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.

Predictive Coverage Placement of Multiple Pinching Antennas with Finite Movement Speed for Multiuser Mobile Indoor Networks

A. Bamila Virgin Louis

bamilavirgin@gmail.com

St. Xavier's Catholic College of Engineering

G Arul Dalton

Saveetha Engineering College

Research Article

Keywords: Pinching antennas, Finite antenna speed, Multiple antennas per waveguide, User mobility, Blockage map, Dynamic programming

Posted Date: November 19th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11363477/v1

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

Additional Declarations: No competing interests reported.

Predictive Coverage Placement of Multiple Pinching Antennas with Finite Movement Speed for Multiuser Mobile Indoor Networks

Dr. A. Bamila Virgin Louis[1*], Dr G Arul Dalton[2]

1. Dept. of Computer Science and Engineering, St. Xavier's Catholic College of Engineering, Nagercoil, India.

Corresponding Author: Email: bamilavirgin@gmail.com

2. Dept. of Computer Science and Engineering, Saveetha Engineering College, Chennai, India.

Email: garuldalton@gmail.com

Abstract

Pinching-antenna systems (PASS) create line-of-sight links on demand by sliding low-cost radiating elements along dielectric waveguides, but existing placement designs assume static users and instantaneous repositioning. This paper studies multiuser, multi-waveguide PASS with finite antenna speed, mobile users, dynamic blockage, and possibly several pinching antennas (PAs) per waveguide. We show that one-to-one user tracking incurs an average travel time of $L/(3v_{\max})$ per re-association, which exceeds the scheduling time scale at practical speeds, and that confining M PAs to separate waveguide segments shortens this time by a factor of M. We propose PRISM, a training-free, two-timescale algorithm that combines Kalman trajectory prediction, a waveguide-aware online blockage map, and speed-constrained coverage placement solved by block-coordinate ascent. For a tractable single-waveguide case, an exact dynamic-programming optimum shows that predictive placement attains 93.2% of the best achievable speed-limited rate. Over 20 realizations of a 60 m × 12 m hall at 28 GHz, PRISM achieves 17.49 bit/s/Hz, 4.2% above reactive tracking and 59.4% above a TD3 agent, and two PAs per waveguide raise the fifth-percentile user rate by 10.7%. In a warehouse with tall shelving and lane-following vehicles, three PAs per waveguide increase the sum rate by 75.4% over a single PA and by 16.5% over reactive tracking with the same hardware. An oracle decomposition shows that the remaining headroom stems from blockage knowledge rather than trajectory prediction.

Keywords

Pinching antennas · Finite antenna speed · Multiple antennas per waveguide · User mobility · Blockage map · Dynamic programming

1 Introduction

1.1 Background and Motivation

Sixth-generation (6G) networks are expected to support immersive, industrial, and massive machine-type services with peak rates of hundreds of gigabits per second and sub-millisecond latency [1], [2]. Meeting these targets requires the large bandwidths available in the millimeter-wave (mmWave) and sub-terahertz bands [3], [4]. Signals in these bands, however, suffer severe free-space path loss and are easily blocked by people, machines, and shelving, so the achievable rate depends critically on the availability of a strong line-of-sight (LoS) path [5], [6]. Indoor industrial environments such as factories, warehouses, and logistics halls are particularly demanding: they combine long, obstacle-rich spaces with mobile workers and automated guided vehicles (AGVs) that require reliable, high-rate connectivity, and dedicated indoor-factory channel models have been standardized for them [7].

Conventional remedies rely on large antenna arrays and beamforming at fixed base station (BS) locations, often combined with learning-based beam combining or beam management [8], [9]. Such arrays increase hardware cost and power consumption, and they cannot restore a blocked LoS path because the antennas remain at fixed positions, far from many users. Flexible-antenna technologies relax this limitation by moving the radiating elements. Fluid antennas switch among the ports of a reconfigurable radiator [10], and movable antennas shift elements within a confined region [11], [12]. Both reshape small-scale fading and spatial correlation, but their movement range spans only a few wavelengths, so they cannot reduce the large-scale path loss or bring the antenna closer to a distant user.

Pinching-antenna systems (PASS), first demonstrated by NTT DOCOMO [13], overcome this limitation. The signal travels along a low-loss dielectric waveguide, and a small dielectric particle, the "pinch", placed anywhere on the waveguide makes it radiate at that point. Because the waveguide can be tens of meters long and the pinch can be repositioned, a pinching antenna (PA) can be brought close to a user, creating a strong LoS link at very low cost [14]–[17]. PASS therefore offers a large-scale, reconfigurable form of antenna flexibility that is particularly attractive for long indoor halls.

1.2 Pinching-Antenna Systems with Mobile Users and Finite Speed

A rapidly growing literature has characterized PASS in terms of array gain [18], rate maximization [19], [20], non-orthogonal multiple access (NOMA) [21], [22], multiple-input multiple-output (MIMO) operation [23], and robustness to blockage [24], [25] and to location uncertainty [26]. Most of these works, however, share two idealizations that limit their applicability to industrial deployments. First, users are static, or the antenna position is re-optimized independently in every slot as if the users had no history. Second, the PA can be moved to any position instantaneously. In practice, a pinch is moved by

a mechanical actuator at a finite speed, typically on the order of a meter per second [27], while workers and vehicles move continuously at comparable speeds [28].

These two facts change the nature of the placement problem. With finite speed, a placement decision constrains all future positions, so the PA must anticipate where users will be rather than react to where they are. With several mobile users sharing a waveguide, a single PA cannot be close to all of them at once, and chasing individual users wastes time in transit. Blockage by moving people and static shelving further changes which positions offer a LoS path, and this information is rarely known in advance. Finally, a waveguide may carry several PAs, which shortens travel distances but splits the transmit power among them. How PAs should be placed under these coupled constraints, and how close practical rules can come to the best achievable performance, are open questions.

1.3 Related Work

Static PASS optimization. Ding *et al.* [14] introduced the PASS system model and showed its advantage over conventional fixed antennas. Placement for rate maximization was solved in closed form or by low-complexity search in [19], [20], [29], and the array gain of PASS was characterized in [18]. Multi-waveguide NOMA and MIMO-PASS designs were studied in [21]–[23], and resource allocation for PASS is surveyed in [30]. All of these works assume static users and instantaneous placement.

Blockage and uncertainty. LoS blockage in PASS was analyzed in [24], [25], where blockage was even shown to help by suppressing inter-user interference. Robust placement under user-location uncertainty was studied in [26]. Measurements show that human bodies cause rapid, deep fades at mmWave frequencies [6], which motivates online rather than offline blockage models.

Mobility and finite speed. A deep deterministic policy gradient (DDPG) agent that jointly optimizes beamforming and placement for one mobile user was proposed in [31]. Trajectory-aware reconfiguration with linear-extrapolation prediction for a single waveguide was analyzed by Monte Carlo simulation in [32]. Antenna roaming, in which a finite-speed PA keeps transmitting while it moves, was introduced in [27], and finite-speed PAs were combined with federated-learning scheduling in [33]. None of these works considers several mobile users on several waveguides together with dynamic blockage.

Learning-based designs. AI-enabled PASS is surveyed in [34], and meta-learning [35] and deep reinforcement learning (DRL) [36] have been applied to static or wireless-powered PASS. DRL has also been used for adaptive beamforming [9] and for proactive handover decisions [37] in 5G and beyond. Actor–critic methods such as DDPG [38] and twin-delayed DDPG (TD3) [39] are natural candidates for continuous PA placement, but they require extensive training and may not transfer across layouts.

Environment awareness and scheduling. Channel knowledge maps (CKMs) store location-indexed channel information to support environment-aware 6G design [40], [41], but PASS has not previously been combined with an online, waveguide-aware blockage map. On the fast time scale, proportional-fair (PF) scheduling with zero-forcing (ZF) precoding [42], [43] is a well-established multiuser transmission layer, which this work adopts.

Table 1. Positioning of this work (✓: explicitly modeled; –: not addressed; entries follow each paper's published summary).

Work	**Multi-user**	**Multi-waveguide**	**User mobility**	**Finite PA speed**	**Dynamic blockage**	**Approach**
[14], [19], [20]	✓/–	–	–	–	–	Analysis / placement
[22], [23]	✓	✓	–	–	–	NOMA / MIMO optimization
[24], [25]	✓	–	–	–	Static	Analysis / optimization
[26]	✓	–	Uncertainty	–	–	Robust optimization
[31]	–	✓	✓	–	✓	DDPG
[32]	✓ (OMA)	–	✓	Update interval	–	Prediction + Monte Carlo
[27]	✓	–	–	✓	–	Dynamic programming
[33]	✓	–	–	✓	–	Coalitional game
This work	✓	✓	✓	✓	✓ (learned)	Predictive coverage, multi-PA, DP optimum

1.4 Research Gap and Questions

Table 1 summarizes the positioning of this work. To the best of our knowledge, no prior work jointly addresses multiple mobile users, multiple waveguides with possibly several PAs each, finite PA speed, and dynamic blockage, or quantifies how far practical placement rules are from the speed-limited optimum. This paper therefore addresses three questions: (i) why do natural extensions of static placement fail when PAs move at a finite speed; (ii) how should PAs be placed, and how many PAs per waveguide are worthwhile, when users move and links are blocked; and (iii) how close can a practical, training-free placement rule come to the best achievable speed-limited performance?

1.5 Contributions

The main contributions of this paper are as follows.

- **System model.** We formulate long-term proportional-fair utility maximization for a multiuser, multi-waveguide PASS with $M \geq 1$ PAs per waveguide. The model includes finite PA speed, random-waypoint user mobility with noisy positioning, geometric blockage by static and moving obstacles, in-waveguide attenuation, Rician fading, and zero-forcing (ZF) multiuser transmission.
- **Analytical insight: coverage versus tracking.** We show that a PA re-associating between uniformly distributed users travels on average $L/3$, so the travel time $L/(3v_{\max})$ exceeds typical PF time constants, and that segment-confined PAs reduce it to $L/(3Mv_{\max})$. This explains why one-to-one tracking performs poorly even with perfect information and motivates a *coverage* objective.
- **PRISM algorithm.** We propose a training-free, two-timescale algorithm that integrates Kalman trajectory prediction, a waveguide-aware online radio blockage map (WA-ORBM) learned from routine channel state information (CSI), and speed-constrained, horizon-based coverage placement over all NM PAs, solved by monotone block-coordinate ascent (BCA) with polynomial complexity.
- **Optimality benchmark.** For a single waveguide with max-rate scheduling, we compute the exact speed-limited optimum by dynamic programming (DP). Predictive placement attains 93.2% of this optimum on average, versus 90.0% for reactive tracking and 72.6% for a fixed PA.
- **Transparent evaluation.** Using an open, seeded simulator, we compare PRISM with genie, oracle, tracking, roaming, fixed, MIMO, DDPG, and TD3 benchmarks over up to 20 realizations with 95% confidence intervals, in a default hall and in a warehouse with tall shelving and lane-following vehicles. An oracle decomposition and an ablation identify which information and which components create the gains, including those that do *not* help.

1.6 Organization and Notation

Section 2 presents the system model and Section 3 the problem formulation, the coverage-versus-tracking analysis, and the DP optimum. Section 4 develops PRISM. Sections 5 and 6 describe the simulation setup and results, Section 7 discusses limitations, and Section 8 concludes. Bold lower-case and upper-case letters denote vectors and matrices; $(\cdot)^{H}$ is the conjugate transpose; $[\mathbf{A}]_{kk}$ is the k-th diagonal entry; $\text{clip}(a, -b, b) = \max(-b, \min(a, b))$; and $\mathbb{1}\{\cdot\}$ is the indicator function.

2 System Model

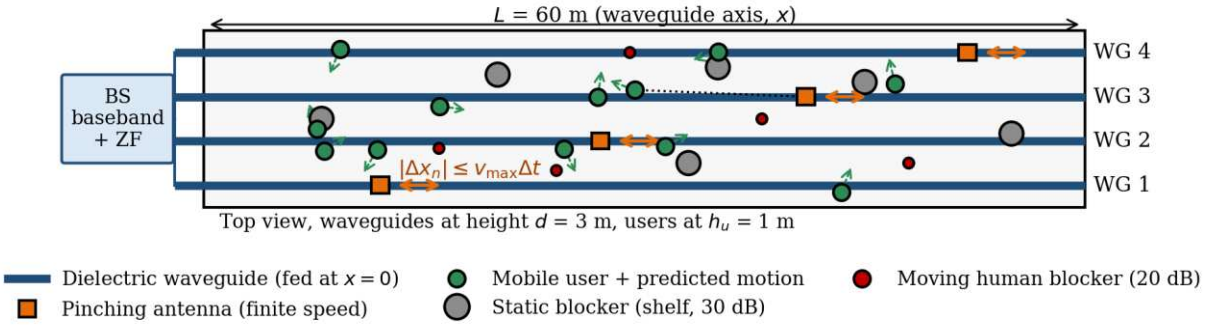


Fig. 1. System model (top view). A base station (BS) feeds N parallel dielectric waveguides mounted at height d along an $L \times W$ hall. Each waveguide carries one pinching antenna whose position can change by at most $v_{\max}\Delta t$ per control interval. Mobile users are served by ZF precoding across waveguides, while static shelves and moving people intermittently block LoS links.

2.1 Network Geometry

Consider the indoor hall of length L and width W shown in Fig. 1. A BS with N radio-frequency chains feeds N parallel dielectric waveguides at height d. Waveguide n runs along the x-axis at lateral offset $y_n = W(2n - 1)/(2N)$ and is fed at $x = 0$. As in multi-waveguide PASS [22], [23], each waveguide carries one PA at position $x_n(t) \in [0, L]$, so the BS can transmit up to N spatial streams. Time is slotted into control intervals of duration Δt. $K \geq N$ single-antenna users at height h_u move in the hall; the position of user k in slot t is $\mathbf{u}_k(t) = [u_{k,x}(t), u_{k,y}(t)]$.

2.2 Finite-Speed Pinching Antenna Movement

A PA is moved by a mechanical actuator with maximum speed $v_{\max}$, which gives the constraints

$$0 \leq x_n(t) \leq L, \quad |x_n(t+1) - x_n(t)| \leq v_{\max}\Delta t, \quad \forall n, t. \tag{1}$$

As in antenna roaming [27], the PA keeps radiating while it moves, so no slot is lost during repositioning. Constraint (1) couples decisions over time: a position that is good now determines which positions are reachable later.

2.3 Channel Model

The distance between the PA on waveguide n and user k is

$$r_{n,k}(t) = [(x_n - u_{k,x})^2 + (y_n - u_{k,y})^2 + (d - h_u)^2]^{1/2}. \quad (2)$$

Following [14], [23], the signal accumulates an in-waveguide phase $2\pi n_{\text{eff}} x_n/\lambda$, where n_{eff} is the effective refractive index, and an in-waveguide attenuation of α dB/m [15]. With free-space gain $\eta = (\lambda/4\pi)^2$, the channel coefficient is

$$h_{k,n} = \sqrt{\eta} \cdot a(x_n) \cdot b_{n,k} \cdot r_{n,k}^{-1} \cdot [\sqrt{(\kappa/(\kappa+1))}\exp(-j\varphi_{n,k}) + \sqrt{(1/(\kappa+1))}\, g_{k,n}], \quad (3)$$

where $a(x_n) = \exp(-\alpha' x_n)$ with $\alpha' = \alpha \ln 10/20$ models the in-waveguide attenuation, $b_{n,k} = \exp(-\beta'_{n,k})$ with $\beta'_{n,k} = \beta_{n,k} \ln 10/20$ models the blockage loss, $\varphi_{n,k} = 2\pi(r_{n,k} + n_{\text{eff}} x_n)/\lambda$ is the total phase, κ is the Rician factor, $g_{k,n} \sim \mathcal{CN}(0, 1)$ captures diffuse scattering, and $\beta_{n,k}$ is the blockage loss in dB defined next. The channel matrix is $\mathbf{H}(t) = [h_{k,n}] \in \mathbb{C}^{K \times N}$.

2.4 Geometric Blockage Model

Obstacles are vertical cylinders $m \in \mathcal{M}$ with center $\mathbf{o}_m$, radius ρ_m, height ζ_m, and penetration loss L_m, with values consistent with mmWave human-blockage measurements [6] and the 3GPP indoor model [7]. They consist of M_s static shelves and M_h people who move by random waypoint. Let $\mathbf{p}(\tau) = (1 - \tau)[x_n, y_n] + \tau \mathbf{u}_k$, $\tau \in [0, 1]$, parameterize the horizontal projection of the link, and let τ_{m} *be the value closest to* o_{m}. *Obstacle* m* blocks the link if the ray passes within ρ_m of the obstacle axis below its top:

$$\beta_{n,k} = \min\{ \Sigma_{m \in \mathcal{M}} L_m \mathbb{1}[\|\mathbf{p}(\tau_{m}) - o_{m}\| < \rho_{m}, d + \tau_m(h_u - d) < \zeta_m], 40 \text{ dB} \}. \quad (4)$$

Because the PA sits above most obstacles, blockage is typically caused near the user end of the link, and it depends on the direction from which the PA illuminates the user. The ORBM in Section 4 exploits this property.

2.5 Mobility, Positioning, and Transmission

Users follow a random-waypoint model [28] with individual speeds drawn uniformly in $[0.5\bar{v}, 1.5\bar{v}]$. The BS observes noisy positions $\tilde{\mathbf{u}}_k(t) = \mathbf{u}_k(t) + \mathbf{e}_k(t)$, with $\mathbf{e}_k \sim \mathcal{N}(\mathbf{0}, \sigma_p^2 \mathbf{I})$, for example from 5G positioning or sensing. In each slot the BS schedules a set $\mathcal{S}(t)$ of at most N users and applies ZF precoding with equal power per stream. User $k \in \mathcal{S}$ then obtains

$$R_k(t) = \log_2(1 + P / (|\mathcal{S}| \sigma^2 [(\mathbf{H}_{\mathcal{S}} \mathbf{H}_{\mathcal{S}}^H)^{-1}]_{kk})), \quad (5)$$

and $R_k(t) = 0$ otherwise. Here P is the total transmit power, σ^2 is the noise power, and $\mathbf{H}_{\mathcal{S}}$ contains the rows of $\mathbf{H}$ indexed by $\mathcal{S}$. CSI at the current PA positions is obtained from uplink pilots in every slot.

2.6 Multiple PAs per Waveguide

A waveguide may carry $M \geq 1$ PAs that radiate the same data stream. As in [14], [22], the input power of waveguide n is split equally among its PAs. To avoid collisions and to bound travel distances, PA m of waveguide n is confined to the segment $[(m-1)L/M, mL/M]$ and moves under the speed constraint (1). The channel from waveguide n to user k is the coherent sum

$$h_{k,n} = \sum_{m=1}^{M} (1/\sqrt{M})\, h_{k,n,m}, \quad (5a)$$

where $h_{k,n,m}$ follows (3) with x_n replaced by the position $x_{n,m}$ of PA m, and each PA has independent diffuse scattering. Additional PAs shorten distances and travel times but split the power by M, a trade-off quantified in Section 6.

3 Problem Formulation and Coverage-versus-Tracking Analysis

3.1 Long-Term Proportional-Fair Utility

Each user maintains an exponentially averaged throughput $\bar{R}_k(t+1) = (1 - 1/t_c)\bar{R}_k(t) + R_k(t)/t_c$. We seek PA trajectories and schedules that maximize long-term proportional fairness:

$$\text{(P1)} \quad \max_{\{\mathbf{x}(t),\, \mathcal{S}(t)\}} \ \lim_{T\to\infty} (1/T) \sum_{t=1}^{T} \mathbb{E}\Big[\sum_{k=1}^{K} \log \bar{R}_k(t) \Big] \quad \text{s.t. (1)}, \ |\mathcal{S}(t)| \leq N. \qquad (6)$$

Problem (P1) is a stochastic control problem with a continuous, speed-coupled action space, a combinatorial scheduling action, an unknown random blockage process, and partially observed user states. Even its single-slot version is non-convex in $\mathbf{x}$ because of the phase terms and the maximum over PAs [14], [22]. We therefore decompose it on two timescales. On the fast scale, PF-ZF scheduling selects $\mathcal{S}(t)$, which is the standard gradient step for the PF utility [42], [43]. On the slow scale, PA roaming shapes the channels that the scheduler can exploit.

3.2 Why One-to-One Tracking Fails under Finite Speed

A natural extension of static PASS designs is to assign each PA to one (high-priority) user and move the PA toward that user; we call this *tracking*. Under PF scheduling, the served user's weight $w_k = 1/\bar{R}_k$ decreases while it is being served, so the assignment changes every few time constants t_c.

Proposition 1 (Re-association travel): If the x-coordinates of the old and new target users are independent and uniform on $[0, L]$, the expected PA travel distance per re-association is $L/3$, and the expected travel time is $T_{tr} = L/(3v_{max})$.

Proof: For independent $X, Y \sim \mathcal{U}[0, L]$, $\mathbb{E}|X - Y| = (2/L^2)\int_0^L\int_0^x (x - y)\mathrm{d}y\mathrm{d}x = (2/L^2)\cdot L^3/6 = L/3$. Dividing by the maximum speed gives T_{tr}. ■

For the default hall (L = 60 m) and $v_{\max}$ = 2 m/s, T_{tr} = 10 s. This is twice the PF averaging window $t_c\Delta t$ = 5 s. A tracking PA therefore spends most of its time in transit toward a target whose priority will have changed before arrival. The simulations confirm this: a *genie-informed* tracker with perfect positions and blockage knowledge achieves only 15.26 bit/s/Hz, below a conventional co-located MIMO BS (15.36 bit/s/Hz). The remedy is to change the objective. Instead of chasing one user, each PA should occupy a position that serves the *set* of users it can reach well, now and over a short prediction horizon, and leave user selection to the fast scheduler. This coverage principle underlies PRISM.

With M segment-confined PAs per waveguide, the same argument applied to a segment of length L/M gives $T_{\mathrm{tr}} = L/(3Mv_{\max})$. Multiple PAs therefore attack the travel-time bottleneck directly, at the cost of the $10\log_{10}M$ dB power split in (5a).

3.3 Speed-Limited Optimum for a Tractable Case

Problem (P1) is intractable in general, so practical schemes cannot be compared with its optimum directly. We therefore isolate the placement layer in a tractable special case: a single waveguide with one PA, K users, and max-rate scheduling, in which the user with the best large-scale channel (including blockage) is served in each slot. Discretizing the waveguide with step $v_{\max}\Delta t$, the PA moves at most one grid point per slot, and the reward $\rho_t(i) = \max_k \log_2(1 + PG_k(x_i, t)/\sigma^2)$ of grid point i at slot t is known offline. The optimal speed-limited trajectory follows from the DP recursion

$$V_t(i) = \rho_t(i) + \max V_{t+1}(j) \text{ over } j \in \{i - 1, i, i + 1\}, \quad (6a)$$

whose complexity is linear in the horizon and the grid size. Because the DP uses the true future positions and blockage, it upper-bounds every causal speed-limited rule in this case, and the ratio between an online rule and the DP value measures its optimality gap.

4 Proposed PRISM Algorithm

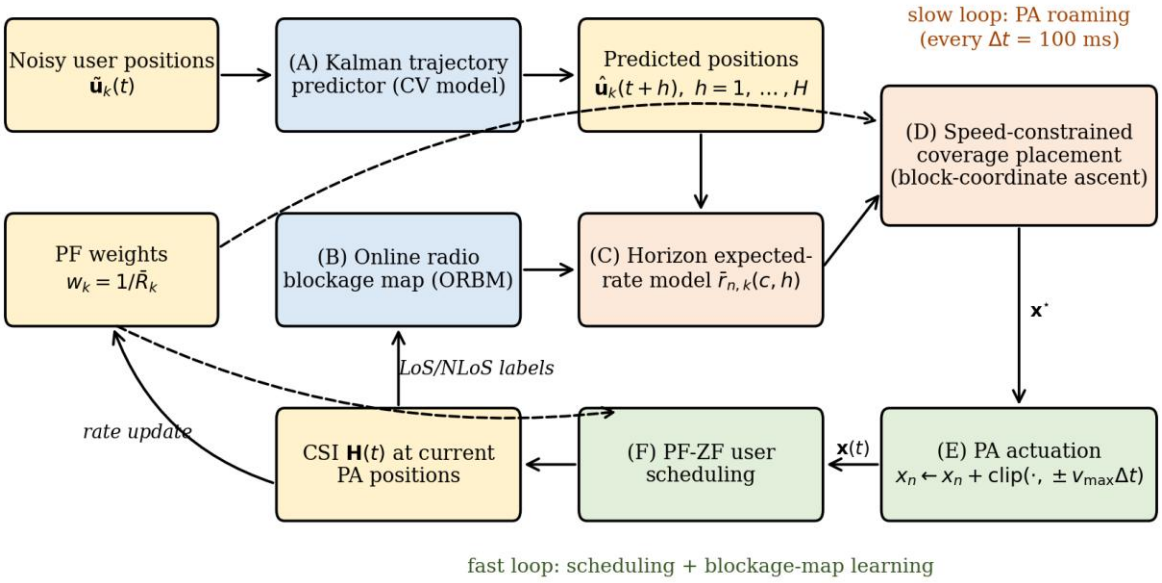


Fig. 2. Block diagram of PRISM. The slow loop (A)–(E) predicts user trajectories, evaluates horizon expected rates using the learned blockage map, and moves each PA toward a speed-reachable coverage position. The fast loop (F) performs PF-ZF scheduling on the resulting channels and feeds LoS/NLoS labels back to the ORBM.

4.1 Overview

PRISM runs once per control interval Δt. Its modules are shown in Fig. 2: (A) a Kalman predictor, (B) the WA-ORBM, (C) a horizon expected-rate model, (D) speed-constrained coverage placement via BCA, (E) PA actuation, and (F) PF-ZF scheduling, whose CSI is reused to update the WA-ORBM. No offline training is required.

4.2 Kalman Trajectory Prediction

Each user is tracked with a constant-velocity Kalman filter [44]. Learning-based mobility prediction has also been applied to proactive handover decisions in 5G [37]; we adopt the Kalman filter because it is training-free and has a negligible computational cost. The filter has state $\mathbf{s}_k = [u_{k,x}, u_{k,y}, \dot{u}_{k,x}, \dot{u}_{k,y}]^T$, transition matrix $\mathbf{F} = [\mathbf{I}_2, \Delta t\mathbf{I}_2; \mathbf{0}, \mathbf{I}_2]$, white-acceleration process noise with intensity q, and measurement noise $\sigma_p^2\mathbf{I}_2$. After the update with $\tilde{\mathbf{u}}_k(t)$, the h-step prediction is

$$\hat{\mathbf{u}}_k(t+h) = [\mathbf{I}_2 \ \ \mathbf{0}]\, \mathbf{F}^h\, \hat{\mathbf{s}}_k(t), \quad h = 1, \ldots, H, \tag{7}$$

clipped to the hall boundaries.

4.3 Waveguide-Aware Online Radio Blockage Map (WA-ORBM)

Section 2.4 shows that blockage depends on where the user is and on the line along which the PA illuminates it. Because every PA is confined to a known waveguide, the WA-ORBM indexes a cell $c = (n, \iota_u, \iota_x)$, where n is the waveguide, ι_u is the user's grid cell (side δ_u), and ι_x is a signed bin of the along-waveguide offset $x_n - u_{k,x}$. Indexing by waveguide separates links that share an azimuth but cross different obstacles, for example shelf rows between aisles, which a purely direction-based map would merge. In every slot, the BS already knows $|h_{k,n}|^2$ and the modeled LoS gain G_0 (3) without blockage, evaluated at the estimated position. It forms a LoS label $\ell = \mathbb{1}[10\log_{10}(|h_{k,n}|^2/G_0) > \theta]$ and updates Beta–Bernoulli statistics with forgetting factor ϱ:

$$S_c \leftarrow \varrho S_c + \ell, \quad C_c \leftarrow \varrho C_c + 1, \quad \hat{\pi}_c = (S_c + c_0) / (C_c + c_0). \tag{8}$$

The prior pseudo-count c_0 encodes an optimistic LoS prior, so unexplored directions remain attractive, in the spirit of optimism under uncertainty. The forgetting factor tracks moving blockers. The NLoS attenuation $\hat{\kappa}$ is estimated as a running mean of the NLoS-labelled observations. An ORBM is a lightweight, online, blockage-specific instance of a channel knowledge map [40], [41].

4.4 Horizon Expected-Rate Model

If PA n heads toward candidate position c from a finite grid $\mathcal{C}$ (spacing δ_c), its trajectory under (1) is

$$x_n^{(c)}(t+h) = x_n(t) + \mathrm{clip}(c - x_n(t), -h \cdot v_{max}\Delta t, h \cdot v_{max}\Delta t). \tag{9}$$

With the LoS SNR $\gamma = PG_0(x_n^{(c)}(t+h), \hat{\mathbf{u}}_k(t+h))/\sigma^2$, the expected rate accounts for the learned blockage mixture:

$$\bar{r}_{n,k}(c, h) = \hat{\pi}\, \log_2(1 + \gamma) + (1 - \hat{\pi})\, \log_2(1 + \hat{\kappa}\gamma). \tag{10}$$

Averaging the rate rather than the SNR is important: blockage is bimodal (roughly 0 dB or 20–30 dB of loss), and averaging the linear attenuation grossly overestimates the rate of partially blocked directions.

4.5 Speed-Constrained Coverage Placement

With PF weights $w_k = 1/\bar{R}_k$ and discount weights $\omega_h = \varsigma^h / \Sigma_{h'} \varsigma^{h'}$, the coverage objective for a joint choice $\mathbf{c} = (c_1, \ldots, c_N)$ is

$$J(\mathbf{c}) = \Sigma_{k=1}^{K} w_k \, \Sigma_{h=1}^{H} \omega_h \, \max_n \bar{r}_{n,k}(c_n, h). \tag{11}$$

The inner maximum reflects the fact that the ZF scheduler can serve each user from whichever PA illuminates it best. Unlike a one-to-one assignment, (11) rewards a PA for covering several users at once and never pays for long trips unless they increase the weighted coverage within the horizon. Because J

is non-concave, we use BCA over waveguides. Starting from the current positions, each PA in turn selects

$$c_n \leftarrow \operatorname{argmax}_{c \in \mathcal{C}} J(c, \mathbf{c}_{-n}), \tag{12}$$

and the procedure stops after I rounds or when no PA changes. Each update does not decrease J, and $\mathcal{C}^N$ is finite, so BCA converges to a coordinate-wise optimum in finitely many steps [45]. Only the first step of each planned trajectory is applied; this receding-horizon principle also absorbs prediction errors.

Multiple PAs per waveguide. With $M > 1$, BCA runs over all NM PAs. Each PA chooses among candidate positions in its own segment, its expected rate uses the power share $1/M$, and the inner maximum in (11) is taken over all other PAs. The receding-horizon actuation and the speed constraint are unchanged.

Algorithm 1: PRISM (executed at every control interval t)

Input: noisy positions $\tilde{\mathbf{u}}_k(t)$, PA positions $\mathbf{x}(t)$, PF averages $\bar{R}_k$, ORBM statistics $\{S_c, C_c\}$

1: Kalman update with $\tilde{\mathbf{u}}_k(t)$; predict $\hat{\mathbf{u}}_k(t+h)$, $h = 1..H$, using (7)

2: for all n, k, $c \in \mathcal{C}$, h: build trajectories (9), look up $\hat{\pi}$ in the ORBM, and compute $\bar{r}_{n,k}(c, h)$ by (10)

3: initialize $c_n \leftarrow$ nearest grid point to $x_n(t)$; compute $w_k = 1/\bar{R}_k$

4: for $i = 1$ to I do

  5: for $n = 1$ to N do $c_n \leftarrow \operatorname{argmax}_c J(c, \mathbf{c}_{-n})$ by (11)–(12)

  6: stop if no c_n changed

7: actuate $x_n(t+1) = x_n(t) + \mathrm{clip}(c_n - x_n(t), \pm v_{\max}\Delta t)$

8: acquire $\mathbf{H}(t+1)$; greedy PF-ZF selection of $\mathcal{S}$ maximizing $\Sigma_{k \in \mathcal{S}} w_k R_k$ using (5)

9: update $\bar{R}_k$; form LoS labels and update the ORBM by (8)

Output: $\mathbf{x}(t+1)$, $\mathcal{S}(t+1)$, rates $R_k(t+1)$

4.6 PF-ZF Scheduling

The fast loop uses the classical greedy user-selection rule [43]. Starting from $\mathcal{S} = \emptyset$, it repeatedly adds the user that most increases $\Sigma_{k \in \mathcal{S}} w_k R_k$, stopping when no addition helps or $|\mathcal{S}| = N$. This scheduler is shared by PRISM and all PASS benchmarks, so performance differences come only from PA positioning.

4.7 Complexity

Step 2 costs $\mathcal{O}(NMKCH)$ with $C = |\mathcal{C}|$ candidates per PA. Each BCA round costs $\mathcal{O}(NM \cdot NMKCH)$ because of the maximum over the other PAs. The Kalman filter costs $\mathcal{O}(K)$, and greedy PF-ZF costs

$\mathcal{O}(N\text{^}2K \cdot N\text{^}3)$ with $N \times N$ inversions. For $N = 4$, $K = 12$, $H = 10$, and $I \leq 3$, our unoptimized single-threaded Python implementation needs 4.7 ms per interval for $M = 1$ and 5.5 ms for $M = 3$, well below $\Delta t = 100$ ms. PRISM requires no training data, unlike DRL-based designs [31], [36].

5 Simulation Setup

All results are produced by an open, seeded Python simulator that implements (1)–(12) exactly; no result is analytical or extrapolated. Each run consists of a 20-s warm-up (200 slots), during which schedulers, filters, and the WA-ORBM initialize, followed by a 60-s measurement period (600 slots). The default-scenario comparison (Table 3) averages 20 independent realizations of the hall layout, mobility, blockers, and fading, the warehouse comparison (Table 4) averages 8, the DP study averages 20, and every sweep point averages 6. We report means with 95% confidence intervals and use paired differences (Student-t) to compare schemes. All schemes experience identical user trajectories, blockers, and small-scale fading in each realization (common random numbers). Table 2 lists the default parameters; $M = 1$ unless stated otherwise.

Table 2. Default simulation parameters.

Parameter	**Value**	**Parameter**	**Value**
Carrier / bandwidth	28 GHz / 100 MHz	Noise figure / noise power	9 dB / −85 dBm
Hall $L \times W$	60 m × 12 m	Waveguides N / users K	4 / 12
Waveguide / user height	3 m / 1 m	Transmit power P	10 dBm
In-waveguide loss α	0.08 dB/m	Effective index n_{eff}	1.4
Rician factor κ	10 dB	Control interval Δt	100 ms
PA speed $v_{\max}$	2 m/s	Mean user speed $\bar{v}$	1.5 m/s (RWP)
Static shelves	6, radius 0.8 m, 2.2 m, 30 dB	Moving people	10, radius 0.3 m, 1.75 m, 20 dB, 1 m/s
Positioning error σ_p	0.2 m	PF time constant t_c	50 slots (5 s)
Horizon H / discount ς	10 slots / 0.9	Candidate spacing δ_c	0.5 m ($C = 121$)

Parameter	Value	Parameter	Value
ORBM cell / sectors / bins	1.5 m / 12 / {2, 5, 10} m	ORBM ϱ / c_0 / θ	0.998 / 3 / −10 dB
BCA rounds I	3	Realizations	20 (Table 3), 12 (Table 4), 8 (sweeps)

Benchmarks. (i) Ideal PASS (genie): infinite PA speed, true positions, and true blockage; each PA jumps to the best position for a Hungarian-assigned user every slot. (ii) Genie tracking (speed-limited): the same genie decision, approached at $v_{\max}$. (iii) Reactive tracking PASS: Hungarian one-to-one assignment [46] using measured positions, each PA moving toward its user at $v_{\max}$; for $M > 1$, each PA tracks the user with the largest PF-weighted rate within its segment. (iv) Fixed PASS: PAs fixed at the segment midpoints. (v) Conventional MIMO BS: an N-element half-wavelength uniform linear array at the ceiling center without waveguide loss. (vi) DDPG PASS [31], [38] and (vii) TD3 PASS [39]: agents that observe normalized PA positions, measured user positions, and log PF weights and output a target position per PA, approached at $v_{\max}$. Both use actors and critics with two hidden layers of 128 ReLU units, Adam, and a sum-rate reward, and are trained on layouts disjoint from the test layouts; TD3 adds twin critics, delayed actor updates, and target-policy smoothing and is trained for 200 episodes of 400 slots. (viii) Oracle-PRISM: PRISM's placement rule fed with the true future positions and the true geometric blockage over the horizon, and two partial oracles that receive only one of the two. For the ablation, we also evaluate PRISM without prediction ($H = 1$) and without the WA-ORBM.

Warehouse-aisle scenario. To test the components of PRISM where they are most likely to matter, we also simulate an intralogistics hall. Here each waveguide runs above an aisle, and three rows of tall shelving (height 2.8 m, i.e., only 0.2 m below the waveguides; modeled as cylinders of radius 0.5 m every 1 m; 30 dB loss) separate adjacent aisles. Users are automated guided vehicles (AGVs) that drive along their aisle at speeds drawn uniformly in $[0.7\bar{v}, 1.3\bar{v}]$, turn at the hall ends, and change to a neighboring aisle only there. Five people also walk through the hall. Cross-aisle links are almost always blocked, and motion is highly predictable. All other parameters follow Table 2.

6 Results and Discussion

6.1 Default Scenario

Table 3. Default scenario, 20 realizations (rates in bit/s/Hz). The last column gives the mean paired difference PRISM ($M = 1$) minus each scheme, its 95% confidence interval, and the number of realizations in which PRISM is better.

Scheme	Sum rate ± CI	5th-pct.	PRISM − scheme (paired)
Ideal PASS (genie, $v_{max} = \infty$)	22.98 ± 0.15	1.77	−5.48 ± 0.28 (0/20)
Oracle-PRISM	17.51 ± 0.28	1.11	−0.02 ± 0.30 (10/20)
PRISM (M = 1)	17.49 ± 0.28	1.13	—
PRISM (M = 2)	17.47 ± 0.20	1.25	+0.02 ± 0.27 (13/20)
PRISM (M = 3)	16.91 ± 0.19	1.26	+0.58 ± 0.29 (17/20)
PRISM w/o WA-ORBM	17.73 ± 0.28	1.13	−0.24 ± 0.14 (3/20)
PRISM w/o prediction	17.29 ± 0.24	1.09	+0.21 ± 0.28 (13/20)
Reactive tracking	16.79 ± 0.25	1.17	+0.70 ± 0.26 (19/20)
Reactive tracking ($M = 3$)	16.87 ± 0.16	1.29	+0.62 ± 0.26 (19/20)
Genie tracking (speed-limited)	15.26 ± 0.22	1.04	+2.23 ± 0.26 (20/20)
TD3 PASS	10.97 ± 0.45	0.65	+6.52 ± 0.55 (20/20)
Conventional MIMO BS	15.36 ± 0.41	0.88	+2.13 ± 0.47 (19/20)
Fixed PASS ($M = 3$)	13.79 ± 0.12	1.05	+3.70 ± 0.27 (20/20)
Fixed PASS ($M = 1$)	12.17 ± 0.31	0.75	+5.32 ± 0.34 (20/20)

Table 3 summarizes the default hall. PRISM with one PA per waveguide achieves 17.49 ± 0.28 bit/s/Hz. Its paired advantage is +0.70 ± 0.26 bit/s/Hz (19/20, significant) over reactive tracking, +2.13 ± 0.47 bit/s/Hz (19/20, significant) over the co-located MIMO BS, and +2.23 ± 0.26 bit/s/Hz (20/20, significant) over the genie-informed tracker, which confirms Proposition 1: under finite speed, perfect information cannot compensate for a tracking objective. In this hall, additional PAs do not raise the sum rate (PRISM with $M = 2$: 17.47 bit/s/Hz; $M = 3$: 16.91 bit/s/Hz), because the $10\log_{10}M$ dB power split offsets the shorter distances. They do, however, improve fairness: the 5th-percentile user rate rises from 1.13 to 1.25 bit/s/Hz with $M = 2$ (+10.7%), because users far from every single PA are now within one segment of a PA.

Learning-based baselines. TD3 reaches 10.97 ± 0.45 bit/s/Hz, a clear improvement over our earlier DDPG agent (9.44 bit/s/Hz with the same architecture), but still 59.4% below PRISM and below fixed PASS. Fig. 3 shows that both agents plateau early. The agents must learn a long-horizon positioning policy from a reward that depends on PF scheduling and blockage, which is hard with 80 000 transitions; larger budgets or graph-structured policies could narrow the gap. PRISM, in contrast, needs no training and transfers across layouts without retuning.

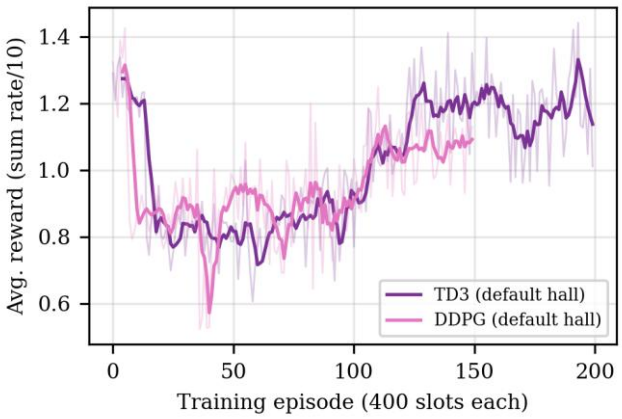


Fig. 3. Training curves of the DRL baselines in the default hall (average per-slot reward, sum rate/10); thin lines show episodes and thick lines a 5-episode moving average.

Ablation and value of information. Relative to PRISM, removing the WA-ORBM changes the sum rate by +0.24 ± 0.14 bit/s/Hz (17/20, significant) and removing prediction by −0.21 ± 0.28 bit/s/Hz (7/20, not significant). The clairvoyant Oracle-PRISM reaches 17.51 bit/s/Hz, only 0.1% above PRISM, so better information would add little in this hall. The infinite-speed genie, in contrast, is 31.2% above the oracle: the dominant loss is PA speed, which motivates the multi-PA results below.

6.2 Gap to the Speed-Limited Optimum

For the tractable case of Section 3.3 (one waveguide, $K = 4$ users, max-rate scheduling, default mobility and blockage, 20 realizations of 60 s), the exact DP optimum attains 7.76 bit/s/Hz. The predictive coverage rule of PRISM, restricted to this case, attains 7.24 bit/s/Hz, i.e., 93.2% of the optimum, compared with 90.0% for reactive tracking and 72.6% for a fixed PA at the waveguide center. The infinite-speed genie reaches 8.17 bit/s/Hz, so finite speed itself costs 5.2% even under optimal clairvoyant control. These numbers show that the placement layer of PRISM operates close to what any speed-limited controller can achieve, and that the residual loss is mostly physical rather than algorithmic.

6.3 Multiple PAs per Waveguide and PA Speed

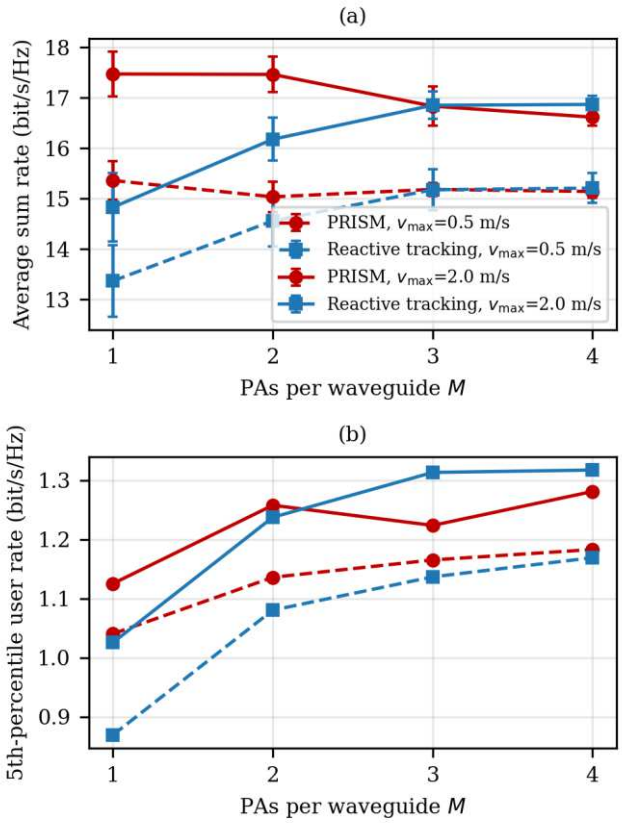


Fig. 4. (a) Average sum rate (95% confidence intervals) and (b) 5th-percentile user rate versus the number of PAs per waveguide M for two PA speeds (default hall, 6 realizations per point). Reactive tracking uses the per-segment rule of benchmark (iii) for all M, so its $M = 1$ value differs from the Hungarian tracker in Table 3.

Fig. 4 varies M at a slow (0.5 m/s) and a nominal (2 m/s) PA speed. Additional PAs mainly help tracking: reactive tracking improves from 13.37 to 15.21 bit/s/Hz at 0.5 m/s and from 14.83 to 16.87 bit/s/Hz at 2 m/s as M grows from 1 to 4, consistent with the $L/(3Mv_{\max})$ travel-time scaling. PRISM, whose coverage objective already avoids long trips, gains little in sum rate (17.47, 17.47, 16.84, and 16.62 bit/s/Hz for $M = 1$ to 4 at 2 m/s), because the $10\log_{10}M$ dB power split offsets the shorter distances. Consequently, the advantage of PRISM over tracking is largest for $M = 1$ and vanishes for $M \geq 3$ in this open hall, where both schemes coincide within the confidence intervals. The 5th-percentile rate in Fig. 4(b) shows the fairness benefit of a second PA. In an open hall, coverage placement and additional PAs are thus two alternative remedies for the same travel-time bottleneck; Section 6.5 shows that they become complementary when obstacles separate the users.

6.4 Effect of Hall Length

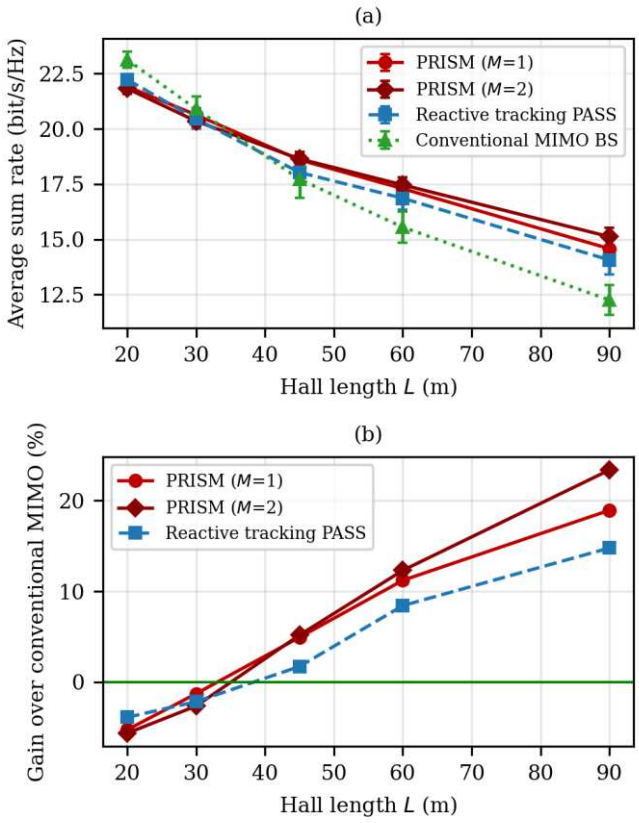


Fig. 5. (a) Average sum rate versus hall length L (95% confidence intervals) and (b) relative gain over the conventional MIMO BS (6 realizations per point).

Fig. 5 maps where PASS is worth deploying. For L = 20 m, the co-located MIMO BS achieves 23.13 bit/s/Hz versus 21.92 for PRISM, because every user is close to the array and its array gain dominates. PRISM overtakes MIMO from L = 45 m onward, and its advantage reaches 18.9% (M = 1) and 23.4% (M = 2) at L = 90 m. PASS degrades much more slowly with L than MIMO because PAs follow the users along the hall, and the benefit of a second PA grows with L as the travel distance increases.

6.5 Warehouse-Aisle Scenario

Table 4. Warehouse-aisle scenario (2.8-m shelving, lane-following AGVs), 8 realizations (rates in bit/s/Hz).

Scheme	Sum rate ± 95% CI	5th-pct. rate	PRISM (M = 3) − scheme (paired)
Oracle-PRISM (M = 1)	14.09 ± 0.77	0.51	+2.25 ± 0.92 (8/8)
Oracle, true blockage only	14.02 ± 0.93	0.49	+2.32 ± 0.92 (8/8)

Scheme	Sum rate ± 95% CI	5th-pct. rate	PRISM ($M = 3$) − scheme (paired)
Oracle, true positions only	9.91 ± 1.27	0.42	+6.43 ± 1.63 (8/8)
PRISM (M = 3)	16.34 ± 0.98	0.74	—
PRISM ($M = 3$) w/o WA-ORBM	16.07 ± 1.02	0.69	+0.27 ± 0.33 (6/8)
Reactive tracking ($M = 3$)	14.02 ± 0.91	0.70	+2.31 ± 0.35 (8/8)
Fixed PASS ($M = 3$)	12.11 ± 0.51	0.62	+4.23 ± 0.68 (8/8)
PRISM (M = 1)	9.32 ± 0.70	0.39	+7.02 ± 1.07 (8/8)
PRISM ($M = 1$) w/o WA-ORBM	9.37 ± 0.90	0.38	+6.97 ± 1.40 (8/8)
PRISM ($M = 1$) w/o prediction	10.95 ± 0.85	0.46	+5.39 ± 1.19 (8/8)
Reactive tracking ($M = 1$)	8.60 ± 0.94	0.40	+7.74 ± 1.27 (8/8)
DDPG PASS	7.77 ± 0.61	0.33	+8.57 ± 1.08 (8/8)
Conventional MIMO BS	3.34 ± 0.55	0.01	+13.00 ± 1.86 (8/8)

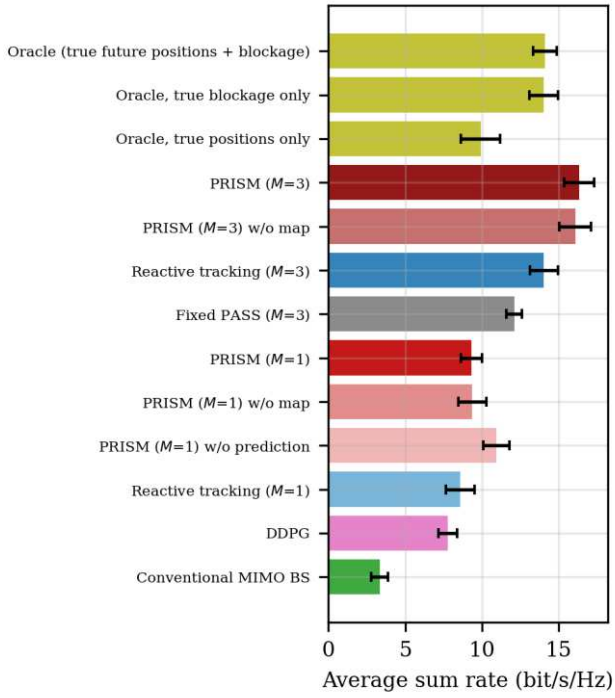


Fig. 6. Warehouse-aisle scenario: average sum rate with 95% confidence intervals, including the oracle decomposition.

Table 4 and Fig. 6 report the warehouse. Here multiple PAs are decisive: PRISM with $M = 3$ achieves 16.34 ± 0.98 bit/s/Hz, 75.4% above PRISM with $M = 1$ and 16.5% above reactive tracking with the same three PAs per waveguide, while the conventional MIMO BS collapses to 3.34 bit/s/Hz because the shelving blocks most of its links. With a single PA, each waveguide must serve several AGVs that are spread along a 52-m aisle, and the PA cannot reach them in time; with three segment-confined PAs, every AGV is always within one segment of a PA in its own aisle. The clairvoyant oracle explains where the remaining headroom of single-PA operation lies. The oracle with true blockage only reaches 14.02 bit/s/Hz, close to the full oracle (14.09), whereas the oracle with true positions only reaches 9.91. The headroom therefore stems from blockage knowledge, not trajectory prediction. Our WA-ORBM does not yet capture it (PRISM $M = 1$ with and without the map: 9.32 vs. 9.37 bit/s/Hz), because the map learns only along the PA trajectories actually visited, whereas the oracle evaluates every candidate position. In addition, Kalman prediction is counterproductive in the warehouse: removing it changes the single-PA sum rate by $+1.63 \pm 1.04$ bit/s/Hz (7/8, significant). A constant-velocity model extrapolates AGVs beyond the aisle ends, where they actually turn, so mobility models that encode the aisle layout are required before prediction can help. Exploiting geometric knowledge of the shelving, which is static and could be surveyed once, is therefore the most promising extension for single-PA deployments, while multiple PAs already recover most of the gap without any such knowledge.

7 Limitations and Future Work

This study has several limitations. First, the results come from a system-level simulator with geometric blockage, Rician fading, and perfect CSI at the current PA positions; hardware effects such as actuator acceleration, pinch-coupling losses, and phase errors during motion [27] are not modeled. Second, the DP optimum covers only a single waveguide with max-rate scheduling; the optimum of the full multiuser PF problem (P1) remains open. Third, the learned WA-ORBM does not capture the blockage headroom revealed by the oracle in the warehouse; surveyed or geometry-aware maps are a natural next step. Fourth, the DRL baselines were trained on a single CPU with a modest budget, and stronger agents may perform better; our released code makes such comparisons straightforward. Finally, experimental validation on a PASS prototype is needed before deployment claims can be made.

8 Conclusion

This paper studied multiuser, multi-waveguide pinching-antenna systems with finite antenna speed, mobile users, dynamic blockage, and possibly several PAs per waveguide. We showed that one-to-one tracking is limited by a re-association travel time of $L/(3v_{\max})$, which segment-confined PAs reduce by a

factor of M, and proposed PRISM, a training-free algorithm that places all PAs to maximize PF-weighted coverage over a prediction horizon under the speed constraint. An exact DP optimum showed that predictive placement attains 93.2% of the best speed-limited rate in a tractable case. In a 60-m hall, PRISM achieved 17.49 bit/s/Hz, 4.2% above reactive tracking and 59.4% above TD3, and two PAs per waveguide raised the 5th-percentile rate by 10.7%. In a warehouse with tall shelving, three PAs per waveguide increased the sum rate by 75.4% over a single PA. Deployment guidance follows directly: use PASS in long halls; in open halls, use either coverage placement or additional PAs to overcome the travel-time bottleneck; where aisles separate the users, combine both; and invest in geometric blockage knowledge where obstacles are nearly as tall as the waveguides.

Statements and Declarations

Funding: The authors did not receive support from any organization for the submitted work.

Competing interests: The authors declare no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Ethics approval: Not applicable. This study is based entirely on computer simulations and involves no human participants, human data, or animals.

Consent to participate and for publication: Not applicable.

Data availability: All data generated in this study (raw simulation results for every table and figure) are included in the public repository listed below.

Code availability: The complete simulator, the trained DDPG and TD3 models, and the scripts that reproduce every figure and table are publicly available at https://github.com/iitrky/pass-prism under the MIT license.

Author contributions: G. Arul Dalton conceived the study, developed the system model and the PRISM algorithm, and carried out the simulations and analysis. A. Bamila Virgin Louis contributed to the analysis and wrote and revised the manuscript. Both authors read and approved the final manuscript.

References

[1]W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," *IEEE Netw.*, vol. 34, no. 3, pp. 134–142, May/Jun. 2020.

[2]ITU-R, "Framework and overall objectives of the future development of IMT for 2030 and beyond," Recommendation ITU-R M.2160-0, Nov. 2023.

[3] T. S. Rappaport *et al.*, "Wireless communications and applications above 100 GHz: Opportunities and challenges for 6G and beyond," *IEEE Access*, vol. 7, pp. 78729–78757, 2019.

[4]T. S. Rappaport *et al.*, "Millimeter wave mobile communications for 5G cellular: It will work!" *IEEE Access*, vol. 1, pp. 335–349, 2013.

[5] T. Bai, R. Vaze, and R. W. Heath, "Analysis of blockage effects on urban cellular networks," *IEEE Trans. Wireless Commun.*, vol. 13, no. 9, pp. 5070–5083, Sep. 2014.

[6] G. R. MacCartney, T. S. Rappaport, and S. Rangan, "Rapid fading due to human blockage in pedestrian crowds at 5G millimeter-wave frequencies," in *Proc. IEEE GLOBECOM*, 2017, pp. 1–7.

[7] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz," TR 38.901, v17.0.0, 2022.

[8] K. Mahendran, H. Sudarsan, S. Rathika, and B. Shankarlal, "Deep adaptive learning-based beam combining framework for 5G millimeter-wave massive 3D-MIMO uplink systems," *Trans. Emerg. Telecommun. Technol.*, vol. 35, no. 12, Art. no. e70030, 2024, doi: 10.1002/ett.70030.

[9] G. Eappen, J. Cosmas, T. Shankar, A. Rajesh, R. Nilavalan, and J. Thomas, "Deep learning integrated reinforcement learning for adaptive beamforming in B5G networks," *IET Commun.*, vol. 16, no. 20, pp. 2454–2466, 2022, doi: 10.1049/cmu2.12501.

[10] K.-K. Wong, A. Shojaeifard, K.-F. Tong, and Y. Zhang, "Fluid antenna systems," *IEEE Trans. Wireless Commun.*, vol. 20, no. 3, pp. 1950–1962, Mar. 2021.

[11] L. Zhu, W. Ma, and R. Zhang, "Modeling and performance analysis for movable antenna enabled wireless communications," *IEEE Trans. Wireless Commun.*, vol. 23, no. 6, pp. 6234–6250, Jun. 2024.

[12]L. Zhu, W. Ma, and R. Zhang, "Movable antennas for wireless communication: Opportunities and challenges," *IEEE Commun. Mag.*, vol. 62, no. 6, pp. 114–120, Jun. 2024.

[13] A. Fukuda, H. Yamamoto, H. Okazaki, Y. Suzuki, and K. Kawai, "Pinching antenna: Using a dielectric waveguide as an antenna," *NTT DOCOMO Tech. J.*, vol. 23, no. 3, pp. 5–12, Jan. 2022.

[14] Z. Ding, R. Schober, and H. V. Poor, "Flexible-antenna systems: A pinching-antenna perspective," *IEEE Trans. Commun.*, vol. 73, no. 10, pp. 9236–9253, Oct. 2025.

[15] Y. Liu *et al.*, "Pinching-antenna systems (PASS): A tutorial," *IEEE Trans. Commun.*, vol. 74, pp. 4881–4918, 2026.

[16] Y. Liu, Z. Wang, X. Mu, C. Ouyang, X. Xu, and Z. Ding, "Pinching-antenna systems: Architecture designs, opportunities, and outlook," *IEEE Commun. Mag.*, vol. 64, no. 1, pp. 190–196, Jan. 2026.

[17] Z. Yang *et al.*, "Pinching antennas: Principles, applications and challenges," arXiv:2501.10753, 2025.

[18] C. Ouyang, Z. Wang, Y. Liu, and Z. Ding, "Array gain for pinching-antenna systems (PASS)," *IEEE Commun. Lett.*, vol. 29, no. 6, pp. 1471–1475, Jun. 2025.

[19] Y. Xu, Z. Ding, and G. K. Karagiannidis, "Rate maximization for downlink pinching-antenna systems," *IEEE Wireless Commun. Lett.*, vol. 14, no. 5, pp. 1431–1435, May 2025.

[20] S. A. Tegos, V. K. Papanikolaou, Z. Ding, and G. K. Karagiannidis, "Minimum data rate maximization for uplink pinching-antenna systems," *IEEE Wireless Commun. Lett.*, vol. 14, no. 5, pp. 1516–1520, May 2025.

[21] K. Wang, Z. Ding, and R. Schober, "Antenna activation for NOMA assisted pinching-antenna systems," *IEEE Wireless Commun. Lett.*, vol. 14, no. 5, pp. 1526–1530, May 2025.

[22] S. Hu, R. Zhao, Y. Liao, D. W. K. Ng, and J. Yuan, "Sum-rate maximization for pinching antenna-assisted NOMA systems with multiple dielectric waveguides," arXiv:2503.10060, 2025.

[23] A. Bereyhi, C. Ouyang, S. Asaad, Z. Ding, and H. V. Poor, "MIMO-PASS: Uplink and downlink transmission via MIMO pinching-antenna systems," *IEEE Trans. Commun.*, vol. 74, pp. 5701–5716, 2026.

[24] K. Wang, C. Ouyang, Y. Liu, and Z. Ding, "Pinching-antenna systems with LoS blockages," *IEEE Wireless Commun. Lett.*, vol. 14, no. 12, pp. 4122–4126, Dec. 2025.

[25] Z. Ding and H. V. Poor, "LoS blockage in pinching-antenna systems: Curse or blessing?" arXiv:2503.08554, 2025.

[26] M. Zeng, X. Wang, Y. Liu, Z. Ding, G. K. Karagiannidis, and H. V. Poor, "Robust resource allocation for pinching-antenna systems under user location uncertainty," *IEEE Trans. Veh. Technol.*, early access, 2026, doi: 10.1109/TVT.2026.3651617.

[27] K. Wang, D. K. C. So, and Z. Ding, "Pinching-antenna systems: From antenna placement to antenna roaming," arXiv:2608.10136, 2026.

[28] C. Bettstetter, H. Hartenstein, and X. Pérez-Costa, "Stochastic properties of the random waypoint mobility model," *Wireless Netw.*, vol. 10, no. 5, pp. 555–567, 2004.

[29] X. Xie, F. Fang, Z. Ding, and X. Wang, "A low-complexity placement design of pinching-antenna systems," *IEEE Commun. Lett.*, vol. 29, no. 8, pp. 1784–1788, Aug. 2025.

[30] M. Zeng *et al.*, "Resource allocation for pinching-antenna systems: State-of-the-art, key techniques and open issues," arXiv:2506.06156, 2025.

[31] A. Amhaz, M. Elhattab, C. Assi, and S. Sharafeddine, "Joint beamforming and antenna placement optimization in pinching antenna systems with user mobility: A deep reinforcement learning approach," arXiv:2605.08039, 2026.

[32] Yüksel Tokur Bozkurt, "Trajectory-aware antenna reconfiguration for pinching antenna systems: Performance analysis under user mobility," *Digital Signal Process.*, 2026, PII: S1051200426002812.

[33] Kaidi Wang, Daniel K C So, Zhiguo Ding"Age-of-information aware federated learning with finite speed pinching antenna," arXiv:2607.23595, 2026.

[34] Fang Fang, Zhiguo Ding, Victor C. M. Leung, Lajos Hanzo, "Pinching antennas meet AI in next-generation wireless networks," arXiv:2511.07442, 2025.

[35] K. Zhou *et al.*, "A gradient meta-learning joint optimization for beamforming and antenna position in pinching-antenna systems," *IEEE Trans. Commun.*, vol. 74, pp. 1099–1112, Feb. 2026.

[36] Saeid Pakravan, Mohsen Ahmadzadeh, Ming Zeng, Xingwang Li, Fang Fang"AI-empowered resource allocation for wirelessly powered pinching-antenna systems," arXiv:2604.11979, 2026.

[37] G. Arul Dalton, A. B. V. Louis, A. Ramachandran, and J. Savija, "Modified deep learning model in proactive decision-making for handover management in 5G," *Concurrency Comput.: Pract. Exper.*, vol. 37, Art. no. e70023, 2025, doi: 10.1002/cpe.70023.

[38] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2016.

[42]S. Fujimoto, H. van Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2018, pp. 1587–1596.

[40] Y. Zeng and X. Xu, "Toward environment-aware 6G communications via channel knowledge map," *IEEE Wireless Commun.*, vol. 28, no. 3, pp. 84–91, Jun. 2021.

[41] Y. Zeng *et al.*, "A tutorial on environment-aware communications via channel knowledge map for 6G," *IEEE Commun. Surveys Tuts.*, vol. 26, no. 3, pp. 1478–1519, 2024.

[42] P. Viswanath, D. N. C. Tse, and R. Laroia, "Opportunistic beamforming using dumb antennas," *IEEE Trans. Inf. Theory*, vol. 48, no. 6, pp. 1277–1294, Jun. 2002.

[43] T. Yoo and A. Goldsmith, "On the optimality of multiantenna broadcast scheduling using zero-forcing beamforming," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 3, pp. 528–541, Mar. 2006.

[44] R. E. Kalman, "A new approach to linear filtering and prediction problems," *J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.

[45] D. P. Bertsekas, *Nonlinear Programming*, 3rd ed. Belmont, MA, USA: Athena Scientific, 2016.

[46] H. W. Kuhn, "The Hungarian method for the assignment problem," *Naval Res. Logist. Quart.*, vol. 2, no. 1–2, pp. 83–97, 1955.