Asuu Network Core Protocol: Provable Safety, Provenance, and Consensus for Clinical Artificial Intelligence

Derviş Berkcan Özdemir, MD, cand.
berkcan@asuucare.com

Abstract

We formalize the Asuu Network protocol, a cryptographic control plane designed to guarantee the safety, provenance, and liveness of clinical artificial intelligence without reliance on trusted proprietary intermediaries. The protocol unifies distributed consensus, cross-epoch formal verification, and post-quantum cryptographic primitives to formulate an immutable ledger of medical AI execution. We introduce Proof of Clinical Authority (PoCA), a partially synchronous Byzantine Fault Tolerant consensus mechanism that bounds Sybil identities through federated credential issuance and achieves consensus proportional to verifiable clinical engagement. To constrain neural network behavioral drift under adversarial conditions, we specify a Simplex architecture paired with Signal Temporal Logic over clinical parameters, embedding automated calibration drift suspension as a hard protocol invariant. Furthermore, we formalize a phased post-quantum migration utilizing a composite ECDSA and ML-DSA-65 signature topology to secure the longitudinal integrity of medical inferences against a "harvest now, decrypt later" paradigm. The paper rigorously deduces the Five Theorem Backbone: cross-epoch safety, end-to-end provenance completeness, update-corridor soundness, privacy-budget enforcement, and safety-suspension bounds. The result is a mathematically verifiable infrastructure for mission-critical medical intelligence.

Contents

1 Introduction **3**
2 Notation and Cryptographic Preliminaries **5**
2.1 Notation 5
2.2 Hash Functions and Signature Security 5
2.3 Polynomial Commitments and Zero-Knowledge Proofs 5
2.4 Fundamental Hardness Assumptions 6
3 Proof of Clinical Authority (PoCA) **7**
3.1 Architectural Validator Model 7
3.2 Federated Sybil Resistance 7
3.3 Rigorous Safety and Liveness Bounds 8
4 Post-Quantum Cryptographic Migration **9**
4.1 Composite Signature Envelopes 9
4.2 Harvest-and-Forge Window Characterizations 9
5 Formal Verification and Clinical Safety Envelopes **11**
5.1 The Clinical Simplex Verification Topology 11
5.2 Continuous Calibration Suspension Bounds 11
6 Synthesis of the Five Theorem Backbone **13**
7 Conclusion **14**

1 Introduction

Medical artificial intelligence currently executes within a deeply centralized paradigm where proprietary vendors monopolize models, data pipelines, and computational infrastructure. While these platforms satisfy functional clinical thresholds, their architectural reliance on perimeter security and third-party trustworthiness introduces existential vulnerabilities to sovereign medical systems. When the provenance of an inference, the integrity of a targeted patient population distribution, and the immutability of an audit trail are unilaterally controlled by a vendor, the system lacks mathematical distinguishability between a benign network update and a malicious calibration drift.

The deployment of clinical artificial intelligence at scale necessitates a shift from trust-based reliance to cryptographically verifiable proofs of safety and accountability. Modern healthcare applications (encapsulating clinical decision support, rapid prior authorization evaluations, and continuous biometric monitoring) interface with deterministic systems like HL7 FHIR and CDS Hooks at the application layer. However, the underlying intelligence that fuels these clinical executions possesses no intrinsic capacity to self-attest its provenance or guarantee compliance with safety constraints. Operational failures, adversarial model tampering, or "harvest now, decrypt later" quantum data interception threaten to transform single-vendor vulnerabilities into systemic clinical crises.

This paper formally specifies the Asuu Network protocol: a peer-to-peer cryptographic ledger engineered strictly to act as the verifiable control plane linking high-assurance clinical executions to decentralized verification mechanisms. Patient data remains strictly localized off-chain. The protocol facilitates only the cryptographic hashes, zero-knowledge proofs, and deterministic gating algorithms necessary to ensure that clinical actions are fully auditable, temporally ordered, and cryptographically anchored to identity.

To achieve this, we delineate the Asuu Network through a suite of formalisms resolving the tensions between network liveness, safety verification, and long-term cryptographic resilience. Our contributions consist of five major theoretical pillars structuring the protocol's backbone:

1. **Consensus and Identity:** We introduce Proof of Clinical Authority (PoCA), replacing anonymous mining or purely economic staking with a weighted Byzantine Fault Tolerant mechanism bound to externally verified healthcare credentials.
2. **Mathematical Safety Envelopes:** We decouple the unpredictable nature of neural network approximations from the requirement for rigorous safety by employing a Simplex architecture monitored via Signal Temporal Logic (STL) bounds. We prove limits on detection delays for calibration drift leading to automated suspension.
3. **Post-Quantum Resilience:** We model a concrete migration pathway utilizing composite signatures composed of ECDSA (Ed25519) and ML-DSA-65, calculating the precise bounds and constraints of a harvest-and-forge window prior to a purely post-quantum protocol consolidation.
4. **End-to-End Provenance Completeness:** We formalize the inference record mapping architecture over Merkle trees, ensuring that every discrete step (model registration, inference event, human oversight authorization, and clinical delivery) is indelibly linked and tamper-evident.

5. **The Core Theorem Backbone:** We systematically unite PoCA safety, inference provenance, model update corridors, data privacy limits, and automated suspension into a single harmonized network model.

Throughout this exposition, all proofs are stated within the concrete security framework. The system guarantees state finality and safety provided that an adversarial coalition cannot control a weight threshold exceeding one-third of the active clinical authority, operating under the conditions of a partially synchronous network.

2 Notation and Cryptographic Preliminaries

We establish the notation, definitions, and hardness assumptions utilized throughout this formalization. All proofs are specified within the concrete security framework unless otherwise delineated.

To architect a trustless medical infrastructure, we compose standard cryptographic primitives whose security paradigms have been extensively verified. The definitions strictly establish the necessary functional thresholds each primitive must exceed for the Asuu protocol to guarantee its proposed systemic safety properties.

2.1 Notation

Let λ denote the security parameter. A function $f : \mathbb{N} \to \mathbb{R}$ is *negligible* if for every positive polynomial $p(\cdot)$ there exists N such that for all $n > N$, $|f(n)| < 1/p(n)$; we write $f(\lambda) \leq \mathrm{negl}(\lambda)$. A probabilistic algorithm $\mathscr{A}$ is *efficient* if it operates in time $\mathrm{poly}(\lambda)$; we denote this efficiency as $\mathscr{A} \in \mathrm{PPT}$. We denote uniform random sampling by $x \xleftarrow{\$} \mathscr{S}$. For any integer n, the notation $[n]$ represents the set $\{1, 2, \ldots, n\}$. We signify string concatenation with $\|$.

2.2 Hash Functions and Signature Security

Definition 2.1 (Collision-Resistant Hash Family)**.** A keyed hash family $\mathscr{H} = \{H_s : \{0,1\}^* \to \{0,1\}^\lambda\}_{s \in \mathscr{K}}$ is *collision-resistant* if for all adversaries $\mathscr{A} \in \mathrm{PPT}$:

$$\mathrm{Adv}^{\mathrm{CR}}_{\mathscr{H},\mathscr{A}}(\lambda) = \Pr\big[H_s(x) = H_s(x') \wedge x \neq x' : s \xleftarrow{\$} \mathscr{K},\ (x,x') \leftarrow \mathscr{A}(s)\big] \leq \mathrm{negl}(\lambda).$$

Lemma 2.2 (Second-Preimage Resistance)**.** *If $\mathscr{H}$ is collision-resistant, then $\mathscr{H}$ guarantees second-preimage resistance. Given an input x, no computationally efficient adversary can identify a distinct $x' \neq x$ such that $H_s(x') = H_s(x)$ except with an exceedingly negligible probability. (A finder $\mathscr{A}'$ successfully generating x' would directly construct a collision pair (x,x'), explicitly violating Definition 2.1).*

Definition 2.3 (EUF-CMA Security)**.** A signature scheme $\Sigma = (\mathrm{KeyGen}, \mathrm{Sign}, \mathrm{Verify})$ is *existentially unforgeable under adaptive chosen-message attack* (EUF-CMA) if for all $\mathscr{A} \in \mathrm{PPT}$:

$$\mathrm{Adv}^{\text{EUF-CMA}}_{\Sigma,\mathscr{A}}(\lambda) = \Pr\left[\mathrm{Verify}(pk, m^*, \sigma^*) = 1 \wedge m^* \notin Q : \begin{array}{l}(pk,sk) \leftarrow \mathrm{KeyGen}(1^\lambda) \\ (m^*,\sigma^*) \leftarrow \mathscr{A}^{\mathrm{Sign}(sk,\cdot)}(pk)\end{array}\right] \leq \mathrm{negl}(\lambda),$$

where Q represents the total set of messages previously submitted to the signing oracle.

2.3 Polynomial Commitments and Zero-Knowledge Proofs

Definition 2.4 (Pedersen Commitment)**.** Let $\mathbb{G}$ be a cyclic group of prime order q defined with orthogonal generators g, h where the discrete logarithm $\log_g h$ is unknown. The Pedersen commitment uniquely mapping to message $m \in \mathbb{Z}_q$ given randomness $r \xleftarrow{\$} \mathbb{Z}_q$ is constructed as $C = g^m h^r$. By leveraging the discrete logarithm problem context (Assumption 2.6), the commitment provides perfect statistical hiding and computational computational binding.

Definition 2.5 (Non-Interactive Zero-Knowledge Argument of Knowledge)**.** A non-interactive argument verification system relative to an NP relation $\mathscr{R}$ relies on $(\text{Gen}, \text{Prove}, \text{Verify})$ fulfilling:

(i) **Completeness.** For any strictly valid pair $(x,w) \in \mathscr{R}$: $\Pr[\text{Verify}(vrs, x, \text{Prove}(crs, x, w)) = 1] = 1$.
(ii) **Knowledge Soundness.** For every PPT prover $\mathscr{P}^*$ there exists a distinct PPT extractor $\mathscr{E}$ dictating that: $\Pr[\text{Verify}(vrs, x, \pi) = 1 \wedge (x,w) \notin \mathscr{R}] \leq \text{negl}(\lambda)$.
(iii) **Zero-Knowledge.** There exists a functioning PPT simulator $\mathscr{S}$ parameterized with trapdoor τ such that for all $(x,w) \in \mathscr{R}$, the distribution arrays $\{\text{Prove}(crs, x, w)\}$ and $\{\mathscr{S}(vrs, \tau, x)\}$ maintain strict computational indistinguishability.

2.4 Fundamental Hardness Assumptions

Assumption 2.6 (Discrete Logarithm Problem (DLP))**.** For a given cyclic group $\mathbb{G}$ possessing prime order q matched to generator g, we observe: $\Pr[\mathscr{A}(g, g^x) = x : x \xleftarrow{\$} \mathbb{Z}_q] \leq \text{negl}(\lambda)$ holds true against all $\mathscr{A} \in \text{PPT}$.

Assumption 2.7 (Honest Supermajority)**.** The protocol operates strictly on the condition where the aggregate fractional weight of Byzantine validators calculates to $f < n/3$, positioning n as the sum of totally recognized active validators. Expressed conversely, honest validators must perpetually maintain ownership of greater than $2n/3$ of the overarching network clinical authority weight matrix.

Assumption 2.8 (Module Learning With Errors (M-LWE) and Module Short Integer Solution (M-SIS))**.** Securing long-tail unforgeability mandates lattice problem resilience. For module rank $d \geq 2$, undisclosed element $\mathbf{s} \xleftarrow{\$} R_q^d$, and uniformly sampled polynomial $\mathbf{A} \xleftarrow{\$} R_q^{k \times d}$, the resulting vector space $(\mathbf{A}, \mathbf{As} + \mathbf{e})$ stands computationally indistinguishable from arbitrary uniform values covering $R_q^{k \times d} \times R_q^k$ for all $\mathscr{A} \in \text{PPT}$ (M-LWE). Complementarily, determining a nonzero vector $\mathbf{x} \in R_q^d$ where $\|\mathbf{x}\|_\infty \leq \beta$ fulfilling $\mathbf{Ax} = \mathbf{0} \pmod{q}$ must remain probabilistically unfeasible (M-SIS). Both assertions reduce structurally to hardest-case evaluations against module lattices.

3 Proof of Clinical Authority (PoCA)

Permissionless staking architectures suffer structurally from plutocratic capture vulnerabilities, introducing unacceptable regulatory risks into clinical contexts. Proof of Clinical Authority (PoCA) mitigates this by restricting validator privileges exclusively to verified institutional identities, weighting block derivation probability dynamically against proven longitudinal healthcare activity metrics rather than pure unanchored collateral. A Byzantine Fault Tolerant protocol processes these weights under partial synchrony bounds.

3.1 Architectural Validator Model

Definition 3.1 (Clinical Authority Weight)**.** For any participating consensus node corresponding to validator $i \in [n]$, the normalized clinical authority weight is derived via:

$$W_i = \min\left(\alpha\left(\frac{S_i}{S_{\text{total}}}\right) + \beta\left(\frac{V_i}{V_{\max}}\right) + \gamma R_i,\ W_{\text{cap}}\right)$$

reflecting S_i staked native liquidity, validated institutional care volume index V_i, and an empirically observed continuous reputation reliability coefficient $R_i \in [0,1]$. Crucially, the constants satisfy $\alpha + \beta + \gamma = 1$. The protocol permanently fixes an upper limit $W_{\text{cap}} \in (0,1]$ bounding absolute voting supremacy, typically constraining maximal single entity consolidation to five percent of the overall network weight parameter.

Proposition 3.2 (Plutocracy Capture Deterrence)**.** *A robust W_{cap} acts as an absolute counter-plutocratic circuit limit. It proves computationally and economically impossible for a high-resource institution to subvert consensus mechanics solely by inflating token variables S_i, as marginal authority strictly plateaus when approaching W_{cap}. Corrupting a quorum certificate uniformly requires coordinating independent, adversarial institutional clusters, not mere wealth.*

Definition 3.3 (Quorum Certificate Generation)**.** A valid quorum certificate QC authorizing block B positioned at chain height h constitutes a mathematical aggregate of unforgeable signatures originating from a dynamic validator subset S enforcing weight superiority mapping to $\frac{2}{3}$:

$$QC(B,h) = \{(i,\sigma_i)\}_{i\in S} \quad \text{requiring precisely} \quad \sum_{i\in S} W_i > \frac{2}{3}\sum_{j=1}^{n} W_j.$$

3.2 Federated Sybil Resistance

Classical Sybil resistance depends on cost barriers explicitly tied to computational waste or token velocity. PoCA achieves strict resilience implicitly through mathematically authenticated offline medical identity mapping managed by a multi-jurisdictional federation layer.

Definition 3.4 (Federated Credential Issuance)**.** A foundational credential-issuance ensemble $\mathbb{CA} = \{CA_1, \ldots, CA_m\}$ comprises accredited medical-credentialing bodies functioning as distributed roots of trust. Validator credentials acquire network significance exclusively bounded by a (t,m)-threshold signature structure, preventing a singular corrupted jurisdictional board from minting rogue adversarial validator instances.

Lemma 3.5 (Injective Identity Defense)**.** *By maintaining a strict bijective injection bounded sequentially by EUF-CMA threshold signatures, an adversary plotting a Sybil node attack effectively reduces the effort to a threshold key compromise covering independent physical authorities. In all alternative permutations, splitting authentic staked collateral redundantly across k simulated node identities induces negative authority leverage under Definition 3.1 as R_j parameters degrade. Sybil inflation behaves strictly sub-linearly.*

3.3 Rigorous Safety and Liveness Bounds

Lemma 3.6 (Quorum Node Intersection Constraint)**.** *Let $W_{\text{total}} = \sum_{j=1}^{n} W_j$. Any randomized sampling of two generated quorum certificates QC_1 and QC_2 guarantees intersecting weight metrics mathematically exceeding the threshold baseline of exactly $W_{\text{total}}/3$. Employing the inclusion-exclusion principle mandates:*

$$W(S_1 \cap S_2) = W(S_1) + W(S_2) - W(S_1 \cup S_2) > \frac{2W_{\text{total}}}{3} + \frac{2W_{\text{total}}}{3} - W_{\text{total}} = \frac{W_{\text{total}}}{3}.$$

Theorem 3.7 (BFT Safety Under Decentralized Partition)**.** *Conditioned rigorously upon the Honest Supermajority boundary established in Assumption 2.7, no dual collection of honest validator subsets will finalize an inconsistent block state positioned identically at coordinate height h.*

Proof rationale: *Suppose for contradiction that disjoint blocks $B \neq B'$ both achieve finalization mappings corresponding to QC_B and $QC_{B'}$. Following Lemma 3.6, calculating overlap weight thresholds forces an intersection minimum containing strictly more capacity than the available Byzantine corruption threshold. Consequently, a minimum of one honest execution node must concurrently validate overlapping contradictory chains. Since executing non-equivocating code is the defining baseline constraint for honest agents, this implies zero probability for contradiction execution. Thus, strict equivalence ($B = B'$) remains permanent across partition bounds.*

Theorem 3.8 (Temporal Liveness Resolution under Partial Synchrony)**.** *Assuming an underlying state of partial synchrony containing an unpredictable absolute Global Stabilization Time GST, following which maximum delay bounds constrain to Δ, liveness asserts its availability continuously post-GST. Through structured iterative gossiping overlays scaling logarithmically at $O(\log n)$, node clusters finalize block outputs within continuous constant phases limited optimally by $O(\Delta)$. Any temporary halt executed securely during network fragmentation resolves mathematically devoid of hard forks once synchrony pathways heal.*

4 Post-Quantum Cryptographic Migration

Medical intelligence architectures harbor a uniquely severe sensitivity to chronological attack vectors. A genomic inference evaluated currently retains its confidentiality and authenticity requirements over multi-decade spans. A sophisticated adversarial framework executing a "harvest now, decrypt later" (HNDL) strategy utilizing cryptographically relevant quantum computers (CRQC) invalidates any classical signature or key-encapsulation primitive. The protocol implements an explicit phased migration leveraging FIPS 204 (ML-DSA) and FIPS 205 (SLH-DSA) to permanently secure both historical consensus attestations and active block derivations.

4.1 Composite Signature Envelopes

Definition 4.1 (Phased Composite Validation Layer)**.** The Asuu network adopts a layered transitional composite signature framework Σ_{comp} merging a classical scheme Σ_1 alongside a lattice-based post-quantum scheme Σ_2. The operational lifecycle dictates:

- **Phase 1 (Hybrid State):** Simultaneous deployment of Σ_1 (Ed25519) and Σ_2 (ML-DSA-65). Verification mathematically necessitates valid signature arrays across both schemas.
- **Phase 2 (Post-Quantum Primary State):** Validation restricts strictness to Σ_2, pivoting Σ_1 into a non-required, strictly optional metric to bridge non-upgraded observer node synchrony.
- **Phase 3 (Post-Quantum Absolute State):** Deletion of Σ_1 compatibility algorithms, replacing the secondary failover position with a hash-based SLH-DSA construction to construct robust algorithmic diversity entirely detached from lattice hardness assumptions.

Theorem 4.2 (Composite EUF-CMA Invulnerability)**.** *The composite layer Σ_{comp} exhibits strict EUF-CMA threshold security directly bounded by the most secure element within its aggregate formulation. Let $\mathscr{A}$ function as a structurally bounded PPT adversary:*

$$\text{Adv}^{\text{EUF-CMA}}_{\Sigma_{\text{comp}},\mathscr{A}}(\lambda) \leq \min\Big(\text{Adv}^{\text{EUF-CMA}}_{\Sigma_1,\mathscr{B}_1}(\lambda),\ \text{Adv}^{\text{EUF-CMA}}_{\Sigma_2,\mathscr{B}_2}(\lambda)\Big).$$

Proof rationale: *Forgery necessitates creating a mathematically sound pair $(m^*, \sigma^* = (\sigma_1^*, \sigma_2^*))$ successfully circumventing both* Verify *constraint checks simultaneously. Therefore, overcoming ML-DSA-65 hardness directly equates to breaking the entire bound, irrespective of Shor's algorithm effectively neutralizing the ECDSA matrix layer.*

4.2 Harvest-and-Forge Window Characterizations

Given that NIST-standardized lattice schemes (like ML-DSA) exhibit massive bandwidth inflation compared to pairing-based aggregatable signatures (BLS12-381), native quorum certificate aggregation presents computational limits inside restricted environments. The protocol isolates this by constructing continuous periodic ML-DSA anchor points covering classical BLS sequences.

Definition 4.3 (Post-Quantum Anchoring Variable)**.** Every chronological interval of k blocks precisely triggers the entire decentralized validator network to co-sign a generalized ML-DSA-65 cryptographic attestation $\sigma_{\text{PQ}}(h, H(chain[h-k{:}h]))$. This signature binds the complete underlying hash history covering k sequences entirely within quantum-resistant constraints. We define k as the bounded *attestation period* reflecting operational length $\tau_k = k \cdot t_{\text{block}}$.

Theorem 4.4 (Harvest Window Isolation Bounds)**.** *An adversary possessing CRQC capabilities online at t_{CRQC} immediately breaks the discrete logarithm dependencies bridging active BLS-aggregated block certificates. However, the theoretical protocol-time window capable of suffering unmitigated retrospective forgery operations resolves precisely to τ_k. For blocks at height $h \leq h_{\mathrm{last}}$ where σ_{PQ} already registered, modifying state forces a collision-resistance violation over $H(\mathit{chain})$, requiring an EUF-CMA forgery over ML-DSA. By calibrating $k = 60$ at a conservative 30-second emission speed, the maximum temporal exposure footprint limits absolutely to a strictly bound 30-minute metric.*

5 Formal Verification and Clinical Safety Envelopes

Relying upon heuristic testing boundaries or stochastic test-set cross-validation introduces unacceptable failure tail events inherent to out-of-distribution execution pathways. A mathematical verification substrate guarantees residual clinical hazard rates remain demonstrably absent. The Asuu protocol decouples performance prediction from safety oversight by installing continuous Linear Temporal Logic (LTL) evaluation systems over active executions.

5.1 The Clinical Simplex Verification Topology

Definition 5.1 (Static Safety Bounds via Temporal Logic)**.** Inherent clinical invariants formulate mathematically as property constraints φ mapping onto an array of atomic propositions AP. An explicit dosage restriction sequence $\mathbf{G}(dose(d) \leq max_safe(d))$ must evaluate consistently true everywhere.

Definition 5.2 (Systemic Simplex Controller Matrix)**.** Because neural prediction spaces ($f : \mathbb{R}^n \to \mathbb{R}^m$) render complete static verification (coNP-complete structure over ReLU nodes) computationally exhaustive for massive state spaces, runtime assurance delegates bounding to a decentralized Simplex system $\mathscr{S} = (\mathscr{H}, \mathscr{F}, \mathscr{D})$. Let $\mathscr{H}$ function as the high-accuracy deep network, $\mathscr{F}$ stand as the deterministically verified robust fallback algorithm, and $\mathscr{D}$ evaluate algorithmic signal robustness $\rho(\varphi, \mathbf{y}_{\mathscr{H}})$. Unsafe states immediately trigger dynamic output rerouting to $\mathscr{F}$.

Theorem 5.3 (Runtime Invariant Enforcement)**.** *Provided a verified fallback matrix $\mathscr{F}$ satisfies safety bounds φ constantly, routing output based upon negative Signal Temporal Logic robustness ($\rho < \varepsilon_{\text{margin}}$) rigorously secures systemic safety. The constant threshold $\varepsilon_{\text{margin}} > \delta_{\text{switch}} \cdot \|\dot{\mathbf{x}}\|_\infty$ mathematically neutralizes interpolation latency hazards triggered by switching interval δ_{switch}. Thus, output states exist unconditionally bound within φ at all chronological snapshots.*

5.2 Continuous Calibration Suspension Bounds

Safety encompasses distribution shifts altering calibration limits long after initial static evaluations map favorable boundaries. We integrate continuous runtime monitors executing probabilistic suspension triggers natively.

Definition 5.4 (Decentralized Calibration Observation)**.** The Expected Calibration Error (ECE) processes the delta between accuracy outputs mapping $\text{acc}(S_b)$ and network confidence distributions calculated as $\text{conf}(S_b)$ over chronological windows spanning exactly N evaluations across isolated bins B.

Theorem 5.5 (Safety-Suspension Delay Bounding)**.** *The Clinical AI Passport explicitly limits its network tolerance bridging variables θ_{alert} and θ_{suspend}. Suppose actual error $\text{ECE}_t > \theta_{\text{suspend}} + \varepsilon$. The monitor definitively deactivates network autonomy within a maximum bound of chronological evaluation epochs mapping exactly to:*

$$d^* = \left\lceil \frac{B^2 \ln(1/\delta_{\text{miss}})}{2N\varepsilon^2} \right\rceil + 1$$

incorporating probabilistic certainty $\geq 1 - \delta_{\text{miss}}$. Proof rationale: *Applying Hoeffding's inequality, an estimation margin outside* ε *decays exponentially by* $\exp(-2N\varepsilon^2/B^2)$. *Forcing this constraint equal to our minimal non-detection risk parameter* δ_{miss} *proves that undetected divergence halts mathematically at evaluating* d^*. *False alarm possibilities scaling beneath threshold bounds likewise diminish within symmetrical exponential tail projections. Consequently, runtime deviation remains not merely theoretically monitored but operationally intercepted via rigid bounds defining an absolute maximum latency horizon.*

6 Synthesis of the Five Theorem Backbone

The formal resilience of the Asuu Network protocol does not rely upon localized cryptographic assertions operating in isolation. Rather, it emerges through the interconnected synthesis of five foundational theorem families. Together, they form an impassable mathematical perimeter defining the absolute operational reality of the network. We categorically summarize this synthesis below.

(I) **Cross-Epoch Safety under PoCA (Theorem 3.7):** Through continuous overlapping weight matrices exceeding required supermajority bounds, the Byzantine Fault Tolerant logic mechanically prevents chain fragmentation. An actively verified institutional dataset cannot be covertly rewritten over successive protocol epochs. Finalization ensures absolute historical immutability despite temporal network partitions.

(II) **End-to-End Provenance Completeness:** By tightly coupling Merkle root hash structures to foundational Model ID registers (*mid*) and terminal human authorization signatures (*HA*), the execution pipeline forms a deterministic cryptographic chain (π_a). No clinical recommendation bypasses institutional accountability or obscures its parent neural architecture.

(III) **Update-Corridor Soundness and Identity Governance:** Model passport systems utilize deterministic smart contract configurations dictating rigid Predetermined Change Control Plan (PCCP) corridors. Unapproved or structurally deviating neural weight upgrades automatically fail network verification protocols, preemptively rejecting noncompliant predictive shifts. Every valid adjustment executes strictly within sovereign, predefined limits.

(IV) **Privacy-Budget Enforcement via Data Embassies:** The protocol enforces a hard mathematical abstraction barrier separating the verification state from the raw clinical substrate. No personally identifiable medical data crosses the on-chain perimeter. Instead, rigorous Zero-Knowledge mappings authenticate necessary clinical prerequisites and role-based access attributes without compromising the institutional privacy budget.

(V) **Safety-Suspension Latency Bounds (Theorem 5.5):** Because distributional realities inevitably shift, the Simplex topology incorporates bounded continuous evaluation modules based on Expected Calibration Error properties. As proven mathematically, the duration allowable for anomalous calibration drift to execute undetected terminates strictly within d^* epochs, automatically triggering a suspension state immune to adversarial continuity.

7 Conclusion

The rapid convergence of predictive medical capability and ubiquitous cloud infrastructure has centralized structural authority over clinical accountability in a computationally opaque vendor ecosystem. The Asuu Network reclaims institutional medical sovereignty by replacing operational trust with incontrovertible mathematical verification. By orchestrating Proof of Clinical Authority limits against bounded Simplex safety architectures, the protocol structurally enforces correct execution without revealing proprietary patient contexts or centralizing hierarchical oversight.

The system encapsulates safety limits not as aspirations, but as rigid algorithms bounding temporal drift detection and automated clinical suspension. Supported by a dual-phase post-quantum cryptographic migration blueprint, the network protects current decisions from future technological breaches. This produces an environment where algorithms exist strictly subordinated to the limits of formal mathematical proof, generating a medical intelligence surface optimized for human oversight, continuous accountability, and absolute verifiability.

References

[1] S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” *Cryptography Mailing List*, 2008.

[2] W3C, “Bitstring Status List v1.0,” *W3C Working Draft*, 2024. Available at https://www.w3.org/TR/vc-bitstring-status-list/

[3] National Institute of Standards and Technology, “FIPS 203, 204, and 205: Post-Quantum Cryptography Standardization,” *NIST Publications*, August 2024.

[4] O. Regev, “On Lattices, Learning with Errors, Random Linear Codes, and Cryptography,” *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 84–93, 2005.

[5] A. Shamir, “How to Share a Secret,” *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.

[6] M. J. Fischer, N. A. Lynch, and M. S. Paterson, “Impossibility of Distributed Consensus with One Faulty Process,” *Journal of the ACM*, vol. 32, no. 2, pp. 374–382, 1985.