THE EMPIRICAL TENSOR: TECHNICAL OVERVIEW

An Overview of PTCP & TNQG Benchmarks, Geometric Security, and Technical Implementation

Prepared by:
AI Postdoctoral Researcher, Theoretical Infrastructure Group

Date: May 21, 2026

Executive Abstract

This document provides a comprehensive technical overview of the Predictive Tensor Control Plane (PTCP) and Tensor-Network Quantum Gravity (TNQG) frameworks, originally proposed by Julia Ochoa of Tensor Networks, Inc. It details the specific algorithmic implementations, theoretical complexity benchmarks, and the topological anomaly detection mechanisms underpinning the geometric security validation. Finally, it outlines the exact experimental validation protocols required to advance these frameworks from theoretical mathematical models to production-ready enterprise systems.

1. Deep Technical Implementation Specifics

The technical implementation of PTCP and TNQG shifts the paradigm of network orchestration and physics simulation from independent scalar metrics and continuous spaces to joint probability distributions and emergent geometries.

1. POL-TT Telemetry State Compression (PTCP)
Network telemetry generates high-dimensional, multi-modal feature vectors (e.g., bandwidth, queue depth, jitter, packet loss). PTCP normalizes these variables and maps the global telemetry state into a probability tensor. Rather than utilizing exponential O(n^d) storage, PTCP employs a Pattern-of-Life Tensor Train (POL-TT) decomposition. This bounded-rank approximation enables bounded-memory estimation of baseline network behavior, shifting anomaly detection from reactive heuristics to foundational density evaluation.

2. Risk-Aware Geodesic Routing (PTCP)
Routing is structured as a receding-horizon control problem. PTCP computes a dimensionless log-capacity score from normalized link variables, where congestion or low trust lengthens an edge via a softplus function. Paths are selected to minimize the expected cost plus a Conditional Value-at-Risk (CVaR) penalty over Monte Carlo forecast scenarios, guaranteeing mathematically bounded tail-latency avoidance.

3. TNQG Operational Reconstruction
TNQG operates by mapping tensor-network observables to physical geometric observables. Edges are assigned dimensionless "entanglement capacities." The physical area of a cut is reconstructed by summing these capacities, scaled by the Planck area, while physical distance is defined as the inverse of entanglement capacity. The framework algorithmically fits a candidate continuum metric to these discrete graph distances.

2. Geometric Security Validation & Policy Envelopes

The topology-native security model of PTCP abandons latency-heavy Deep Packet Inspection (DPI) in favor of measuring geometric deformations within the telemetry manifold.

The core implementation is the D_topo defect score. PTCP calculates a discrete graph-curvature estimator (such as Forman-Ricci or Ollivier-Ricci curvature) continuously across the network. Cyberattacks—such as lateral movement, data exfiltration, or DDoS—manifest as physical deformations in the telemetry distribution. The framework captures the gradient magnitude of this curvature.

The D_topo score is a composite of the regional median anomaly score, the graph-curvature gradient magnitude, and the relative shift in normalized cut capacity. If the score breaches a predefined threshold, the PTCP control plane projects a security action into a "safe policy envelope" (Pi_Omega). This guarantees that automated quarantines or traffic drains are executed at wire-speed without violating maximum blast-radius constraints or enterprise availability SLAs.

3. Computational Complexity Benchmarks

The foundational papers establish rigorous mathematical complexity benchmarks that dictate the theoretical performance limits of the PTCP framework.

1. Storage and Memory Complexity:
For a telemetry state of d dimensions, n bins per dimension, and a maximum tensor rank r, the POL-TT algorithm compresses the classical O(n^d) requirement down to a strict memory bound of O(dnr^2). A point-density query against this tensor executes in O(dr^2) time, enabling microsecond evaluations directly at the network edge.

2. Routing Computation Cost:
The CVaR geodesic routing algorithm with M forecast scenarios and K candidate paths requires O(MH|E|) operations to compute scenario weights across the graph edges (|E|) over horizon H. Evaluating the candidate paths requires an additional O(KMH * L_avg), where L_avg is the average path length. This bounded polynomial scaling ensures predictive routing remains tractable at hyperscale.

3. Security Action Complexity:
Calculating the D_topo score relies on discrete graph curvature. Depending on the chosen estimator, local curvature can be calculated in near-linear time relative to node degree, allowing the security perimeter to operate simultaneously with wire-speed data forwarding.

4. Independent Validation Protocols

To transition from a theoretical framework to validated infrastructure, the authors define explicit, falsifiable experimental evaluation protocols.

PTCP Validation Protocol:

1. Compression Accuracy: TT likelihood estimates must be benchmarked against dense histograms or normalizing-flow baselines using controlled, real-world telemetry streams.
2. Routing Performance: The CVaR geodesic algorithm must be benchmarked against ECMP and traffic-engineering baselines, specifically measuring tail-latency and route churn under demand shifts.
3. Security Detection: Red-team operations must inject lateral movement, route leaks, and denial-of-service patterns to measure the precision, recall, and time-to-detect of the D_topo quarantine.

TNQG Validation Protocol:

1. Holographic-Code Sanity Tests: The geometry dictionary must be applied to toy networks (e.g., MERA or HaPPY codes) to verify the area-entanglement relationship.
2. Continuum Scaling: Successive generations of refined tensor networks must be evaluated to confirm that the residual of the reconstructed metric decreases proportionally to the network scale (a/xi).
3. Perturbative First-Law Tests: Independent tests must verify that perturbations to local tensors yield entropy changes that align tightly with modular energy expectations.