Decentralized Identity for Student Credentials: Simulation-Based Comparison of Revocation in Verifiable Credential, AnonCreds, and Polygon ID Stacks

Measured cryptographic costs and a state-scale event simulation of two million academic credentials

[1]Pravat Satpathy, [2]Sonia Mohapatra, [3]Saswat Swain
[1]Researcher, [2]Researcher, [3]Ph.D Scholar
[1,2]Research and Development Division,
Tekstein Solutions Private Limited, Bhubaneswar, Odisha, India
[3]KSOM — KIIT Deemed to be University,
Bhubaneswar, Odisha, India
[1]analytics.pravat@gmail.com, [2]smctc98@gmail.com, [3]saswat.mck@gmail.com

***Abstract*—Academic credentials are issued once, verified many times, and revoked rarely, yet the revocation mechanism decides most of the cost, latency, and privacy of a decentralized-identity deployment for students. This paper compares three deployable stacks: a W3C Verifiable Credential with did:web issuer keys, SD-JWT selective disclosure, and a Bitstring Status List; an AnonCreds-style accumulator registry on a permissioned ledger; and an iden3 (Polygon ID) sparse-Merkle identity state anchored on a public chain. A lifecycle model is derived with four propositions: an entropy bound on status-list size, the holder-side update burden of accumulators, a closed form for anchoring transactions under batching, and a linear law for false acceptance of revoked credentials under caching. Cryptographic primitives are measured in a reference implementation, and the three stacks are replayed against one simulated population of two million credentials, 17.6 thousand revocations, and 3.7 million verification requests over five years. The status-list stack costs 2.9 kilobytes and 0.2 milliseconds per verification, but every presentation is linkable and its herd privacy collapses for small issuers: one third of list fetches expose anonymity sets below 1000, far from the 131 072-entry list length the specification assumes. Federating lists across issuers restores a median anonymity set of 130 thousand for 2.7 times more list bytes per fetch. The accumulator stack is unlinkable but shifts up to 120 witness updates per presentation onto the holder, and the Merkle-state stack is unlinkable at an anchoring cost governed by whether issuance is on-chain rather than by batching. All code and result logs are released.**

***Index Terms*—Decentralized identifiers, verifiable credentials, credential revocation, status list, cryptographic accumulator, sparse Merkle tree, zero-knowledge proof, academic records, self-sovereign identity.**

I. Introduction

India now holds more than 1.1 billion academic records on government digital platforms, and roughly 263 million learners carry an APAAR identifier [1]. These records are verifiable because a central platform vouches for them at the moment of lookup. A decentralized identity (DID) design moves the proof into the credential itself: the institution signs a W3C Verifiable Credential (VC) [2], the student holds it in a wallet, and an employer checks the signature against the issuer's public key without calling anyone. Signature verification is the easy part. The part that decides whether such a system is cheap, timely, and private is revocation, because a degree withdrawn for malpractice or a mark-sheet replaced after a correction must stop verifying, and the mechanism that makes it stop is consulted at every presentation.

Three mechanisms are deployable today. The Bitstring Status List [3] publishes one bit per credential in a compressed, signed list that verifiers fetch over HTTPS. AnonCreds, the credential format of the Hyperledger Aries and Indy ecosystem, uses a cryptographic accumulator whose current value lives on a permissioned ledger and whose non-membership is proved in zero knowledge [4], [5]. The iden3 protocol behind Polygon ID (now Privado ID) keeps revocation nonces in a sparse Merkle tree (SMT) whose root is part of an identity state committed on a public chain through a Groth16-verified transition [6], [7]. Each is well described in isolation, and education-sector proposals adopt one of them, often with a token or a bespoke chain attached [8], [9]. What is missing is a like-for-like quantification under one population of students, issuers, and verifiers with the arrival pattern of academic credentials: bursty issuance at convocation, revocation near one percent, and verifications that trail issuance by months.

This paper fills that gap with a lifecycle model, measured primitives, and an event simulation. No new cryptography is proposed; the contribution is quantitative. Specifically:

1) A lifecycle cost model for credential issuance, presentation, revocation, and verification across the

three stacks, with four propositions that are checked numerically: a Shannon bound on status-list size, the holder-side witness burden of accumulators, a closed form for anchoring transactions under batching, and a linear law for false acceptance under verifier caching.

2) Measured costs of the primitives each stack depends on, from a reference implementation with a fixed seed.
3) A replay of the three stacks against one simulated state-scale population of two million credentials over five years, reporting false-accept rate, propagation latency, bytes per verification, holder and verifier compute, anonymity set, cross-verifier linkability, and issuer cost.
4) Two findings that run against common assumptions: the herd privacy of status lists fails for most issuers in a real institution-size distribution, and per-issuer batching barely reduces anchoring cost when revocations are sparse.

II. Related Work

Table I summarises the work this study builds on. Camenisch and Lysyanskaya introduced dynamic accumulators for anonymous-credential revocation [10]; the bilinear variant of Camenisch, Kohlweiss, and Soriente [5] is the construction behind AnonCreds revocation registries and their public tails file [4]. Batched witness updates [11] exist but deployed wallets still update one revocation at a time. The W3C status list [3] inherits the certificate-revocation-list idea and defends holder privacy by herd size, with a recommended minimum of 131 072 entries. Hoops, Gebele, and Matthes [12] show that both status lists and accumulators leak issuer metadata and propose CRSet, a padded Bloom-filter cascade; their comparison is qualitative on the accumulator side. The iden3 specification [6] combines an SMT per identity with Groth16 proofs [7], and the Polygon ID documentation describes the on-chain state contract but publishes no cost figures. Surveys [13]–[15] classify these systems without measuring them, and the education-sector work that established the use case [8], [9] treats revocation as secondary. The gap addressed here is a measured, population-level comparison of the three mechanisms that production wallets already implement.

Table I. Prior work on credential revocation and decentralized identity relevant to this study

Work (year)	Mechanism / method	Strength	Limitation or open question
Camenisch and Lysyanskaya [10] (2002)	Dynamic RSA accumulator; ZK non-membership	Unlinkable presentations; constant-size proof	Holder must update witness per revocation; no system-scale evaluation
Camenisch et al. [5] (2009)	Bilinear accumulator; basis of AnonCreds [4]	Pairing-friendly; public tails file	Tails file grows linearly with registry; wallet burden unquantified
W3C [3] (2025)	Bitstring Status List, GZIP, HTTPS	Simple; cacheable; offline-capable	Herd privacy argued from a 131 072 minimum; linkable index
Hoops et al. [12] (2025)	CRSet: padded Bloom-filter cascade	Hides issuer metadata	Accumulator comparison qualitative; education arrival pattern not studied
iden3 [6], Groth [7]	SMT identity state; Groth16 on-chain transition	Unlinkable; public auditability	No published anchoring or wallet cost model
Grech and Camilleri [8], Blockcerts [9]	Blockchain-anchored diplomas	Established the education use case	Revocation and privacy treated as secondary

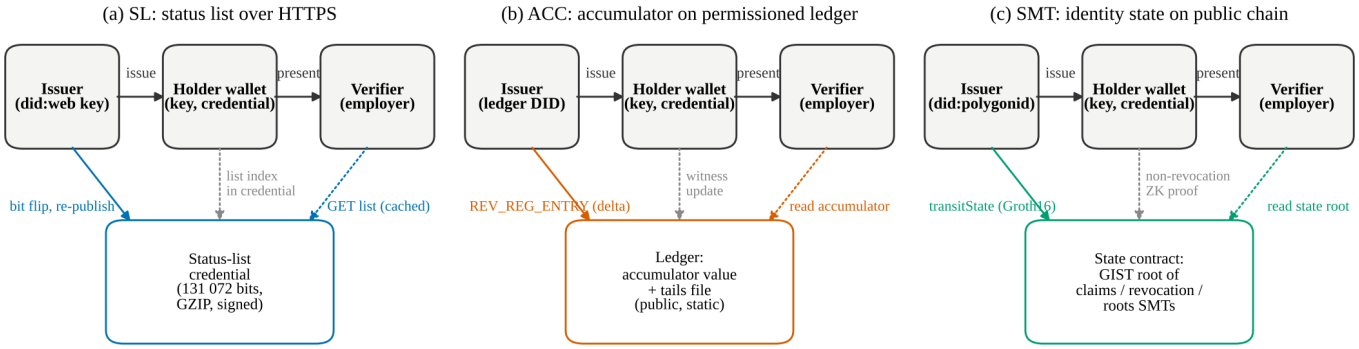


Fig. 1. Lifecycle of the three stacks. Solid arrows are issuer writes, dotted arrows are verifier reads; the grey arrow is what the holder must carry or compute to prove non-revocation.

III. System Model and Problem Statement

Three roles interact. An *issuer* (an institution) signs credentials and later revokes some of them. A *holder* (a student) stores credentials in a wallet and presents them. A *verifier* (an employer, a university, a scholarship board) checks a presentation. A trust registry publishes issuer identifiers and keys; its design is orthogonal to revocation and is not modelled. Fig. 1 shows the three stacks compared.

Stack SL. The issuer publishes a did:web document with an Ed25519 key, issues an SD-JWT VC [16] with a status-list entry, and maintains one or more Bitstring Status Lists of L = 131 072 bits [3]. Revocation flips a bit; the list is re-signed and re-published at most once per interval T. Verifiers fetch the

list over HTTPS and may cache it for a lifetime τ. A state-level federation may instead host *shared* lists in which indices from all issuers are interleaved; this variant is denoted SL-F.

Stack ACC. Credentials are assigned to revocation registries of capacity R. Each registry has a public tails file of R group elements and an accumulator value on a permissioned ledger. Revocation removes an element; the issuer writes a delta entry at most once per T. Before presenting, a holder updates its witness against every revocation published in its registry since its last presentation and proves non-revocation in zero knowledge.

Stack SMT. Each issuer keeps claims, revocation, and roots SMTs; revocation inserts a nonce into the revocation tree and the new state is committed on-chain once per T through a transition that verifies a Groth16 proof. Credentials are issued by signature (no transition) unless stated otherwise. The holder proves non-revocation against the committed state; the verifier reads the global state root, which it may cache for τ.

Scale scenario and metrics. One state's higher-education sector is modelled with the order-of-magnitude parameters of Table II: 1000 issuers of log-normal size, 400 000 credentials per year, ρ = 1 %, verifications that trail issuance by months, and Zipf-distributed verifier demand. For each stack the simulation reports the false-accept rate (share of verifications of already-revoked credentials that the verifier accepts), publication latency, verifier-side staleness, bytes per verification, holder and verifier computation per presentation, the anonymity set k against the issuer at each status fetch, the fraction of cross-verifier presentation pairs that are linkable, and issuer cost. The question is which mechanism meets a false-accept target below one percent at the lowest holder burden and an acceptable privacy level, and where each mechanism's cost actually sits.

IV. Lifecycle Model and Propositions

Let a credential be revoked at time t_r. The issuer makes the revocation visible at $t_v = T\lceil t_r/T \rceil + \delta$, where δ is the write latency of the publication medium. A verifier that cached the status object at time $t_c \le t$ sees the revocation only if $t_v \le t_c$. The effective window during which a revoked credential still verifies is therefore

$$W = (t_v - t_r) + (t - t_c), \tag{1}$$

the sum of publication latency and cache staleness. Under uniform revocation times the first term averages $T/2 + \delta$; the second depends on τ and on how often the same verifier revisits the same status object, which is why it is simulated rather than assumed.

Proposition 1 (status-list size). For L entries revoked independently with density p, any lossless encoding of the list needs at least $L\,H_2(p)/8$ bytes on average, where H_2 is the binary entropy function [17]. GZIP, which the specification mandates, approaches this bound only for p above a few percent and exceeds the raw 16 384 bytes at $p = 0.5$. Consequence: a list is small because revocations are rare, not because of compression, and clustered revocations (a cohort withdrawal) compress far better than scattered ones.

Proposition 2 (holder burden under accumulators). If revocations in a registry arrive at rate λ_R and a holder presents after an idle time Δt, the expected number of witness updates is $\lambda_R \Delta t$, each costing one group operation. Summed over the H non-revoked holders of the registry, every revocation induces $\Theta(H)$ work, against $\Theta(1)$ for a status-list bit flip; the work is deferred to presentation time and lands on the device least able to do it.

Proposition 3 (anchoring under batching). If issuer i generates state-changing events at rate λ_i and commits at most once per interval T, the expected number of transitions per year is

$$N(T) = \Sigma_i\,(365/T)[1 - \exp(-\lambda_i T)], \tag{2}$$

with cost $c_{tx}N(T)$ and mean publication latency $T/2 + \delta$. For $\lambda_i T \ll 1$ the exponential linearises and $N \to 365\,\Sigma_i\lambda_i$: batching buys nothing, cost is proportional to the number of events, and latency is paid for free. For $\lambda_i T \gg 1$, $N \to 365n/T$ and cost and latency trade off as a constant product. Equation 2 assumes Poisson arrivals; for clustered arrivals it is an upper bound, because clustering raises the probability that several events share a batch.

Proposition 4 (false acceptance). Let $f(g)$ be the density of the gap $g = t - t_r$ between a revocation and a subsequent verification of the same credential, conditioned on $g > 0$. For small W,

$$P_{fa} \approx f(0^+) \cdot \mathrm{E}[W] \equiv f_0 W, \tag{3}$$

so the false-accept rate is linear in the effective window with a slope set by the application, not by the mechanism. In the academic setting f_0 is small because employers check credentials months after issuance; the same mechanisms deployed for access badges would face a much steeper slope.

V. Simulation Setup

All experiments are reproducible from a single seed (20261025) with Python 3.13, NumPy, and the cryptography package on one 2.1 GHz Xeon core; no hardware wallet or production ledger was used. Table II lists every parameter and whether it was measured or assumed.

Table II. Simulation parameters (M = measured here, A = assumed or modelled; see README for sources)

Parameter	Value	Src.
Issuers / credentials per year / horizon	1000 / 400 000 / 5 yr	A
Revocation ρ, delay; verifications, lag	1 %, Exp(180 d); 3, Exp(730 d)	A
Verifiers; demand skew	5000; Zipf 1.1	A
List length; publish interval T; cache τ	131 072; 24 h; 24 h	A
SD-JWT presentation (6 of 12 claims); list envelope	2189 B; 646 B	M
Ed25519 sign / verify	0.035 / 0.086 ms	M
Registry size; tails element; ledger write δ	10 000; 128 B; 5 s	A
Witness update / accumulator verify (RSA proxy)	6.3 / 3.1 ms	M
AnonCreds presentation; ledger read	4200 B; 900 B	A
SMT depth; non-membership proof	32; 548 B	M
Gas per transition; gas price; POL price	300 000; 30 gwei; 0.20 USD	A
Groth16 prove (phone) / verify; ZK presentation	1.5 s / 8 ms; 1100 B	A

Measured primitives. Ed25519 and ECDSA P-256 signing and verification; the byte size of an SD-JWT VC with twelve claims (eight semester grades, identity fields, programme, and award) at five disclosure levels; the GZIP size of a 131 072-bit list at ten revocation densities with random and with contiguous placement, five trials each; a 32-level SMT over SHA-256 with 10^3 to 10^5 inserted nonces (iden3 uses Poseidon, which is slower natively but proof-friendly); and a 2048-bit RSA accumulator with hash-to-prime elements, used as a measured proxy for the asymptotic cost of accumulator operations, since the bilinear accumulator's per-revocation witness update is also one group operation.

Event simulation. The population generator draws issuers, issuance times (30 % uniform, 70 % in two convocation peaks), revocations, and verification requests, then sorts requests in time. For each stack a cache replay walks every (verifier, status object) pair in time order and re-fetches when the cached copy is older than τ; the status object is a list for SL, a registry accumulator for ACC, and the global state root for SMT. Holder witness updates are counted per presentation from the revocations published in the holder's registry since its previous presentation. Linkability is computed combinatorially from the verification count of each credential and the number of batch-issued copies used round-robin. Sweeps vary T, τ, ρ, registry size, copies, and gas price.

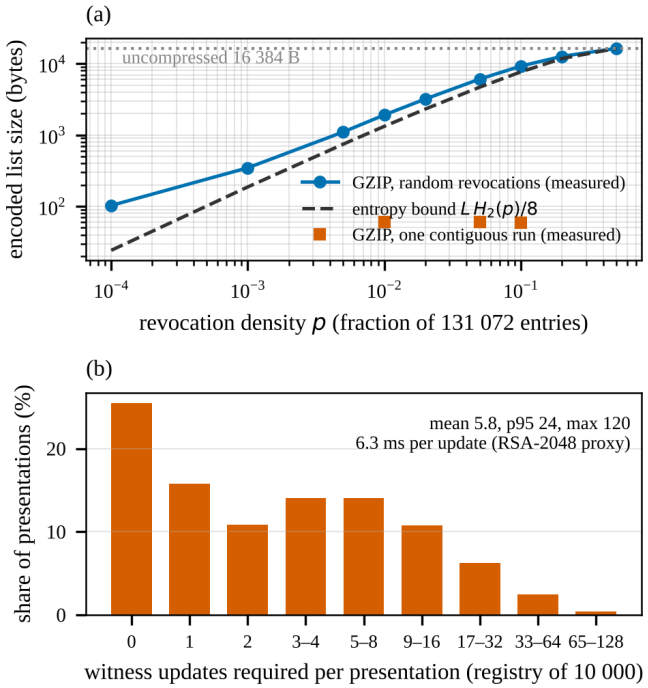


Fig. 2. (a) Encoded size of a 131 072-bit status list against revocation density, measured GZIP output versus the Shannon bound; (b) distribution of witness updates a holder must perform before a presentation under the accumulator stack with registries of 10 000.

VI. Results and Discussion

A. Primitive Costs

Ed25519 signing took 0.035 ms and verification 0.086 ms; P-256 took 0.020 and 0.058 ms. The twelve-claim SD-JWT VC is 2573 bytes as issued and 2189 bytes when six claims are disclosed with a key-binding JWT, against 1118 bytes for the same content as a JSON-LD VC with full disclosure: selective disclosure roughly doubles presentation size. SMT membership and non-membership proofs were 292 to 548 bytes with empty-sibling compression and verified in 0.012 ms. RSA-accumulator verification took 3.1 ms and a holder witness update 6.3 ms per revocation; the issuer-side deletion with trapdoor took 23 ms.

B. Status-List Compression

Fig. 2(a) confirms Proposition 1. GZIP output is 1.44 times the entropy bound at $p = 0.01$, 1.21 times at $p = 0.1$, and 1.00 at $p = 0.5$, where it is 16 408 bytes, larger than the raw list. A single contiguous run of 13 107 revoked bits compressed to 59 bytes. Status lists are therefore cheap in this domain because ρ ≈ 1 % makes them cheap; a list that reaches 10 % density costs 9.3 kilobytes per fetch, and the W3C minimum length amplifies that cost for every verifier.

C. Baseline Comparison

Table III. Baseline comparison at T = τ = 24 h, ρ = 1 %, publication latency 12.0 h for all (SL-F = federated lists)

Metric	SL	SL-F	ACC	SMT
False-accept rate (%)	0.11	0.14	0.11	0.19
Verifier staleness, mean (h)	1.4	6.5	1.3	8.8
Bytes per verification	2866	3142	5000	1233
Holder compute per presentation	0.03 ms	0.03 ms	36 ms	1.5 s (A)
Verifier compute per presentation	0.21 ms	0.21 ms	4.1 ms	8 ms (A)
Anonymity set k, median	1598	129 861	10 000	n/a
Linkable cross-verifier pairs	1.00	1.00	0.00	0.00
Issuer writes per year	3462	1295	3465	3462
Issuer cost per year (USD)	≈0	≈0	ledger ops.	6.2 (A)

Table III compares the stacks at $T = \tau = 24$ h. All three deliver the same false-accept rate of 0.11 %, because at this setting the window is dominated by the shared publication latency of 12.0 h. The differences are elsewhere. SL moves 2.9 kilobytes per verification and asks 0.03 ms of the holder, but every presentation of a credential carries the same issuer signature and status index, so all cross-verifier pairs are linkable. ACC and SMT are unlinkable; ACC costs the holder a mean of 36 ms and a 95th percentile of 151 ms per presentation in the RSA proxy, with 25 % of presentations needing no update and 2.8 % needing more than 32 (Fig. 2(b)), plus a 1.28 megabyte tails download per registry. SMT's holder cost is a modelled 1.5 s Groth16 proof. The verifier's cached object also matters: the global state root of SMT is revisited by every request a verifier makes, so at τ = 24 h its mean staleness is 8.8 h against 1.4 h for a per-list cache, and its false-accept rate is 0.19 % rather than 0.11 %.

D. Effective Window and Offline Operation

Fig. 3(a) plots the false-accept rate of all three stacks over the interval and cache sweeps against the mean effective window. The points collapse onto the line of Proposition 4 with $f_0 = 0.0021$ per day, measured directly from the population as the share of post-revocation verifications that fall within one day of the revocation. At $T = 1$ h no revoked credential was accepted; at $T = 30$ d the rate is 3.2 %. Fig. 3(b) shows the cache sweep: a seven-day offline cache keeps SL and ACC at 0.32 to 0.33 % and SMT at 0.82 %, and a thirty-day cache raises them to 1.45 % and 3.23 %. For field verification with intermittent connectivity, a weekly refresh satisfies a one-percent target with any stack, and the stack with per-object caching degrades most gracefully.

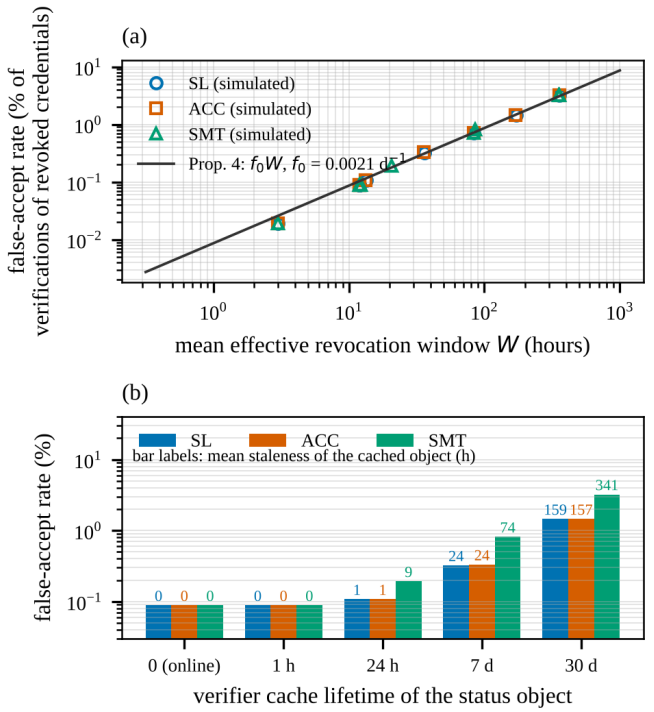


Fig. 3. (a) False-accept rate against mean effective window for all stacks and sweeps, with the prediction of Proposition 4; (b) false-accept rate against verifier cache lifetime, labelled with the mean staleness of the cached status object.

E. Anchoring Cost

Fig. 4(a) tests Proposition 3. With signature-based issuance, the 17 564 revocations over five years produced 3462 transitions per year at $T = 24$ h and still 2562 at $T = 30$ d, within 1.1 % of Equation 2; batching cannot amortise events that arrive 3.5 times per issuer per year. At the modelled gas and prices this is 6.2 USD per year for the whole state, or 312 USD at 300 gwei and 1.00 USD per POL. Moving issuance on-chain changes the picture: 148 416 transitions per year at $T = 24$ h (267 USD base, 13 357 USD high), and here batching works because issuance is dense, cutting transitions thirtyfold between $T = 1$ h and $T = 30$ d. The Poisson form overestimates the on-chain-issuance count by 26 % at $T = 24$ h, as convocation clustering predicts. Anchoring is not where the money goes in this application; the holder's proving time and the operator's infrastructure are.

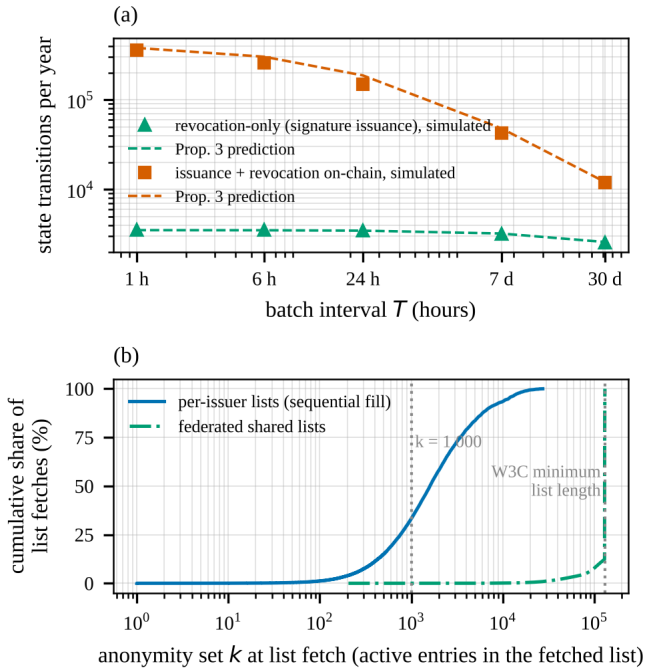


Fig. 4. (a) On-chain state transitions per year against batch interval, simulated and predicted by Equation 2, for revocation-only and for on-chain issuance; (b) cumulative distribution of the anonymity set at each status-list fetch for per-issuer lists and for federated shared lists.

F. Herd Privacy and Linkability

Fig. 4(b) is the central negative result. With one list per issuer filled in issuance order, the anonymity set at a fetch is the number of active entries that issuer has placed so far, not the list length. In the simulated size distribution, 42 % of issuers issue fewer than 1000 credentials in five years and the largest issues 27 631, so no list ever fills: the median anonymity set is 1598, the 5th percentile 233, and 33 % of fetches expose a set below 1000. The 131 072-entry minimum protects nobody at an institution with 400 graduates a year. Federating lists at state level (SL-F), with indices from all issuers interleaved, raises the median to 129 861 and the 5th percentile to 81 716, at the cost of a list that is 2.7 times larger per fetch (2092 against 766 bytes) because it is fuller, and of a federation signature on the list. Sequential index assignment also reveals issuance order and hence cohort; random assignment within the list removes this at no cost. Linkability is a separate problem. SD-JWT presentations of one credential share the issuer signature; batch-issuing n copies with distinct salts and indices, used round-robin, reduced the linkable fraction of cross-verifier pairs from 1.00 to 0.33 (n = 2), 0.13 (n = 3), 0.02 (n = 5), and 0.00 (n = 10), with a 1.6 % increase in list consumption at n = 10. Full unlinkability without copies requires the ZK stacks or a BBS signature suite [18].

G. Sensitivity

Table IV reports the remaining sweeps. Raising ρ from 0.1 % to 5 % increases SL list bytes per fetch from 707 to 944 and ACC holder time from 3.7 to 187 ms, while SMT anchoring scales linearly with revocations. Registry size trades tails download against update burden: registries of 1000 cut the mean holder time to 13 ms and the tails file to 0.13 megabytes but require 2508 registries and 502 registry definitions per year; registries of 100 000 need 12.8 megabyte tails files per holder. Gas sensitivity spans two orders of magnitude between plausible price regimes and is the least stable number in this paper.

Table IV. Sensitivity of the stacks to revocation rate, registry size, and gas price (T = τ = 24 h unless stated)

Sweep	SL B/fetch	ACC ms (p95)	Tails MB	SMT USD/yr
ρ = 0.1 %	707	3.7	1.28	0.6
ρ = 1 % (base)	766	37.4 (151)	1.28	6.3
ρ = 5 %	944	186.9	1.28	29.6
Registry 1000	—	13.1 (50)	0.13	—
Registry 100 000	—	41.3 (176)	12.8	—
Gas 300 gwei, POL 1 USD	—	—	—	312
On-chain issuance, base gas	—	—	—	267
On-chain issuance, high gas	—	—	—	13 357

VII. Recommendation and Limitations

For student credentials the evidence points to the status-list stack with three corrections: federated lists so that small issuers inherit the herd of the state, random index assignment, and batch-issued copies for holders who present to many verifiers. It verifies in 0.2 ms, works offline for a week within a one-percent false-accept target, and costs the holder one signature. The accumulator and Merkle-state stacks are the right choice only where unlinkability is a legal requirement rather than a preference, and then the budget item to plan for is the wallet, not the chain.

Four limitations bound these claims. Timings are from a pure-Python reference on a server core; absolute values would shift by one to two orders of magnitude on native code and on phones, although the ratios between stacks are structural. The BBS, Groth16, and AnonCreds proof costs and the gas figures are modelled, not measured, and gas prices are volatile. The population is synthetic; institution sizes, revocation rates, and verification lags are plausible orders of magnitude, and the sensitivity sweep shows which conclusions depend on them. Usability, key recovery, and trust-registry governance were out of scope.

VIII. Conclusion

A lifecycle model with four numerically verified propositions, measured primitives, and a two-million-credential event simulation were used to compare status-list, accumulator, and Merkle-state revocation for student credentials. False acceptance is governed by the effective window and the application's verification lag, not by the mechanism; holder burden and privacy are where the mechanisms differ. Status lists are the cheapest by a wide margin but, as deployed per issuer, fail to deliver the herd privacy their specification assumes, a failure that federation repairs. Accumulators push revocation work onto wallets, and on-chain anchoring is inexpensive for revocation but not for issuance. Future work should replace the modelled proof costs with measurements on mid-range handsets and test federated lists against real institution-size data from a national depository.

Acknowledgment

The authors thank the Research and Development Division of Tekstein Solutions Private Limited for computing resources. Author designations are placeholders pending confirmation.

References

[1] Press Information Bureau, Government of India, "Eleven years of Digital India," New Delhi, Jul. 2026. [Online]. Available: https://www.pib.gov.in

[2] M. Sporny, T. Thibodeau, I. Herman, M. B. Jones, and G. Cohen, Eds., "Verifiable credentials data model v2.0," W3C Recommendation, May 2025. [Online]. Available: https://www.w3.org/TR/vc-data-model-2.0/

[3] M. Sporny, D. Longley, and others, Eds., "Bitstring status list v1.0," W3C Recommendation, May 15, 2025. [Online]. Available: https://www.w3.org/TR/vc-bitstring-status-list/

[4] Hyperledger Foundation, "AnonCreds specification v1.0," 2023. [Online]. Available: https://hyperledger.github.io/anoncreds-spec/

[5] J. Camenisch, M. Kohlweiss, and C. Soriente, "An accumulator based on bilinear maps and efficient revocation for anonymous credentials," in *Public Key Cryptography (PKC 2009)*, LNCS vol. 5443, Berlin: Springer, 2009, pp. 481–500.

[6] iden3, "Iden3 protocol specification," 2024. [Online]. Available: https://docs.iden3.io

[7] J. Groth, "On the size of pairing-based non-interactive arguments," in *Advances in Cryptology – EUROCRYPT 2016*, LNCS vol. 9666, Berlin: Springer, 2016, pp. 305–326.

[8] A. Grech and A. F. Camilleri, *Blockchain in Education*, JRC Science for Policy Report EUR 28778 EN, Luxembourg: Publications Office of the European Union, 2017.

[9] MIT Media Lab Learning Initiative, "Blockcerts: an open standard for blockchain credentials," 2016. [Online]. Available: https://www.blockcerts.org

[10] J. Camenisch and A. Lysyanskaya, "Dynamic accumulators and application to efficient revocation of anonymous credentials," in *Advances in Cryptology – CRYPTO 2002*, LNCS vol. 2442, Berlin: Springer, 2002, pp. 61–76.

[11] D. Boneh, B. Bünz, and B. Fisch, "Batching techniques for accumulators with applications to IOPs and stateless blockchains," in *Advances in Cryptology – CRYPTO 2019*, LNCS vol. 11692, Cham: Springer, 2019, pp. 561–586.

[12] F. Hoops, J. Gebele, and F. Matthes, "CRSet: non-interactive verifiable credential revocation with metadata privacy for issuers and everyone else," arXiv:2501.17089, 2025.

[13] A. Mühle, A. Grüner, T. Gayvoronskaya, and C. Meinel, "A survey on essential components of a self-sovereign identity," *Computer Science Review*, vol. 30, pp. 80–86, 2018.

[14] M. Kuperberg, "Blockchain-based identity management: a survey from the enterprise and ecosystem perspective," *IEEE Transactions on Engineering Management*, vol. 67, no. 4, pp. 1008–1027, Nov. 2020.

[15] Y. Liu, D. He, M. S. Obaidat, N. Kumar, M. K. Khan, and K.-K. R. Choo, "Blockchain-based identity management systems: a review," *Journal of Network and Computer Applications*, vol. 166, p. 102731, 2020.

[16] D. Fett, K. Yasuda, and B. Campbell, "Selective disclosure for JWTs (SD-JWT)," RFC 9901, Nov. 2025.

[17] C. E. Shannon, "A mathematical theory of communication," *Bell System Technical Journal*, vol. 27, no. 3, pp. 379–423, Jul. 1948.

[18] T. Looker, V. Kalos, A. Whitehead, and M. Lodder, "The BBS signature scheme," IETF Internet-Draft draft-irtf-cfrg-bbs-signatures, work in progress, 2025.