Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Physics Informed Adaptive Antenna Control for High Speed Railway Links with Paired Monte Carlo Validation

PRABIR BANERJEE

prabir.banerjee@gmail.com

Heritage Institute of Technology https://orcid.org/0000-0001-8198-1397

Research Article

Keywords: railway communications, rooftop antenna, antenna mode selection, trajectory prediction, link margin, paired Monte Carlo simulation

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11385202/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Physics Informed Adaptive Antenna Control for High Speed Railway Links with Paired Monte Carlo Validation

Prabir Banerjee
Electronics and Communication Engineering Department
Heritage Institute of Technology, Kolkata, India
prabir.banerjee@heritageit.edu
ORCID ID: 0000-0001-8198-1397

Abstract

Reliable railway connectivity depends on both the electromagnetic characteristics of the rooftop antenna and its response to changing propagation geometry. This paper integrates analytical antenna design with a trajectory-aware adaptive low-profile mobility controller (ALMC) and a paired numerical evaluation. Rectangular-patch synthesis, impedance matching, conductor loss and realized gain establish the physical basis for a low-profile radiator; slot loading is considered as a route to wider bandwidth. These relations motivate, but do not independently validate, the discrete gain, beamwidth and interference-rejection states used in a 3.6 GHz railway channel surrogate. Five controllers are compared over a 6 km route using 50 paired stochastic trials. Relative to fixed control, the risk-aware controller improves mean SINR by 3.989 dB, reduces outage by 13.009 percentage points and increases the fraction of route samples with at least 6 dB link margin by 12.160 percentage points. The corresponding approximate 95% confidence intervals are 3.934–4.044 dB, 12.297–13.720 percentage points and 11.449–12.871 percentage points. Greedy instantaneous SINR selection nevertheless achieves the best mean values of all five reported RF metrics. Recovered source settings and direct paired analysis show that risk-aware control also incurs 27.58 additional mode switches and 2.206 mJ more modeled control energy per route than greedy selection. All five controllers are included in the stress comparison. A leaky-feeder extension describes tunnel coverage without claiming validated handover performance. The contribution is an integrated physical rationale and controlled comparison of adaptive antenna policies, conditional on the assumed channel and antenna states; the predictor uses oracle look-ahead, and full-wave, hardware and railway field validation remain necessary. The paper has concentrated on upper bound analysis.

Keywords: railway communications; rooftop antenna; antenna mode selection; trajectory prediction; link margin; paired Monte Carlo simulation.

1 Introduction

A high-speed train experiences changing base-station direction, shadowing, interference and coverage transitions along its route. These effects make a static peak-gain specification insufficient for evaluating the availability of a train-to-ground link [1]. Rapid channel variation also places demands on channel estimation and prediction [2]. An antenna intended for this environment must therefore be assessed both as an electromagnetic structure and as part of a mobile radio system.

Railway antenna integration introduces constraints beyond those of an isolated radiator. The 5GRAIL antenna study addresses system requirements and installation considerations [3], while Bai et al. examine roof-position optimization to improve antenna isolation and interference performance [4]. Shi et al. demonstrate a compact

wideband end-fire design for a different 5G-R frequency range [5]. That design provides useful evidence of the gain–bandwidth–size tradeoff, but its measured performance cannot be transferred to the 3.6 GHz radiator considered here. These studies motivate attention to mounting and directional response rather than treating nominal gain as a geometry-independent constant.

Coverage inside a tunnel is also an infrastructure problem. Published system-level work evaluates leaky coaxial cable channels for high-speed trains [6]. ITU-R descriptions distinguish railway communication services and train-to-trackside architectures [7,8]. Accordingly, this paper treats an outdoor base station and a tunnel distribution system as complementary radio sources. It does not claim that a single rooftop radiator alone solves tunnel penetration loss or that the proposed controller implements certified train-protection functions.

The research question is whether a physically motivated, trajectory-aware choice among antenna states improves link reliability under a common stochastic channel model. The study connects analytical antenna relations to the quantities used by the controller, compares five policies on paired channel realizations, and examines whether gains persist beyond a favourable high-SNR example. The work contributes a system-level integration and evaluation framework. It does not introduce a fabricated wideband antenna, a learned policy or a proven globally optimal controller.

The remainder of the paper proceeds from antenna physics in Section 2 to the propagation and reliability model in Section 3. Section 4 describes the decision logic and comparator policies. Section 5 defines the paired evaluation, Section 6 interprets its results, and Section 7 identifies the limits of the evidence and the next validation requirements.

2 Physical basis of the rooftop antenna model

2.1 Reference geometry and frequency conventions

The physical starting point is a rectangular microstrip patch on a low-loss substrate, mounted above a conductive ground structure. Standard patch synthesis provides initial dimensions rather than a substitute for a full-wave solution [9]. The controller benchmark uses a nominal frequency of 3.6 GHz, relative permittivity of approximately 2.2 and substrate thickness of 0.787 mm. The permittivity is consistent with the nominal RT/duroid 5880 material specification [10]. The manufacturer's laminate data do not validate the complete antenna or its assumed adaptive states.

The earlier antenna screening used 3.5 GHz and a target interval of 3.3–3.8 GHz. Both frequencies lie in this interval, but they refer to distinct numerical configurations. The paired controller results throughout this paper belong to the 3.6 GHz configuration. Illustrative mobility calculations at 3.5 GHz are retained only where identified explicitly; antenna screening values are not presented as measurements of the benchmark states.

For design frequency f_0, substrate relative permittivity ε_r, thickness h and speed of light c, the initial width and effective permittivity are

$$W = \frac{c}{2f_0}\sqrt{\frac{2}{\varepsilon_r + 1}}. \qquad (1)$$

$$\varepsilon_{\text{eff}} = \frac{\varepsilon_r + 1}{2} + \frac{\varepsilon_r - 1}{2}\left(1 + \frac{12h}{W}\right)^{-1/2}. \qquad (2)$$

The fringing-field extension and physical length are

$$\frac{\Delta L}{h} = 0.412\frac{(\varepsilon_{\text{eff}} + 0.3)(W/h + 0.264)}{(\varepsilon_{\text{eff}} - 0.258)(W/h + 0.8)}, \qquad (3)$$

$$L = \frac{c}{2f_0\sqrt{\varepsilon_{\text{eff}}}} - 2\Delta L. \qquad (4)$$

At 3.6 GHz these relations give approximately $W = 32.918$ mm and $L = 27.706$ mm for the stated substrate. They define an un-slotted reference geometry. Feed position, finite ground plane, radome and nearby roof structures require further electromagnetic modelling before fabrication.

2.2 Bandwidth and slot loading

The reference thin patch has an estimated fractional bandwidth near 1.73%, approximately 62 MHz at 3.6 GHz. Covering 3.3–3.8 GHz requires approximately 14.1% fractional bandwidth when normalized to the 3.55 GHz band centre. A conventional narrowband patch therefore does not establish full-band operation.

A U-slot is a candidate bandwidth-enhancement mechanism. Characteristic-mode and coupled-mode analysis explains the interaction between patch and slot-related resonances [11]. The number and placement of useful resonances depend on the geometry and excitation; introducing a U-slot does not automatically establish three usable resonances or a specified impedance bandwidth. This distinction is important because the current controller study uses assumed mode characteristics rather than full-wave pattern libraries derived from a finalized U-slot structure.

A consistent matching metric is the target-band coverage

$$C_B = \frac{1}{N_f}\sum_{k=1}^{N_f} \mathbf{1}\{20\log_{10}|S_{11}(f_k)| \le -10\}. \qquad (5)$$

Here N_f is the number of frequency samples in the target band and $\mathbf{1}$ is the indicator function. Coverage fraction, contiguous matched bandwidth and the span between outermost matched points are different quantities. Earlier screening reported approximately 88% coverage and a 488 MHz matched span. These quantities cannot be equated: 88% of a uniformly sampled 500 MHz interval is approximately 440 MHz of covered bandwidth. Without the underlying trace and interval definitions, those preliminary values are excluded from the paper's validated performance claims.

2.3 Loss and realized gain

Conductor and dielectric losses connect material selection to useful antenna gain. For a good conductor with conductivity σ, permeability approximately μ_0 and angular frequency $\omega = 2\pi f$, the skin depth and thick-conductor surface resistance are [12]

$$\delta = \sqrt{\frac{2}{\omega\mu_0\sigma}}, \qquad R_s \simeq \sqrt{\frac{\pi f\mu_0}{\sigma}}. \qquad (6)$$

The surface-resistance approximation applies when the metal thickness is several skin depths. Thinner films require an appropriate finite-thickness surface-impedance model. No measured conductivity, film stack or fabrication tolerance is available here to substantiate a hybrid-conductor advantage; copper is retained as the fabrication reference.

Using radiation, conductor and dielectric quality factors Q_r, Q_c and Q_d, a reduced loss model is

$$Q_L^{-1} = Q_r^{-1} + Q_c^{-1} + Q_d^{-1}, \qquad \eta_{\mathrm{rad}} = \frac{Q_r^{-1}}{Q_r^{-1} + Q_c^{-1} + Q_d^{-1}}. \qquad (7)$$

These expressions omit additional loss channels unless incorporated in the corresponding quality factors. If D is linear directivity and Γ is the input reflection coefficient, realized gain is

$$G_{\mathrm{real}} = D\eta_{\mathrm{rad}}(1 - |\Gamma|^2), \qquad G_{\mathrm{real,dBi}} = 10\log_{10} G_{\mathrm{real}}. \qquad (8)$$

The reference surrogate gives approximately 4.51 dBi realized gain and 89.81% radiation efficiency at 3.6 GHz. Earlier material screening near 3.5 GHz gave approximately 89.3% efficiency. These are configuration-dependent analytical estimates and are not averaged or treated as independent confirmation. Roof loading, polarization mismatch and off-axis gain must be considered separately from broadside realized gain. In particular, upward broadside gain is not necessarily the gain toward a wayside base station at low elevation [3,4].

2.4 Connection to discrete antenna states

The control model comprises four candidate states: M1 Coverage, M2 Balanced, M3 High-SINR and M4 Recovery. Each state is represented by assumed gain, pattern width and interference rejection. The physical tradeoff is that a narrower pattern may improve desired-direction gain or interference rejection but may lose coverage as the train-to-base-station angle changes. A broad state sacrifices some directional selectivity for angular tolerance.

This mapping is the interface between antenna design and control. The patch equations constrain plausible size and loss, but do not prove that a single patch with unspecified switching circuitry realizes all four states. Mode-specific patterns, gain tables, switching losses and hardware realization are required before the control results can be attributed to a built antenna. In the present work, the states form a conditional design space for numerical evaluation.

2.5 Analytic consistency and physical realization

The implemented angular gain is $G_m(\varphi, \varepsilon) = g_m - \min\{35, 12[(90 - \varepsilon)^2 + (0.35|\varphi|)^2]/h_m^2\}$ dBi, with wrapped train-relative azimuth φ and elevation ε in degrees. It is a broadside-centred surrogate: the boresight is not steered toward the selected base station. Therefore h_m is best called a pattern-width parameter. At fixed elevation in the unclipped region, a 3 dB azimuth drop from $\varphi = 0$ occurs at $|\varphi| = h_m/0.7$, giving full azimuth width $2h_m/0.7$ when that crossing lies within ±180°. For M3 this is about 165.7°, rather than the nominal 58° label; M2 gives 308.6°, while M1 and M4 have no such crossing within ±180°. This audit prevents the nominal labels from being misrepresented as validated azimuth HPBWs.

For a hypothetical low-elevation direction $(\varphi, \varepsilon) = (0°, 0°)$, the formula gives gains of −0.12, −3.03, −21.79 and 0.56 dBi for M1–M4 respectively, versus −0.82 dBi for FIXED. These are analytic evaluations of the surrogate, not roof-mounted measurements. M3's 7.1 dBi peak therefore does not imply improved wayside gain. The extra interference rejection is applied only to non-reference base stations whose azimuth separation exceeds $h_m/2$; it is a separate heuristic attenuation and is not derived from the gain pattern or polarization.

A physically realizable replacement can use a switched parasitic radiator or a compact switched-feed subarray, but its dimensions, feed network, switch impedances and bias routing must be specified and solved. For each state, the required evidence is input matching, radiation efficiency and co/cross-polarized realized gain over direction and frequency with the train roof and radome included. One consistent pattern library must determine both desired and interfering gains; adding an unrelated rejection bonus would double-count discrimination. The current study retains the original surrogate to preserve comparability, while treating physical realizability as an unclosed validation requirement [3,4,9,11].

3 Railway channel and reliability model

3.1 Route geometry and open-track link

The controller benchmark represents a 6 km route with multiple wayside base stations. A log-distance model describes nominal large-scale attenuation, while correlated shadowing, intercell interference and terrain-zone losses introduce route-dependent variation. The nominal path-loss exponent is 2.45. Standardized channel descriptions provide context for channel modelling [13], but the simplified surrogate is not claimed to implement the complete 3GPP TR 38.901 model.

For base-station separation $d \geq d_0$, reference distance $d_0 = 1$ m, path-loss exponent n, shadowing term X_σ in decibels and additional zone loss L_z,

$$PL(d) = PL(d_0) + 10n\log_{10}(d/d_0) + X_\sigma + L_z. \qquad (9)$$

A free-space reference uses $PL(d_0) = 20\log_{10}(4\pi d_0 f_c/c)$. Received desired power for antenna state m is

$$P_{r,m} = P_t + G_t + G_{r,m}(\theta, f_c) - PL(d) - L_f. \qquad (10)$$

Powers are in dBm, gains in dBi and losses in dB. $G_{r,m}$ is the gain toward the serving transmitter, including the modeled angular response; L_f represents feeder and other implementation losses. At a nominal thermal noise density of −174 dBm/Hz,

$$N_{\mathrm{dBm}} = -174 + 10\log_{10}(B) + NF. \qquad (11)$$

For desired power S_m, interfering powers $I_{j,m}$ and noise N expressed consistently in linear units,

$$\gamma_m = \frac{S_m}{N + \sum_j I_{j,m}}, \qquad \gamma_{m,\mathrm{dB}} = 10\log_{10}\gamma_m. \qquad (12)$$

SNR applies when interference is omitted; SINR is used for the controller benchmark. These quantities are not interchangeable. For example, the earlier illustrative link used $B = 100$ MHz and $NF = 7$ dB, giving a noise floor of −87 dBm. A received power near −58 dBm then corresponds to approximately 29 dB SNR. That strong-link example checks arithmetic, but cannot by itself explain the interference-inclusive trial results.

3.2 Mobility and control timescales

For train speed v and propagation angle θ_v relative to the direction of travel,

$$f_D = \frac{v f_c}{c}\cos\theta_v, \qquad f_{D,\max} = \frac{v f_c}{c}, \qquad T_c \approx \frac{0.423}{f_{D,\max}}. \qquad (13)$$

At 300 km/h and 3.5 GHz, the maximum Doppler shift is approximately 973 Hz and the illustrative coherence time is 0.435 ms. At 3.6 GHz the corresponding values are approximately 1001 Hz and 0.423 ms. The coherence-time approximation depends on the scattering model; a maximum line-of-sight Doppler shift is not itself a measured Doppler spread.

Sub-millisecond channel variation motivates appropriate receiver estimation and tracking [2]. It does not imply that an antenna mode must switch every coherence interval. Mode selection can follow slower geometry and large-scale interference changes, with prediction horizon and dwell time chosen independently of the waveform tracking interval. The stored route samples therefore support system-level comparison rather than validation of sub-millisecond modem behaviour.

3.3 Tunnel extension and network functions

For a repeater-fed leaky feeder, a reduced loss model can be written as

$$PL_{\text{tun}}(x, r) = L_c(r_0) + \alpha x + \Delta L_r(r; r_0) + L_m, \qquad (14)$$

$$P_{r,\text{tun}} = P_{\text{rep}} + G_r - PL_{\text{tun}}(x, r). \qquad (15)$$

Here x is distance from the feed point, α is longitudinal attenuation in dB/m, $L_c(r_0)$ is coupling loss at a stated reference separation, and ΔL_r is only the incremental correction for actual separation r. This convention prevents double-counting radial coupling loss. If a measured coupling-loss specification already includes the reference receive-antenna gain, the gain term must also be handled consistently. L_m contains remaining implementation losses [6].

Figure 1 retains the source manuscript's illustrative received-power envelopes to explain why coverage overlap is desirable at a tunnel portal. Its crossing point is not a simulated or measured handover event. A complete implementation would require network procedures consistent with the NR architecture and RRC functions [14,15]. The paired controller experiment below does not establish handover interruption time, packet continuity or tunnel-specific outage improvement.

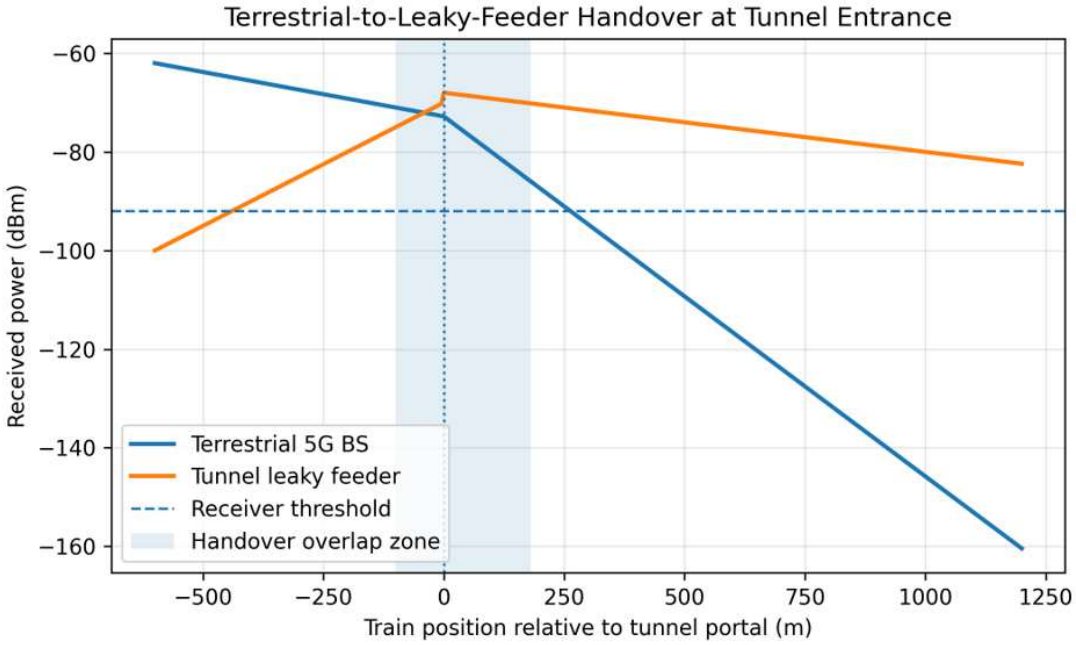


Figure 1. Illustrative terrestrial and leaky-feeder received-power envelopes from the analytical study. The overlap and receiver threshold are scenario assumptions, not measured coverage or executed NR handover.

3.4 Common performance measures

The conservative link margin combines a received-power requirement and a SINR requirement:

$$LM_i = \min\{P_{r,i} - P_{\text{sens}},\ \gamma_{i,\text{dB}} - \gamma_{\text{req,dB}}\}. \qquad (16)$$

Each difference is a margin in decibels. The frozen benchmark uses receiver sensitivity −92 dBm and required SINR 8 dB. Thus LM is the smaller of received power plus 92 dB and SINR minus 8 dB. The earlier illustrative 5 dB SNR threshold is not used for this benchmark.

For N_s samples along one route,

$$O = \frac{100}{N_s}\sum_{i=1}^{N_s} \mathbf{1}\{LM_i < 0\}, \qquad A_q = \frac{100}{N_s}\sum_{i=1}^{N_s} \mathbf{1}\{LM_i \geq q\}. \qquad (17)$$

O is outage percentage and A_q is the route fraction with at least q dB margin. The retained label ALMC6 in the result files denotes A_6; ALMC otherwise names the adaptive low-profile mobility controller. The fifth percentile of LM characterizes the lower tail. For uniformly spaced route samples, these fractions describe sampled route coverage; interpreting them as time availability also requires a consistent speed and sampling convention.

The reported rate is a capped Shannon-like proxy. Its general form is

$$R_{\text{proxy}} = \min\{R_{\max},\ \eta B \log_2(1+\gamma)\}. \qquad (18)$$

The implemented rate proxy is R = 20 min{7.4, log₂[1 + 10^(SINR/10)]} Mbit/s, corresponding to a 20 MHz scaling, unit efficiency factor and 148 Mbit/s cap. Noise is independently fixed at −95 dBm. If interpreted with −174 dBm/Hz thermal density and 20 MHz bandwidth, this implies approximately 6 dB noise figure. The proxy remains positive during negative-margin outage; it is not outage-gated goodput. Neither this expression nor the earlier fixed-modulation illustration resolves NR scheduling, coding, retransmission or resource allocation [16].

4 Trajectory aware antenna control

4.1 Decision logic

The risk-aware policy uses route and channel information with ideal future-sample access, evaluates candidate link behaviour, scores the four antenna modes and applies a state subject to a switching penalty and minimum dwell. This is an offline simulation of predictive mode selection. The term agentic describes this closed-loop mode selection; the evidence does not establish autonomous learning, a large-language-model agent or online discovery of a new antenna geometry.

The supplied controller description includes predicted SINR, margin and rate benefits, with penalties for outage risk, insufficient angular coverage, switching, inappropriate recovery selection and unsafe narrow-beam operation. A conceptual representation is

$$J_m = U_m - C_{\text{out},m} - C_{\text{ang},m} - C_{\text{sw},m} - C_{\text{rec},m} - C_{\text{guard},m}, \qquad (19)$$

$$m^* = \underset{m\in\mathcal{M}}{\operatorname{argmax}}\ J_m. \qquad (20)$$

Equations (19)–(20) give the general scoring structure. Section 4.3 specifies the archived implementation, including numerical weights and penalty functions. Finite penalties do not impose a hard outage constraint, and the policy therefore has no certified risk bound.

The implementation uses h_req = clip{2[ωH + 2° + 6°], 20°, 175°}, where ω is the absolute finite-difference derivative of the unwrapped serving-base-station azimuth and H = 0.25 s. This is a full-width requirement. Serving-cell changes can contribute to the computed angular derivative. Moreover, the pattern width parameter does not equal a measured azimuth HPBW, as shown in Section 2.5; this coverage test is a heuristic constraint on surrogate modes.

After scoring, a 0.6 utility penalty discourages mode changes and dwell logic limits successive changes. There is no additional mode-utility hysteresis threshold in the archived code. The separate 3 dB hysteresis selects the serving base station. Mode switching and serving-cell changes are distinct events; neither includes simulated interruption time.

4.2 Comparator policies

Table 1 identifies the five policies and preserves the implementation labels used in the supplied plots. The V1.3-style comparator approximates earlier predictive rule logic and is not claimed to reproduce every historical implementation detail.

Table 1. Controller definitions used in the paired benchmark.

Label	Decision basis	Interpretation
FIXED	Retains the reference antenna state	Nonadaptive baseline
THRESHOLD	Uses simple link and beamwidth thresholds	Reactive rule comparator
GREEDY_SINR	Maximizes instantaneous modeled SINR	Instantaneous RF benchmark
V13_STYLE	Uses earlier predictive rule logic	Approximate predictive comparator
V14_RISK_AWARE	Scores predicted utility and risk penalties with switching safeguards	Evaluated risk-aware policy

GREEDY_SINR evaluates all four modeled candidate SINRs at the current sample. V14_RISK_AWARE evaluates the same candidate model at a future sample, using the realized future shadowing as well as future geometry. The paper does a comparative analysis of look-ahead horizon penalties vs. instantaneous adaptivity. Thus the benchmark provides ideal candidate-state information and an oracle look-ahead to the predictive controller; it does not test an implementable channel predictor. All policies share each seed's channel realization, but their time-indexed information differs. The comparison quantifies the archived policies under these explicit assumptions.

4.3 Recovered scoring and timing implementation

Let $[x]_+ = \max(0,x)$, γ denote candidate SINR in dB, L the conservative margin in dB, R the rate proxy in Mbit/s, h the mode width and h_req the required width. The exact score is $U = \gamma + 1.35L + 0.018R - 7q - 0.09a - 0.6s - 2b - 0.5r$. Here $q = [2 - L]_+ + 0.7[9 - \gamma]_+$; $a = [h_req - h]_+$; s is one when the candidate differs from the previous mode; $b = [7 - L]_+ + 0.7[13 - \gamma]_+$ for M3 and zero otherwise; and r is one for M4 when $L > -2$ and $\gamma > 5$, and zero otherwise. The score mixes engineering quantities with empirical weights; it is not a normalized probability of failure.

At sample i, candidates are scored at $j = \min(N - 1, i + \max\{1, \text{round}(H/\Delta t)\})$. After selecting the highest score, M4 is forced if that candidate has $L < -2$ dB or $\gamma < 5$ dB. Otherwise, h_req > 95° excludes M3. A proposed change is held until the previous mode has lasted 0.18 s, or 0.12 s if the previous mode was M4. The initial mode is M2. The balanced-pattern reference serving sequence uses 3 dB received-power hysteresis; the final serving sequence is recomputed after mode selection. This two-stage procedure may choose a final serving cell different from the reference used during scoring.

With 700 uniformly spaced route-parameter samples, $\Delta s = 6000/699 = 8.584$ m and $\Delta t = \Delta s/(320/3.6) = 0.09657$ s. The 0.25 s horizon rounds to three samples, or 0.2897 s; both dwell settings effectively require two sample intervals, or 0.1931 s. At 420 samples, $\Delta t = 0.1611$ s, the horizon rounds to two samples, and the two dwell settings require two and one intervals respectively. The S-curve uses $x = s$ and $y = 120 \sin(2\pi s/6000)$, so

s is a longitudinal route parameter, not exact arc length. The nominal traversal time is 67.5 s; geometric speed varies slightly along the curve. The code's Doppler quantity is diagnostic and does not drive a small-scale fading waveform.

5 Numerical evaluation and statistical treatment

5.1 Paired trial design

The principal evaluation comprises 50 nominal trials using seeds 1001–1050 and 700 route samples per trial. All five controllers use the same stochastic channel realization within each seed. Pairing reduces the effect of differences in random environments when estimating a controller contrast. Correlated samples within one route are not treated as independent trials: the seed is the statistical unit.

Table 2 consolidates the settings that are explicitly available. Antenna screening and numerical control validation are separate evidence layers. Evaluation guidelines such as ITU-R M.2412 emphasize explicit scenarios and assumptions [17]; citing them does not make this simplified experiment a compliant IMT evaluation.

Table 2. Reported configurations and evidence scope.

Item	Setting or scope
Controller carrier	3.6 GHz
Physical screening	Earlier 3.5 GHz illustrations; 3.3–3.8 GHz target band
Patch reference	Relative permittivity 2.2; thickness 0.787 mm; 32.918 × 27.706 mm at 3.6 GHz
Nominal route	6 km; multiple wayside base stations; path-loss exponent 2.45
Nominal trials	50 paired seeds, 1001–1050; 700 samples per route
Antenna states	Four assumed gain, beamwidth and rejection states
Ablation study	20 trials per variant
Stress study	Six trials per scenario; 420 route samples
External validation	No full-wave, prototype or railway field measurements supplied

The numerical benchmark is now specified by the archived V1.4 engine, the V1.5 validation wrapper and the study configuration. Tables 3–5 report its geometry, channel, controller and mode settings. Coordinates are synthetic Cartesian positions in metres, not surveyed geographical locations. The nominal study uses Config(samples=700) and seeds 1001–1050. The archived run_study.py retains an earlier 12-trial stress command, whereas study_config.json and the stored stress CSV identify six trials per scenario with 420 samples; the latter configuration defines the stress results reported here.

Table 3. Synthetic wayside geometry and conducted transmit power. All base-station antenna gains are 15 dBi; EIRP excludes any unmodelled transmitter feeder loss.

Station	x (m)	y (m)	Height (m)	Power (dBm)	EIRP (dBm)
BS1	650	55	30	43	58
BS2	1750	-70	32	42	57
BS3	2950	60	30	43.5	58.5
BS4	4150	-80	34	43	58

Station	x (m)	y (m)	Height (m)	Power (dBm)	EIRP (dBm)
BS5	5350	65	30	42.5	57.5

Table 4. Recovered benchmark settings and model conventions.

Quantity	Value or definition
Route and receiver	S-curve; x = s, y = 120 sin(2πs/6000) m; train height 4.2 m; nominal speed 320 km/h
Carrier and path loss	3.6 GHz; exponent 2.45; reference distance 1 m; $PL_0 = 32.44 + 20\log_{10}(3600) - 60 = 43.566$ dB
Shadowing	Independent sequence per base station; σ = 3.5 dB, variance 12.25 dB²; correlation distance 90 m
Shadow recursion	$X_0 = 0$; $X_i = \rho X_{i-1} + \sqrt{(1-\rho^2)}e_i$; $e_i \sim N(0,\sigma^2)$, $\rho = \exp(-\Delta s/90)$; NumPy default_rng(seed)
Noise and thresholds	Noise −95 dBm; receiver sensitivity −92 dBm; required SINR 8 dB; cable/miscellaneous loss 2 dB
Zone losses	s = 1150–1700 m: 12 dB; 2550–3100 m: 3 dB; 4300–4850 m: 4.5 dB; zero elsewhere
Interference	All four non-serving stations transmit cochannel continuously; powers summed in linear units; no load or scheduling model
Sampling and seeds	Main: 700 samples, 50 seeds 1001–1050; ablation: 700 samples, 20 seeds 3001–3020; stress: 420 samples, six seeds 5001–5006
Prediction and angular terms	H = 0.25 s; uncertainty allowance 2°; guard 6°; width clip 20°–175°; exact scoring in Section 4.3
Control cost	0.08 mJ per counted mode switch; 0.006 mJ per sample for each adaptive policy; fixed control cost zero

The zone labelled TUNNEL in the code is only a 12 dB outdoor-link attenuation segment applied to every base station. It does not implement the leaky-feeder model in Section 3.3. Shadowing starts at zero rather than a stationary random draw, producing a short initial variance transient. These conventions are retained for reproduction and should not be interpreted as measured railway conditions.

Table 5. Assumed state library. Width is the h parameter in Section 2.5; rejection is the separate off-reference attenuation. PIN codes are logical labels without a verified circuit mapping.

State	Peak gain (dBi)	Width (°)	Rejection (dB)	PIN label
FIXED	4.51	135	0	—
M1	4.2	150	0.8	1010
M2	5.3	108	1.8	1100
M3	7.1	58	5	1111
M4	4	168	0.2	0000

The source additionally lists mode bandwidth/efficiency pairs of 4.5%/86%, 3.2%/89%, 2.2%/88% and 6.0%/82% for M1–M4. These metadata do not enter the RF comparison and are not validated performance. The standalone varactor calculation uses f_0 = 3.6 GHz, temperature 35°C, −0.25 MHz/°C relative to 25°C, radome/roof shifts −12/−8 MHz, C_0 = 2.2 pF, Vj = 0.75 V, exponent 0.48, Cref = 1 pF, coupling 0.085 and an 8-bit 0–12 V sweep. Its selected tuning value is recorded but does not modify link gain, matching or outage in this benchmark; tuning improvement is therefore not an established cause of the reported RF gains.

5.2 Confidence intervals and paired changes

For a metric y_s evaluated once per seed, the reported interval is

$$\overline{y} = \frac{1}{n}\sum_{s=1}^{n} y_s\,, \qquad CI_{95} \approx \overline{y} \pm 1.96\frac{s_y}{\sqrt{n}}, \quad n = 50. \qquad (21)$$

For metrics in which larger values are preferable, the paired improvement is $d_s = y_{s,\mathrm{risk}} - y_{s,\mathrm{fixed}}$; for outage its sign is reversed. The same mean and standard-error calculation is applied to d_s. This pairing must be retained when estimating uncertainty. The source intervals are preserved as approximate normal intervals,

rather than relabelled as exact finite-sample intervals. A Student-t interval is the usual alternative when population variance is estimated [18]; recalculating it requires the trial data or sufficiently precise summary statistics.

The original intervals quantify seed-to-seed variation under the chosen model and are retained for continuity. Newly computed V14-minus-greedy contrasts use paired Student-t intervals, mean difference ± t(0.975,49) times its standard error, with t(0.975,49) approximately 2.010. These are pointwise intervals without multiplicity adjustment. They quantify simulation variability, not uncertainty in physical calibration. Seed-level CSV records have been reanalysed for the new contrasts. An independent rerun of all 50 nominal seeds with the archived engine reproduced all 14 reported summary metrics for all five controllers to a maximum absolute difference of 1.43×10^{-14} in their respective numerical units.

5.3 Ablation and stress scope

The ablation study uses 20 paired seeds 3001–3020 per variant with 700 samples. The stress archive includes all five controllers, six paired seeds 5001–5006 per scenario and 420 samples. Changes are speed 350 km/h; shadowing standard deviation 6 dB; path-loss exponent 2.2 for the dense-interference label; prediction horizon 0.5 s; or nominal transmit power 40 dBm. Dense interference therefore means a changed path-loss exponent, not extra base stations. Figure 7 now displays all five policies. The six-trial stress results remain exploratory.

6 Results and discussion

6.1 Nominal controller comparison

Table 6 reports the supplied across-seed means. Risk-aware control increases mean SINR from 12.59 to 16.58 dB and reduces outage from 35.89% to 22.89% relative to fixed control. It also improves ALMC6 and the fifth-percentile link margin. These changes show that mode adaptation can mitigate some of the geometric and interference losses represented in the channel surrogate.

Table 6. Across-seed means for the 50-trial nominal comparison.

Controller	SINR (dB)	Outage (%)	ALMC6 (%)	LM 5th percentile (dB)	Rate proxy (Mbit/s)
GREEDY_SINR	16.95	21.40	51.66	−5.17	97.65
V14_RISK_AWARE	16.58	22.89	51.34	−6.46	96.33
V13_STYLE	14.86	34.77	44.54	−7.96	87.95
THRESHOLD	14.71	32.59	42.68	−7.97	87.77
FIXED	12.59	35.89	39.18	−8.07	82.17

Figures 2 and 3 show the mean SINR and outage comparisons with the supplied approximate confidence intervals. The risk-aware policy improves on both simple rule comparators in these nominal results. However, GREEDY_SINR has the best mean value for every metric in Table 6. The data therefore support effectiveness relative to fixed and selected rule-based policies, not universal superiority of the risk-aware approach.

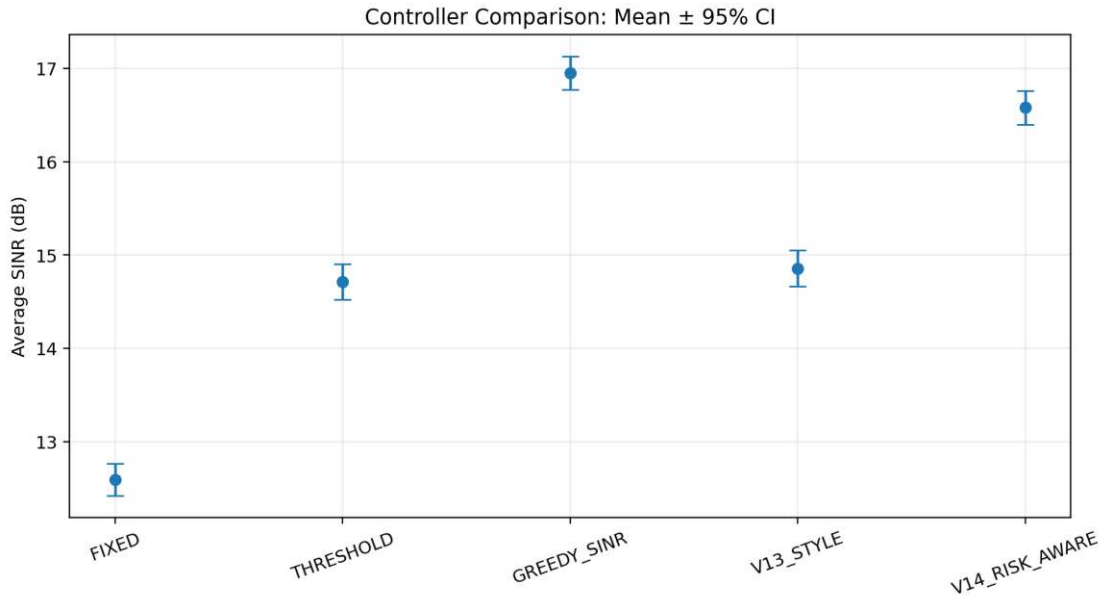


Figure 2. Across-seed mean SINR and approximate 95% confidence intervals over 50 paired trials. Source implementation labels are retained.

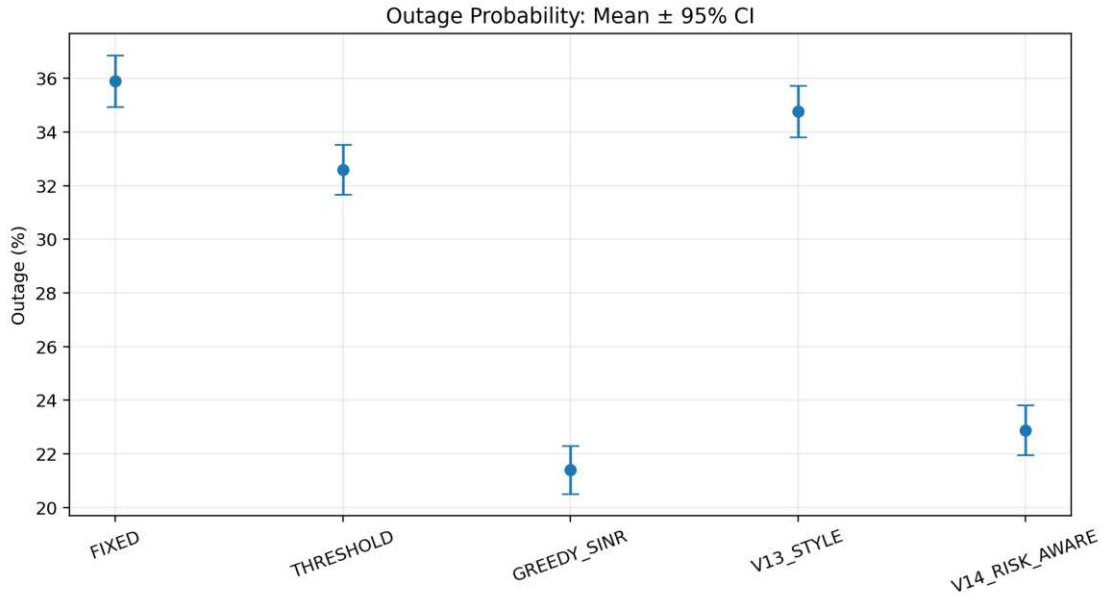


Figure 3. Across-seed outage percentage and approximate 95% confidence intervals over 50 paired trials. Outage is defined by negative conservative link margin.

6.2 Paired improvements and lower-tail performance

Table 7 gives the paired risk-aware improvements over FIXED. Positive entries denote improvement, so outage is reported as a reduction. Percentage-based differences are expressed in percentage points, not relative percent. The difference between rounded Table 6 entries may differ slightly from the paired means computed before rounding.

Table 7. Paired V14_RISK_AWARE improvements over FIXED across 50 seeds.

Metric and direction	Mean improvement	Approximate 95% CI
Mean SINR increase (dB)	3.989	3.934 to 4.044
Fifth-percentile margin increase (dB)	1.611	1.436 to 1.785
Outage reduction (percentage points)	13.009	12.297 to 13.720
ALMC6 increase (percentage points)	12.160	11.449 to 12.871

The paired intervals exclude zero for these four contrasts under the reported approximation. This provides evidence of consistent improvement over the fixed state within the numerical model. It does not establish transferability to a measured railway channel.

Lower-tail margin remains negative for every controller, including the greedy policy (Figure 4). Risk-aware control improves the fifth percentile from −8.07 to −6.46 dB, but does not eliminate the vulnerable route segments. Its nominal outage remains about 23%, while ALMC6 is only about 51% (Figure 5). Consequently, the results should not be interpreted as achievement of stringent railway service availability or URLLC targets.

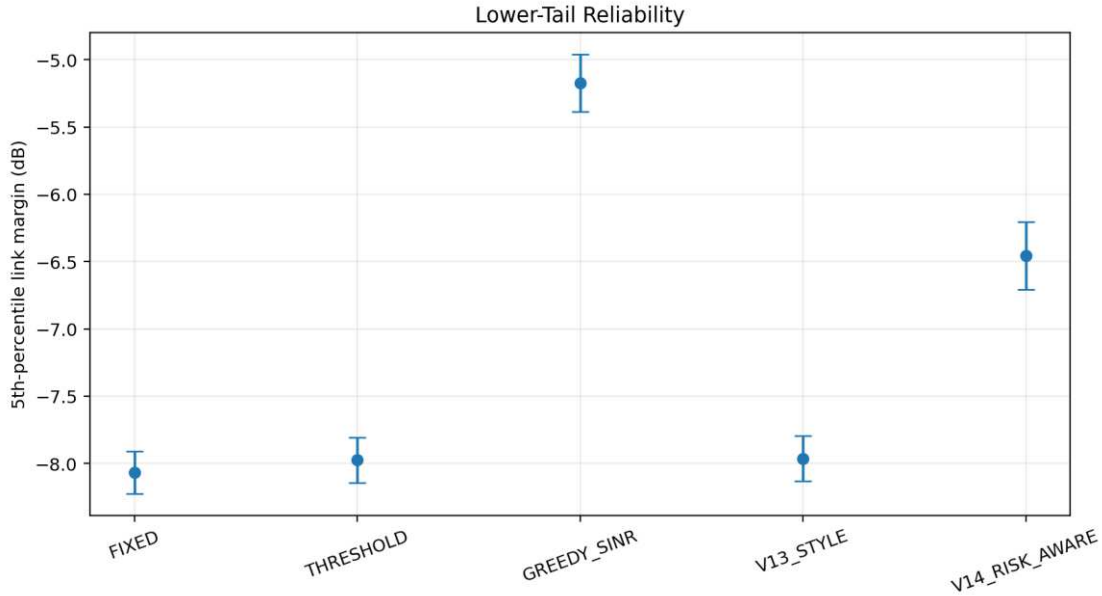


Figure 4. Mean per-trial fifth-percentile link margin with approximate 95% confidence intervals. A negative value indicates that the lower tail remains below the operating requirement.

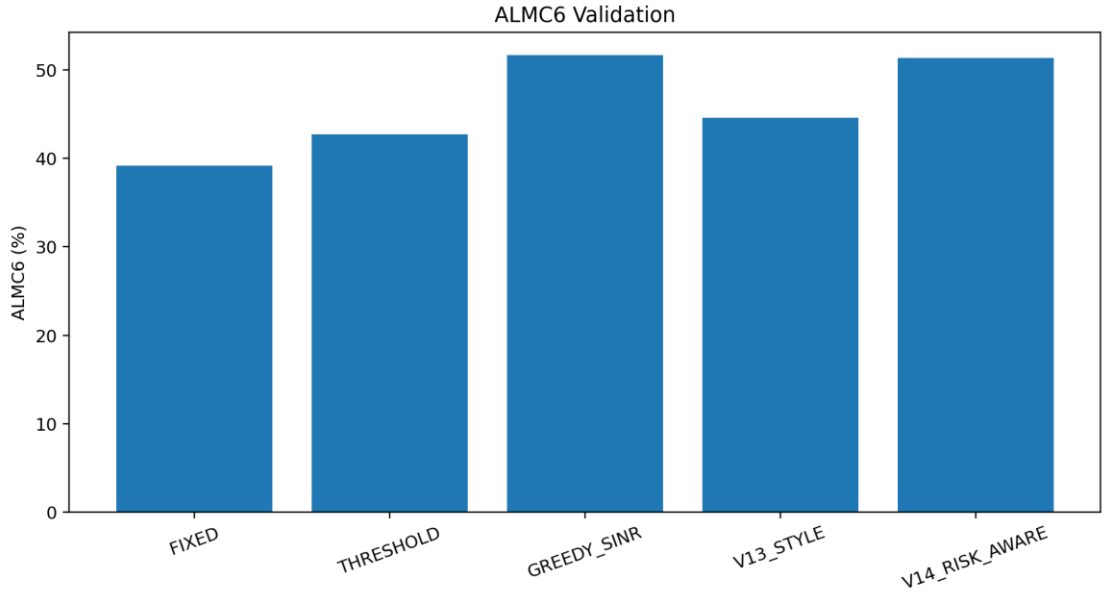


Figure 5. Across-seed mean ALMC6, the percentage of route samples with at least 6 dB conservative margin. The supplied figure does not show confidence intervals.

6.3 What the greedy comparison reveals

The direct paired comparison in Table 8 shows that V14_RISK_AWARE has lower SINR, ALMC6, lower-tail margin and rate proxy, and higher outage than GREEDY_SINR in the nominal benchmark. Its mode-switch count is also higher: 58.54 versus 30.96 switches per route, an increase of 89.08%. The mean control-energy proxy is 8.8832 versus 6.6768 mJ, an increase of 33.05%. These findings rule out a switch-count or modeled control-energy advantage for the archived risk-aware policy.

Control energy is $E = 0.08\ N_{sw} + 0.006\ N$ mJ for every adaptive policy, where $N = 700$ and the initial state assignment is excluded from N_{sw}. Consequently, the energy contrast is exactly 0.08 times the switch-count contrast; it is not independent hardware evidence. The proxy excludes static bias power, RF-chain consumption and measurement/computation cost. The paired handover-count difference is −0.04 with 95% interval −0.492 to 0.412, providing no clear reduction. No switching interruption or packet-latency benefit has been measured.

Table 8. Direct paired comparison over 50 seeds. Differences are V14_RISK_AWARE minus GREEDY_SINR; negative is favourable only for outage, switches, energy and serving-cell changes. Intervals are pointwise paired Student-t 95% intervals.

Metric	Greedy	Risk-aware	Difference	95% interval
Mean SINR (dB)	16.950	16.581	-0.369	-0.393 to -0.345
Outage (percentage points)	21.400	22.886	1.486	1.325 to 1.646
5th-percentile margin (dB)	-5.173	-6.457	-1.284	-1.474 to -1.095
ALMC6 (percentage points)	51.663	51.340	-0.323	-0.398 to -0.247
Rate proxy (Mbit/s)	97.650	96.326	-1.323	-1.427 to -1.219
Mode switches per route	30.960	58.540	27.580	25.153 to 30.007
Control-energy proxy (mJ)	6.677	8.883	2.206	2.012 to 2.401
Serving-cell changes	6.400	6.360	-0.040	-0.492 to 0.412

The policy penalties and overrides do not guarantee a monotonic decrease in switching or outage. Oracle look-ahead can select a future-favoured state too early, recovery overrides can select lower immediate SINR, and dwell constraints can retain a suboptimal state. These mechanisms explain plausible performance costs, but their separate causal contributions require controlled ablation. A switching-delay or measurement-noise study would be a new experiment, not a justification inferred from the present results.

6.4 Ablation and exploratory stress findings

Figure 6 shows the supplied ablation means. Removing the outage penalty gives visibly higher mean outage, consistent with its intended role in the objective. Removing the M3 guard gives a lower outage mean than the full controller in this plot. The latter observation is important: not every safeguard improves the selected RF metric, and a guard may carry a measurable cost. Without paired intervals and the companion switching or safety metrics, these trends do not establish the necessity of each term.

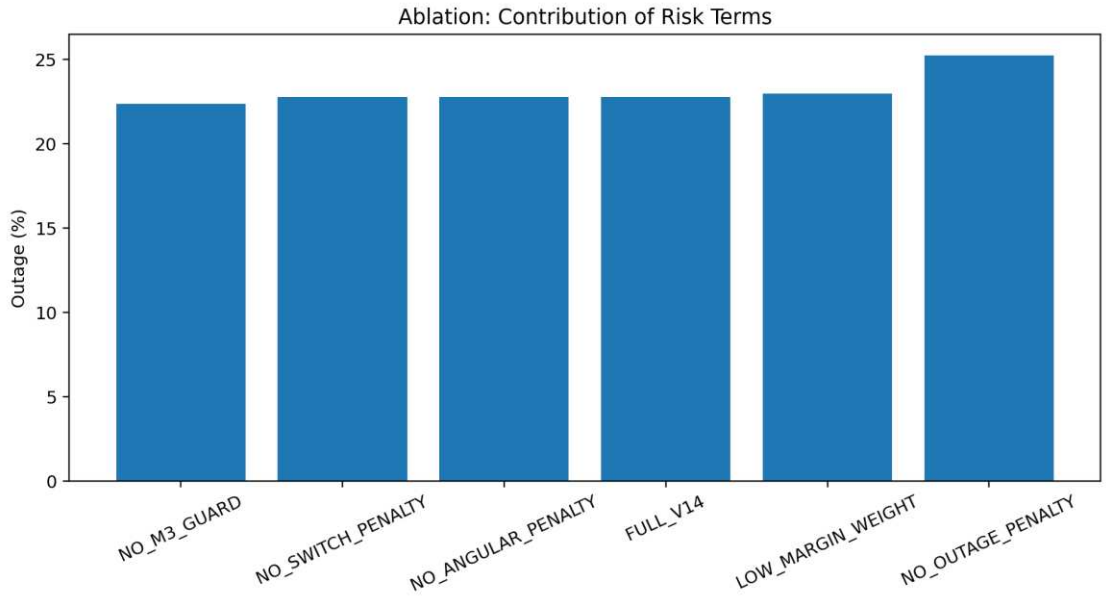


Figure 6. Mean outage for the risk-control ablations, with 20 trials per variant. No confidence intervals are available in the supplied plot; differences are interpreted descriptively.

Figure 7 shows risk-aware outage below FIXED and V13_STYLE in all six archived scenarios, but above GREEDY_SINR throughout. The risk-aware-minus-greedy mean outage gap ranges from 1.310 percentage points at 350 km/h and the stress nominal setting to 2.897 percentage points under high shadowing. At 420 samples, the speed change does not change the rounded prediction lead or effective dwell in sample counts, explaining the identical nominal and 350 km/h outage means. This discretization limits conclusions about speed robustness. These stress comparisons do not establish worst-case reliability or risk-aware dominance.

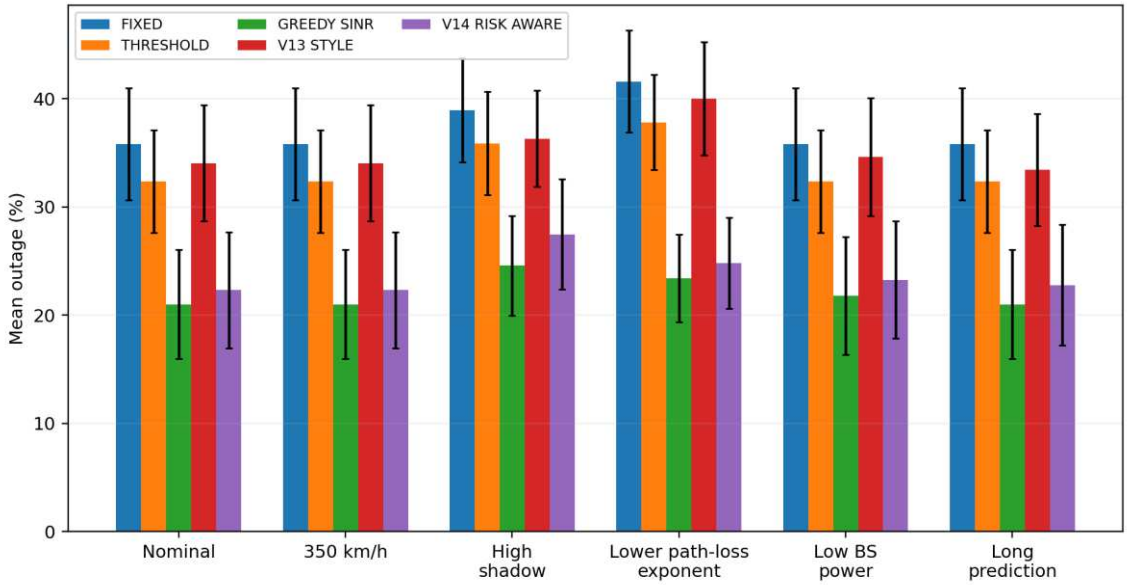


Figure 7. Archived mean outage for all five controllers in six stress scenarios, with six paired trials per scenario and 420 samples per route. Error bars show pointwise Student-t 95% intervals across seeds. The nominal stress configuration has different sampling from the 700-sample main study.

6.5 Physical interpretation

The analytical and numerical components answer complementary questions. Patch synthesis and loss relations show why a compact low-loss radiator is plausible and why bandwidth, mounting and off-axis response need attention. The paired trials show the benefit of selecting among assumed directional states when geometry and interference vary. The trials do not validate the electromagnetic realizability of those states, and the physical calculations do not predict the reported outage reduction without the channel and control model.

This distinction gives the integrated study a defensible scope. A conventional patch provides an interpretable baseline; slot loading is a proposed route toward broader matching; and the controller experiment quantifies conditional system benefits. Tunnel coverage remains a complementary infrastructure extension, with portal overlap illustrating the need for a coordinated network transition rather than proving one.

7 Limitations and reproducibility requirements

The principal limitation is model fidelity. The mode table and analytic angular response are now explicit, and Section 2.5 separates the surrogate width parameter from physical HPBW. However, no independent full-wave or measured mode library establishes that one roof-mounted structure realizes these states. PIN bit labels, efficiency and bandwidth entries alone do not specify a realizable RF circuit or geometry. Roof equipment, radome loading, coupling and polarization require electromagnetic evaluation before hardware-level claims can be made.

The propagation model is a stylized railway surrogate. Site-specific shadowing, interference and tunnel coupling remain uncalibrated. The numerical precision of the paired intervals should therefore not be confused

with accuracy of the physical model. The current evidence establishes repeatability across selected random realizations, not demonstrated railway deployment performance.

The recovered source and seed-level records provide the exact numerical policy and configuration. Tables 3–8 remove ambiguity about parameters and operational metrics. Remaining reproducibility concerns are the difference between the archived stress driver and stored stress settings, and dependency-version sensitivity in future reruns. More fundamentally, the predictive policy accesses realized future shadowing and does not estimate it from past observations. A deployment-oriented comparison must replace this oracle with a causal estimator and give policies a common measurement-noise and delay model.

No protocol-level reliability or latency certification is claimed. ITU-R M.2410 defines minimum technical requirements for IMT-2020 radio interfaces [19], and M.2411 describes associated requirements and evaluation criteria [20]. The simplified outage and rate proxies used here do not demonstrate compliance with these requirements. Likewise, a 3.6 GHz study frequency is a research configuration, not evidence of a universally allocated railway operating band.

Further validation should replace the surrogate with full-wave or measured three-dimensional realized-gain patterns, substitute a causal predictor for oracle look-ahead, and calibrate the propagation and switching models. These requirements are distinct from numerical reproducibility. The present comparison establishes a benefit of adaptation relative to fixed control, while identifying the current greedy policy as the stronger nominal RF and control-cost baseline.

8 Conclusions

This paper integrates the electromagnetic basis of a low-profile railway rooftop radiator with a trajectory-aware mode-selection model and paired numerical validation. The physical analysis motivates compact geometry, low material loss, broader impedance matching and control of off-axis response. It also identifies the gap between an analytically dimensioned patch and the assumed adaptive states used in the channel surrogate.

Across 50 paired trials at 3.6 GHz, risk-aware control improves mean SINR by 3.989 dB, reduces outage by 13.009 percentage points and improves ALMC6 by 12.160 percentage points relative to fixed control. These results demonstrate consistent conditional benefit under the selected model. Greedy instantaneous SINR selection nevertheless leads all five reported RF metrics, and substantial residual outage remains. The study therefore supports adaptive mode selection and an evidence-based comparison of control policies, while leaving full-wave antenna validation, measured mode switching and railway field performance as essential next steps. The recovered switch and energy records favour greedy selection as well: risk-aware control adds 27.58 switches per route (95% paired t interval 25.153–30.007) and 2.206 mJ of control-energy proxy. Explicit source-level parameters make the numerical comparison reproducible, while the angular-pattern audit and oracle-prediction disclosure define its physical and causal limits.

References

[1] B. Ai, X. Cheng, T. Kürner, Z.-D. Zhong, K. Guan, R.-S. He, L. Xiong, D. W. Matolak, D. G. Michelson, and C. Briso-Rodriguez, "Challenges toward wireless communications for high-speed railway," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, no. 5, pp. 2143–2158, 2014. https://doi.org/10.1109/TITS.2014.2310771

[2] X. Chen, W. Chen, X. Gong, B. Ai, and I. Wassell, "Wireless channel estimation for high-speed rail communications: Challenges, solutions and future directions," *High-speed Railway*, vol. 1, no. 1, pp. 18–22, 2023. https://doi.org/10.1016/j.hspr.2022.11.004

[3] N. Salhab, A. Haidar, J. J. Muñoz, and C. Reboul, "Train antenna requirements, design and integration for 5GRAIL project," in *Proceedings of the IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, 2023, pp. 1–5. https://doi.org/10.1109/VTC2023-Spring57618.2023.10199389

[4] Y. Bai, J. Ren, and Y. Wen, "Fast optimization of the installation position of 5G-R antenna on the train roof," *Applied Sciences*, vol. 14, no. 16, article 6954, 2024. https://doi.org/10.3390/app14166954

[5] Y. Shi, Y. Wu, C. Hu, Z. Ma, and W. Luo, "A miniaturized wideband high-gain end-fire antenna for 5G-R communication applications," *High-speed Railway*, vol. 2, no. 4, pp. 259–264, 2024. https://doi.org/10.1016/j.hspr.2024.11.004

[6] P. Liu, J. Feng, W. Ge, H. Wang, X. Liu, D. Wang, T. Song, and J. Chen, "A system-level performance evaluation for a 5G system under a leaky coaxial cable MIMO channel for high-speed trains in the railway tunnel," *Electronics*, vol. 11, no. 8, article 1185, 2022. https://doi.org/10.3390/electronics11081185

[7] ITU-R, *Introduction to railway communication systems*, Report M.2395-0, November 2016. https://www.itu.int/pub/R-REP-M.2395-2016

[8] ITU-R, *Description of railway radiocommunication systems between train and trackside (RSTT)*, Report M.2418-0, November 2017. https://www.itu.int/pub/R-REP-M.2418

[9] C. A. Balanis, *Antenna Theory: Analysis and Design*, 4th ed. Wiley, 2016. ISBN 978-1-118-64206-1.

[10] Rogers Corporation, *RT/duroid 5870/5880 High Frequency Laminates*, data sheet, 2022. https://www.rogerscorp.com/advanced-electronics-solutions/rt-duroid-laminates/rt-duroid-5880-laminates

[11] J. J. Borchardt and T. C. LaPointe, "U-slot patch antenna principle and design methodology using characteristic mode analysis and coupled mode theory," *IEEE Access*, vol. 7, pp. 109375–109385, 2019. https://doi.org/10.1109/ACCESS.2019.2933175

[12] D. M. Pozar, *Microwave Engineering*, 4th ed. Wiley, 2012. ISBN 978-0-470-63155-3.

[13] 3GPP, *Study on channel model for frequencies from 0.5 to 100 GHz*, TR 38.901, version 17.1.0, Release 17, ETSI TR 138 901, 2024. https://portal.etsi.org/webapp/workprogram/Report_WorkItem.asp?WKI_ID=69925

[14] 3GPP, *NR and NG-RAN overall description; Stage 2*, TS 38.300, version 18.4.0, Release 18, ETSI TS 138 300, January 2025. https://www.etsi.org/deliver/etsi_ts/138300_138399/138300/18.04.00_60/ts_138300v180400p.pdf

[15] 3GPP, *NR; Radio Resource Control (RRC); Protocol specification*, TS 38.331, version 18.4.0, Release 18, ETSI TS 138 331, January 2025. https://www.etsi.org/deliver/etsi_ts/138300_138399/138331/18.04.00_60/ts_138331v180400p.pdf

[16] 3GPP, *NR; Physical layer procedures for data*, TS 38.214, version 18.4.0, Release 18, ETSI TS 138 214, October 2024. https://www.etsi.org/deliver/etsi_ts/138200_138299/138214/18.04.00_60/ts_138214v180400p.pdf

[17] ITU-R, *Guidelines for evaluation of radio interface technologies for IMT-2020*, Report M.2412-0, November 2017. https://www.itu.int/pub/R-REP-M.2412-2017

[18] NIST/SEMATECH, "Confidence limits for the mean," *e-Handbook of Statistical Methods*, Section 1.3.5.2, accessed October 13, 2026. https://www.itl.nist.gov/div898/handbook/eda/section3/eda352.htm

[19] ITU-R, *Minimum requirements related to technical performance for IMT-2020 radio interface(s)*, Report M.2410-0, November 2017. https://www.itu.int/pub/R-REP-M.2410-2017

[20] ITU-R, *Requirements, evaluation criteria and submission templates for the development of IMT-2020*, Report M.2411-0, November 2017. https://www.itu.int/pub/R-REP-M.2411-2017