

Elliptic Curve Cryptography for Lightweight IoT and Industrial IoT Security: A Comparative Review

Krishiyana Bhakta
M2 Data Science
Polytech Nantes, Nantes Université
Nantes, France
krishiyana.bhakta@etu.univ-nantes.fr

Sulakshyana Ghimire
M2 Smart Computing
Faculty of Sciences and Techniques
Nantes Université, Nantes, France
sulakshyana.ghimire@etu.univ-nantes.fr

Sabil Shrestha
M2 Data Science
Polytech Nantes, Nantes Université
Nantes, France
sabil.shrestha@etu.univ-nantes.fr

Abstract— Elliptic curve cryptography (ECC) is a public-key crypto primitive that is widely used because it achieves security levels that are roughly equivalent to those of the conventional public-key schemes with smaller key sizes; thus, a 256-bit ECC key can provide security levels comparable to those of a 3072-bit RSA modulus. This paper presents a selective review rather than a comprehensive survey. Three recent peer-reviewed studies have been chosen following a well-defined procedure, given that they come at different points in the process of proposing, breaking and systematizing lightweight security schemes: propose-break-systematize. The first provides a protocol that integrates ECC with a physically unclonable function (PUF) and applies a password retrieval mechanism based on the Chinese Remainder Theorem, with security analysis carried out with BAN logic and ProVerif. The second one tests a formally proven anonymous ECC authentication and key-agreement system and shows an impersonation attack on the end-device of a system in which ephemeral randomness can be leaked under an adversarial model, which is then confirmed by AVISPA. The third proposes a federated learning framework for the industrial IoT with the adoption of certificateless ECC, similarity-based update screening, and blockchain-based checkpoints. The three methods operate on different architectures, threat models, platforms and workloads, so the results of the measured performance are not directly comparable. The review shows that the security claims are critically dependent on what adversary model is assumed; that none of the three reviewed studies reports any failure of the underlying elliptic-curve arithmetic; and that the reviewed literature does not provide the same reference configuration to evaluate efficiency claims. A minimum reporting standard is proposed to fill this void.

Keywords— *Authenticated key agreement, eCK adversary model, elliptic curve cryptography, federated learning, formal verification, industrial IoT, physically unclonable function.*

I. INTRODUCTION

There are tens of billions of connected devices deployed and accessing them is more difficult than protecting conventional computers. A typical sensor node has kilobytes of memory, a low-power microcontroller and a battery that is expected to last for years; it is typically installed where no one goes back to read it and communicates over an open wireless channel, so that an adversary can listen, forge messages, or capture the device and read its memory. These limits preclude traditional PK cryptography, a 3072-bit RSA key just can't be managed on a node with a few kilobytes of RAM. The problem is avoided by elliptic curve cryptography: A 256-bit elliptic curve key is generally believed to be equivalent to a 3072-bit RSA key, and the best-known attacks on the elliptic curve discrete logarithm problem (ECDLP) are much more expensive, relative to key size, than the best-known attacks on integer factorization [1, 2].

So, the question isn't whether to use elliptic curves, it's how to create protocols around them. There are numerous published schemes for which formal security proof exists and that are based on ECC. A number of these have since been cracked by adversarial third parties and the literature then follow a cycle: a scheme is proposed, it is broken by a tougher adversary, an architectural solution sums up the lesson. The paper is not general and complete in scope review of ECC in IoT and IIoT, but rather a comparative review of three papers purposively selected from different stages of ECC in IoT and IIoT cycle, as described in Section II. Each is called by its methodology not its authors:

- An ECC-PUF authenticated key agreement (AKA) scheme with Chinese Remainder Theorem based password recovery [3];
- A published provably secure anonymous ECC AKA scheme with a security evaluation using eCK [4]; and
- Leash-FL is a certificateless ECC self-healing federated learning framework for industrial IoT [5].

II. REVIEW METHODOLOGY

The question that springs to mind upon reading a review based on three studies is "Why three studies?". The questions that guided the review, how candidate studies were identified, the number that were accepted and those that were rejected, the reasons that three studies were selected for retention, and the cost of generality in making the decision are therefore described. The questions that guided the work throughout were what security properties are new schemes based on ECC actually claiming under which adversary model; when an adversary fails, does it fail in the elliptic curve arithmetic or in the composition of the primitives to a protocol; what should be provided once an adversary is authenticated and thus indistinguishable from the legitimate participant.

The search terms "elliptic curve cryptography", "lightweight authentication", "key agreement", "IoT", "industrial IoT", "eCK", "physically unclonable function", "certificateless" and "federated learning" were used to search the Internet for candidate studies, in IEEE Xplore, MDPI, ScienceDirect and Google Scholar databases. For the discussion, the window was limited to 2023–2026 to ensure the assumptions made are hardware-constrained and can thus be assumed to be deployable. Classical literature on elliptic curves and adversary models and verification tools was introduced separately, regardless of their publication date—these

literatures do not compete with the recent literature but rather form the language that the recent studies use.

The study was valid if it was peer-reviewed and indexed, if it referenced ECC as its main public-key primitive, if it explicitly mentioned an adversarial model, and if it validated some variant of the model, formal, automated, or experimental. Studies that were not reviewed pre-prints would be left out, if ECC was not present as an explicit element in the primitive that was created and reported in the study but was not present in the security argument, and if an ECC was not even reported at all, as the core question in the review is exactly what is taken for granted, and what is demonstrated.

Three studies were selected using purposive sampling from the list of eligible studies. The sampling frame consisted of the one exemplar kept for each of the three stages (propose, break, systematize) discussed above, selected for exemplarity of the stage by the clarity with which it does so, for being recent, and because it has a complete report of its own validation. Here, the appropriateness of purposive sampling is justified because the object of study is the cycle and not the content of the field. A sample to maximize coverage would dilute exactly the contrast that this review discusses; transitioning from design, through independent evaluation, to architectural response. The outcome is a comparative case series and is presented as a case series.

The findings have three implications following directly from the design. The sample is small and purposive, and no claim is made for statistical generalization to the ECC literature; what the review provides is an analytic generalization about the cycle, not a claim of frequency for the field. The review also depends on the results as reported by the source studies, because re-implementation and re-measurement of neither of the schemes was attempted, and this is the reason for limiting the comparison to the dimensions which do not rely on measurement in Section V. Lastly, no single sample will represent its stage's internal diversity: a different pair may alter the emphasis but, as in Section VIII, the comparison of adversary models is based on one that can be duplicated by any pair.

III. RELATED WORK

Early IoT security relied on symmetric keys shared in advance. Symmetric primitives are fast and cheap, but they provide no forward secrecy, so that compromise of a key also compromises historical traffic, and key distribution scales poorly across large networks. For the first time, public-key methods were made available on small hardware with the introduction of elliptic curve cryptography, independently in [1] and [2].

ECC is performed on a set of points on an elliptic curve over a finite field, denoted P . It's based on the ECDLP: for an appropriate choice of parameters, it's hard to compute the value of k from a point $Q = kP$. The Diffie–Hellman key exchange protocol is based on elliptic curve Diffie–Hellman; elliptic curve DSA is used to create digital signatures, but in this case no private value is sent over the network. Nearly all the IoT authentication protocols presented in the literature are based on a Diffie–Hellman style exchange with a standardized curve, with subsequent protocols incorporating smart cards,

biometrics, and multi-factor authentication designs over the basic core.

Two models of adversary are found throughout this review and the difference between them is of great import to the argument. The Dolev–Yao model [6] gives the adversary total control of the channel (messages can be intercepted, modified, replayed and forged) and the cryptography is assumed to be unbreakable; usually, IoT papers add physical capture of an IoT device or smart card. The Canetti–Krawczyk model [7] models session-based key exchange and allows for long-term key compromise and session compromise. The extended form of the model, the eCK model [8] lets the adversary learn the ephemeral random values used in a session. This is not a theoretical refinement: cheap IoT devices often don't possess a hardware entropy source, and thus the randomness of their session may be weak or predictable, and that's exactly what eCK detects.

Three instruments of validation are used in the reviewed studies. BAN logic [9] is a belief logic to argue by hand that the intended parties end a run believing in a shared key; the automated symbolic verification tool ProVerif [10] works in the Dolev–Yao setting; and the automated attack trace searcher AVISPA [11] has a backend of CL-AtSe and a backend of OFMC. In all three cases, what is established symbolically is bound by the adversary capabilities that the specification grants. There are also three enabling technologies present. Physically unclonable functions exploit the fact that uncontrollable variation in semiconductor manufacturing causes each chip to respond differently to the same challenge [12]; because real responses are noisy, a fuzzy extractor is required to derive a stable key from them [13], and no usable key material remains on the device once it is powered down. Certificateless public-key cryptography [14] eliminates certificate management, key escrow, and distributes the private key between the key generation center and the device. In federated learning [15], a shared model is trained across devices that share model updates instead of data, with the update channel being a potential attack surface as well, as updates can be crafted to inject backdoors [16].

IV. MODEL ARCHITECTURE

All three techniques can be implemented on low-end machines and are distinct in the adversaries they defend against, the assets they defend, and how their claims are verified.

A. ECC–PUF AKA with CRT-based password recovery [3]

The scheme focuses on the wireless sensor network consisting of three roles: A user, sensor nodes and a trusted authority (gateway). The gateway is removed from the data path after registering the user and nodes and is only used for authentication, the session key is directly generated between the user and the sensor node (Fig. 1).




Fig. 1. System model of Approach A, showing users, the gateway node and sensor nodes. Reproduced from [3], licensed under CC BY 4.0.

The mutual authentication phase takes four messages (Fig. 2). The terminal compares the password against a local verifier and aborts if they do not match, thus avoiding the possibility of using the protocol as a password-guessing oracle. It then creates new random values and disguises the identity with a pseudonym replacing the real identity; the real identity never passes through the channel in clear, and the pseudonym changes in each session, only the gateway knows what the real identity is, having its own private key. The session key is pulled from both random contributions of the user and the sensor node, and secrets are stored at the node and tied to a PUF response; once the sensor node is powered down, no key material is of any use.




Fig. 2. The four-message mutual authentication exchange of Approach A. Reproduced from [3], licensed under CC BY 4.0.

What sets this scheme apart from similar schemes is how forgotten passwords are handled. Most schemes will have to ask a user to go through the entire registration process if he forgets his password, which is inconvenient, and is potentially problematic when implemented on a large scale as a denial-of-service vector. At enrollment the user answers a set of security questions, and a recovery secret is derived from those answers using the Chinese Remainder Theorem. That secret is embedded in the material held on the user's terminal and can be reconstructed locally, without contacting the gateway, which never stores it. The larger point is that there is no competition between hardware-protected and mathematical-protected keys – a key in flash can be recovered from a captured device, but a

key created with a PUF can't be read when power is shut down, and there is no purely algebraic construction of such a key.

B. eCK evaluation of a provably secure anonymous ECC AKA scheme [4]

This work is the opposite of what is normally done, in that it is an evaluation, rather than a proposal, of a scheme. The target is an anonymous authentication and key agreement protocol published by the protocol authors which they claim has provable session-key security, in a system of trusted authority, service provider and end devices. The assessment is done in the eCK model. Specifically, in the original proof, no mistake is claimed, but rather the assumptions used are overly favorable to the defender. If the adversary can learn ephemeral session values, then the original argument no longer covers the capabilities the adversary has.

A session key, shared with the SP, can be reconstructed analytically and algebraically (Fig. 3) by an adversary impersonating an end device. This key is fundamental for confidentiality and integrity, and is therefore a total rather than partial break, and is reproduced independently in AVISPA with both the CL-AtSe and OFMC backends.




Fig. 3. The end-device impersonation attack against the target scheme of Approach B. Reproduced from [4], licensed under CC BY 4.0.

The lesson is more than the one target. A security proof is a claim for one adversary and does not apply to an adversary more powerful than the model anticipates. This is a practical rather than a theoretical warning, as cheap IoT devices often have poor randomness.

C. Leash-FL, certificateless ECC with self-healing federated learning [5]

With this approach you will be able to handle what's left after the channel is secured. The difference between conventional authentication and federated learning is that the adversary is a participant in federated learning and can thus submit maliciously corrupted updates, inject a backdoor into the shared model, or free ride by contributing nothing and enjoying the rewards. Authentication alone addresses only the first of the relevant threat classes, which also include model poisoning, update injection, free-riding and Sybil behavior alongside

identity spoofing. The structure is built in three tiers: at the edge layer, cryptography and screening are taking place, above it, aggregation, anomaly detection and rollback (Fig. 4).




Fig. 4. Layered architecture of Approach C, showing the device, edge and aggregation tiers. Reproduced from [5], licensed under CC BY 4.0.

Three mechanisms are used. Certificateless ECC with rotating pseudonyms can authenticate a device in milliseconds with very low overheads in metadata, and the contributions of a device become unlinkable even if it participates in many rounds. Similarity-based screening checks each new update before it is aggregated and discards any new updates that are too different from the honest majority. With Blockchain based checkpoints, contamination can be detected after the fact, and the system can be rolled back to a known good state without interrupting the system's operation. Combined, these signal a paradigm change in the objective of the design from prevention to recovery: The first two methods try to prevent the adversary from getting in, the latter asks how quickly the system recovers when an adversary is already in the system. A reorientation is a requirement, and not refinement in an industrial environment, where a compromise cannot be afforded to until the next maintenance period.

V. DATASET AND EVALUATION SETTING

These three approaches have widely divergent numerical outcomes which would not methodologically be comparable, because they use fundamentally different methodologies: protocol design validated symbolically, adversarial re-analysis, and empirical machine-learning experimentation. The assessment criteria, then, are set beforehand.

Five dimensions are compared across the three approaches: (i) the adversary model assumed; (ii) the method by which the claims are validated; (iii) the class of threats addressed; (iv) the position occupied on the prevention–detection–recovery continuum; and (v) the deployment domain and class of device targeted. These five are treated as commensurable because each is stated explicitly by all three source studies, each is independent of the implementation choices made within an approach, and together they span the prevention–detection–recovery continuum. They are the only dimensions on which cross-approach statements are made in Sections VIII and IX.

Execution time, energy, communication cost and model accuracy are not compared across approaches, because no common basis for them exists: the reported figures depend on the choice of curve, the library, the processor, the compiler and the workload, none of which is shared by the three studies. Where such figures appear in Table II, they are in-study claims measured against baselines chosen by the study itself, and no cross-approach ranking is drawn from them.

Table I summarizes the evaluation setting of each approach. Two of the three are protocol studies whose claims are verified formally with symbolic tools, while the third is evaluated on datasets in the conventional machine-learning manner. That heterogeneity is itself a finding, and it is taken up in Section VIII.

TABLE I
EVALUATION SETTING OF THE THREE APPROACHES

Approach	Type of work	Application domain	Adversary assumed	Validation method
A: ECC–PUF AKA [3] (Entropy, 2026)	Protocol design	Wireless sensor networks	Dolev–Yao extended with physical device capture	BAN logic and ProVerif
B: eCK evaluation [4] (Sensors, 2025)	Security evaluation	General-purpose IoT	eCK, including leakage of ephemeral session randomness	Analytic eCK proof plus AVISPA (CL-AtSe, OFMC)
C: Leash-FL [5] (Sensors, 2025)	System architecture	Industrial IoT with federated learning	Authenticated but malicious participant (insider)	Experiments on intrusion, anomaly and vision datasets

Unlike machine learning, IoT authentication based on ECC has no consensus benchmark suite. Baselines and hardware are self-selected, curve parameters are often unspecified, and efficiency claims therefore cannot be verified or reproduced by a third party.

VI. EXPERIMENTAL RESULT

Each study is reported in Table II. As in Section V, these are within-study claims, and they are not normalized nor ranked against one another.

TABLE II
RESULTS REPORTED BY EACH APPROACH

Approach	Outcome as reported by the source study, within its own setting
A [3]	Reduction of computation and communication cost when compared to the schemes. Implements resistance to impersonation, replay, offline password guessing, node cloning and stolen smart card attacks. The identities are not leaked when they are anonymous and different each session.
B [4]	An end device impersonation attack is shown which is followed by a compromise of the session key. Under eCK, the analyzed scheme does not achieve session-key security. Confirmed by AVISPA with both the CL-AtSe and the OFMC backends.
C [5]	Even with 50% malicious clients, accuracy is maintained at 85% and higher. The success rate for a backdoor attack is less than 5% after four recovery rounds. Accurate up to 3 times quicker than screening only methods. Joins and leaves in less than 50ms, re-admission in less than 60ms, forward and backward secrecy maintained.

The most notable of these is that Approach C remains largely resistant at up to 50% malicious clients, a level at which

many robust aggregation schemes have already degraded. The figure is, however, dependent on the data sets used in that study, the number of clients and attack strategy used, and should be interpreted as proof of the potential of the mechanism, and not as a constant to be transferred to other situations.

VII. STRENGTH AND LIMITATION

The strengths and weaknesses of each approach, like architecture, are outlined in Table III without referring to the numerical results above.

TABLE III
STRENGTHS AND LIMITATIONS OF THE THREE APPROACHES

Approach	Strengths	Limitations
A [3]	- Mathematical / hardware protection combination - Online guessing is prevented by local password check. - It is possible to recover a password without re-registering. - Verified using two separate procedures - Gateway is not on the data path.	- Only tested with Dolev-Yao, not the stronger eCK model - Such PUF responses are noisy in real hardware and require additional correction circuits which are not fully costed. - The answers to security questions can be easily guessed
B [4]	- Implements a more realistic and powerful attack model - Verifies the attack using a separate program - The lesson is transferable to other than one scheme - Provides positive design suggestions	- Makes no alternative protocol available - Applies only to one target scheme - Does not consider attacks on a hardware level
C [5]	- Detects and responds to an authenticated attacker - Recovers automatically, without shutting down. - No need to deal with certificate management and key escrow. - The blockchain offers regulators an audit trail	- Blockchain is too intensive to run on the smallest sensors and takes on a gateway / edge tier. - A conscientious intruder could create an update that gets through. - May incorrectly label honest devices that contain different data to most devices

VIII. CRITICAL ANALYSIS AND DISCUSSION

When taken together, and in the commensurable dimensions of Section V, the three approaches allow for several observations which none of them makes alone. The outcome depends on the adversary, not on the proof. Approach A and the scheme targeted by Approach B were each validated by two independent methods, and the scheme that Approach B breaks had already been formally proven secure before the attack was found. The difference in outcome is not a matter of rigor or competence in the analysis, but of the capabilities granted to the adversary in each case. One adversary controls the channel (Dolev-Yao); the other additionally learns the ephemeral randomness of a session (eCK); everything else follows from that single choice. As a result, with commodity IoT hardware, with weak or no entropy sources, it makes sense to use eCK as a default option instead of an optional strengthening. It also suggests a different way of presenting: the threat model should not be contained in the preliminaries, but rather in the title in terms of queries to the oracle that are allowed and those that are not. A reader otherwise cannot understand what the word "secure" in the title means. In Approach A the gap is therefore

one of evaluation rather than of construction: it has been proven to be secure in the face of the adversary model specified in the design, but its security in the face of the stronger eCK model is not demonstrated.

Errors do not occur due to curve arithmetic but due to composition. There are no reported defects in elliptic curve arithmetic in any of the three studies. Each failure mentioned is related to composing primitives into a protocol: which values are to be put into a hash, which are not, which quantity to which identity is bound, and which freshness guarantee(s) to enforce. So, what is not failing now is something that is better optimized for when it's ready, rather than being done by the micro-optimization of scalar multiplication.

It becomes necessary, but not sufficient, to use authentication. Both of approach A and B have a common ground: the enemy is the external, and the goal is to keep the enemy out. Approach C does away with that premise. A successfully authenticated participant may still cause damage to the system by interacting with a shared artefact, like a learned model. The series of prevention, detection and recovery processes across the three therefore is not an accident of selection, but a response to what collaborative IoT workloads need.

A definite remedy is required for the benchmark gap. The heterogeneity shown in Table I is not simply inconvenient: it leaves the efficiency arguments that motivate lightweight cryptography unverifiable. If a scheme reports a lower cost than baselines chosen by the scheme, hardware on which the scheme ran, and curve parameters that the scheme refuses to disclose, then an independent reader cannot confirm, refute or reuse the result. Two concrete steps are proposed to close the gap. The first is a minimum reporting requirement: Table IV outlines eight items that a lightweight ECC authentication paper should report, items that are inexpensive for authors to provide and sufficient for readers to reconstruct the comparison. Two of those items (R3 and R4) are deliberately hardware-independent, so that they remain comparable across different boards.

TABLE IV
PROPOSED MINIMUM REPORTING STANDARD FOR LIGHTWEIGHT ECC AUTHENTICATION SCHEMES

Reporting item	What should be stated
R1 Adversary model	Model (Dolev-Yao, CK, eCK) declared in the abstract, with oracle queries granted and withheld listed explicitly.
R2 Curve and parameters	Named curve and field size, point representation, and the hash function used.
R3 Operation counts	Scalar multiplications, point additions, hash and symmetric operations, counted per role and independent of hardware.
R4 Message cost	Number of messages, bytes per message and direction, and total handshake bytes.
R5 Platform	Processor and clock, RAM and flash footprint, cryptographic library and version, compiler and optimization flags.
R6 Energy	Measured or modelled energy per authentication session, which governs battery life more directly than latency.
R7 Baselines	Why each comparison scheme was chosen, and whether its figures were re-measured or quoted from its own paper.
R8 Artefacts	Public release of the ProVerif, AVISPA or Tamarin scripts and of the implementation used for timing.

The second step is a common reference configuration (CRC) consisting of a small set of reference platforms (representing the realistic deployment range, that is, 8- or 16-bit microcontroller class, ARM Cortex-M class device, single-board Linux gateway), a fixed curve (that is, secp256r1, a named library version, and published reference protocols to which new schemes report). A shared harness, like the common datasets used to ensure comparability of machine-learning results could transform claims of efficiency to measurable quantities, and venues could help drive its adoption by mandating that the artefacts for item R8 be provided as a prerequisite for publication.

It is for all three that the post quantum horizon is relevant. If a quantum computer could be built to run Shor's algorithm, it would solve the ECDLP [17] and traffic that is recorded today could be decrypted later. This is a present procurement matter, not a future one, since the equipment used for industrial IoT can be in use for 10 years or more. There are now standardized post-quantum mechanisms available [18] but their keys and signatures are significantly larger than for elliptic curve mechanisms, just what ECC was designed to prevent. While these security outcomes are possible, the specific case of achieving post-quantum security within a sensor-class energy and bandwidth budget or staging a hybrid migration for devices already deployed remains underexplored and is the most impactful open problem for this line of work.

IX. CONCLUSION

For constrained IoT and industrial IoT devices, the efficiency of elliptic curve cryptography relative to RSA is both settled and important, since every transmitted byte costs battery life. The remaining difficulties lie elsewhere. The design approach of Approach A demonstrates that it is possible to combine the features of mathematical protection with hardware-based protection and simultaneously address a usability issue that is usually overlooked by most scheme designs, namely password recovery, albeit under a less threatening adversary than the deployment conditions warrant. Approach B proposes no replacement protocol, but establishes that a security proof does not hold against every adversary; it holds only against the adversary the proof assumes. Approach C shows that once IoT systems learn collectively, authentication ceases to be sufficient, since a correctly identified device can still corrupt the shared model; the price of its answer is an additional blockchain tier that the smallest hardware cannot carry. For the sake of conciseness, the comparison is summarized in Table V along the commensurable dimensions given in Section V.

TABLE V
SYNTHESIS OF THE THREE REVIEWED APPROACHES

Aspect	A: ECC-PUF AKA [3]	B: eCK evaluation [4]	C: Leash-FL [5]
Approach	ECC with PUF binding for mutual authentication and key agreement	Re-analysis of a published ECC scheme under a stronger adversary	Certificateless ECC with update screening and blockchain checkpoints

Aspect	A: ECC-PUF AKA [3]	B: eCK evaluation [4]	C: Leash-FL [5]
Key focus	Lightweight mutual authentication with usable password recovery	Showing that the assumed adversary, not the proof, decides security	Resilience and recovery once an adversary is already inside
Adversary assumed	Dolev-Yao plus physical capture	eCK with leaked session randomness	Authenticated but malicious participant
Innovation	CRT-based password reset; verifier resistant to offline guessing	Concrete impersonation attack on a "provably secure" scheme	Similarity screening plus blockchain rollback for automatic recovery
Threats addressed	Impersonation, replay, offline guessing, node cloning, verifier theft	Session-key compromise via end-device impersonation	Identity spoofing, model poisoning, update injection, free-riding, Sybil
How checked	BAN logic and ProVerif	eCK analysis and AVISPA	Experiments on three dataset families
Lifecycle stage	Prevention	Detection of failure	Recovery after compromise

The three approaches are related and logically follow one another from prevention to failure detection to success recovery. Four directions follow. For constrained systems, the stronger adversary models, and eCK in particular, should be the default assumption rather than an optional strengthening. The minimum reporting standard of Table IV, in conjunction with a reference configuration agreed by the parties, should be followed, to make efficiency claims verifiable. Analysis should extend below the protocol level to the hardware and side-channel behavior of real deployments, including the error-correcting logic that PUF-based designs require. Lastly, migration routes to post-quantum primitives will have to be designed now for equipment that is expected to last well beyond the time for which the ECDLP is assumed to be hard. Mathematics is not the difficult part; the composition and the evaluation are.

REFERENCES

- [1] N. Kobitz, "Elliptic curve cryptosystems," *Mathematics of Computation*, vol. 48, no. 177, pp. 203–209, 1987.
- [2] V. S. Miller, "Use of elliptic curves in cryptography," in *Advances in Cryptology — CRYPTO '85*, LNCS vol. 218, Springer, 1986, pp. 417–426.
- [3] Y. Yu, K. Wei, K. Qi and W. Wu, "Privacy-preserving ECC-based AKA for resource-constrained IoT sensor networks with forgotten password reset," *Entropy*, vol. 28, no. 2, art. 185, 2026.
- [4] K. Park, M. Kim and Y. Park, "Security evaluation of provably secure ECC-based anonymous authentication and key agreement scheme for IoT," *Sensors*, vol. 25, no. 1, art. 237, 2025.
- [5] M. M. Salim, F. Naaz and K. Choi, "Lightweight ECC-based self-healing federated learning framework for secure IIoT networks," *Sensors*, vol. 25, no. 22, art. 6867, 2025.
- [6] D. Dolev and A. Yao, "On the security of public key protocols," *IEEE Transactions on Information Theory*, vol. 29, no. 2, pp. 198–208, 1983.
- [7] R. Canetti and H. Krawczyk, "Analysis of key-exchange protocols and their use for building secure channels," in *Advances in Cryptology — EUROCRYPT 2001*, LNCS vol. 2045, Springer, 2001, pp. 453–474.
- [8] B. LaMacchia, K. Lauter and A. Mityagin, "Stronger security of authenticated key exchange," in *ProvSec 2007*, LNCS vol. 4784, Springer, 2007, pp. 1–16.
- [9] M. Burrows, M. Abadi and R. Needham, "A logic of authentication," *ACM Transactions on Computer Systems*, vol. 8, no. 1, pp. 18–36, 1990.
- [10] B. Blanchet, "An efficient cryptographic protocol verifier based on Prolog rules," in *Proc. 14th IEEE Computer Security Foundations Workshop (CSFW)*, 2001, pp. 82–96.

- [11] A. Armando et al., "The AVISPA tool for the automated validation of internet security protocols and applications," in *Computer Aided Verification (CAV 2005)*, LNCS vol. 3576, Springer, 2005, pp. 281–285.
- [12] B. Gassend, D. Clarke, M. van Dijk and S. Devadas, "Silicon physical random functions," in *Proc. 9th ACM Conf. on Computer and Communications Security (CCS)*, 2002, pp. 148–160.
- [13] Y. Dodis, R. Ostrovsky, L. Reyzin and A. Smith, "Fuzzy extractors: How to generate strong keys from biometrics and other noisy data," *SIAM Journal on Computing*, vol. 38, no. 1, pp. 97–139, 2008.
- [14] S. S. Al-Riyami and K. G. Paterson, "Certificateless public key cryptography," in *Advances in Cryptology — ASIACRYPT 2003*, LNCS vol. 2894, Springer, 2003, pp. 452–473.
- [15] B. McMahan, E. Moore, D. Ramage, S. Hampson and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. AISTATS*, 2017, pp. 1273–1282.
- [16] E. Bagdasaryan, A. Veit, Y. Hua, D. Estrin and V. Shmatikov, "How to backdoor federated learning," in *Proc. AISTATS*, 2020, pp. 2938–2948.
- [17] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM Journal on Computing*, vol. 26, no. 5, pp. 1484–1509, 1997.
- [18] National Institute of Standards and Technology, "Module-lattice-based key-encapsulation mechanism standard," *FIPS 203*, 2024.