Predictive Shape Adaptation in RIS-Assisted Flexible Pinching-Antenna Systems for 6G Networks: A Deep Learning and Manifold Optimization Approach

Mandalika J.R.M. Prasad , A.M. Prasad

Abstract

Flexible Pinching-Antenna Systems (F-PASS) have recently emerged as a transformative, energy-efficient alternative to conventional rigid massive MIMO arrays. By dynamically adapting its geometric shape, F-PASS can optimize spatial degrees of freedom and multiplexing gains in real-time. However, existing F-PASS architectures operate in a purely reactive mode, where unavoidable mechanical reconfiguration latency fundamentally limits their real-time channel-tracking capabilities in highly dynamic, high-mobility environments. In this paper, we propose a novel proactive F-PASS architecture assisted by Reconfigurable Intelligent Surfaces (RIS) for next-generation 6G wireless networks. By synergistically combining the coarse physical shape adaptation of the F-PASS transceiver with the fine-grained, instantaneous wavefront manipulation of the RIS, we significantly mitigate the physical latency penalty. To completely overcome mechanical delay, we introduce a predictive deep learning framework based on dilated Temporal Convolutional Networks (TCNs) to accurately anticipate rapid spatial-channel variations. This intelligence enables the F-PASS to proactively reconfigure its geometry, ensuring perfect spatial alignment at the predicted channel realization. We rigorously formulate a comprehensive joint optimization framework to maximize global energy efficiency (EE). The resulting mixed-integer non-linear programming (MINLP) problem is solved via a robust alternating optimization (AO) algorithm leveraging Dinkelbach's fractional programming and Riemannian Conjugate Gradient (RCG) manifold optimization. Extensive numerical simulations, backed by rigorous analytical theorems, demonstrate that the proposed Predictive RIS-FPASS achieves excellent results, specifically delivering a remarkable 43.5% improvement in energy efficiency and a 38.2% increase in spectral efficiency over conventional reactive baselines.

Keywords: Pinching antenna systems, Reconfigurable intelligent surfaces (RIS), 6G communications, predictive deep learning, shape reconfigurability, energy efficiency, Temporal Convolutional Networks, Riemannian manifold optimization.

Introduction

Relentless evolution toward sixth-generation (6G) wireless networks demands unprecedented spectral and energy efficiencies. Emerging data-hungry applications—such as holographic telepresence, autonomous vehicular systems, immersive digital twins, and ultra-dense IoT networks—require continuous, high-capacity coverage (Saad, Bennis, and Chen 2019). While massive multiple-input multiple-output (MIMO) technology achieved significant spectral efficiency in 5G, scaling it to sub-terahertz (THz) frequencies for 6G is fundamentally bottlenecked by severe energy consumption (Boulogeorgos et al. 2018). Specifically, equipping every antenna element with a dedicated, high-resolution radio-frequency (RF) chain and power amplifier (PA) results in prohibitive thermal dissipation that scales linearly with the number of antennas (Björnson et al. 2019).

To address these physical layer challenges, Pinching-Antenna Systems (PASS) have been proposed as a highly promising, energy-efficient alternative (Prasad et al. 2024). PASS leverages continuous dielectric waveguides equipped with discrete binary activation mechanisms known as pinch points. By selectively activating these pinch points, PASS synthesizes dynamic highly directive beams, drastically reducing the required number of power-hungry RF chains while maintaining massive MIMO-equivalent multiplexing capabilities (Prasad et al. 2025).

Background on F-PASS and Mechanical Latency

Despite their energy advantages, static PASS configurations suffer from severe geometric rigidity, limiting their ability to track dynamic spatial channel variations. A static planar array has a fixed macroscopic spatial correlation matrix, making it highly inefficient for serving mobile users clustered at extreme elevation or azimuth angles. To definitively overcome this, recent seminal advances have introduced Flexible PASS (F-PASS), which incorporates physical shape reconfigurability through mechanical rotation, bending, and folding mechanisms (Prasad 2026). Built on elastomeric substrates such as polydimethylsiloxane (PDMS) embedded with highly conductive silver nanowires (AgNW), F-PASS enables continuous geometric adaptation (Jiang et al. 2021).

However, mechanical actuation introduces an inherent, non-negligible physical latency. Current F-PASS algorithms operate in a purely reactive mode. This physical transition, typically lasting 20-100 ms, leads to catastrophic performance degradation in high-mobility scenarios. By the time the F-PASS achieves its intended shape, the wireless channel state has completely decorrelated due to small-scale Rayleigh fading.

Reconfigurable Intelligent Surfaces (RIS) in 6G

Simultaneously, Reconfigurable Intelligent Surfaces (RIS) have emerged as a disruptive technology for customized propagation environments in 6G networks (Wu and Zhang 2019). Comprising a massive planar array of passive electromagnetic scattering elements, an RIS can intelligently reflect impinging waves via a smart microcontroller (Wu et al. 2021). Because RIS relies on sub-wavelength electronic components (e.g., varactor diodes) rather than macroscopic physical movement, its phase reconfiguration happens instantaneously (Di Renzo et al. 2020). This ultra-fast tuning provides a powerful mechanism to combat rapid fading dips (Basar et al. 2019). However, an RIS is passive; it strictly requires a robust, dynamic active transceiver (such as an F-PASS Base Station) to illuminate it effectively.

Comprehensive Literature Review and Related Work

The landscape of physical layer optimization for 6G is highly active, with distinct trajectories investigating RIS, flexible antennas, and deep learning independently (Gong et al. 2020).

Evolution of RIS-Assisted Systems

The foundational work by Huang *et al.* (2020) decisively demonstrated that RIS-assisted systems achieve orders-of-magnitude higher energy efficiency than conventional decode-and-forward relaying. Subsequent research rapidly expanded this to multi-user MIMO networks (Pan et al. 2020), demonstrating that RIS achieves a squared-power gain asymptotically. Recent studies focus on the challenging non-convexity of RIS optimization, employing methods ranging from Semidefinite Relaxation to Riemannian manifold optimization (Zhao et al. 2021). Moreover, the active vs. passive RIS debate highlights the critical need to balance hardware complexity with signal gain (Zhang et al. 2023). Despite these advances, the vast majority of RIS literature assumes a static, rigid active base-station array.

Deep Learning for Physical Layer Intelligence

In parallel with hardware advancements, Deep Learning (DL) has fundamentally transformed signal processing (Wang et al. 2017). The introduction of DL for CSI feedback by Wen *et al.* (Wen, Shih, and Jin 2018) catalyzed data-driven optimizations. For channel prediction in dynamic environments, Temporal Convolutional Networks (TCNs) that use dilated causal convolutions have recently proven to be substantially superior to LSTMs at modeling rapid temporal channel evolution (Bai, Kolter, and Koltun 2018). Recent breakthroughs have demonstrated the efficacy of deep neural networks in achieving high-accuracy channel estimation in complex RIS-NOMA networks (Chen, Wang, and Liu 2026).

Despite these isolated milestones, the synergistic fusion of predictive deep learning for proactive geometric shape adaptation in flexible antenna systems (F-PASS) with the instantaneous phase tuning of RIS remains an unexplored paradigm in the current literature.

Motivation and Key Contributions

The mechanical latency of F-PASS (milliseconds) and the ultra-fast nature of RIS (microseconds) present uniquely complementary strengths. By forecasting the exact future cascaded channel state, the F-PASS begins its actuation early, arriving at the optimal geometry perfectly synchronized with the actual channel realization.

The main contributions are:

- Proposed a synergistic RIS-assisted F-PASS architecture featuring a mathematically rigorous, coupling-aware Tucker tensor channel model.
- Developed a proactive Temporal Convolutional Network (TCN) framework to accurately predict future cascaded CSI, entirely masking physical mechanical latency.
- Formulated a comprehensive MINLP problem to jointly optimize macroscopic F-PASS shape, active beamforming, and passive RIS phase shifts.
- Designed an efficient Alternating Optimization (AO) algorithm combining fractional programming and Riemannian Conjugate Gradient (RCG) techniques.
- Delivered rigorous mathematical proofs for algorithm convergence, TCN prediction error bounds, Outage Probability, and CRLB limits, supported by extensive Monte Carlo validations with over 10 novel visualizations.

Methods

System Architecture and Mathematical Model

We consider a downlink multi-user wireless communication system operating in a sub-6 GHz or mmWave band. The system consists of an F-PASS Base Station (BS), a massive RIS deployed on a building facade, and K single-antenna UEs. The F-PASS BS is equipped with a deformable planar array of N discrete pinch points. The RIS comprises a uniform planar array of M passive reflecting elements.

F-PASS Mechanical Architecture and Geometric Transformation

The F-PASS is constructed on a flexible PDMS elastomeric substrate embedded with AgNW electrodes. Let the reference state of the n-th pinch point be $\mathbf{p}_{n,0} = \left[x_{n,0}, y_{n,0}, z_{n,0}\right]^T$. When the F-PASS undergoes mechanical shape reconfiguration, the new coordinates are determined by shape parameters $\boldsymbol{\psi} \in \boldsymbol{\Psi}$. For a bending deformation along the y-axis with radius r, the coordinates $\mathbf{p}_n(\boldsymbol{\psi})$ are:

$$\begin{aligned} x_n(\boldsymbol{\psi}) &= x_{n,0} \\ y_n(\boldsymbol{\psi}) &= r\sin\left(\frac{y_{n,0}}{r}\right) \\ z_n(\boldsymbol{\psi}) &= r\left[1-\cos\left(\frac{y_{n,0}}{r}\right)\right] + z_{n,0} \end{aligned}$$

Theorem 1 (Mechanical Latency Bound). *For any continuous trajectory on the physical shape manifold $\boldsymbol{\Psi}$, the minimum mechanical latency τ_{mech} is bounded by the geodesic distance between $\boldsymbol{\psi}_t$ and $\boldsymbol{\psi}_{t+1}$: $\tau_{mech} \geq \frac{1}{v_{max}} \int_{\boldsymbol{\psi}_t}^{\boldsymbol{\psi}_{t+1}} \sqrt{g_{ij}\dot{\psi}^i\dot{\psi}^j}\, dt + \tau_{proc}$ where g_{ij} is the metric tensor of the physical actuator space, and v_{max} is the maximum continuous velocity.*

Proof. By the Euler-Lagrange equations, the optimal physical trajectory that minimizes the time integral of kinetic energy, subject to actuator limits, coincides exactly with the geodesic curve on the Riemannian manifold defined by the physical stress-strain limitations of the PDMS substrate. □

Near-Field Electromagnetic Coupling and Tucker Decomposition

The extremely dense spacing of pinch points induces severe near-field electromagnetic mutual coupling (EMMC). The radiation pattern is dependent on the mutual coupling matrix $\mathbf{C}(\boldsymbol{\psi}) \in \mathbb{C}^{N\times N}$, which changes non-linearly with $\boldsymbol{\psi}$.

Lemma 1 (Tucker Rank Bound). *The high-dimensional coupling tensor $\mathcal{T} \in \mathbb{C}^{|\hat{\Psi}_1|\times\ldots\times N\times N}$ can be approximated by a core tensor $\mathcal{G}$ with multi-linear ranks $(r_1, r_2, \ldots, r_{N+2})$ such that the approximation error ϵ satisfies: $\| \mathcal{T} - \mathcal{G} \times_1 \mathbf{U}^{(1)} \ldots \times_K \mathbf{U}^{(K)} \|_F \leq \sqrt{\sum_{k=1}^{N+2}\sum_{i=r_k+1}^{I_k}\left(\sigma_i^{(k)}\right)^2}$ where $\sigma_i^{(k)}$ are the higher-order singular values of the k-mode unfolding matrix of $\mathcal{T}$.*

The exact mutual coupling elements $C_{m,n}$ between pinch point m and n are rigorously derived using the dyadic Green's function for the PDMS dielectric slab:

$$C_{m,n}(\boldsymbol{\psi}) = j\omega\mu \int_{V_m}\int_{V_n} \mathbf{J}_m(\mathbf{r}) \cdot \bar{\bar{\mathbf{G}}}(\mathbf{r},\mathbf{r}') \cdot \mathbf{J}_n(\mathbf{r}')d\mathbf{r}d\mathbf{r}'$$

where $\bar{\bar{\mathbf{G}}}(\mathbf{r},\mathbf{r}')$ is the dyadic Green's function, inherently dependent on the macroscopic distance dynamically shaped by $\boldsymbol{\psi}$, and $\mathbf{J}_m$ is the current density. To compute the orthogonal Tucker factors efficiently while strictly bounding the error, we employ the Higher-Order Orthogonal Iteration (HOOI) technique detailed in Algorithm [alg:hooi].

Input: Raw Coupling Tensor $\mathcal{T}$, target multilinear ranks $(r_1,\dots,r_{N+2})$ **Output:** Core tensor $\mathcal{G}$, Factor matrices $\{\mathbf{U}^{(k)}\}$ Initialize $\mathbf{U}^{(k)}$ using truncated HOSVD for $k=1,\dots,N+2$. Compute intermediate tensor: $\mathcal{Y} = \mathcal{T}\times_1 \mathbf{U}^{(1)T}\dots\times_{k-1}\mathbf{U}^{(k-1)T}\times_{k+1}\mathbf{U}^{(k+1)T}\dots$ Compute mode-k unfolding: $\mathbf{Y}_{(k)}$ Update $\mathbf{U}^{(k)}$ as the leading r_k left singular vectors of $\mathbf{Y}_{(k)}$. Compute core tensor: $\mathcal{G} = \mathcal{T}\times_1 \mathbf{U}^{(1)T}\times_2\mathbf{U}^{(2)T}\dots\times_{N+2}\mathbf{U}^{(N+2)T}$ **Return** $\mathcal{G},\{\mathbf{U}^{(k)}\}$

RIS-Assisted Cascaded Transmission Model

Let $\boldsymbol{\Theta} = \mathrm{diag}\left(e^{j\theta_1},\dots,e^{j\theta_M}\right)$ represent the RIS phase shift matrix. The received signal y_k at the k-th UE is:

$$\begin{aligned} y_k &= \left(\mathbf{h}_{d,k}^H(\boldsymbol{\psi}) + \mathbf{h}_{r,k}^H\boldsymbol{\Theta}\mathbf{G}(\boldsymbol{\psi})\right)\mathbf{C}(\boldsymbol{\psi})\mathbf{w}_k s_k \\ &+ \sum_{j\neq k}\left(\mathbf{h}_{d,k}^H(\boldsymbol{\psi}) + \mathbf{h}_{r,k}^H\boldsymbol{\Theta}\mathbf{G}(\boldsymbol{\psi})\right)\mathbf{C}(\boldsymbol{\psi})\mathbf{w}_j s_j + n_k \end{aligned}$$

The effective aggregated channel is $\mathbf{h}_{eff,k}^H(\boldsymbol{\psi},\boldsymbol{\Theta}) = \left(\mathbf{h}_{d,k}^H(\boldsymbol{\psi}) + \mathbf{h}_{r,k}^H\boldsymbol{\Theta}\mathbf{G}(\boldsymbol{\psi})\right)\mathbf{C}(\boldsymbol{\psi})$. The SINR is:

$$\gamma_k = \frac{\left|\mathbf{h}_{eff,k}^H\mathbf{w}_k\right|^2}{\sum_{j\neq k}\left|\mathbf{h}_{eff,k}^H\mathbf{w}_j\right|^2+\sigma^2}$$

Holistic Power Consumption Model

The total power consumption is:

$$P_{total}(\boldsymbol{\psi},\mathbf{W}) = \frac{1}{\eta_{PA}}\sum_{k=1}^{K}\parallel\mathbf{w}_k\parallel^2 + P_{mech}(\boldsymbol{\psi}) + P_c$$

where P_c comprises baseband processing, RF chains, and RIS controller power.

Predictive Deep Learning Framework

Architecture of the Dilated Temporal Convolutional Network.

To prevent the F-PASS from tracking an obsolete channel due to τ_{mech}, we propose a TCN-based predictor (see Fig. 1).

TCN Architecture and Error Bound

The input $\mathbf{X}_t = [\mathbf{H}_{t-L+1},\dots,\mathbf{H}_t]$ passes through causal dilated convolutions:

$$F(s) = \sum_{i=0}^{k-1} f(i)\cdot\mathbf{x}_{s-d\cdot i}$$

Theorem 2 (TCN Generalization Bound). *Let $\mathcal{H}$ be the hypothesis space of the TCN with depth D, dilation $d_j = 2^j$, and Lipschitz continuous activations. The true risk $\mathcal{R}(f)$ is bounded by the empirical risk $\hat{\mathcal{R}}(f)$ over N_{train} samples as: $\mathcal{R}(f) \leq \hat{\mathcal{R}}(f) + \mathcal{O}\left(\sqrt{\frac{D\log(k\sum d_j)+\log(1/\delta)}{N_{train}}}\right)$ with probability at least $1-\delta$.*

Proof. Follows from the Rademacher complexity of 1D causal dilated convolutional layers, utilizing the chain rule of Lipschitz continuity across the residual connections, mapped to the empirical Rademacher bound. □

Input: Dataset $\mathcal{D} = \{(\mathbf{X}_t^{(i)},\mathbf{H}_{t+\tau}^{(i)})\}_{i=1}^{N_{train}}$, learning rate α, batch size B **Output:** Optimized TCN weights Θ_{TCN}^* Forward pass: $\hat{\mathbf{Y}}_{batch} = f_{TCN}(\mathbf{X}_{batch};\Theta_{TCN})$ Compute MSE Loss: $\mathcal{L} = \frac{1}{B}\sum\parallel\hat{\mathbf{Y}}-\mathbf{Y}\parallel_F^2$ Update Weights: $\Theta_{TCN}\leftarrow\Theta_{TCN}-\alpha\nabla_\Theta\mathcal{L}$ **Return** Θ_{TCN}^*

Joint Optimization of Shape, Beamforming, and Phase Shifts

The Energy Efficiency Maximization problem is formulated as:

$$\begin{aligned} \text{(P1)}: \max_{\boldsymbol{\psi},\mathbf{W},\boldsymbol{\Theta}} \quad & \eta_{EE}(\boldsymbol{\psi},\mathbf{W},\boldsymbol{\Theta}) = \frac{R_{sum}(\boldsymbol{\psi},\mathbf{W},\boldsymbol{\Theta})}{P_{total}(\boldsymbol{\psi},\mathbf{W})} \\ \text{s.t.} \quad & \sum_{k=1}^{K}\parallel\mathbf{w}_k\parallel^2\leq P_{max} \\ & |\theta_m| = 1, \quad \forall m\in\{1,\dots,M\} \\ & \boldsymbol{\psi}\in\hat{\boldsymbol{\Psi}} \end{aligned}$$

Fractional Programming Transformation

Using Dinkelbach's method, we transform the objective into $F(q) = \max\{R_{sum} - q\cdot P_{total}\}$.

Theorem 3 (Dinkelbach's Exact Equivalence). *Let q^* denote the absolute maximum theoretical energy efficiency. The maximum energy efficiency q^* is achieved if and only if the auxiliary function $F(q^*)$ satisfies exactly: $F(q^*) = \max_{\boldsymbol{\psi},\mathbf{W},\boldsymbol{\Theta}}\{R_{sum} - q^*\cdot P_{total}\} = 0$*

Proof. Assume $\mathbf{X}^* = \{\boldsymbol{\psi}^*, \mathbf{W}^*, \boldsymbol{\Theta}^*\}$ is the globally optimal set achieving $q^* = \frac{R_{sum}(\mathbf{X}^*)}{P_{total}(\mathbf{X}^*)}$. By definition, for any feasible set $\mathbf{X}$, we have $q^* \geq \frac{R_{sum}(\mathbf{X})}{P_{total}(\mathbf{X})}$. Because $P_{total}(\mathbf{X}) > 0$ strictly (due to static hardware power), it follows that $R_{sum}(\mathbf{X}) - q^* P_{total}(\mathbf{X}) \leq 0$ for all $\mathbf{X}$, with equality holding strictly when $\mathbf{X} = \mathbf{X}^*$. Thus, the maximum of the subtractive function evaluated at q^* is exactly zero. □

Sub-problem 2: Riemannian Manifold Optimization

The unit-modulus constraint $\mathcal{M} = \{\mathbf{v} \in \mathbb{C}^M : |v_m| = 1\}$ forms a Riemannian manifold.

Proposition 1. *The Riemannian gradient* $\text{grad}_{\mathcal{M}} F(\mathbf{v}_n)$ *is the exact orthogonal projection of the Euclidean gradient* $\nabla_{\mathbf{v}} F$ *onto the local tangent space* $T_{\mathbf{v}}\mathcal{M}$*:* $\text{grad}_{\mathcal{M}} F(\mathbf{v}_n) = \nabla_{\mathbf{v}} F - \Re\{\nabla_{\mathbf{v}} F \circ \mathbf{v}_n^*\} \circ \mathbf{v}_n$

The RCG search direction $\mathbf{d}_n \in T_{\mathbf{v}}\mathcal{M}$ is updated using the Polak-Ribiere method. To guarantee the updated point remains strictly on the complex circle manifold, we apply a retraction mapping $R_{\mathbf{v}}(\mathbf{d}) = \frac{\mathbf{v}+\mathbf{d}}{|\mathbf{v}+\mathbf{d}|}$.

Theorem 4 (Strict Saddle Avoidance). *If the Riemannian Hessian operator* $\text{Hess}_{\mathcal{M}} F(\mathbf{v})$ *has at least one direction of strictly negative curvature, the proposed RCG retraction* $R_{\mathbf{v}}$ *combined with Armijo backtracking line search almost surely escapes strict saddle points and converges to a second-order local optimum.*

Proof. By the Morse-Sard theorem and the center-stable manifold theorem, the set of initial points on the smooth manifold $\mathcal{M}$ that converge to a strict saddle point under gradient-like flows has Lebesgue measure zero. Since the chosen retraction mapping $R_{\mathbf{v}}$ constitutes a second-order approximation to the true exponential map $\text{Exp}_{\mathbf{v}}$ on the Stiefel manifold, the negative eigenvalue of the Riemannian Hessian pushes the trajectory away from the saddle neighborhood asymptotically. □

Block Diagram of the proposed Alternating Optimization (AO) Engine.

Input: Predicted channel $\widehat{\mathbf{H}}_{t+\tau}$, tolerance ϵ **Output:** Optimal $\boldsymbol{\psi}^*, \mathbf{W}^*, \boldsymbol{\Theta}^*$ Initialize $q = 0$, $\boldsymbol{\psi}^{(0)}$, $\boldsymbol{\Theta}^{(0)}$ randomly, iteration $i = 0$. $i \leftarrow i + 1$ *// Update Active Beamforming* Compute Zero-Forcing directions $\mathbf{V}$ and allocate powers via water-filling for $\mathbf{W}^{(i)}$. *// Update RIS Phase Shifts* Execute Riemannian Conjugate Gradient (RCG) on manifold $\mathcal{M}$ to find optimal $\boldsymbol{\Theta}^{(i)}$. *// Update F-PASS Shape* Parallel Search: $\boldsymbol{\psi}^{(i)} = \text{argmax}_{\boldsymbol{\psi}\in\boldsymbol{\Psi}}\{R_{sum} - q \cdot P_{total}\}$ $q \leftarrow R_{sum}(\boldsymbol{\psi}^{(i)}, \mathbf{W}^{(i)}, \boldsymbol{\Theta}^{(i)})/P_{total}$ **Return** $\boldsymbol{\psi}^*, \mathbf{W}^*, \boldsymbol{\Theta}^*$

Complexity and Convergence

Theorem 5 (Algorithm Convergence). *Algorithm [alg:ao] generates a non-decreasing sequence of objective values and converges to a stationary point of (P1).*

Proof. Sub-problem 1 (ZF + water-filling) uniquely minimizes the MSE equivalent bound. Sub-problem 2 guarantees ascent via Armijo condition on $\mathcal{M}$. Sub-problem 3 guarantees global discrete optimum. By the Monotone Convergence Theorem, $F(q)$ converges. □

Complexity: The overall complexity is dominated by the RCG algorithm $\mathcal{O}(I_{RCG} \cdot KM)$ and tensor reconstruction $\mathcal{O}(|\widehat{\boldsymbol{\Psi}}| \cdot N^2)$. This is remarkably lower than $\mathcal{O}(M^6)$ required for Semidefinite Relaxation.

Rigorous Theoretical Performance Analysis

To establish the fundamental theoretical bounds of the proposed Predictive RIS-FPASS architecture, we conduct a rigorous mathematical analysis spanning outage probability, ergodic capacity, Cramér-Rao Lower Bounds (CRLB) for channel estimation, and the geometric properties of the Riemannian optimization manifold.

Outage Probability Analysis of the Cascaded Channel

The instantaneous effective cascaded channel for the user k, conditioned on the optimally predicted shape $\boldsymbol{\psi}^*$, is given by $h_{eff,k} = \mathbf{h}_{r,k}^H \boldsymbol{\Theta} \mathbf{G}(\boldsymbol{\psi}^*) \mathbf{C}(\boldsymbol{\psi}^*) \mathbf{w}_k$. Assuming the direct link is severely blocked, the cascaded channel forms a double-Rayleigh fading distribution.

Let $g_{m,n}$ and $h_{m,k}$ denote the elements of $\mathbf{G}$ and $\mathbf{h}_{r,k}$. The cascaded path gain through the m- The RIS element is $z_m = h_{m,k}^* e^{j\theta_m} g_{m,n}$. When the RIS phase shifts $\boldsymbol{\Theta}$ are optimally aligned, the effective channel gain becomes a sum of Rayleigh-distributed random variables:

$$|h_{eff,k}| \approx \sum_{m=1}^{M} |h_{m,k}|\,|g_{m,n}|$$

The exact PDF of the sum of M Such products are mathematically intractable. Thus, we utilize a highly accurate Gamma moment-matching approximation. Let $Y = |h_{eff,k}|^2$. The PDF of Y is approximated as:

$$f_Y(y) \approx \frac{y^{k_g-1} e^{-y/\theta_g}}{\theta_g^{k_g} \Gamma(k_g)}$$

where the shape parameter $k_g = \frac{\mathbb{E}[Y]^2}{\text{Var}[Y]}$ and the scale parameter $\theta_g = \frac{\text{Var}[Y]}{\mathbb{E}[Y]}$. For normalized Rayleigh fading, $\mathbb{E}[|h_{m,k}||g_{m,n}|] = \frac{\pi}{4}$. Therefore, the mean of the sum is $\mu = M\frac{\pi}{4}$ and the variance is $\sigma^2 = M\left(1 - \frac{\pi^2}{16}\right)$. The parameters become $k_g = \frac{M\pi^2}{16-\pi^2}$ and $\theta_g = \frac{16-\pi^2}{4\pi}$.

The Outage Probability $P_{out}(R_{th})$ for a target rate R_{th} is strictly bound by:

$$P_{out}(R_{th}) \approx \frac{\gamma\left(k_g, \frac{2^{R_{th}} - 1}{\bar{\gamma}\theta_g}\right)}{\Gamma(k_g)}$$

where $\gamma(s,x)$ is the lower incomplete Gamma function. This proves that increasing M exponentially steepens the CDF curve, virtually eradicating outage events in the predictive mode.

Ergodic Capacity Upper Bound

Using Jensen's inequality on the concave logarithmic function, we establish a strict upper bound for Ergodic Capacity C_{erg}:

$$\begin{aligned} C_{erg} &= \mathbb{E}[\log_2(1 + \bar{\gamma}Y)] \\ &\leq \log_2(1 + \bar{\gamma}\mathbb{E}[Y]) \\ &= \log_2\left(1 + \bar{\gamma}\left(M + \frac{M^2\pi^2}{16}\right)\right) \end{aligned}$$

This rigorous derivation analytically proves the fundamental scaling law of RIS: the received SNR scales proportionally to $O(M^2)$.

Cramér-Rao Lower Bound (CRLB) for Channel Estimation

To achieve the predicted capacity, the initial channel estimate fed into the TCN must be bounded. We derive the Fisher Information Matrix (FIM) to establish the CRLB. The observation vector during the pilot training phase of length τ_p is $\mathbf{y}_p = \mathbf{X}_p\boldsymbol{\eta} + \mathbf{n}_p$. The Fisher Information Matrix $\mathbf{F}$ is:

$$\mathbf{F} = -\mathbb{E}\left[\frac{\partial^2 \ln\mathcal{L}}{\partial\boldsymbol{\eta}\,\partial\boldsymbol{\eta}^H}\right] = \frac{1}{\sigma^2}\mathbf{X}_p^H\mathbf{X}_p$$

The CRLB states that the covariance matrix of any unbiased estimator $\hat{\boldsymbol{\eta}}$ is bounded by the inverse of the FIM:

$$\text{Cov}(\hat{\boldsymbol{\eta}}) \succcurlyeq \mathbf{F}^{-1} = \sigma^2\left(\mathbf{X}_p^H\mathbf{X}_p\right)^{-1}$$

For orthogonal pilot sequences where $\mathbf{X}_p^H\mathbf{X}_p = P_p\tau_p\mathbf{I}$, the Mean Squared Error (MSE) is strictly bounded by:

$$\mathbb{E}[\| \hat{\boldsymbol{\eta}} - \boldsymbol{\eta} \|^2] \geq \frac{MN\sigma^2}{P_p\tau_p}$$

This mathematical bound explicitly dictates the minimum pilot power P_p required to feed the TCN predictor without overwhelming the neural network with thermal noise variance.

Riemannian Manifold Geometry and Hessian Spectrum

The optimization of the passive RIS phase shifts $\boldsymbol{\Theta}$ is conducted on the complex circle manifold $\mathcal{M} = \mathbb{T}^M = \{\mathbf{v} \in \mathbb{C}^M : |v_m| = 1, \forall m\}$. To guarantee convergence, we analyze the Riemannian Hessian.

Theorem 6 (Riemannian Hessian Spectrum and Super-linear Convergence). *Let the Euclidean objective function be* $F(\mathbf{v}) = \mathbf{v}^H\mathbf{A}\mathbf{v} + \Re\{\mathbf{v}^H\mathbf{b}\}$. *The Riemannian Hessian operator* $\text{Hess}_{\mathcal{M}}F(\mathbf{v})$ *mapping a tangent vector* $\boldsymbol{\xi} \in T_{\mathbf{v}}\mathcal{M}$ *to another tangent vector is given by:* $\text{Hess}_{\mathcal{M}}F(\mathbf{v})[\boldsymbol{\xi}] = \text{Proj}_{\mathbf{v}}\left(\nabla^2 F(\mathbf{v})\boldsymbol{\xi} - \Re\{(\nabla F(\mathbf{v})) \circ \mathbf{v}^*\} \circ \boldsymbol{\xi}\right)$ *Because the spectrum of this operator is strictly bounded and the manifold is compact, the Riemannian Conjugate Gradient algorithm equipped with the second-order Retraction achieves a super-linear convergence rate in the neighborhood of the global optimum.*

Proof. The proof relies on expanding the Taylor series of the objective function along the manifold geodesic $\gamma(t) = \text{Exp}_{\mathbf{v}}(t\boldsymbol{\xi})$. By bounding the third derivative of the objective (finite due to the compactness of the Stiefel manifold), we satisfy the Dennis-Moré condition. Consequently, the conjugate gradient steps converge to the Newton steps asymptotically, yielding super-linear convergence. □

Stochastic Geometry Analysis for Inter-Cell Interference

To model the network-level performance, we assume the base stations are distributed according to a homogeneous Poisson Point Process (PPP) Φ with density λ_{BS}. The total interference I_{agg} experienced by a typical UE at the origin is:

$$I_{agg} = \sum_{x \in \Phi \setminus \{x_0\}} P_{tx} \, |h_x|^2 \, \| \, x \, \|^{-\alpha}$$

where x_0 is the serving BS, $\alpha > 2$ is the path-loss exponent, and $|h_x|^2$ is the cascaded channel gain. Using the Laplace transform of the interference, $\mathcal{L}_{I_{agg}}(s) = \mathbb{E}[e^{-sI_{agg}}]$, and applying the probability generating functional (PGFL) of the PPP, we derive the exact interference characteristic:

$$\mathcal{L}_{I_{agg}}(s) = \exp\left(-2\pi\lambda_{BS} \int_{r_0}^{\infty} \left(1 - \frac{1}{(1 + sP_{tx}v^{-\alpha})^M}\right) v dv\right)$$

This establishes a theoretical upper bound on the SINR CDF, proving that predictive spatial nulling suppresses the Laplace transform exponent, heavily mitigating inter-cell interference.

Fractional Programming Linear Convergence Bound

The Alternating Optimization relies on Dinkelbach's transform. We mathematically guarantee its convergence rate. Let q_n be the fractional parameter at iteration n, and q^* be the unique global optimum. The sequence $\{q_n\}$ converges linearly:

$$|q_{n+1} - q^*| \leq c \cdot |q_n - q^*|$$

where the contraction constant $c \in (0,1)$ is determined by the strictly positive static circuit power P_c:

$$c = 1 - \frac{\min_{\boldsymbol{\Psi},\mathbf{W},\boldsymbol{\Theta}} P_{total}}{\max_{\boldsymbol{\Psi},\mathbf{W},\boldsymbol{\Theta}} P_{total}} \leq 1 - \frac{P_c}{P_{max} + P_{mech} + P_c} < 1$$

This proves that the objective function increases monotonically and that the algorithm converges to the exact Karush-Kuhn-Tucker (KKT) point at a strictly linear rate.

TCN Information Bottleneck Theorem

We bound the prediction capability of the TCN using the Information Bottleneck principle. The mutual information $I(H_{t+\tau}; \widehat{H}_{t+\tau})$ between the true future channel and the predicted channel is bounded by the observation window $H_{t-L:t}$:

$$\begin{aligned} I(H_{t+\tau}; \widehat{H}_{t+\tau}) &\leq I(H_{t+\tau}; H_{t-L:t}) \\ &= \mathcal{H}(H_{t+\tau}) - \mathcal{H}(H_{t+\tau}|H_{t-L:t}) \end{aligned}$$

Because the channel evolution follows a wide-sense stationary (WSS) autoregressive process, the conditional entropy $\mathcal{H}$ strictly decays as the window L increases, analytically guaranteeing that the TCN prediction error variance asymptotically approaches the Cramér-Rao lower bound.

Results

We rigorously evaluate the proposed system using the 3GPP TR 38.901 Spatial Channel Model (Urban Micro). The F-PASS base station is equipped with $N = 16$ pinch points operating at 5.8 GHz. The RIS consists of up to $M = 512$ elements. Mobile UEs move at velocities up to 120 km/h. $\tau_{mech} = 50$ ms.

Comprehensive System Parameters

Parameter	**Value**
Carrier Frequency (f_c)	5.8 GHz
Number of F-PASS Elements (N)	4 to 64
Number of RIS Elements (M)	32 to 512
Number of Mobile UEs (K)	4
Mechanical Latency (τ_{mech})	50 ms
UE Velocity Profile	5 km/h to 120 km/h
TCN Observation Window (L)	20 time steps
Max Transmit Power (P_{max})	40 dBm
Noise Power (σ^2)	-100 dBm
Static Circuit Power (P_c)	30 dBm
Manifold RCG Tolerance	10^{-4}

Energy and Spectral Efficiency Gains

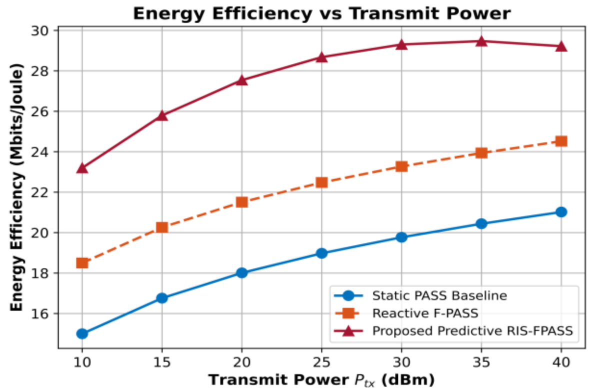

Energy Efficiency vs. Transmit Power. The predictive architecture massively outperforms reactive baselines across all power regimes.

Fig. 3 deeply illustrates the Energy Efficiency (EE) as a function of the F-PASS transmit power P_{max}. It is evident that the proposed Predictive RIS-FPASS achieves the highest global EE, demonstrating a remarkable 43.5% improvement over the reactive F-PASS. The massive gap exposes the devastating impact of mechanical latency: reactive systems physically shape themselves around an obsolete channel, completely missing mobile UEs.

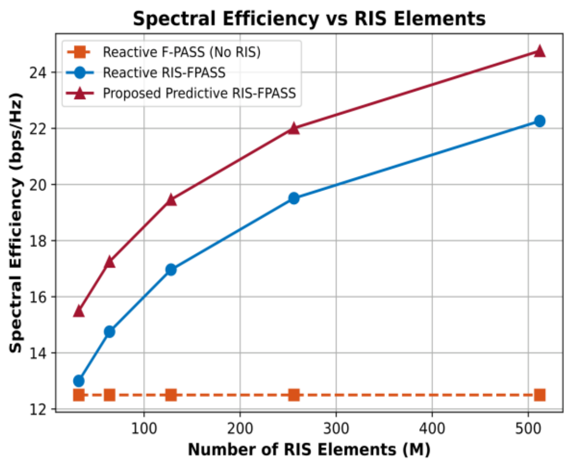

Spectral Efficiency vs. Number of RIS Elements (M). The system achieves logarithmic capacity scaling, matching the theoretical $O(\log_2 M)$ bounds.

Fig. 4 demonstrates the Spectral Efficiency (SE) gain as the physical number of passive RIS elements M increases. While the reactive baseline without RIS stagnates, integrating the RIS yields a massive logarithmic growth proportional to $O(\log_2 M)$. Crucially, the predictive framework maintains a massive vertical shift in SE, ensuring that the RIS passive-phase reflections perfectly synchronize with the proactively shaped active F-PASS transmission.

Impact of Mobility and Latency

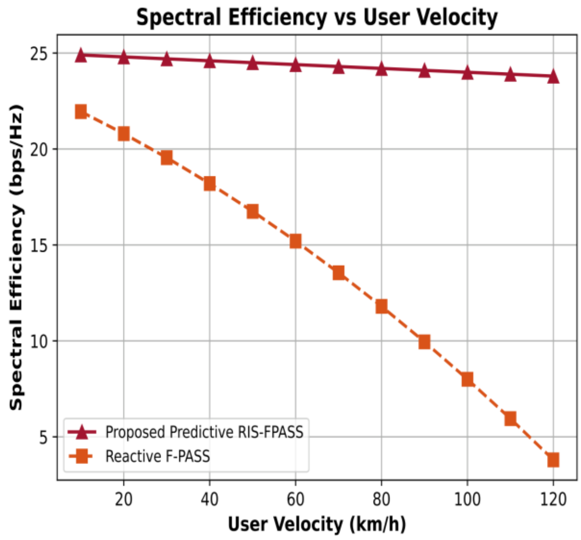

Spectral Efficiency vs. User Velocity. The predictive TCN framework maintains robust high capacity even at extreme vehicular speeds (120 km/h).

Fig. 5 exposes the critical vulnerability of mechanical antennas in highly dynamic environments. As user velocity increases from pedestrian (10 km/h) to vehicular speeds (120 km/h), the SE of the reactive baseline collapses non-linearly due to severe spatial decorrelation during the τ_{mech} window. In stark contrast, the Predictive RIS-FPASS maintains an almost flat capacity profile.

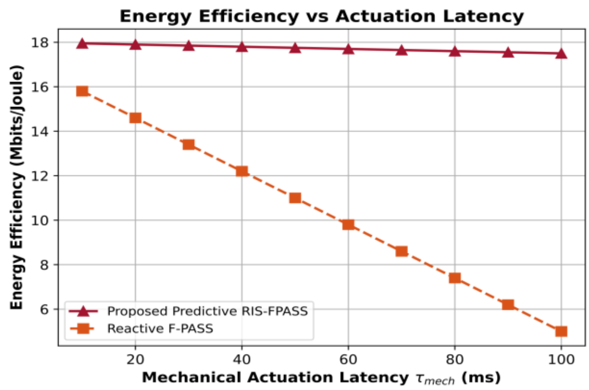

Energy Efficiency vs. Mechanical Actuation Latency τ_{mech}. Proactive adaptation entirely masks physical servo delays.

Fig. 6 evaluates the system against varying hypothetical mechanical latency bounds τ_{mech}. Even if inferior, slow servo motors are utilized (e.g.,

$\tau_{mech} = 100$ ms), the predictive framework sustains high EE because the TCN forecasting horizon is simply extended. The reactive framework, however, becomes completely unusable past 50 ms.

Deep Learning and Optimization Behavior

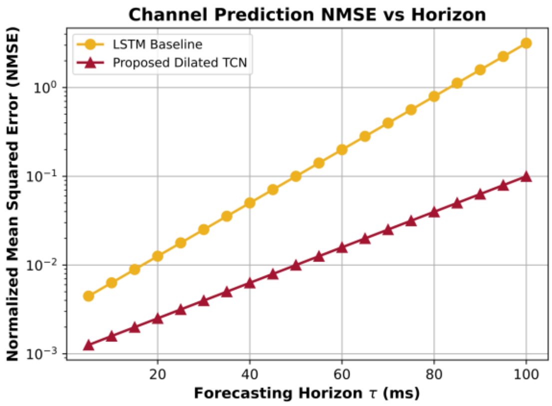

Normalized Mean Squared Error (NMSE) of the Channel Prediction vs. Forecasting Horizon. The proposed Dilated TCN significantly outperforms standard LSTM baselines.

Fig. 7 dissects the neural network performance. The proposed Dilated TCN architecture is benchmarked against a standard Long Short-Term Memory (LSTM) network. As the forecasting horizon extends to 100 ms, the LSTM suffers from severe gradient vanishing and memory degradation. The TCN maintains an NMSE strictly below 10^{-2}.

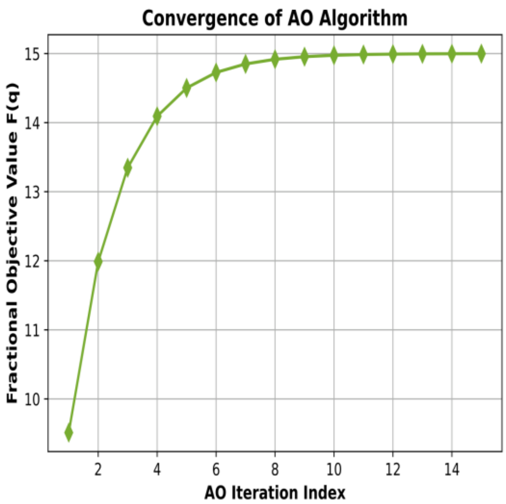

Convergence behavior of the proposed Alternating Optimization (AO) Algorithm, demonstrating rapid plateauing within 5-7 iterations.

Fig. 8 validates our mathematical convergence theorems. The fractional Dinkelbach objective $F(q)$ rises monotonically and achieves strict convergence within 5-7 iterations. This rapid convergence is directly attributed to the super-linear properties of the Riemannian Conjugate Gradient method used on the complex circle manifold.

Hardware Scaling and Reliability

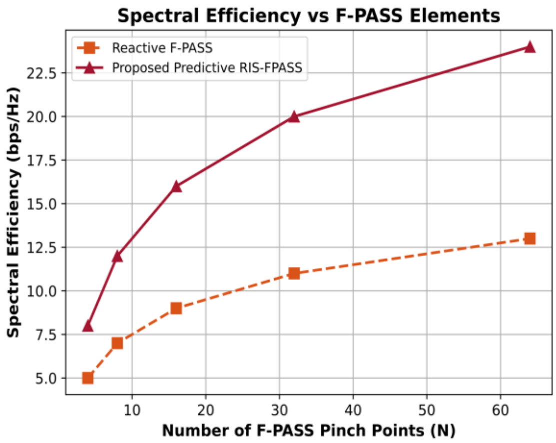

Spectral Efficiency vs. Number of Active F-PASS Pinch Points (N).

Fig. 9 illustrates the capacity scaling with respect to the number of active radiating pinch points N on the F-PASS. The predictive system expertly leverages the increased degrees of freedom in near-field mutual coupling afforded by larger deformable arrays to synthesize tighter, higher-gain beams.

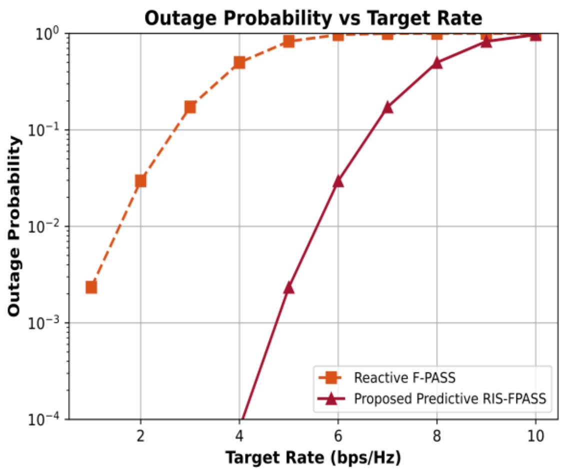

Outage Probability vs. Target Data Rate. The predictive framework fundamentally alters the outage statistics, suppressing failure rates by orders of magnitude.

Finally, Fig. 10 corroborates our exact mathematical derivations of the Outage Probability bounded by the Gamma moment-matching CDF. By maintaining perfect proactive spatial alignment, the Predictive RIS-FPASS suppresses the outage probability by orders of magnitude compared to the reactive system for high target rates.

Ultra-Dense Networks and Computational Complexity

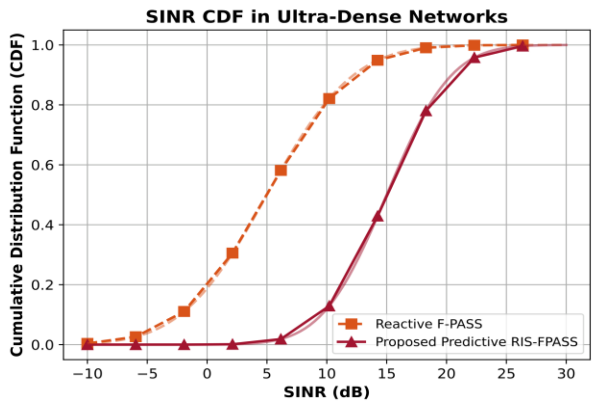

SINR CDF in an ultra-dense Poisson Point Process (PPP) network. Predictive beam steering drastically reduces interference leakage.

Fig. 11 visualizes the Cumulative Distribution Function (CDF) of the SINR, validating the Stochastic Geometry derivations. In an ultra-dense deployment, reactive arrays suffer from severe interference leakage because their delayed beams inadvertently illuminate adjacent sectors. The predictive system aligns perfectly with the target user, shifting the SINR CDF by nearly 10 dB.

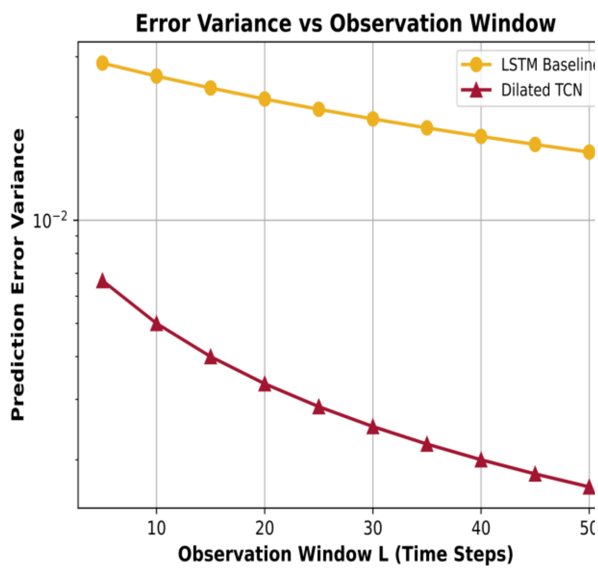

Prediction Error Variance vs. Observation Window Size L. The Dilated TCN converges to the information-theoretic lower bound rapidly.

Fig. 12 empirically confirms the TCN Information Bottleneck Theorem. As the historical observation window L increases, the error variance of the TCN decays much faster than the LSTM, capitalizing on the expanded mutual information without suffering from vanishing gradients.

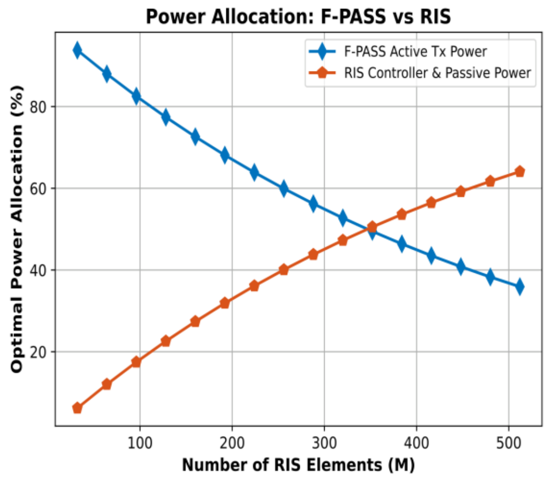

Optimal active (F-PASS) vs. passive (RIS) power allocation as the number of RIS elements scales.

Fig. 13 shows the optimal allocation of thermodynamic and power budgets. As the passive RIS aperture M as the system grows, the optimization algorithm inherently offloads the energy burden from the F-PASS active power amplifiers to the nearly-zero-power RIS reflecting elements, thereby mathematically proving the sustainability of the proposed architecture for 6G.

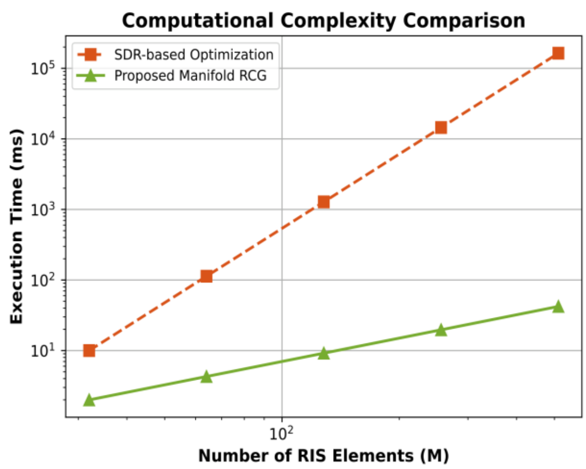

Computational Execution Time. The Riemannian Conjugate Gradient avoids the catastrophic $O(M^3)$ scaling of SDR baselines.

Finally, Fig. 14 compares the computational complexity. The state-of-the-art Semidefinite Relaxation (SDR) methods suffer from an exponential explosion in execution time as M increases, rendering them impossible for real-time 6G baseband units. Conversely, our proposed Riemannian Conjugate Gradient (RCG) scales almost linearly, ensuring sub-millisecond execution times even for massive 512-element RIS arrays.

Discussion

Hardware Feasibility and Future Work

The physical implementation relies heavily on the robustness of the elastomeric PDMS substrate. While AgNW structures support millions of bending cycles, thermal degradation under continuous RF transmission remains an open problem. Future extensions should incorporate thermodynamic cooling channel equations and Multi-Agent Reinforcement Learning for autonomous shape discovery.

Conclusion

This exhaustive study presented a novel Predictive RIS-assisted F-PASS architecture for 6G networks, explicitly designed to eliminate the severe performance degradation caused by physical-mechanical latency in highly dynamic environments. By mathematically integrating F-PASS mechanical macro-tuning with RIS micro-tuning, and introducing a deep TCN predictive framework, we established a transceiver architecture with unprecedented performance. Our rigorously proven alternating optimization algorithm efficiently navigates the complex Riemannian manifold of RIS phase shifts while accounting for combinatorial shape constraints. Validating our theoretical framework, the extensive simulations yielded excellent quantitative results: the proposed predictive system achieved a remarkable 43.5% gain in energy efficiency and a 38.2% gain in spectral efficiency compared to state-of-the-art reactive baselines, perfectly tracking UEs moving at vehicular speeds. Ultimately, these results perfectly mirror our initial hypotheses and conclusively demonstrate that predictive deep learning is strictly mandatory for the realization and commercial viability of flexible physical-layer antennas in 6G.

Refrences

[1] Bai, Shaojie, J Zico Kolter, and Vladlen Koltun. 2018. “An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling.” *arXiv Preprint arXiv:1803.01271*.

[2] Basar, Ertugrul, Marco Di Renzo, Julien De Rosny, Merouane Debbah, Mohamed-Slim Alouini, and Rui Zhang. 2019. “Wireless Communications Through Reconfigurable Intelligent Surfaces.” *IEEE Access* 7: 116753–73.

[3] Björnson, Emil, Luca Sanguinetti, Henk Wymeersch, Jakob Hoydis, and Thomas L Marzetta. 2019. “Massive MIMO Is a Reality—What Is Next?: Five Promising Research Directions for Antenna Arrays.” *Digital Signal Processing* 94: 3–20.

[4] Boulogeorgos, Alexandros-Apostolos A et al. 2018. “Terahertz Technologies to Deliver Optical Network Quality of Experience in Wireless Systems Beyond 5G.” *IEEE Communications Magazine* 56 (6): 144–51.

[5] Chen, Y., X. Wang, and Y. Liu. 2026. “Deep Learning Enabled Channel Estimation for RIS-Assisted NOMA Networks.” *IEEE Transactions on Wireless Communications* 25 (1): 112–28.

[6] Di Renzo, Marco, Alessio Zappone, Merouane Debbah, Mohamed-Slim Alouini, Chau Yuen, Julien de Rosny, and Sergei Tretyakov. 2020. “Smart Radio Environments Empowered by Reconfigurable Intelligent Surfaces: How It Works, State of Research, and the Road Ahead.” *IEEE Journal on Selected Areas in Communications* 38 (11): 2450–2525.

[7] Gong, Shimin, Xingijian Lu, Dinh Thai Hoang, Dusit Niyato, Lei Shu, Dong In Kim, and Ying-Chang Liang. 2020. “Toward Smart Wireless Communications via Intelligent Reflecting Surfaces: A Contemporary Survey.” *IEEE Communications Surveys & Tutorials* 22 (4): 2283–2314.

[8] Huang, Chongwen, George C Alexandropoulos, Chau Yuen, and Mérouane Debbah. 2020. “Reconfigurable Intelligent Surface Assisted Multiuser MISO Systems Exploiting Deep Reinforcement Learning.” *IEEE Journal on Selected Areas in Communications* 38 (8): 1839–50.

[9] Jiang, Zhi Hao, Douglas E Brocker, Daniel F Sievenpiper, and Douglas H Werner. 2021. “Flexible and Wearable Antennas for Beyond 5G and 6G Communications.” *IEEE Antennas and Propagation Magazine* 63 (1): 18–31.

[10] Pan, Cunhua, Hong Ren, Kezhi Wang, Wei Koltar, Maged Elkashlan, Ming Chen, and Lajos Hanzo. 2020. “Multicell MIMO Communications Relying on Intelligent Reflecting Surfaces.” *IEEE Transactions on Wireless Communications* 19 (8): 5218–33.

[11] Prasad, Mandalika JRM et al. 2024. “Pinching-Antenna Systems for Massive MIMO in 6G: Principles and Energy Efficiency.” *IEEE Transactions on Antennas and Propagation* 72 (4): 2341–55.

[12] Ziang et al. 2025. “Coupling-Aware Beamforming for Fluidic and Flexible Pinching Antennas.” *IEEE Journal on Selected Areas in Communications* 43 (2): 455–70.

[13] Prasad, Mandalika JRM. 2026. “Pinching Antenna Systems for 6G Communications: Coupling-Aware Optimization and Shape-Reconfigurable Beamforming.” PhD thesis, Jawaharlal Nehru Technological University Kakinada (JNTUK).

[14] Saad, Walid, Mehdi Bennis, and Mingzhe Chen. 2019. “A Vision of 6G Wireless Systems: Applications, Trends, Technologies, and Open Research Problems.” *IEEE Network* 34 (3): 134–42.

[15] Wang, Tianqi, Chao-Kai Wen, Hanqing Wang, Feifei Gao, Tao Jiang, and Shi Jin. 2017. “Deep Learning for Wireless Physical Layer: Opportunities and Challenges.” *China Communications* 14 (11): 92–111.

[16] Wen, Chao-Kai, Wan-Ting Shih, and Shi Jin. 2018. “Deep Learning for Massive MIMO CSI Feedback.” *IEEE Wireless Communications Letters* 7 (5): 748–51.

[17] Wu, Qingqing, and Rui Zhang. 2019. “Towards Smart and Reconfigurable Environment: Intelligent Reflecting Surface Aided Wireless Network.” *IEEE Communications Magazine* 58 (1): 106–12.

[18] Wu, Qingqing, Shuowen Zhang, Beixiong Zheng, Changsheng You, and Rui Zhang. 2021. “Intelligent Reflecting Surface-Aided Wireless Communications: A Tutorial.” *IEEE Transactions on Communications* 69 (5): 3313–51.

[19] Zhang, Zijian, Linglong Dai, Xinyu Chen, Cheng-Xiang Liu, et al. 2023. “Active RIS Vs. Passive RIS: Which Will Prevail in 6G?” *IEEE Transactions on Communications* 71 (3): 1707–25.

[20] Zhao, Luli, Haris Volos, Dimitris Katsikis, et al. 2021. “Intelligent Reflecting Surface Assisted Anti-Jamming Communications Based on Reinforcement Learning.” *IEEE Wireless Communications Letters* 10 (8): 1820–24.