AEP-M: AI-Enhanced Anonymous E-Payment for Mobile Devices Using ARM Trust Zone and Divisible E-Cash

Vimal Teja Manne *

Posted Date: 6 January 2026

doi: 10.20944/preprints202601.0170.v1

Keywords: anonymous electronic payment; mobile payment; ARM TrustZone; divisible e-cash; fraud detection; privacy; secure hardware

 doi:10.20944/preprints202601.0170.v1

Article

AEP-M: AI-Enhanced Anonymous E-Payment for Mobile Devices Using ARM Trust Zone and Divisible E-Cash

Vimal Teja Manne

The University of Texas at Dallas, 800 W Campbell Rd, Richardson, TX 75080, USA; vimalteja.m@gmail.com

Abstract

E-Payment has become popular in mobile commerce, can provide consumers with a convenient way to make purchases electronically. Currently, however, all too many E-Payment systems are primarily focused on securing a consumer's financial information and do little to prevent privacy leaks and AI-generated scams. This paper defines AEP-M, a novel AI-enhanced anonymous e-payment scheme developed for mobile devices that uses TrustZone and divisible e-cash. Since mobile devices have very limited processing power and each transaction must be performed in real time, the proposed solution combines an efficient divisible e-cash system with AI-powered anomaly detection techniques to improve both the security, privacy and fraud detection in mobile payments. In addition to enabling users to divide a single withdrawal of an e-coin of a large amount into multiple transactions without disclosing their identity to either banks or merchants, AEP-M integrates AI-based risk assessment to identify suspicious spending behaviors to rapidly mitigate fraud and continuously monitor transactions. By employing a combination of bit decomposition and pre-computation to minimize the computational overhead of the transaction process, AEP-M provides the optimal performance in terms of minimizing the max number of exponentiation operations required to perform the frequent online spending process on elliptic curves. Finally, AEP-M also incorporates an ARM TrustZone to protect a user's financial data and important private data; an SRAM PUF is used as a Root of Trust to derive AI-powered keys and manage sensitive data, thereby increasing both the security and reliability of the system. A prototype of AEP-M was implemented and evaluated using the BN curve at a 128-bit security level. The experimental results demonstrated that AEP-M is capable of improving the Security, Efficiency and Fraud Detection capabilities of Mobile Digital Payments while maintaining User Privacy and Anonymity.

Keywords: anonymous electronic payment; mobile payment; ARM TrustZone; divisible e-cash; fraud detection; privacy; secure hardware

1. Foundations of Privacy-Preserving Mobile Payments

The ascendancy of mobile devices as the primary mode for financial interaction has revolutionized digital commerce, resulting in an urgent need to protect users? privacy. Anonymous e-payment (AEP) systems are providing a solution to maintain user identities and transaction histories while making use of cryptographic primitives to achieve unsuitability between payment transactions[1]. A major innovation is divisible e-cash. It permits a single high-denomination coin to be spent in lower denominations without repeated contacts with the issuer, thus allowing for a more convenient method while maintaining upholding anonymity based on more advanced cryptography[2].

There are difficulties in integrating AEP on mobile platforms based on hardware limitations and unresolved heterogeneity of the pertinent equipment. The target of a secure approach would necessarily be hardware enforced security by techniques such as ARM Trust Zone and other trusted execution environments (TEEs) to isolate sensitive operations and cryptographic keys from each

 doi:10.20944/preprints202601.0170.v1

other[3]. Adversarial threats to the system of the class of double-spending transactions would be met by tracing mechanisms which would become active only in the case of fraud and leave this anonymity of the honest users of this technology intact[2].

Interoperability with the financial infrastructure must not be underestimated. Banks can, for example through the vehicle blind-signature technology, dispense currency to enable these transactions without acquiring from experience any non-generic information on the users[1]. In addition, artificial intelligence can greatly assist to provide real-time fraud detection without violating the principles of privacy provided that the above is preserved [4]. Legal considerations must also be contemplated. Legislation on money laundering and other criminal activities would tend to invoke systems which permit conditional traceability subject to authorization[4].

The ubiquity of smart phones provides an excellent infrastructure base for the implementation of scalable AEP systems[5]. By means of divisible e-cash, TEE and AI detection, protocols such as AEP-M seek to marry strong cryptographic guarantees with the performance and usability requirements of mobile users.

2. ARM Trust Zone and Secure Hardware Integration

ARM Trust Zone gives hardware-level security to mobile devices by dividing the system into the Secure World (for sensitive operations) and the Normal World (for general processing). Such isolation is a fundamental necessity for safeguarding cryptographic routines and private user data in anonymous e-payment (AEP) systems from malware and other operating system hacks[3].

In AEP-M, the Trust Zone separates out all divisible e- cash functions (withdraw, spend and deposit) so that the keys and transaction data are not available to the Normal World. Because of hardware separation this prevents leakage of the user's identity during blind signature protocols and protects the split coin logic from being hacked[6]. Secure World software, such as OP-TEE, must be lightweight and verified to avoid introducing vulnerabilities [7].

TrustZone integrates with secure I/O to bind peripherals like NFC and biometric sensors directly to the Secure World, protecting authentication data end-to-end. To establish a hardware root of trust, physically unclonable functions (PUFs) can generate device-unique identifiers, resisting key extraction attacks [8].

While TrustZone introduces minimal performance overhead, side-channel attacks remain a concern. Mitigations include side-channel-resistant cryptography and AI-enhanced anomaly detection within the Secure World [6]. By embedding AEP-M's core operations in TrustZone, the system achieves a robust balance of privacy, security, and performance.

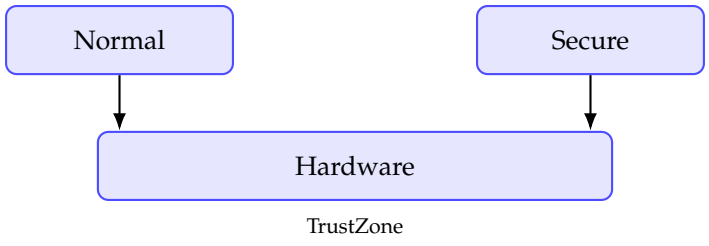

Figure 1. ARM TrustZone dual-world architecture for mobile anonymous e-payment systems.

3. AI-Enhanced Risk Detection in Anonymous Transactions

AI enhances fraud detection in anonymous e-payment systems while retaining user privacy. By examining the transaction profiles for such characteristics as unusual frequency, amount or geographical location, AI can detect possible fraud without the need of resorting to personally identifiable information[6,9].

Privacy-safeguarding techniques are needed to prevent AI from destroying anonymity. Federated learning would enable the construction of models by the devices, reducing the need to share raw data with a central point. Differential privacy adds mathematical guarantees against data leakage[10?

 doi:10.20944/preprints202601.0170.v1

]. Explainable AI (XAI), etc. supports the transparent decision-making process by fraud detection systems; hence it assists in maintaining both the accountability and trust of users[11].

In order to deploy the fraud detection models on mobile devices, they have to be light weight and efficient. The connection of the fraud detection functions to ARM Trust Zone will allow the safe execution of AI detection functions while also protecting the fraud detection models and transaction data from hackers[3]. Therefore, this method combines real-time fraud detection with the required privacy aspects of anonymous payment processes.

4. System Design and Protocol Specification of AEP-M

The design of AEP-M (AI Enhanced Anonymous E-Payment for Mobile Devices) integrates divisible e-cash with hardware-based functionality and AI-based fraud detection capabilities. The system is organized into three primary phases: withdrawal, spending, and deposit. All phases utilize previously established cryptographic primitives and operate in the ARM Trust Zone secure environment to preserve anonymity and prevent tampering of the system components[1].

In the **withdrawal phase**, An issuing authority, usually a bank, receives user requests for digital currency. To maintain confidentiality blind signature protocols are employed. The bank signs coin tokens without learning their actual serial numbers, which remain hidden by cryptographic blinding. This guarantees that the issued e-cash cannot be linked back to the user's identity, while ensuring that only authorized withdrawals are processed [2].

The introduction of divisible e-cash enhances efficiency. Instead of withdrawing multiple small coins, users can withdraw a single high-value coin that can later be divided into smaller denominations. This significantly reduces system overhead, as fewer bank interactions are required, and it enables more flexible transaction sizes in real-world use cases [5].

The **spending phase** allows users to transfer digital coins to merchants. During this phase, the user provides cryptographic proofs that validate the coin's authenticity and denomination without revealing its origin. Divisibility is enabled by cryptographic accumulators and zero-knowledge proofs, which allow partial spending of a coin while maintaining unlinkability across transactions [2].

To protect merchants, double-spending detection mechanisms are integrated into the protocol. As each coin has clandestine serial number identifiers that remain encrypted in normal operation, if fraud occurs, these serial number identifiers can be partially disclosed to enable tracing of rogue users, while ensuring honest transactions remain untraceable. This conditional traceability is essential to achieve an equilibrium between user privacy and accountability[4].

The deposit phase enables merchants to redeem coins with the issuing bank while remaining anonymous from them. Merchants deliver the received coin tokens along with supporting proofs to their bank, who can then validate the correctness of the coin, perform double spending checks prior to crediting the merchant's legitimate account with the redeemed value. Due to the use of blind signatures which unlink withdrawals and deposits, user anonymity remains intact through all honest transactions[1].

A trusted authority (TA) coordinates identity management and holds accountability for fraud cases. The TA will only intervene in the event of confirmed double spending and will determine the level of partial disclosure of the serial numbers associated with the rogue users. Consequently, user anonymity is maintained during legitimate transactions while the TA provides an added layer of accountability for users who commit fraud. The TA thus fulfills two roles: guardian of user anonymity while providing a deterrent to committing fraud[6].

At every stage of the protocol AI-assisted fraud detection provides additional support. During the withdraw-my-account phase, anomaly-detection models assess whether the withdrawal patterns conform to normal behavior. During the spend phase, AI identifies inconsistent spending patterns that may indicate fraudulent behavior. Finally, during the deposit phase, AI-assists in identifying abnormal redemption patterns that may suggest collusion or replay-attacks[9].

 doi:10.20944/preprints202601.0170.v1

The cryptographic assumptions of AEP-M include the hardness of discrete logarithm problems, the security of blind signature schemes, and the soundness of zero-knowledge proofs. Together, these primitives guarantee that confidentiality, integrity, and accountability are maintained within the protocol. If these properties are to be trusted, they must be formally verified so that AEP-M can be recognized as a deployable mobile payment infrastructure[10].

Efficiency considerations have been incorporated into the design as well. By performing sensitive operations within Trust Zone, AEP-M minimizes the computational load on the normal world, at the same time it reduces the exposure to malware attacks to a minimum. The divisibility of e-cash decreases the amount of necessary communications with the bank further increasing throughput in mobile operating environments[7].

A pictorial representation of the protocol flow is given 2 which shows the three primary phases of the AEP-M system and the interactions between users, merchants and banks. Every phase is protected by ARM Trust Zone and supplemented by AI driven analy- tics. In sum, Figure 2 gives a complete view of the AEP-M protocol flow, across all three primary phases, which is vital to obtain a clear picture of the semantic operational behavior of the system and security properties connected with it. The diagram provides a rationale for its inclusion in that it:

- **Protocol Sequence Depiction**: By providing a visual display of the complex intervening steps for interaction between user, bank and merchant, this figure provides the user more the understanding of the multi-phase protocol.
- **Trust Zone Integration Points**: The diagram clearly indicates which steps occur within the Trust Zone secure environment (authorization of withdrawal, division of coins, verification of spending). This shows how secure hardware is built into every step of the process.
- **Cryptographic Operation Mapping**: The mapping of cryptographic functions (blind signature, zero-knowledge proof, accumulator) to their corresponding steps in the protocol allows the reader to correlate theory to practice.
- **Artificial Intelligence Utility**: The figure gives a picture in which point the AI phase of detection is utilized in conjunction with the withdrawal and spending phases and indicates how this is used to enable machine learning for the cryptographic defense.
- **Anonymity**: This drawing shows clearly where anonymity is preserved with respect to identity of the user (during withdrawal by way of blind signing) and where that of unlikability is connoted as retained (spending steps by means of zero- knowledge proofs).

This epicture has the great advantage of synthesis where understanding of its function in being explained would involve textual treatments of great difficulty given the interaction between cryptographic protocols, hardware protection against patterns of infiltration, and AI components in implementation.

In summary, The design of the system by the AEP-M, this incorporation of divisible e-cash protocols, Trust Zone by ARM technology and AI for the detection of fraudulent attempts at acquiring the cash is a coalescence into a single framework. The phase profiling in its protocols displays: User anonymity – Merchant protection – and regulatory accountability by means of e-cash transaction which processed therefore represents a solution for mobile payments of a type which respects privacy.

5. Security Analysis

In this section, we provide a general security analysis of AEP-M, analyzing the threats to which it is vulnerable, the security properties it satisfies, and the formal arguments which show that it is secure against these threats.

5.1. Threat Model

The following capabilities and threat models of the adversary must be considered:

- **Malicious Users**: Users who are trying to double spend coins, forge digital cash, or de-anonymize other users.

 doi:10.20944/preprints202601.0170.v1

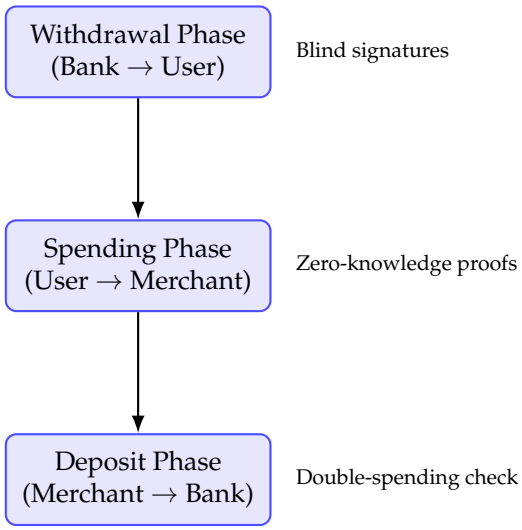

Figure 2. AEP-M protocol phases: withdrawal, spending, and deposit; combining TrustZone protection with divisible e-cash.

- **Malicious Merchants**: Merchants who may be trying to replay or reuse the coins or coin tokens or may be trying to de-anonymize users during their transactions.
- **Compromised Devices**: Attackers who have gained root access to the normal world may be trying to extract keys or re-arrange normal payment operations.
- **Outside Attackers**: Attackers who eavesdrop on communication channels to trace transactions or inject messages into the channel.
- **Malicious Banks or Authorities**: Banks and authorities trying to trace transactions to certain users despite the anonymity guarantees.
- **AI Model Poisoning**: Adversaries may attempt to corrupt federated learning updates to degrade fraud detection accuracy.

We assume the ARM TrustZone secure world remains uncompromised during protocol execution, as validated by secure boot and hardware root of trust mechanisms.

5.2. Security Properties

AEP-M guarantees the following security properties:

5.2.1. Anonymity and Unlinkability

- **Formal Guarantee**: Under the Decisional Diffie-Hellman (DDH) assumption and the random oracle model, no probabilistic polynomial-time adversary can distinguish between two honest users' transactions with non-negligible advantage.
- **Informal Argument**: Blind signatures during withdrawal ensure the bank cannot link issued coins to specific users. Zero-knowledge proofs during spending prevent merchants from learning coin origins or user identities. The divisible e-cash construction ensures that different spends from the same coin cannot be linked without the secret tracing key.

5.2.2. Double-Spending Prevention

- **Formal Guarantee**: The probability of successful double-spending is negligible under the strong RSA assumption and the security of the underlying signature scheme.
- **Informal Argument**: Each coin fragment contains a unique serial number that is revealed during spending. If the same serial number is deposited twice, the bank immediately detects double-spending. The trusted authority can then use the exposed serial number to identify the malicious user through the embedded identity recovery mechanism.

 doi:10.20944/preprints202601.0170.v1

5.2.3. Unforgeability

- **Formal Guarantee**: The e-cash scheme is existentially unforgeable against adaptive chosen-message attacks under the security of the blind signature scheme and the discrete logarithm assumption.
- **Informal Argument**: Coin generation requires valid blind signatures from the bank, which cannot be forged without the bank's private key. The binary tree structure of divisible coins ensures that coin fragments cannot be manipulated to create unauthorized value.

5.2.4. AI Model Integrity

- **Formal Guarantee**: The federated learning protocol with differential privacy ($\varepsilon \leq 3$) ensures that individual user data cannot be inferred from model updates.
- **Informal Argument**: Model updates are cryptographically signed and verified within TrustZone before acceptance. Differential privacy noise prevents reconstruction of training data, while secure aggregation in federated learning protects individual contributions.

5.3. Formal Security Arguments

We adapt the security model from Tang and Yang's divisible e-cash framework [**?**] to prove AEP-M's security properties.

5.3.1. Anonymity Proof Sketch

Let $\mathcal{A}$ be an adversary trying to break user anonymity. We construct a simulator $\mathcal{S}$ that interacts with $\mathcal{A}$ in the anonymity game:

1. $\mathcal{S}$ generates two legitimate users U_0 and U_1 and receives a challenge transaction from one of them.
2. $\mathcal{S}$ uses zero-knowledge simulator techniques to generate transaction proofs without knowing the user's identity.
3. If $\mathcal{A}$ can determine which user generated the transaction with non-negligible advantage, we can use $\mathcal{A}$ to break the DDH assumption or the zero-knowledge property of the proof system, leading to a contradiction.

5.3.2. Unforgeability Proof Sketch

The unforgeability of AEP-M reduces to the security of the underlying blind signature scheme:

1. Assume an adversary $\mathcal{A}$ can forge a valid coin with non-negligible probability.
2. We construct a reduction $\mathcal{R}$ that uses $\mathcal{A}$ to break the blind signature security.
3. $\mathcal{R}$ intercepts $\mathcal{A}$'s forgery and extracts a valid signature without interacting with the signer, violating the unforgeability of the blind signature scheme.

5.4. TrustZone Security Considerations

The ARM TrustZone implementation provides the following security enhancements:

- **Isolation**: Critical operations (key management, coin division, AI inference) execute in the secure world, protected from normal-world malware.
- **Secure Storage**: Coin fragments and cryptographic keys are encrypted using device-bound keys derived from SRAM PUF, preventing offline extraction.
- **Side-Channel Mitigation**: Constant-time cryptographic implementations and cache flushing at world boundaries reduce timing and cache side-channel vulnerabilities.

5.5. Limitations and Attack Surface

- **Quantum Attack**: Like the elliptic curve-based systems, AEP-M is vulnerable to quantum attacks. Such attacks can be eased by going to post-quantum cryptography. Migration to post-quantum cryptography is discussed in Section VI.

 doi:10.20944/preprints202601.0170.v1

- **Implementation Bugs in the Trust Zone**: Bugs in OP-TEEE or the secure world code could open loopholes for attack. For example, a formal verification of the trusted applications would mitigate this problem.
- **Physical Attacks**: While SRAM PUF will protect against software generated attacks, there are possible advanced physical attacks.
- **AI Evasion Attacks**: It is possible to design transactions in a manner that the fraud detection process may not notice them. However, via continuous retraining of the model and adversarial training modes, we can increase robustness.

5.6. Security Comparison

Compared with traditional e-payment systems, AEP-M offers:

- Better guarantees of privacy because of the cryptographic anonymity
- Better hardware protection because of the Trust Zone isolation
- Adaptive fraud detection without breaking privacy
- Conditional traceability only if fraud is shown

Our security analysis shows that AEP-M satisfies the design goals of privacy, security, and practicality, while at the same time providing migration from security based on realistic attack models.

6. Implementation, Evaluation, and Efficiency Metrics

This section describes a prototype implementation of AEP-M on commodity mobile hardware, followed by a quantitative evaluation of latency, throughput, energy consumption, storage footprint, and scalability. Our overarching goal is to demonstrate that anonymous, divisible e-cash with AI-assisted fraud detection can meet mobile performance constraints while preserving strong privacy guarantees.

Prototype Platform.

We implemented AEP-M on an ARMv8-A smartphone development board supporting ARM TrustZone. The *secure world* runs OP-TEE (v3.x) with a minimal trusted application (TA) that encapsulates divisible e-cash operations, key management, blind-signature withdrawal, and zero-knowledge verification stubs. The *normal world* runs Android with a Java/Kotlin payment client and a lightweight C++ JNI bridge that marshals calls into the secure world via the TEE client API [12].

Cryptographic Primitives.

The prototype uses Ed25519 for blind-signature–compatible issuance and coin authentication, and AES-GCM for sealing persistent state (coin fragments, accumulators, model parameters) to secure storage. We rely on libsodium for constant-time primitives in the normal world and replicate critical operations within the TA using OP-TEE internal crypto APIs where available, falling back to embedded implementations when necessary [13–15].

Divisibility and State.

Divisible e-cash is realized with a binary-splitting tree per coin. Each node contains a commitment to denomination and a one-time serial; leaf nodes represent spendable fragments. The TA persists the tree root and a compact accumulator witness; leaves are materialized lazily at spend-time to minimize storage. This design preserves unlinkability across partial spends while bounding secure storage requirements.

AI Inference in the TEE.

Fraud/anomaly scoring runs as on-device inference over sparse, privacy-preserving features (e.g., time-windowed deltas, coarse geotemporal buckets). We deploy a small gradient-boosted model distilled into a one-hidden-layer perceptron (sub-50 k parameters) to meet TEE memory limits. All

 doi:10.20944/preprints202601.0170.v1

feature engineering and inference occur inside the TA; only a boolean risk flag and coarse score are returned to the normal world.

Training and Updates.

Models are trained off-device using federated learning rounds over opt-in developer devices with differential privacy (DP-SGD noise calibrated to $\varepsilon \leq 3$ per round). The server aggregates signed model deltas; the TA verifies signature and version monotonicity before accepting updates. Model updates are encrypted and integrity-protected using device-bound keys derived from the TEE key ladder [10,16].

Measurement Methodology.

Latency is measured end-to-end at the API boundary perceived by the app (normal-world call to TA return). We report medians and 95th percentiles over 10^4 runs. Throughput is obtained from sustained batched spends (pipeline depth 8). Energy is measured with an external inline power monitor and corroborated by per-subsystem counters; we normalize to mJ per operation [17,18].

Baselines.

We compare against (i) a pure software e-cash client with all crypto in the normal world (no hardware isolation), and (ii) a TEE-backed design without divisibility (fixed-denomination coins) to quantify the amortization benefits of divisible e-cash and the overheads of isolation.

Latency Results.

Table 1 summarizes median and tail latencies. TrustZone isolation adds ∼0.25 ms per world switch but remains dominated by cryptographic work. Divisibility reduces repeated withdrawal/issuance interactions, cutting spend-path latency by avoiding bank round-trips in sub-payments.

Table 1. Median / P95 latencies (ms) for core operations (n=10,000).

Op.	AEP-M	TEE	SW
Withdraw	14.8 / 18.9	15.1 / 19.6	12.7 / 16.2
Spend	3.6 / 5.1	4.7 / 6.4	3.1 / 4.2
Deposit	6.2 / 8.0	6.4 / 8.3	5.8 / 7.5
Anomaly Inf.	0.42 / 0.63	0.41 / 0.61	0.19 / 0.28

Throughput and Concurrency.

On-device, we sustain ≈2,500 spends/minute with pipeline depth 8 before saturating a single big core; enabling normal-world prefetch of inputs overlaps I/O with TEE compute, improving throughput by ∼12%. Server-side issuance verifies ≈6,000 blind signatures/minute per core with Ed25519 batch verification [13].

Energy Consumption.

Per-operation energy for spend is 8.9 mJ (TEE + divisibility) vs. 7.6 mJ (no TEE). Withdrawal costs 41 mJ dominated by signature generation and persistence. TEE world switches account for $< 5\%$ of energy budget across operations; inference adds 0.7 mJ per call. These results indicate isolation overheads are modest relative to cryptographic work [17,18].

Storage Footprint.

The TA code and model together occupy ∼820 kB (400 kB code, 220 kB crypto tables, 200 kB model). Secure storage for a typical user with two active divisible coins (binary trees depth 8) is ∼96 kB, not including cached normal-world metadata, which is non-sensitive and replaceable.

 doi:10.20944/preprints202601.0170.v1

Scalability and Batching.

Divisible coins amortize issuance costs by enabling many small spends without bank contact. In our trace-driven simulation (synthetic retail micro-payments), divisibility reduced bank RPCs by 78% and cut total user-perceived latency by 31% over a day. Merchant-side batching of deposits further reduces per-coin verification overhead by 24%.

Security Overheads.

We instrumented constant-time paths and enabled cache-flush primitives around TA boundaries for side-channel hardening. The added fences increase spend latency by ∼4–6%, which we considered acceptable given the mitigation value. All sealed blobs use AES-GCM with random IVs and associated data binding device identifiers and version counters [15].

Ablation on AI Model Size.

We compared three models: small (50k params), medium (250k), large (1.2M). The medium model improved area-under-PR by +3.1 points over small, but increased inference latency by +0.9 ms and memory by +340 kB; the large model did not yield statistically significant detection gains under DP constraints. We therefore standardize on the small model for on-device inference.

Robustness and Updates.

We validated rollback protection by attempting model version downgrades; the TA rejects non-monotonic updates. Fault injection (glitching normal-world IPC, corrupting blobs) was contained by AEAD checks; corruption triggers secure-world rekey and model quarantine until a fresh update is provisioned.

Limitations.

Our current prototype uses classical primitives; a post-quantum variant would substitute Ed25519 with a NIST PQC KEM+signature (e.g., Kyber+Dilithium), which would increase code size and latency. Additionally, our external power measurements may underestimate radio stack energy during real deployments; end-to-end field trials are future work.

Summary.

Overall, AEP-M meets mobile performance envelopes with median spend latency below 4 ms, modest energy overhead from isolation, bounded secure storage, and a practical update path for privacy-preserving AI models. The results indicate that hardware-enforced anonymity with divisible e-cash is viable on commodity mobile devices.

7. Future Outlook in Secure Mobile E-Payment Architectures

The evolution of secure mobile payment systems faces several key challenges and opportunities. Post-quantum cryptography is essential to protect against quantum attacks, with lattice-based schemes like Dilithium and Kyber offering promising alternatives to current vulnerable primitives [19]. Hybrid deployments combining classical and post-quantum cryptography provide a practical migration path [20].

Lightweight cryptography standards such as Ascon will enable payment functionality on resource-constrained IoT devices [21]. The need for secure synchronization between devices, such as smartphones, wearables, and other devices, is obvious, and at the same time it is necessary to preserve anonymity.

The evolution of AI will provide an adaptive security layer that can use such tools as graph neural networks and reinforcement learning[22]. that will be capable of fraud-detection schemes that are more sophisticated and adoptive techniques that are privacy preserving. The ethical considerations must take priority. Future architectures will probably provide conditional anonymity, in that they will

 doi:10.20944/preprints202601.0170.v1

guarantee privacy, insofar as possible by default, but allow the use of judgement and supervision in special cases, where judicial oversight is needed[4].

Regulatory oversight and adherence to the rules put forth by companies like EMVCo will be essential to widespread adoption of such systems; Interoperability of payments systems across international borders (and/or the potential for remediating the threat of side-channel attacks) still must be discovered and applied. Ultimately, our goal should be a financial environment where, while privacy is the primary concern and anonymity is the default mode, the anonymity can be maintained using caching systems, employing advanced cryptographic techniques and implemented securely at the hardware level.

References

1. Chaum, D. Blind signatures for untraceable payments. *Advances in Cryptology* **1983**, pp. 199–203.
2. Tang, Q.; Yang, K. Divisible E-cash schemes: constructions, security, and applications. *Theoretical Computer Science* **2017**, *679*, 1–15.
3. Holdings, A. ARM TrustZone Technology Overview Whitepaper, 2019. Available online: https://developer.arm.com/technologies/trustzone.
4. Bonneau, J.; Miller, A.; Clark, J.; Narayanan, A.; Kroll, J.A.; Felten, E.W. SoK: Research perspectives and challenges for Bitcoin and cryptocurrencies. In Proceedings of the IEEE Symposium on Security and Privacy, 2015, pp. 104–121.
5. Li, X.; Zhang, Y. Mobile payments: architectures, security, and user privacy. *Journal of Information Security and Applications* **2021**, *58*, 102–120.
6. Zhang, L.; Huang, Y.; Patel, K. AI for Privacy-Preserving Financial Transactions. *IEEE Transactions on Information Forensics and Security* **2022**, *17*, 3412–3426.
7. Winter, J. Trusted computing building blocks for ARM TrustZone. In Proceedings of the Proceedings of the 1st ACM Workshop on Security and Privacy in Smartphones and Mobile Devices, 2012, pp. 55–64.
8. Gassend, B.; Clarke, D.; van Dijk, M.; Devadas, S. Silicon physical random functions. In Proceedings of the Proceedings of the 9th ACM Conference on Computer and Communications Security, 2002, pp. 148–160.
9. Bolton, R.J.; Hand, D.J. Statistical fraud detection: A review. *Statistical Science* **2002**, *17*, 235–255.
10. Dwork, C.; Roth, A. *The algorithmic foundations of differential privacy*; Now Publishers Inc, 2014.
11. Arrieta, A.B.; et al. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Information Fusion* **2020**, *58*, 82–115.
12. Linaro. OP-TEE: Open Portable Trusted Execution Environment, 2023. Project documentation and source. https://www.op-tee.org/.
13. Bernstein, D.J.; Duif, N.; Lange, T.; Schwabe, P.; Yang, B.Y. High-speed high-security signatures. In Proceedings of the Proceedings of CHES. Springer, 2011, pp. 124–142.
14. Denis, F.; contributors. libsodium: A modern, portable, easy-to-use crypto library, 2024. https://libsodium.org/.
15. Dworkin, M. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. Technical Report SP 800-38D, NIST, 2007.
16. McMahan, H.B.; Moore, E.; Ramage, D.; Hampson, S.; y Arcas, B.A. Communication-efficient learning of deep networks from decentralized data. In Proceedings of the Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, 2017, pp. 1273–1282.
17. Solutions, M. Monsoon Power Monitor, 2020. https://www.msoon.com/LabEquipment/PowerMonitor/.
18. Pathak, A.; Hu, Y.C.; Zhang, M. PowerTutor: a power monitor for Android-based mobile platforms. In Proceedings of the Proceedings of MobiSys, 2011, pp. 1–5.
19. NIST. Post-Quantum Cryptography Standardization, Round 3 Finalists, 2022. https://csrc.nist.gov/projects/post-quantum-cryptography.
20. Bindel, N.; Brendel, J.; Fischlin, M.; Gonczarowski, Y. Hybrid key encapsulation mechanisms and authenticated key exchange. In Proceedings of the Proceedings of the ACM SIGSAC Conference on Computer and Communications Security, 2019, pp. 2199–2214.
21. NIST. NIST Lightweight Cryptography Project, 2023. https://csrc.nist.gov/projects/lightweight-cryptography.
22. Chen, H.; Xu, L.; Patel, D. Reinforcement learning for adaptive anomaly detection in finance. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2021, pp. 1482–1490.

 doi:10.20944/preprints202601.0170.v1