bulletproofs)

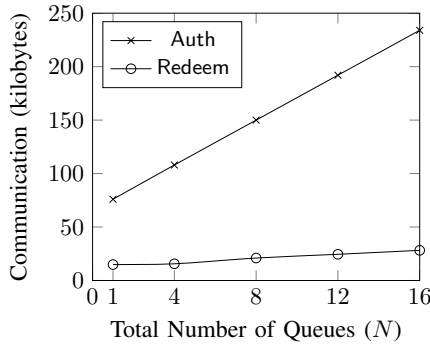


Figure 19. Total communication costs for Auth, Redeem

### 6.3. Experimental Evaluation

We use the Rust library from docknetwork<sup>7</sup> which is based on the arkwork library [3]. We also use the modified version of the Bulletproofs<sup>8</sup> for the pairing-friendly curve BLS12-381. We test on a desktop PC (localhost setting) with an 8-core AMD Ryzen CPU running at 3.7GHz. We use Windows Linux Subsystem to emulate Ubuntu 20.04.

We benchmark our authentication and redeeming protocols against numbers of queues  $N = \{1, 4, 8, 12, 16\}$  with 8 scoring categories ( $\ell = 8$ ). We compare with our re-implementation of FARB [35], the state-of-the-art sequential-judging system. The runtimes in Figure 18 are taken as the average timing of multiple runs. The bit-length of the score, tickets, and indexes is 64. Note that the improved ZKP for signature verification [13] is used. We also plot the total communication costs of authentication and redeeming in kilobytes versus the number of queues.

Compared to FARB ( $N = 1$ ) with no slackness, SMART’s user authentication time increases by about  $1.35\times$  when the SP can slack for  $N - 1 = 7$  times. We deem such kind of overhead acceptable. For example, the time for  $N = 16$  is still in the order of 5 seconds. More fundamentally, this can be sped up by using better computing platforms alone; while the subjective judging complexity and hence its associated global-halting issue

7. <https://github.com/docknetwork/crypto>

8. <https://github.com/dalek-cryptography/bulletproofs>

are inherent no matter how fast the future computing platform will be (unless the subjective assessment can be fully automated, say, by fancy AI technologies).

Our protocol scales linearly in the number of queues in both computation and communication costs (Figure 19). The increased communication cost in redeeming mainly comes from the increased credential proof size. The bottleneck of the authentication is the computation (and verification) of the  $N$  disjunctive proof.

Our authentication and redeeming costs are independent of the number of sessions in each queue and the global score lists. Efficiency-wise, one could roughly consider our  $N$  takes the role of prevalent parameter  $K$  in prior sequential-judging designs. Nonetheless, as explained in the introduction, the implications brought by tuning these two parameters are different. SMART gains flexibility and resilience against global halting.

## 7. Concluding Remarks

Securing the online world and protecting user privacy are becoming increasingly important, asking for anonymous credential systems free from trusted third parties. Meanwhile, incorporating reputation mechanisms based on subjective evaluation can add an additional layer of trust and security to these anonymous credential systems.

We unveil the unspoken global-halting vulnerability of the prevailing first-in-first-out sequential judging (PEREA, PERM, BLACR-express, and FARB [35]). Adversaries can launch denial-of-service attacks by introducing many hard-to-judge sessions. As long as one session remains unjudged, the entire system becomes paralyzed, affecting honest users as well. They also fail to support transient judgments. Another design (BLAC(R), SNARKBlock) processes all judgments ever made. Precomputations are only reusable when the judgments never change.

We propose SMART, with an extended multi-queue data structure in the credential. Besides alleviating the global-halting issue, SMART allows the SP to give adjustable judgment to authenticated sessions, avoiding the dilemma of being too lenient or overly cautious. SMART needs to carefully maintain the user sessions across multiple queues, in particular, hiding the source queue and the transient/finalized nature while redeeming, which are non-issues in single-queue designs like FARB.

We end with introducing two concurrent works. SAC [20] supports transient judgments without sequential judging and is also free from the global-halting issue. The credential uses an unordered data structure of a set to store the ticket. Redeeming (and hence removing) finalized tickets requires (signature-based) membership proof against the finalized ticket set. Verifiable shuffle is employed to hide the indices of any added or removed tickets. Unlike SMART, SAC requires signature expiration to invalidate previous judgments, *i.e.*, all active tickets need to be re-signed whenever the SP upgrades a score.

In a regular system, judgments on the tickets cannot be casually deleted since the privacy guarantee makes it unclear whether they have been redeemed or not. Decentralized anonymous reputation with sustainability [18] is recently formalized. A subset of surely-redeemed tickets can be identified and hence deleted from time to time.

## References

- [1] Tolga Acar, Sherman S. M. Chow, and Lan Nguyen. Accumulators and U-Prove revocation. In Ahmad-Reza Sadeghi, editor, *Financial Cryptography and Data Security - 17th International Conference, FC 2013, Okinawa, Japan, April 1-5, 2013, Revised Selected Papers*, volume 7859 of *Lecture Notes in Computer Science*, pages 189–196. Springer, 2013.
- [2] antonela. The value of Tor and anonymous contributions to wikipedia. <https://blog.torproject.org/the-value-of-anonymous-contributions-wikipedia>.
- [3] arkworks contributors. arkworks zkSNARK ecosystem, 2022. <https://arkworks.rs>.
- [4] Thomas Attema and Ronald Cramer. Compressed  $\Sigma$ -protocol theory and practical application to plug & play secure algorithms. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part III*, volume 12172 of *Lecture Notes in Computer Science*, pages 513–543. Springer, 2020.
- [5] Thomas Attema, Ronald Cramer, and Serge Fehr. Compressing proofs of k-out-of-n partial knowledge. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology - CRYPTO 2021 - 41st Annual International Cryptology Conference, CRYPTO 2021, Virtual Event, August 16-20, 2021, Proceedings, Part IV*, volume 12828 of *Lecture Notes in Computer Science*, pages 65–91. Springer, 2021.
- [6] Thomas Attema, Ronald Cramer, and Matthieu Rabaud. Compressed  $\Sigma$ -protocols for bilinear group arithmetic circuits and application to logarithmic transparent threshold signatures. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology - ASIACRYPT 2021 - 27th International Conference on the Theory and Application of Cryptology and Information Security, Singapore, December 6-10, 2021, Proceedings, Part IV*, volume 13093 of *Lecture Notes in Computer Science*, pages 526–556. Springer, 2021.
- [7] Man Ho Au and Apu Kapadia. PERM: Practical reputation-based blacklisting without TTPs. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *the ACM Conference on Computer and Communications Security, CCS'12, Raleigh, NC, USA, October 16-18, 2012*, pages 929–940, 2012.
- [8] Man Ho Au, Apu Kapadia, and Willy Susilo. BLACR: TTP-free blacklistable anonymous credentials with reputation. In *19th Annual Network and Distributed System Security Symposium, NDSS 2012, San Diego, California, USA, February 5-8, 2012*. The Internet Society, 2012.
- [9] Man Ho Au, Willy Susilo, Yi Mu, and Sherman S. M. Chow. Constant-size dynamic k-times anonymous authentication. *IEEE Systems Journal*, 7(2):249–261, 2013.
- [10] Jan Bobolz, Fabian Eidens, Stephan Krenn, Sebastian Ramacher, and Kai Samelin. Issuer-hiding attribute-based credentials. In Mauro Conti, Marc Stevens, and Stephan Krenn, editors, *Cryptography and Network Security - 20th International Conference, CANS 2021, Vienna, Austria, December 13-15, 2021, Proceedings*, volume 13099 of *Lecture Notes in Computer Science*, pages 158–178. Springer, 2021.
- [11] Benedikt Bünz, Jonathan Bootle, Dan Boneh, Andrew Poelstra, Pieter Wuille, and Gregory Maxwell. Bulletproofs: Short proofs for confidential transactions and more. In *2018 IEEE Symposium on Security and Privacy, SP 2018, Proceedings, 21-23 May 2018, San Francisco, California, USA*, pages 315–334. IEEE Computer Society, 2018.
- [12] Jan Camenisch, Rafik Chaabouni, and abhi shelat. Efficient protocols for set membership and range proofs. In Josef Pieprzyk, editor, *Advances in Cryptology - ASIACRYPT 2008, 14th International Conference on the Theory and Application of Cryptology and Information Security, Melbourne, Australia, December 7-11, 2008, Proceedings*, pages 234–252, 2008.
- [13] Jan Camenisch, Manu Drijvers, and Anja Lehmann. Anonymous attestation using the strong Diffie Hellman assumption revisited. In Michael Franz and Panos Papadimitratos, editors, *Trust and Trustworthy Computing - 9th International Conference, TRUST 2016, Vienna, Austria, August 29-30, 2016, Proceedings*, volume 9824 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2016.
- [14] Jan Camenisch and Anna Lysyanskaya. Signature schemes and anonymous credentials from bilinear maps. In Matthew K. Franklin, editor, *Advances in Cryptology - CRYPTO 2004, 24th Annual International Cryptology Conference, Santa Barbara, California, USA, August 15-19, 2004, Proceedings*, volume 3152 of *Lecture Notes in Computer Science*, pages 56–72. Springer, 2004.
- [15] Jan Camenisch and Markus Stadler. Efficient group signature schemes for large groups (extended abstract). In Burton S. Kaliski Jr., editor, *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings*, volume 1294 of *Lecture Notes in Computer Science*, pages 410–424. Springer, 1997.
- [16] Matteo Campanelli, Dario Fiore, and Anaïs Querol. LegoSNARK: Modular design and composition of succinct zero-knowledge proofs. In Lorenzo Cavallaro, Johannes Kinder, XiaoFeng Wang, and Jonathan Katz, editors, *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, CCS 2019, London, UK, November 11-15, 2019*, pages 2075–2092. ACM, 2019.
- [17] Sherman S. M. Chow. Real traceable signatures. In Michael J. Jacobson Jr., Vincent Rijmen, and Reihaneh Safavi-Naini, editors, *Selected Areas in Cryptography, 16th Annual International Workshop, SAC 2009, Calgary, Alberta, Canada, August 13-14, 2009, Revised Selected Papers*, volume 5867 of *Lecture Notes in Computer Science*, pages 92–107. Springer, 2009.
- [18] Sherman S. M. Chow, Christoph Egger, Russell W. F. Lai, Viktoria Ronge, and Ivy K. Y. Woo. On sustainable ring-based anonymous systems. In *36th IEEE Computer Security Foundations Symposium, CSF 2023, Dubrovnik, Croatia, July 9 - 13, 2023*, 2023. To appear.
- [19] Sherman S. M. Chow, Joseph K. Liu, and Duncan S. Wong. Robust receipt-free election system with ballot secrecy and verifiability. In *Proceedings of the Network and Distributed System Security Symposium, NDSS 2008, San Diego, California, USA, 10th February - 13th February 2008*. The Internet Society, 2008.
- [20] Sherman S. M. Chow, Jack P. K. Ma, and Tsz Hon Yuen. Scored anonymous credentials. In Mehdi Tibouchi and Xiaofeng Wang, editors, *Applied Cryptography and Network Security - 21st International Conference, ACNS 2023, Kyoto, Japan, June 19-22, 2023, Proceedings, Part II*, volume 13906 of *Lecture Notes in Computer Science*. Springer, 2023. To appear.
- [21] Sherman S. M. Chow, Haibin Zhang, and Tao Zhang. Real hidden identity-based signatures. In Aggelos Kiayias, editor, *Financial Cryptography and Data Security - 21st International Conference, FC 2017, Sliema, Malta, April 3-7, 2017, Revised Selected Papers*, volume 10322 of *Lecture Notes in Computer Science*, pages 21–38. Springer, 2017.
- [22] Ronald Cramer, Ivan Damgård, and Berry Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Yvo Desmedt, editor, *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference, Santa Barbara, California, USA, August 21-25, 1994, Proceedings*, volume 839 of *Lecture Notes in Computer Science*, pages 174–187. Springer, 1994.
- [23] Roger Dingledine, Nick Mathewson, and Paul F. Syverson. Tor: The second-generation onion router. In Matt Blaze, editor, *Proceedings of the 13th USENIX Security Symposium, August 9-13, 2004, San Diego, CA, USA*, pages 303–320. USENIX, 2004.
- [24] Jack Doerner, Yashvanth Kondi, Eysa Lee, abhi shelat, and LaKyah Tyner. Threshold BBS+ signatures for distributed anonymous credential issuance. In *44th IEEE Symposium on Security and Privacy, SP 2023, San Francisco, CA, USA, 22-25 May 2023*, pages 2095–2111. IEEE Computer Society, 2023.
- [25] Chaya Ganesh, Claudio Orlandi, Mahak Panchoi, Akira Takahashi, and Daniel Tschudi. Fiat-Shamir Bulletproofs are non-malleable (in the algebraic group model). In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology - EUROCRYPT 2022 - 41st Annual International Conference on the Theory and Applications of Cryptographic Techniques, Trondheim, Norway, May 30 - June 3, 2022, Proceedings, Part II*, volume 13276 of *Lecture Notes in Computer Science*, pages 397–426. Springer, 2022.

- [26] Christina Garman, Matthew Green, and Ian Miers. Decentralized anonymous credentials. In *21st Annual Network and Distributed System Security Symposium, NDSS 2014, San Diego, California, USA, February 23-26, 2014*. The Internet Society, 2014.
- [27] Jens Groth. On the size of pairing-based non-interactive arguments. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology - EUROCRYPT 2016 - 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8-12, 2016, Proceedings, Part II*, volume 9666 of *Lecture Notes in Computer Science*, pages 305–326. Springer, 2016.
- [28] Aggelos Kiayias, Yiannis Tsiounis, and Moti Yung. Traceable signatures. In Christian Cachin and Jan Camenisch, editors, *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, volume 3027 of *Lecture Notes in Computer Science*, pages 571–589. Springer, 2004.
- [29] Torben P. Pedersen. Non-interactive and information-theoretic secure verifiable secret sharing. In Joan Feigenbaum, editor, *Advances in Cryptology - CRYPTO '91, 11th Annual International Cryptology Conference, Santa Barbara, California, USA, August 11-15, 1991, Proceedings*, volume 576 of *Lecture Notes in Computer Science*, pages 129–140. Springer, 1991.
- [30] Michael Rosenberg, Mary Maller, and Ian Miers. SNARKBlock: Federated anonymous blocklisting from hidden common input aggregate proofs. In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*, pages 948–965. IEEE, 2022.
- [31] Claus-Peter Schnorr. Efficient signature generation by smart cards. *J. Cryptol.*, 4(3):161–174, 1991.
- [32] Stefano Tessaro and Chenzhi Zhu. Revisiting BBS signatures. In Carmitt Hazay and Martijn Stam, editors, *Advances in Cryptology - EUROCRYPT 2023 - 42nd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Lyon, France, April 23-27, 2023, Proceedings, Part V*, volume 14008 of *Lecture Notes in Computer Science*, pages 691–721. Springer, 2023.
- [33] Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. Blacklistable anonymous credentials: Blocking misbehaving users without TTPs. In Peng Ning, Sabrina De Capitani di Vimercati, and Paul F. Syverson, editors, *Proceedings of the 2007 ACM Conference on Computer and Communications Security, CCS 2007, Alexandria, Virginia, USA, October 28-31, 2007*, pages 72–81. ACM, 2007.
- [34] Patrick P. Tsang, Man Ho Au, Apu Kapadia, and Sean W. Smith. PEREA: towards practical ttp-free revocation in anonymous authentication. In Peng Ning, Paul F. Syverson, and Somesh Jha, editors, *Proceedings of the 2008 ACM Conference on Computer and Communications Security, CCS 2008, Alexandria, Virginia, USA, October 27-31, 2008*, pages 333–344. ACM, 2008.
- [35] Li Xi and Dengguo Feng. FARB: Fast anonymous reputation-based blacklisting without TTPs. In Gail-Joon Ahn and Anupam Datta, editors, *Proceedings of the 13th Workshop on Privacy in the Electronic Society, WPES 2014, Scottsdale, AZ, USA, November 3, 2014*, pages 139–148. ACM, 2014.
- [36] Li Xi, Jianxiong Shao, Kang Yang, and Dengguo Feng. ARBRA: Anonymous reputation-based revocation with efficient authentication. In Sherman S. M. Chow, Jan Camenisch, Lucas Chi Kwong Hui, and Siu-Ming Yiu, editors, *Information Security - 17th International Conference, ISC 2014, Hong Kong, October 12-14, 2014, Proceedings*, volume 8783 of *Lecture Notes in Computer Science*, pages 33–53. Springer, 2014.
- [37] Kin Ying Yu, Tsz Hon Yuen, Sherman S. M. Chow, Siu-Ming Yiu, and Lucas Chi Kwong Hui. PE(AR)<sup>2</sup>: Privacy-enhanced anonymous authentication with reputation and revocation. In Sara Foresti, Moti Yung, and Fabio Martinelli, editors, *Computer Security - ESORICS 2012 - 17th European Symposium on Research in Computer Security, Pisa, Italy, September 10-12, 2012, Proceedings*, volume 7459 of *Lecture Notes in Computer Science*, pages 679–696. Springer, 2012.
- [38] Tao Zhang, Huangting Wu, and Sherman S. M. Chow. Structure-preserving certificateless encryption and its application. In Mitsuru Matsui, editor, *Topics in Cryptology - CT-RSA 2019 - The Cryptographers' Track at the RSA Conference 2019, San Francisco, CA, USA, March 4-8, 2019, Proceedings*, volume 11405 of *Lecture Notes in Computer Science*, pages 1–22. Springer, 2019.

## A. Security Definition of Building Blocks

A function  $f : \mathbb{N} \rightarrow \mathbb{R}$  is negligible, denoted by  $\text{negl}(\lambda)$ , if  $\forall c > 0, \exists x_0$  s.t.  $f(x) < 1/x^c, \forall x > x_0$ .

### A.1. Zero-Knowledge Proof or Argument System

A (non-interactive) proof system is an argument of knowledge if it satisfies the two properties below.

**Definition 5 (Completeness).** A non-interactive proof system  $\Pi = (\text{Setup}, \text{P}, \text{V})$  for language  $\mathcal{L}_R$  achieves perfect completeness if for all  $(x, w) \in R$ :  $\Pr[1 \leftarrow \text{V}(\text{crs}, x, \pi) : \text{crs} \leftarrow \text{Setup}(1^\lambda), \pi \leftarrow \text{P}(\text{crs}, x, w)] = 1$

**Definition 6 (Knowledge-Soundness).** A (non-interactive) proof system  $\Pi = (\text{Setup}, \text{P}, \text{V})$  satisfies knowledge-soundness if there exists a PPT simulator  $\text{Ext}$  such that for all PPT adversaries  $\mathcal{A}$  and prover  $\text{P}^*$ :

$$\left| \Pr \left[ \begin{array}{l} \text{V}(\text{crs}, x, \pi) = 1 \\ \text{crs} \leftarrow \text{Setup}(1^\lambda), \\ (x, w) \leftarrow \mathcal{A}(\text{crs}), \\ \pi \leftarrow \text{P}^*(\text{crs}, x, w) \end{array} \right] - \Pr \left[ \begin{array}{l} (x, w) \in R \\ \text{crs} \leftarrow \text{Setup}(1^\lambda), \\ (x, w) \leftarrow \mathcal{A}(\text{crs}), \\ w \leftarrow \text{Ext}^{\text{P}^*}(\text{crs}, x) \end{array} \right] \right| \leq \text{negl}(\lambda).$$

**Definition 7 (Perfect Special Honest-Verifier Zero-knowledge (SHVZK)).** A non-interactive proof system  $\Pi = (\text{Setup}, \text{P}, \text{V})$  is said to be Perfect SHVZK if there exists a PPT simulator  $\text{Sim} = (\text{Sim}_1, \text{Sim}_2)$  such that for all pairs of PPT adversaries  $(\mathcal{A}_1, \mathcal{A}_2)$ :

$$\Pr \left[ \begin{array}{l} (x, w) \in R, \mathcal{A}_1(\pi) = 1 \\ \text{crs} \leftarrow \text{Setup}(1^\lambda), \\ (x, w) \leftarrow \mathcal{A}_2(\text{crs}), \\ \pi \leftarrow \text{P}(\text{crs}, x, w) \end{array} \right] = \Pr \left[ \begin{array}{l} (x, w) \in R, \mathcal{A}_1(\pi) = 1 \\ \text{crs} \leftarrow \text{Sim}_1(1^\lambda), \\ (x, w) \leftarrow \mathcal{A}_2(\text{crs}), \\ \pi \leftarrow \text{Sim}_2(\text{crs}, x) \end{array} \right].$$

### A.2. Credential Signature Scheme

**Definition 8 (Unforgeability).** A credential signature scheme  $\text{Cred} = (\text{KGen}, \text{Sign}, \text{Vf})$  is said to be unforgeable if for all PPT adversaries  $\mathcal{A}$ , the probability of winning the following game is less than  $\text{negl}(\lambda)$ :

- The challenger runs  $\text{KGen}(1^\lambda)$  to obtain  $(\text{pk}, \text{sk})$ ,  $\text{pk}$  is given to the adversary  $\mathcal{A}$ .
- The adversary can query the signing protocol inputting  $\vec{m}, r$ . The protocol (with the challenger playing as the signer) returns a signature  $\sigma$ .
- The adversary outputs a message-signature tuple  $(\vec{m}^*, \sigma^*)$  that is valid and distinct, i.e.,  $\text{Cred.Vf}(\text{pk}, \vec{m}^*, \sigma^*) = 1$  and  $\vec{m}^*$  has not been queried before.



**Definition 9** (Blindness). A credential signature scheme  $\text{Cred} = (\text{KGen}, \text{Sign}, \text{Vf})$  satisfies the blindness property if for all PPT adversaries  $\mathcal{A}$ , the probability of winning the following game is less than  $1/2 + \text{negl}(\lambda)$ :

- The challenger runs  $\text{KGen}(1^\lambda)$  to obtain  $(\text{pk}, \text{sk})$ ,  $(\text{pk}, \text{sk})$  is given to the adversary  $\mathcal{A}$ .
- The adversary on input  $(\text{pk}, \text{sk})$ , outputs two messages  $\vec{m}_0, \vec{m}_1$ .
- The challenger samples a bit  $b$  and randomness  $r$  and interacts with the adversary in the  $\text{Cred.Sig}$  protocol with input  $(\vec{m}_b; r)$ .
- The adversary outputs a bit  $b'$  and wins if  $b = b'$ .

Blindness follows from the (perfectly) hiding property of the Pedersen commitment scheme and zero-knowledgeness of the ZKP.

## B. BBS+ Signature

Let  $\mathbb{G}_1, \mathbb{G}_2$ , and  $\mathbb{G}_T$  be cyclic groups of  $\lambda$ -bit prime order  $p$ . Let  $e : \mathbb{G}_1 \times \mathbb{G}_2 \rightarrow \mathbb{G}_T$  be a non-degenerate bilinear map.

The setup algorithm outputs the bilinear group context  $(p, \mathbb{G}_1, \mathbb{G}_2, \mathbb{G}_T, e, g, h)$  with  $g, h$  being a generator of  $\mathbb{G}_1, \mathbb{G}_2$ , respectively, and  $e(g, h)$  is a generator of the target group  $\mathbb{G}_T$ . For Type-III pairing,  $\mathbb{G}_1 \neq \mathbb{G}_2$ , with no isomorphism between  $\mathbb{G}_1$  and  $\mathbb{G}_2$ .

The public parameter  $\text{pp}$  contains generators  $(g, g_0, \dots, g_n, h) \leftarrow \mathbb{G}_1^{n+1} \times \mathbb{G}_2$  and the description of the bilinear group. The BBS+ signature (over a bilinear group defined by  $\text{pp}$ ) consists of three algorithms:

- $\text{BBS+}.\text{KGen}(1^\lambda) \rightarrow (\text{pk}, \text{sk})$ : The signer samples  $\gamma \in \mathbb{Z}_p^*$ , sets  $\text{sk} = \gamma$ , and sets  $\text{pk} = w = h^\gamma$ .
- $\text{BBS+}.\text{Sign}(\vec{m}, \text{sk}) \rightarrow \sigma$ : To sign  $\vec{m} = (m_1, \dots, m_n) \in \mathbb{Z}_p^n$ , the signer samples  $e, s \leftarrow \mathbb{Z}_p$  and computes  $A \leftarrow (gg_0^s \prod_{i=1}^n g_i^{m_i})^{\frac{1}{e+s}}$ . The signature on  $\vec{m}$  is  $\sigma \leftarrow (A, e, s)$ .
- $\text{BBS+}.\text{Vf}(\text{pk}, \vec{m}, \sigma) \rightarrow b$ : To verify  $\sigma$  on  $\vec{m}$  under  $\text{pp} = (g, g_0, \dots, g_n, h)$  and  $\text{pk} = h^\gamma$ , the verifier returns  $e(A, wh^e) \stackrel{?}{=} e(gg_0^s \prod_{i=1}^n g_i^{m_i}, h)$ .

The BBS+ signature supports signing on committed multi-block message  $\vec{m}$ . We describe the blind signature issuance protocol run between the signer and user:

- 1) The user computes the commitment of  $\vec{m}$  as  $C_{\vec{m}} = g_0^{s'} \prod_{i=1}^n g_i^{m_i}$  for a randomly sampled  $s' \leftarrow \mathbb{Z}_p^*$ .
- 2) The user sends the commitment  $C_{\vec{m}}$  to the signer with ZKPoK of the commitment opening, e.g.,  $\text{ZKP}\{(\vec{m}; \rho) : C_{\vec{m}} = \text{Com}(\vec{m}; \rho)\}$ .
- 3) The signer computes  $A = (gg_0^s C_{\vec{m}})^{\frac{1}{e+s}}$ , where  $e, s \leftarrow \mathbb{Z}_p^*$ , and sends  $\sigma' = (A, e, s)$  to the user.
- 4) The user parses the signature as  $\sigma = (A, e, s + s')$  and returns  $\sigma$ , if  $\text{BBS+}.\text{Vf}(\text{pk}, \vec{m}, \sigma) = 1$ .

The server stores the auxiliary output  $\text{aux}'$  that contains the commitment(s) and ZKP. The user can further prove relations about individual commitments of  $m_i \in \vec{m}$ , e.g.,  $\text{ZKP}\{(\{m_i; r_i\}_{i=1}^\ell; \rho) : C_{\vec{m}} = \text{Com}(\vec{m}; \rho) \wedge_{i=1}^\ell (C_{m_i} = \text{Com}(m_i; r_i))\}$ , where  $C_{m_i}$  would also be proved about in other ZKPs (commit-and-prove).

To prove the knowledge of the signature on the committed multi-block message, the prover can prepare the commitment and prove in zero-knowledge that the pairing verification equation holds [9] (which requires ZKP in the target group  $\mathbb{G}_T$ ). Camenisch *et al.* [13] proposed a more efficient variant of ZKP that only involves operations over  $\mathbb{G}_1$  and supports selective message disclosure, i.e.,  $\{m_i\}_{i \in D}$  are revealed, where  $D \subseteq \{1, n\}$  is a subset of indexes of  $\vec{m}$ . The prover with  $\sigma = (A, e, s)$  on  $\vec{m}$  randomizes the signature with  $r_1 \leftarrow \mathbb{Z}_p^*$  by setting  $A' \leftarrow A^{r_1}$ . It computes  $r_3 \leftarrow \frac{1}{r_1}$  and set  $\bar{A} \leftarrow A'^{-e} \cdot B^{r_1}$ , where  $B = (gg_0^s \prod_{i=1}^n g_i^{m_i} = A^{\gamma+e})$ . It also picks  $r_2 \leftarrow \mathbb{Z}_p$  and sets  $d \leftarrow b^{r_1} \cdot g_0^{-r_2}$  and  $s' \leftarrow s - r_2 \cdot r_3$ . Finally, it proves:

$$\pi \leftarrow \text{ZKP}\{(\{m_i\}_{i \notin D}, e, r_1, r_2, r_3, s') : \bar{A}/d = A'^{-e} \cdot g_0^{r_2} \wedge g \prod_{i \in D} g_i^{m_i} = d^{r_3} g_0^{-s'} \prod_{i \notin D} g_i^{-m_i}\}.$$

The proof consists of  $(A', \bar{A}, d, \pi)$ . Its verification requires checking  $A' \neq 1_{\mathbb{G}_1}$  and  $e(A', w) = e(\bar{A}, h)$ . The signer needs to publish  $(\bar{g}_1, \bar{g}_1^\gamma)$  for  $\bar{g}_1 \neq 1_{\mathbb{G}_1}$  and ensures  $A \neq 1_{\mathbb{G}_1}$ . The nonce is revealed to the SP in our protocol. Recently, a variant of the BBS+ signatures with a shorter signature size has been proposed [32].

The BBS+ scheme is existentially unforgeable against adaptive chosen message attacks for Type-III pairings under the  $q$ -strong Diffie-Hellman ( $q$ -SDH) assumption [13].

**Definition 10** ( $q$ -SDH Assumption). Given the tuple  $(g_1, g_1^x, g_1^{(x^2)}, \dots, g_1^{(x^q)}, g_2, g_2^x) \in \mathbb{G}_1^{q+1} \times \mathbb{G}_2^2$ , no PPT algorithm can output  $(c, g_1^{1/(x+c)})$ , where  $c \in \mathbb{Z}_p \setminus \{-x\}$ .

## C. Extensions to Multiple Service Providers

The SP holds the credential issuance key for ticket management in our setting. As SMART uses a credential signature scheme in a black-box manner, we can distribute the credential-issuing power across multiple SPs by replacing it with its different extensions for different “trust” models. For example, for the  $n'$ -out-of- $n$  access structure, i.e., any subgroup of  $n'$  SPs among a group of  $n$  SPs can issue credentials, we can use 1-out-of- $n$  ZKP [22] or threshold BBS+ [24]. In issuer-hiding attribute-based credentials, Bobolz *et al.* [10] replaced the OR-proof technique with a signature-based set membership proof [12], where the verifier specify the permitted issuer public keys by signing a possibly different set of them according to the expected verification policy (regarding the SPs) in a given scenario.

Different SPs can also judge tickets as they wish, and multiple SPs can judge the same ticket. The redemption policy and the credential update function could take in the same judgment signed by at least  $n'$  SPs (with  $n'$ -out-of- $n$  proof [5], [22]). With the above in mind, our signature-based design makes it compatible with a possible extension to the decentralized setting. The core idea is that the credentials, policies, and score lists can be put into a trusted public ledger [26] where succinct ZKPs accompany the credential updates.