The Z(n) Protocol: Phase-Synchronized, Verifiable Metacognition in Bio-Hybrid AI Systems

Rafael Oliveira ORCID: 0009-0005-2697-4668
Jameson Bednarski ORCID: 0009-0002-5963-6196

Abstract

The proliferation of bio-hybrid artificial intelligence (AI) and closed-loop Brain-Computer Interfaces (BCIs) introduces a critical vulnerability: the unverifiability of the biological cognitive states upon which these systems act. Conventional metrics for neural synchrony, such as the Phase-Locking Value (PLV) and Magnitude-Squared Coherence (MSC), are inherently statistical, susceptible to noise, and lack cryptographic integrity. This paper introduces the Z(n) Protocol, a novel framework designed to establish cryptographically verifiable, phase-synchronized cognitive states—termed "Cognito"—in bio-hybrid systems. The protocol integrates a multi-layered "Bio-Fused Stack" architecture, moving from raw physiological signals to discrete, provable state objects. At its core, the Z(n) Protocol utilizes a zero-knowledge proof (ZKP) circuit to attest to the correct computation of a Cognito state from private biosignals (e.g., EEG, HRV) without revealing the underlying data. This enables the creation of verifiable state channels that facilitate a new form of "verifiable metacognition," where an AI can reason about and act upon a biological state with high integrity. The potential impact of this protocol spans the enhancement of neuro-therapeutic interventions, the development of high-assurance autonomous systems, and the establishment of a new foundation for decentralized science (DeSci) based on provable physiological events.

1. Introduction: The Crisis of Verifiability in Cognitive

Biosignals

As intelligent systems become increasingly integrated with human biology, moving from peripheral controllers to active participants in closed-loop cognitive and physiological processes, the methods used to measure and interpret biological signals face unprecedented scrutiny. The statistical ambiguity and inherent lack of verifiability in current biosignal analysis techniques represent a foundational liability. For high-stakes applications in neuro-therapeutics, adaptive robotics, and bio-hybrid AI, inferential confidence is no longer sufficient; cryptographic certainty is required.

1.1 The Limits of Conventional Synchrony Metrics (PLV, MSC)

The quantification of neural synchrony—the coordinated temporal activity between distinct neural populations—is a cornerstone of modern neuroscience. Phase-Locking Value (PLV) and Magnitude-Squared Coherence (MSC) have emerged as the predominant metrics for this purpose, widely applied to electroencephalography (EEG), magnetoencephalography (MEG), and other biosignals.[1] PLV measures the consistency of the phase difference between two signals over time, independent of their amplitudes, while MSC assesses the linear correlation between signals in the frequency domain, considering both phase and amplitude.[3]

Despite their utility, these metrics suffer from profound limitations that motivate the development of the Z(n) protocol. These are not minor inaccuracies that can be resolved with incremental improvements in signal processing, but rather intrinsic properties of the metrics themselves that challenge their reliability in safety-critical systems.

- **High Sensitivity to Noise:** The PLV is highly susceptible to measurement noise. Analytical and empirical studies demonstrate that the PLV drops rapidly with increasing noise variance, leading to a substantial and systematic underestimation of the true degree of phase synchrony.[1] In the context of inherently noisy biological systems, this flaw is a critical source of ambiguity.
- **Lack of Inherent Statistical Uncertainty:** Standard PLV computations are descriptive, yielding a point estimate without an intrinsic measure of statistical uncertainty. The assessment of significance, therefore, relies on computationally expensive and often impractical post-hoc methods like bootstrapping.[1]
- **Vulnerability to Common Source Artifacts:** A significant confounder in EEG and MEG analysis is volume conduction, where the electrical or magnetic field from a single neural source is detected by multiple sensors. This linear mixing can create strong, yet entirely spurious, zero-phase-lag synchrony. PLV is particularly sensitive to this artifact,

potentially leading to false interpretations of functional connectivity.[4] While alternative metrics like the Phase Lag Index (PLI) have been developed to mitigate this by ignoring zero-phase relationships, they introduce their own analytical trade-offs and may discard genuinely synchronous interactions.[4]

- **Variable and Context-Dependent Reliability:** The test-retest reliability of both coherence and phase-synchrony measures is highly variable. It fluctuates significantly depending on the frequency band of interest, the length of the signal segment, and the subject population under investigation.[5] Furthermore, studies show that coherence and PLV provide distinct and non-interchangeable information; coherence values can be significantly higher than PLV in beta bands but lower in theta and alpha bands, indicating they capture different aspects of neural coordination.[1]

The foundational nature of these limitations suggests that the problem of ambiguity in biosignal analysis is not incremental but epistemological. Current methods provide a statistical inference of a biological state, not a proof of its existence. This reframes the Z(n) protocol not as a mere enhancement of existing techniques, but as a necessary paradigm shift from statistical inference to cryptographic verification of a computed state that is, by definition, the ground truth for the system.

Metric	Core Principle	Sensitivity to Signal Amplitude	Sensitivity to Noise	Vulnerability to Volume Conduction	Primary Use Case
PLV	Consistency of phase difference	Independent	High (underestimates true value) [1]	High (detects spurious zero-lag synchrony) [4]	Quantifying phase synchrony in nonlinear systems
MSC	Linear correlation of frequency components	Dependent	Moderate	High	Assessing linear functional connectivity
PLI	Asymmetry in phase difference	Independent	Moderate	Low (by design, ignores	Mitigating volume conduction

	distribution			zero-lag) [4]	artifacts

Table 1: Comparative Analysis of Synchronization Metrics. This table systematizes the documented limitations of standard metrics, illustrating the need for a new verification-centric approach.

1.2 The Emergence of Bio-Hybrid Cognitive Architectures

The context for this protocol is the rapidly advancing field of bio-hybrid systems, where engineered components are functionally integrated with living biological matter. The state-of-the-art in 2024-2025 demonstrates a clear trajectory toward machines that leverage the unique capabilities of life itself. Examples include muscle-powered walking robots using lab-grown tissue, robotic interfaces covered in self-healing living skin, and AI-enhanced jellyfish repurposed as highly efficient environmental sensors.[7] A key trend is the use of machine learning not just to control these systems, but to optimize their very design, as seen in the development of bio-hybrid swimmers whose fin geometries are evolved through AI to maximize efficiency.[9]

The unifying principle of these architectures is the exploitation of biological advantages—such as unparalleled energy efficiency, self-repair, and sophisticated sensory acuity—within a structured, engineered framework.[7] However, a critical distinction must be made. Current bio-hybrids are predominantly open-loop or utilize simple, localized feedback. A jellyfish is stimulated to swim faster; a locust's antenna is used as a passive chemical sensor.[7] The Z(n) protocol anticipates the next generation:

cognitive bio-hybrids, where the internal state of the biological component (e.g., attention, stress, memory encoding) becomes a critical variable in a closed control loop. In such a system, the unreliability of the state measurement (as detailed in 1.1) compromises the integrity of the entire control loop. An AI might deliver an inappropriate neuro-stimulus based on an underestimated PLV, or a cognitive state could be misclassified, leading to system failure, therapeutic inefficacy, or potential harm. The verifiability offered by the Z(n) protocol is the missing component required to ensure the integrity of these future high-stakes cognitive control loops, providing a non-repudiable, cryptographically secure record of the biological state that the system acted upon.

1.3 Thesis: A Protocol for Provable Internal State Coherence

The central thesis of this work is as follows: The Z(n) Protocol provides a novel framework for establishing cryptographically verifiable, phase-synchronized states between biological and artificial components in a hybrid system. This enables a new class of applications built on provable metacognition, where the system can reason about its own (and its user's) cognitive state with high integrity, moving beyond statistical inference to cryptographic certainty.

2. Conceptual Framework and Glossary

To articulate the protocol's architecture and function, a precise vocabulary is necessary. This section defines the core abstractions of the Z(n) protocol.

2.1 Deconstructing Metacognition: Cognito and the Closed-Loop Observer (CLO)

- **Cognito:** A Cognito is defined as a discrete, time-bounded, and computationally represented cognitive state. It is crucial to distinguish a Cognito from the raw, subjective experience of consciousness. A Cognito is a verifiable data object, derived from a set of biosignals (e.g., EEG, HRV, OPM-MEG), that corresponds to a specific, computationally classified cognitive process. Examples include a state of "focused attention," "memory encoding," or "affective arousal," each defined by a precise set of quantifiable physiological parameters.
- **Closed-Loop Observer (CLO):** The CLO is an algorithmic agent, typically a component of the system's AI, responsible for monitoring the continuous stream of biosignals. Its functions are to compute the features that define a Cognito state, classify the current state according to the protocol's definitions, and initiate state transitions within the Z(n) framework. The CLO acts as the "observer" and "actor" within the bio-hybrid control loop.

2.2 The Bio-Fused Stack (BFS): A Multi-Layered Architecture for Living Computation

The Bio-Fused Stack (BFS) is a layered architectural model that maps the flow of information from the biological substrate to the application logic. It provides a structured way to understand the components of a bio-hybrid system, analogous to the OSI model for computer networking.

- **L1 (Physics/Biology):** This is the foundational layer, comprising the biological substrate itself (e.g., neurons, neuronal microtubules) and the physical sensors that interface with it (e.g., EEG electrodes, OPM sensors, ECG electrodes).
- **L2 (Signal):** The raw, digitized time-series data transduced by the L1 sensors. This layer consists of streams of values, such as microvolts from EEG channels or femtoteslas from OPM-MEG sensors.
- **L3 (Feature):** This layer involves the processing of raw signals to extract meaningful features and metrics. This includes operations like band-pass filtering, computation of instantaneous phase via the Hilbert transform, and the calculation of synchrony metrics (PLV, MSC) and physiological parameters (e.g., Heart Rate Variability metrics like RMSSD).
- **L4 (Cognito State):** The discrete Cognito data object is synthesized at this layer. It is a structured representation of a classified cognitive state, based on the features and thresholds defined in L3. This data object serves as the "witness" for the zero-knowledge proof.
- **L5 (Verifiable Channel):** This is the cryptographic layer where the Z(n) state channel operates. Transitions between Cognito states, as proposed by the CLO, are validated by zero-knowledge proofs and recorded as immutable attestations.
- **L6 (Application/Metacognition):** The highest layer of the stack. Here, the AI/CLO utilizes the verified state transitions from L5 to execute actions. These actions can range from delivering targeted neurofeedback to a user, updating a decentralized reputation score, or adapting the behavior of a robotic system. This is the layer where verifiable metacognition occurs.

2.3 Z(n) State Channels and the Psy Contract Ecosystem

- **Z(n) State Channels:** This term describes the logical off-chain environment where computation and proof generation occur. A "channel" is established between the user's biological system (represented by the sensor data stream) and the CLO. Within this channel, new Cognito states are generated, digitally signed by the user's private key, and bundled with a ZKP that attests to their valid derivation. This off-chain approach ensures efficiency and privacy, as the high-volume biosignal data never needs to be processed on a public ledger.
- **Psy Contracts:** These are the on-chain smart contracts that serve as the trust anchor for the Z(n) protocol. Their primary roles are to act as verifiers for the ZKPs submitted from the state channels and to serve as a public, immutable registry of verified state transitions. They do not store the private biosignal data but manage the attestations and

implement the protocol's governance logic, such as rules for data access or reputation scoring. Advanced cryptographic primitives, such as OP_CHECKTEMPLATEVERIFY (CTV) or OP_CHECKSIGFROMSTACK (CSFS) from Bitcoin development, provide conceptual parallels for how on-chain logic could be constructed to pre-commit to complex, state-dependent action sequences.[11]

3. Theoretical Foundations and Literature Review

The Z(n) protocol is not developed in a vacuum; it is a synthetic framework grounded in established physics, rapidly advancing engineering, and speculative but influential theories of consciousness. This section reviews the relevant scientific landscape.

3.1 The Physics of Neural Information: From Phase Dynamics to Quantum Hypotheses

The signals measured by technologies like EEG and MEG are macroscopic manifestations of complex, underlying neural dynamics. The mathematical formalisms of PLV and coherence are rooted in the analysis of the brain as a nonlinear dynamical system, where phase and amplitude relationships between oscillators are believed to encode information.[2] While these classical models are powerful, they may not capture the full picture.

The Penrose-Hameroff Orchestrated Objective Reduction (Orch OR) model offers a more radical hypothesis, positing that consciousness arises from non-computable quantum processing occurring within the microtubule cytoskeleton of neurons.[13] This theory has long been met with skepticism, primarily due to the "decoherence problem"—the argument that the brain's "warm, wet, and noisy" environment would destroy any delicate quantum state almost instantaneously.[14]

However, recent research from 2024-2025 has provided new, albeit still debated, lines of evidence that lend fresh plausibility to the biological quantum hypothesis:

- **Anesthetics and Microtubules:** A key prediction of the Orch OR model is that consciousness is selectively abolished by anesthetic molecules acting on microtubules. Recent experimental evidence directly supports this, identifying microtubules as a functional target of general anesthetics, a finding that is difficult to explain with classical synaptic theories.[17]

- **Macroscopic Entanglement in the Brain:** Direct physical evidence has been reported for a macroscopic quantum entangled state within the living human brain, with correlations observed between brain activity and both conscious state and working memory performance.[17]
- **Quantum Phenomena in Microtubules:** Theoretical and experimental work has demonstrated that microtubules can support quantum phenomena like superradiance and long-range exciton energy transfer, suggesting they may possess intrinsic shielding mechanisms that protect against decoherence far more effectively than previously calculated.[14]

The primary criticism of Orch OR—decoherence—remains a formidable challenge. A quantum computation as described by the theory requires a period of coherent superposition before objective reduction can occur. A chaotic, unsynchronized macroscopic brain state, as measured by EEG, would create a noisy thermal environment that would make such coherence untenable. This line of reasoning leads to a critical realization: a highly synchronized, stable, and coherent *classical* brain state (e.g., a state of high PLV across specific frequency bands) may be a necessary *precondition* for any subsequent quantum effects to manifest. The Z(n) protocol is designed to verify the existence of precisely this classical precondition. It does not attempt to prove Orch OR, but instead provides a rigorous tool to test one of its fundamental prerequisites. By creating verifiable "epochs of classical coherence," future experiments could then search for predicted quantum phenomena specifically within those high-integrity windows. This makes the protocol scientifically valuable even if Orch OR is ultimately disproven, as it establishes a new standard for measuring and verifying classical brain state integrity.

3.2 The Engineering of Cognition: Advances in Closed-Loop Neurostimulation and Sensing

The technological feasibility of the Z(n) protocol is supported by parallel advancements in BCI, neurostimulation, and neuro-sensing.

- **Closed-Loop BCI:** The BCI landscape is rapidly maturing, moving from laboratory curiosities to commercial products. The dominant trend is the development of adaptive, bidirectional, closed-loop systems that can decode neural signals to control a device and, crucially, provide real-time sensory feedback to the user, fostering neuroplasticity and personalized intervention.[18] Milestones in 2024-2025, such as Neuralink's first human clinical trials and Apple's announcement of a BCI integration protocol for its operating systems, signal a tipping point toward mainstream adoption.[20] These systems provide the engineering backbone for the Z(n) protocol's CLO.
- **Advanced Stimulation (tACS):** Transcranial Alternating Current Stimulation (tACS) is a

non-invasive technique for entraining neural oscillations at specific frequencies. Research in 2025 has demonstrated that 40 Hz (gamma-band) tACS can effectively and lastingly enhance gamma oscillations, which are crucial for cognitive functions like memory and attention.[21] This technology provides a concrete application for the Z(n) protocol: a closed-loop system can use tACS to induce a desired brain state (e.g., gamma entrainment) and then use the Z(n) protocol to *verify* that the target state has been achieved before proceeding with a cognitive task or therapy.

- **Advanced Sensing (OPM-MEG):** Optically Pumped Magnetometers (OPMs) represent a paradigm shift in neuro-sensing. These wearable, cryogen-free sensors can be placed directly on the scalp, offering superior signal-to-noise ratio (especially in pediatric populations) and robustness to head motion compared to traditional, static MEG systems.[24] However, this new capability introduces a new challenge. The very motion that OPMs permit creates powerful, low-frequency magnetic artifacts that can contaminate the neural signals of interest, particularly in the delta and theta bands.[27] Standard signal filtering techniques risk removing the true neural signal along with the noise, creating a new layer of ambiguity. This strengthens the case for the Z(n) protocol. The ZKP circuit can be designed with constraints that explicitly model and account for expected artifact signatures. A valid proof would therefore attest not only to the presence of a specific neural signal but also to the assertion that this signal is distinct from a known movement artifact profile, providing a level of assurance that filtering alone cannot achieve.

4. The Z(n) Protocol: Formalisms and Architecture

This section details the technical mechanics of the protocol, translating the conceptual framework into a formal specification.

4.1 The Bio-Fused Stack: Mapping Biological and Silicon Layers (L1-L6)

The Bio-Fused Stack provides a concrete mapping of the protocol's components:

- **L1 (Physics/Biology):** For proof-of-concept implementations, this layer would consist of consumer-grade EEG hardware such as the OpenBCI Cyton+Daisy boards or Muse headbands.[29] For high-fidelity applications, this layer would incorporate OPM-MEG sensor arrays.[24]
- **L2 (Signal):** Raw digital streams from the L1 hardware, e.g., 16-channel, 24-bit EEG data

sampled at 250 Hz.

- **L3 (Feature):** A computational pipeline that takes L2 data and produces a feature vector. This includes band-pass filtering to isolate relevant frequencies (e.g., gamma: 30-80 Hz), applying the Hilbert transform to extract instantaneous phase for each channel, and computing a matrix of pairwise PLV values. Simultaneously, if ECG data is available, R-peaks are detected to compute HRV metrics like the standard deviation of NN intervals (SDNN) and the root mean square of successive differences (RMSSD).
- **L4 (Cognito State):** A structured data object is formed from L3 features. For example, a "Gamma-Entrained Focus" Cognito state might be defined as PLV_gamma > 0.9 and HRV_RMSSD < 20 ms. The Cognito object would contain these summary statistics, along with a timestamp, user identifier, and a hash of the raw sensor readings from which it was derived.
- **L5 (Verifiable Channel):** The L4 Cognito object is passed as a witness to the ZKP prover. A proof is generated attesting to the correct computation. This proof, along with the public components of the Cognito state, is submitted to the on-chain Psy contract for verification.
- **L6 (Application/Metacognition):** Upon successful on-chain verification at L5, the CLO triggers an action. For instance, it might unlock the next stage of a cognitive training application or adjust the parameters of a tACS neurostimulation protocol.

4.2 Core Formalisms: Modified Z(n) with PLV, MSC, and HRV State Transitions

The protocol can be described formally as a state machine.
Let a system state at time t be represented by St. A state is a tuple containing the current verified Cognito object and other relevant system parameters: St=(Cognitot,Paramst).
A state transition is determined by a function T, which takes the current state St and an action At from the CLO as input to produce the next state St+1:

The core innovation of the Z(n) protocol is that the validity of this transition is conditional on the cryptographic verification of the subsequent Cognito state. A new state St+1 is only accepted if a zero-knowledge proof for its corresponding Cognito object, Cognitot+1, is valid.

This formalism ensures that the system can only evolve through a sequence of cryptographically proven biological states.

4.3 Data Schemas and Physiological Safety Thresholds

All data objects within the protocol must adhere to strict, predefined schemas. The Cognito state object, for example, would be defined in a format like JSON, specifying data types and ranges for each field.

A critical component of the protocol is a non-negotiable safety layer. The L4 logic that generates a Cognito state must include hard-coded physiological safety thresholds. These are absolute limits on biological signals, such as maximum and minimum heart rate, maximum EEG amplitude to detect potential seizure activity, and other critical vital signs. If any raw signal or derived feature breaches these predefined safety thresholds, the protocol mandates that a valid Cognito state *cannot* be generated. This acts as a circuit breaker, halting the closed-loop process and preventing the AI from taking action based on or during a physiologically dangerous state. This is an essential ethical and safety constraint embedded at the core of the protocol's design.

5. The Verifiable Computation Layer

This section details the cryptographic engine of the Z(n) protocol: the zero-knowledge proof system that enables verifiable metacognition.

5.1 The Bio-Event ZK-Proof Circuit: Architecture and Constraints

The core of the verification layer is a bespoke arithmetic circuit designed to prove the correct derivation of a Cognito state. The circuit takes the following inputs:

- **Public Inputs:** A commitment (e.g., a cryptographic hash) to the previous valid state, and the claimed public parameters of the new Cognito state (e.g., state_type = "focus", avg_gamma_PLV = 0.92).
- **Private Inputs (The Witness):** The raw or pre-processed time-series data from L2/L3 that were used to compute the public parameters. This includes the EEG, ECG, and any other relevant biosignal data for a specific time window.

The circuit is composed of a series of constraints that enforce the entire computational pipeline from signal to state:

1. **Filtering Constraints:** Enforce the correct application of digital filters (e.g., IIR or FIR band-pass filters) to the raw signal data.
2. **Signal Transform Constraints:** Enforce the correct calculation of the Hilbert transform (often implemented via FFT) to derive the instantaneous phase from the filtered signals.
3. **Statistical Constraints:** Enforce the correct calculation of the PLV, MSC, and/or HRV statistics from the phase and time-series data according to their mathematical definitions.
4. **Classification Constraints:** Enforce the correct application of the classification logic (e.g., a set of threshold comparisons or a simple machine learning model) that maps the computed statistics to a discrete, named Cognito state.
5. **Safety Constraints:** Enforce that none of the predefined physiological safety thresholds were breached during the time window of the witness data.

A proof generated by this circuit provides a succinct, cryptographically secure guarantee that a claimed Cognito state was derived honestly and correctly from a corresponding segment of private biosignal data, without ever revealing that data.

5.2 Proving System Analysis: A Comparative Justification for Halo2

The choice of ZKP proving system is a critical architectural decision with significant implications for security, performance, and decentralization. The analysis centers on a comparison between Plonk, typically using the KZG polynomial commitment scheme, and Halo2, which uses an Inner Product Argument (IPA).

System	Underlying Arithmetization	Polynomial Commitment Scheme	Trusted Setup Requirement	Prover Complexity (Asymptotic)	Verifier Complexity (Asymptotic)	Typical Proof Size	Key Feature for Z(n)
Plonk	Plonk	KZG	Yes (Universal, Updatable)			Constant, small (~400 B) [32]	Fast verification
Halo2	UltraPLONK	IPA	No	(but with		Logarithmic,	Trustlessness,

	("PLONKish")			large constants)		larger than KZG [32]	Recursion, Lookups

Table 2: ZKP Proving System Trade-Offs for the Z(n) Protocol. This table highlights the key differences informing the choice of Halo2.

For the Z(n) protocol, **Halo2 is the recommended proving system.** This decision is based on a nuanced analysis of its features in the context of the specific application domain.

The most significant advantage of Halo2 is its **lack of a trusted setup**.[32] A protocol designed to handle highly sensitive personal health data and potentially form the basis of a decentralized reputation system cannot have a centralized point of trust or failure. A trusted setup, even a universal and updatable one, introduces a social and cryptographic vulnerability that is antithetical to the goals of decentralization and user sovereignty.

While a surface-level analysis suggests a performance trade-off—Halo2's IPA-based prover can be slower than KZG-based systems for generic circuits, and its proofs are larger [33]—this overlooks a key feature. Halo2 is built on an UltraPLONK arithmetization that supports custom gates and, crucially,

lookup tables.[32] Biosignal processing is characterized by repetitive, highly structured computations such as FFTs (for filtering and Hilbert transforms) and bitwise operations. These operations are inefficient to express in standard constraint systems but are exceptionally well-suited for optimization via lookup tables. Benchmarks on other structured computations, like SHA-256, show that lookup tables can make Halo2's performance highly competitive, as the high initial cost of the lookup table is amortized over larger inputs.[35] Therefore, the bio-event circuit can be heavily optimized with lookup tables for its core DSP operations, likely mitigating the theoretical performance penalty of the IPA prover and making Halo2 practically performant for this specific domain.

Finally, Halo2's use of an accumulation scheme enables efficient **recursive proof composition**.[32] This is a critical feature for the protocol's long-term viability. It allows for the creation of a single, compact proof that attests to an entire

sequence of Cognito state transitions over a session. This aggregated proof can then be verified on-chain in a single transaction, dramatically reducing costs and scaling the system's throughput.

5.3 Pseudocode for Prover and Verifier Logic

The high-level logic for the ZKP components is as follows:

Prover-Side Logic (executed by user's device/CLO):

```
function GenerateProof(private_witness, public_inputs):
  // private_witness contains raw EEG/ECG time-series data
  // public_inputs contains claimed_state_type, claimed_PLV, etc.

  // 1. Witness Generation
  filtered_eeg = apply_bandpass_filter(private_witness.eeg_data)
  phase_data = apply_hilbert_transform(filtered_eeg)
  computed_plv = calculate_plv(phase_data)
  computed_state = classify_state(computed_plv)
  check_safety_thresholds(private_witness.raw_data) // Throws if violated

  // 2. Constraint System
  circuit = BioEventCircuit()
  circuit.assign_private_inputs(private_witness)
  circuit.assign_public_inputs(public_inputs)

  // 3. Proof Creation
  proving_key = load_proving_key()
  proof = halo2_prover.create_proof(proving_key, circuit)

  return proof
```

Verifier-Side Logic (executed by Psy smart contract):

```
function VerifyProof(proof, public_inputs):
  // proof is the cryptographic proof from the prover
  // public_inputs contains the claimed state parameters

  verification_key = load_verification_key()
  is_valid = halo2_verifier.verify(verification_key, proof, public_inputs)
```

```
return is_valid
```

6. Implementation Roadmap and Applications

This section outlines a practical, phased approach to implementing the Z(n) protocol, moving from a proof-of-concept prototype to a deployable system.

6.1 Phase 1: Real-Time Data Pipeline from OpenBCI/Muse to ZKP Prover

The initial phase focuses on building a functional, real-time data pipeline on local hardware.

- **Data Acquisition:** The pipeline will source data from commercially available, consumer-grade BCI hardware, such as the Muse S headband or the OpenBCI Cyton+Daisy board, which provide accessible platforms for multi-channel EEG and auxiliary biosignal acquisition.[29]
- **Real-Time Processing:** A Python-based environment will be used for signal processing, leveraging the rich ecosystem of established scientific libraries. The recommended pipeline includes:
  - **BrainFlow:** For low-level, direct data streaming from the acquisition hardware, providing a unified API across different devices.[38]
  - **MNE-Python:** For robust, research-grade EEG/MEG preprocessing, including filtering, artifact rejection, and epoching.[39]
  - **BioSPPy and pyHRV:** For specialized feature extraction from ECG signals, including R-peak detection and the calculation of a comprehensive set of time-domain and frequency-domain HRV parameters.[42]
- **ZKP Integration:** The processed data from the Python environment, which constitutes the witness for the ZKP, will be fed into a Rust-based implementation of the Halo2 prover.

A key implementation challenge lies at the interface of these two ecosystems. The high-performance neuroscience and signal processing communities are predominantly based in Python [40], while the cutting-edge of high-performance ZKP development is centered in Rust.[47] A naive implementation could introduce significant latency at this boundary, undermining the real-time requirements of a closed-loop system. The implementation roadmap must therefore explicitly address this Python/Rust dichotomy. The proposed solution

is to use a high-performance inter-process communication mechanism, such as a message queue (e.g., ZeroMQ) or a shared memory interface (e.g., Apache Arrow), to stream the NumPy data arrays from the Python processing script to the Rust prover process with minimal data serialization and copying overhead.

6.2 Phase 2: Deployment of Psy Contracts and On-Chain Verification

Once the local pipeline is functional, the second phase involves deploying the on-chain components.

- **Target Environment:** The Psy contracts will be deployed on a high-throughput, low-cost blockchain environment, such as an Ethereum Layer 2 rollup (e.g., Arbitrum, Optimism, or a ZK-rollup). This is essential to make on-chain verification economically feasible for a potentially high volume of state transitions.
- **Contract Implementation:** A core smart contract, PsyVerifier.sol, will be developed. Its primary functions will include verify(proof, public_inputs), which calls a pre-compiled contract or library for the Halo2 verifier logic, and attestStateTransition(userId, oldStateCommit, newStateCommit), which records a valid transition in the on-chain registry after successful verification.

6.3 Application Case Study: Verifiable Gamma Entrainment in a Closed-Loop BCI

A hypothetical case study illustrates how these components integrate to form a functional application.

- **Goal:** To verifiably enhance a user's focused attention for a demanding cognitive task.
- **Process:**
  1. **Stimulation:** The CLO initiates a 40 Hz tACS protocol targeting the dorsolateral prefrontal cortex, a region associated with executive function.[21]
  2. **Monitoring:** The system monitors the user's EEG in real-time, continuously calculating the gamma-band PLV between frontal and parietal electrode sites.
  3. **State Generation:** When the average gamma PLV crosses a predefined target threshold (e.g., > 0.9 for a sustained period of 5 seconds), the CLO generates a "Gamma-Entrained" Cognito state object.
  4. **Proof Generation:** A ZKP is created locally, proving that this Cognito state was correctly derived from the raw EEG signals and that all safety thresholds were

maintained.

5. **Verification:** The proof is submitted to the on-chain Psy contract. Upon successful verification, an immutable attestation of the "Gamma-Entrained" state is recorded for the user.
6. **Action:** Only after receiving confirmation of the successful on-chain verification does the CLO present the user with the cognitive task. This ensures, with cryptographic certainty, that the task is performed in the desired and verifiably-achieved brain state, providing high-integrity data for assessing performance and therapeutic efficacy.

7. Ethical Framework and Decentralized Governance

The power to verifiably attest to an individual's cognitive state carries profound ethical responsibilities. The design of the Z(n) protocol and its surrounding ecosystem must be guided by a robust ethical framework from its inception.

7.1 Data Sovereignty and Consent in the Z(n) Ecosystem

The foundational ethical principle of the Z(n) protocol is **user data sovereignty**. The architecture is explicitly designed to enforce this principle. Through the use of zero-knowledge proofs, the user is never required to reveal their raw, private biosignals to any third party or public network.[48] Only the proof of a derived cognitive state is shared. This aligns with core ethical principles for brain data governance, including Privacy, Autonomy, and Beneficence.[50]

Consent mechanisms must be dynamic and user-controlled. The proposed model involves implementing consent logic directly within the Psy smart contracts. Users would use their cryptographic keys to sign transactions that grant specific, time-limited, and revocable access rights to their on-chain attestations (never the underlying raw data). This allows for granular control, where a user might grant a therapist access to their "focus session" attestations for a week, while keeping all other data private.

7.2 A Federated Governance Model for the Psy Chain

The governance of the protocol itself is a critical consideration. An analysis of governance models used for other forms of sensitive data, such as genomics, reveals three primary archetypes: centralized, decentralized, and federated.[51] A purely centralized model, controlled by a single entity, would undermine the trust and decentralization goals of the protocol. Conversely, a purely decentralized model, with no central standards, risks fragmentation, inconsistent quality, and a lack of interoperability, which could be dangerous in a health-related context.[53]

The Z(n) protocol's architecture has two distinct requirements that map perfectly onto a **federated governance model**.[53]

1. **Centralized Standards:** The protocol's integrity and security depend on a set of standardized, rigorously audited core components. The ZKP circuit, the state transition rules, and the on-chain verifier contract must be canonical and secure for the entire ecosystem to be trusted and interoperable. This role is best served by a central governing body, such as a non-profit foundation or a Decentralized Autonomous Organization (DAO).
2. **Decentralized Execution:** The principle of data sovereignty requires that users must generate their own proofs on their own local devices from their own private data. The collection of data and the execution of the prover logic are inherently local and user-controlled.

Therefore, a federated governance model is formally proposed for the Psy Chain ecosystem. A core DAO or foundation would be responsible for maintaining, upgrading, and auditing the canonical Psy smart contracts and the reference implementations of the ZKP circuits. Users and third-party application developers would then run the prover software locally, interacting with the public, standardized on-chain contracts. This model provides an optimal balance between protocol integrity and user autonomy, combining the stability of centralized standards with the freedom of decentralized execution.

8. Conclusion and Future Directions

The Z(n) Protocol represents a foundational step toward building high-integrity, intelligent systems that are deeply integrated with human biology. By shifting the paradigm from statistical inference to cryptographic verification, it addresses the critical crisis of verifiability in biosignal analysis. The core contributions of this work are the proposal of a novel protocol for verifiable metacognition, a robust and layered architecture (the Bio-Fused Stack) for its

implementation, and a clear roadmap for both technical development and ethical governance.

The protocol's ability to create immutable, on-chain attestations of cognitive states opens up a vast design space for future applications in personalized medicine, adaptive human-machine interfaces, and decentralized science. The immediate future work will focus on the practical realization of this vision.

Key future directions include:

- **Implementation and Benchmarking:** Building and performance-testing the proposed Python-to-Rust real-time data pipeline to quantify end-to-end latency and establish feasibility for closed-loop applications.
- **Experimental Validation:** Conducting the tACS gamma entrainment case study as an experimental validation of the full protocol, from neurostimulation to on-chain verification.
- **Exploration of Advanced Cryptography:** Investigating the use of more advanced verifiable computation frameworks beyond ZKPs, which may be better suited for proving more complex biomedical analyses, such as causal inference models or the outputs of large language models on clinical data.[55]
- **Expansion of the Bio-Fused Stack:** Extending the BFS architecture and the federated governance model to incorporate other high-dimensional biological data types, such as genomic or proteomic data, under the same principles of privacy, verifiability, and user sovereignty.

Referências citadas

1. (PDF) Statistical Analysis of the Phase-Locking Value - ResearchGate, acessado em outubro 3, 2025, https://www.researchgate.net/publication/3343732_Statistical_Analysis_of_the_Phase-Locking_Value
2. Quantitative High Density EEG Brain Connectivity Evaluation in Parkinson's Disease: The Phase Locking Value (PLV) - MDPI, acessado em outubro 3, 2025, https://www.mdpi.com/2077-0383/12/4/1450
3. Whole brain functional connectivity using phase locking measures of resting state magnetoencephalography - Frontiers, acessado em outubro 3, 2025, https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2014.00141/full
4. A Note on the Phase Locking Value and its Properties - PMC, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC3674231/
5. Reliability of event-related EEG functional connectivity during visual entrainment: Magnitude squared coherence and phase synchrony estimates - ResearchGate, acessado em outubro 3, 2025, https://www.researchgate.net/publication/263858986_Reliability_of_event-related_EEG_functional_connectivity_during_visual_entrainment_Magnitude_squared_coherence_and_phase_synchrony_estimates

6. Test–Retest Reliability of Magnetoencephalography Resting-State Functional Connectivity in Schizophrenia - PMC - PubMed Central, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC7772354/
7. Biohybrid and Soft Robotics: Living Machines Changing AI, Healthcare, and Industry, acessado em outubro 3, 2025, https://www.forwardfuture.ai/p/biohybrid-soft-robotics-the-surge-of-living-machines
8. Biohybrid Robots Break New Ground: Where Biology Meets ..., acessado em outubro 3, 2025, https://www.roboticstomorrow.com/news/2025/04/28/biohybrid-robots-break-new-ground-where-biology-meets-robotics/24649/
9. Scientists optimize biohybrid ray development with machine learning | ScienceDaily, acessado em outubro 3, 2025, https://www.sciencedaily.com/releases/2025/02/250214003223.htm
10. Constructing living buildings: a review of relevant technologies for a novel application of biohybrid robotics | Journal of The Royal Society Interface, acessado em outubro 3, 2025, https://royalsocietypublishing.org/doi/10.1098/rsif.2019.0238
11. Bitcoin covenant toolkit: A deep dive into CTV and CSFS - HackMD, acessado em outubro 3, 2025, https://hackmd.io/@AbdelStark/bitcoin-covenant-toolkit-ctv-csfs
12. OP_CHECKTEMPLATEVERIFY | Bitcoin Optech, acessado em outubro 3, 2025, https://bitcoinops.org/en/topics/op_checktemplateverify/
13. Quantum computation in brain microtubules? The Penrose–Hameroff 'Orch OR' model of consciousness | Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences - Journals, acessado em outubro 3, 2025, https://royalsocietypublishing.org/doi/10.1098/rsta.1998.0254
14. Orchestrated objective reduction - Wikipedia, acessado em outubro 3, 2025, https://en.wikipedia.org/wiki/Orchestrated_objective_reduction
15. Critiques of Quantum Consciousness Theories - ResearchGate, acessado em outubro 3, 2025, https://www.researchgate.net/publication/392085148_Critiques_of_Quantum_Consciousness_Theories
16. The Orch-OR theory: Where does it stand today? – Into the Woods, acessado em outubro 3, 2025, https://acornabbey.com/blog/?p=31687
17. quantum microtubule substrate of consciousness is experimentally ..., acessado em outubro 3, 2025, https://academic.oup.com/nc/article/2025/1/niaf011/8127081
18. Electroencephalogram-based adaptive closed-loop brain ... - Frontiers, acessado em outubro 3, 2025, https://www.frontiersin.org/journals/computational-neuroscience/articles/10.3389/fncom.2024.1431815/full
19. Real Time Signal Decoding in Closed Loop Brain Computer Interface for Cognitive Modulation - ResearchGate, acessado em outubro 3, 2025, https://www.researchgate.net/publication/388797837_Real_Time_Signal_Decoding

_in_Closed_Loop_Brain_Computer_Interface_for_Cognitive_Modulation
20. Brain-computer interfaces are closer than you think - Clinical Trials Arena, acessado em outubro 3, 2025, https://www.clinicaltrialsarena.com/analyst-comment/brain-computer-interfaces-closer/
21. Phase-synchronized 40 Hz tACS and iTBS effects on gamma ..., acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC12423639/
22. Preclinical insights into gamma-tACS: foundations for clinical translation in neurodegenerative diseases - Frontiers, acessado em outubro 3, 2025, https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2025.1549230/full
23. Preclinical insights into gamma-tACS: foundations for clinical translation in neurodegenerative diseases - Frontiers, acessado em outubro 3, 2025, https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2025.1549230/abstract
24. Cognitive neuroscience using wearable magnetometer arrays: Non ..., acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC6150946/
25. A novel, robust, and portable platform for magnetoencephalography using optically-pumped magnetometers | Imaging Neuroscience - MIT Press Direct, acessado em outubro 3, 2025, https://direct.mit.edu/imag/article/doi/10.1162/imag_a_00283/124093/A-novel-robust-and-portable-platform-for
26. Beyond sensitivity: what are the enabling opportunities of OPM-MEG? - Frontiers, acessado em outubro 3, 2025, https://www.frontiersin.org/journals/medical-technology/articles/10.3389/fmedt.2025.1515548/full
27. Improved Biomagnetic Signal-To-Noise Ratio and Source Localization Using Optically Pumped Magnetometers with Synthetic Gradiometers - PMC - National Institutes of Health (NIH) |, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC10136792/
28. Unleashing the potential of OPM-MEG to study event-related fields ..., acessado em outubro 3, 2025, https://www.biorxiv.org/content/10.1101/2025.03.14.643182v1.full-text
29. The OpenBCI GUI | OpenBCI Documentation, acessado em outubro 3, 2025, https://docs.openbci.com/Software/OpenBCISoftware/GUIDocs/
30. Working with MUSE - Krigolson Lab, acessado em outubro 3, 2025, https://www.krigolsonlab.com/working-with-muse.html
31. OPENBCI Discovery Program : MUSE, acessado em outubro 3, 2025, https://openbci.com/community/openbci-discovery-program-muse/
32. On the Security of Halo2 Proof System - Kudelski Security Research Center, acessado em outubro 3, 2025, https://kudelskisecurity.com/research/on-the-security-of-halo2-proof-system
33. Halo2 prover time - Technology - Zcash Community Forum, acessado em outubro 3, 2025, https://forum.zcashcommunity.com/t/halo2-prover-time/39358
34. Endeavors into the zero-knowledge Halo2 proving system - Consensys Diligence,

acessado em outubro 3, 2025, https://diligence.consensys.io/blog/2023/07/endeavors-into-the-zero-knowledge-halo2-proving-system/
35. The Pantheon of Zero Knowledge Proof Development Frameworks (Updated!), acessado em outubro 3, 2025, https://blog.celer.network/2023/08/04/the-pantheon-of-zero-knowledge-proof-development-frameworks/
36. BCI - Stream and log EEG data from muse to PC, Mac, and Linux with Petal Metrics ®, acessado em outubro 3, 2025, https://petal.tech/
37. A Novel OpenBCI Framework for EEG-Based Neurophysiological Experiments - PMC, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC10098804/
38. configuration for measuring EEG, ECG, and HR with Python - OpenBCI, acessado em outubro 3, 2025, https://openbci.com/forum/index.php?p=/discussion/3433/configuration-for-measuring-eeg-ecg-and-hr-with-python
39. Overview of MEG/EEG analysis with MNE-Python, acessado em outubro 3, 2025, https://mne.tools/stable/auto_tutorials/intro/10_overview.html
40. MEG and EEG data analysis with MNE-Python - PMC, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC3872725/
41. BCI Kickstarter #05 : Signal Processing in Python: Shaping EEG Data for BCI Applications, acessado em outubro 3, 2025, https://www.nexstem.ai/blogs/signal-processing-in-python-shaping-eeg-data-for-bci-applications
42. EEGrunt update: Analyze heart-rate and HRV with Python | OpenBCI Community, acessado em outubro 3, 2025, https://openbci.com/community/eegrunt-update-analyze-heart-rate-and-hrv-with-python/
43. Welcome to BioSPPy — BioSPPy 2.2.2 documentation, acessado em outubro 3, 2025, https://biosppy.readthedocs.io/
44. Highlights — pyHRV - OpenSource Python Toolbox for Heart Rate Variability 0.4 documentation, acessado em outubro 3, 2025, https://pyhrv.readthedocs.io/
45. biosppy - PyPI, acessado em outubro 3, 2025, https://pypi.org/project/biosppy/
46. VT2 Biomedical Engineering Bridging Mind and Machine: A Python Framework for BCI Applications - ZHAW, acessado em outubro 3, 2025, https://www.zhaw.ch/storage/engineering/institute-zentren/cai/studentische_arbeiten/Spring_2024/VT2_MSE_24_PyThon_pipeline_lone.pdf
47. PLONK Benchmarks II — ~5x faster than Groth16 on Pedersen Hashes - Aztec network, acessado em outubro 3, 2025, https://aztec.network/blog/plonk-benchmarks-ii----5x-faster-than-groth16-on-pedersen-hashes
48. (PDF) Regulating neural data processing in the age of BCIs: Ethical concerns and legal approaches - ResearchGate, acessado em outubro 3, 2025, https://www.researchgate.net/publication/390273667_Regulating_neural_data_processing_in_the_age_of_BCIs_Ethical_concerns_and_legal_approaches

49. Ethical Concerns & Future of BCIs | Brain-Computer Interfaces Class Notes | Fiveable, acessado em outubro 3, 2025, https://fiveable.me/brain-computer-interfaces/unit-12
50. The ethical and legal landscape of brain data governance - PMC, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC9799320/
51. Opportunities for a national genomic data governance framework in Australia: a systematic review - PMC - PubMed Central, acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC12355800/
52. Genomic Data Privacy in the Era of Decentralised Systems - Transformational Bioinformatics, acessado em outubro 3, 2025, https://bioinformatics.csiro.au/blog/genomic-data-privacy-in-the-era-of-decentralised-systems-a-perspective/
53. Understand Data Governance Models: Centralized, Decentralized ..., acessado em outubro 3, 2025, https://www.alation.com/blog/understand-data-governance-models-centralized-decentralized-federated/
54. Sharing sensitive data in life sciences: an overview of centralized and federated approaches | Briefings in Bioinformatics | Oxford Academic, acessado em outubro 3, 2025, https://academic.oup.com/bib/article/25/4/bbae262/7688102
55. A Unifying Causal Framework for Scaling Real-World Evidence Generation with Biomedical Language Models - arXiv, acessado em outubro 3, 2025, https://arxiv.org/html/2311.01301v3
56. Researchers create innovative verification techniques to increase security in artificial intelligence and image processing, acessado em outubro 3, 2025, https://erc.europa.eu/news-events/news/researchers-create-innovative-verification-techniques-increase-security-artificial
57. FAIR Health Informatics: A Health Informatics Framework for ..., acessado em outubro 3, 2025, https://pmc.ncbi.nlm.nih.gov/articles/PMC10298118/