Channel Disparity Engineering for NOMA Leveraging Pinching Antenna Systems

Derek Chui†, Collin Fiske†, Wei Wang‡ and Krishna Ramamoorthy†

†*Department of Computer Science and Engineering, Santa Clara University,* California, USA.

‡*Department of Computer Science, San Diego State University,* California, USA.

dchui2@scu.edu, cfiske@scu.edu, wwang@sdsu.edu, kkattiyanramamoorthy@scu.edu

***Abstract*—Successive interference cancellation in non-orthogonal multiple access (NOMA) relies on channel disparity between paired users, which in fixed-antenna deployments is determined by user geometry and can only be exploited as found. Pinching antenna systems (PASS) create radiating points at selected locations along a waveguide, turning that disparity into a quantity the network can engineer. In turn this paper studies the resource allocation problem. We formulate the joint selection of user pairing, pinching antenna (PA) assignment, and intra-group power allocation, and establish two structural results. The power split separates across groups for any fixed pairing and assignment, so it can be folded into the reward of a learning agent without loss of optimality. Pairing and assignment, in contrast, do not separate, because the effective channel and hence the weak-strong ordering within a pair are themselves functions of the assignment. This non-separability is a consequence of the reconfigurability that PASS introduces and does not arise under fixed antennas. Guided by these results we develop a Q-learning agent that co-optimizes PA assignment and power allocation, priced against the exactly enumerated joint optimum, and we evaluate the channel disparity it engineers rather than throughput alone. Across free-space and Sionna CDL-C NLoS channels, the proposed design attains higher disparity than proximity-based, random, and fixed-antenna assignment at equal radiated power, and converts it into consistently higher semantic utility while attaining 98.9% of the utility of exhaustive search.**

***Index Terms*—Next-Generation Multiple Access (NGMA), Non-Orthogonal Multiple Access (NOMA), Pinching Antenna Systems (PASS), Channel Disparity, Reinforcement Learning, User Grouping, Power Allocation.**

I. Introduction

Next-generation multiple access (NGMA) paradigms aim to move beyond conventional orthogonal designs toward more flexible, scalable, and intelligent resource sharing mechanisms in emerging 6G networks. Non-Orthogonal Multiple Access (NOMA) is a key enabler in this direction, allowing multiple users to share the same spectral resource through Successive Interference Cancellation (SIC) at the receiver [1], [2]. By pairing users with different channel gains, NOMA achieves higher spectral efficiency than orthogonal schemes [3], [10]. In parallel, Pinching Antenna Systems (PASS) address the need for spatial adaptability by attaching small, low-cost dielectric clips along a waveguide, which effectively create controllable radiating points wherever pinched [5], [6], [9]. This enables near-optimal proximity to served users, fundamentally reducing path loss compared to fixed deployments [4].

The integration of NOMA and PASS provides a promising framework for NGMA systems, jointly enabling efficient multi-user access and spatially adaptive transmissions, as illustrated in Fig. 1. In particular, PASS can intentionally shape channel disparities through pinch placement, which is favorable for NOMA operation. However, jointly optimizing which users to pair, where to activate pinches, and how to allocate power among NOMA users remains a combinatorially hard problem [8].

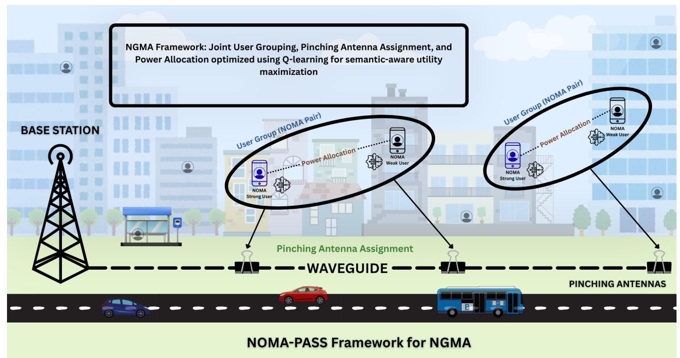

Fig. 1: Illustration of an NGMA system integrating NOMA and PASS

Beyond spectral and spatial efficiency, NGMA systems are expected to incorporate application-layer semantics into the communication process. Optimizing throughput alone is insufficient for next-generation applications, which impose diverse fidelity and quality requirements. Semantic communications [11] address this by embedding application-level importance and fidelity directly into the utility metric, rather than treating all bits equally.

To address the resulting high-dimensional and non-convex optimization problem, learning-based approaches provide a scalable alternative to conventional methods. In particular, reinforcement learning (RL), and Q-learning in particular, enables efficient exploration of large decision spaces without requiring explicit modeling of system dynamics [7]. By framing user grouping, pinching antenna assignment, and power allocation as a sequential decision-making process, the learning agent can capture the strong coupling between these decisions and learn near-optimal policies without resorting to exhaustive search.

Several recent works have investigated the integration of NOMA and PASS, each focusing on specific aspects of the resource allocation problem. Prior work in [1] proposes a matching-based antenna activation algorithm to maximize sum rate. The study in [2] addresses QoS-aware design to satisfy per-user rate requirements under PASS deployments. Joint optimization of antenna placement and power allocation is explored in [3], while total transmit power minimization across multiple waveguides using iterative methods is considered in [10]. Joint user grouping and power allocation has been studied in NOMA mmWave systems [8] and reinforcement learning has shown promise for NOMA resource management in vehicular networks [7]. However, existing PASS-NOMA works generally decouple user grouping from antenna assignment, rely on conventional iterative methods, and optimize throughput metrics alone. None jointly address user grouping, antenna assignment, and power allocation in NOMA-enabled PASS within a learning-based and semantic-aware framework.

A key structural observation motivates this work. In conventional NOMA, the channel disparity that SIC depends on is fixed by user geometry and can only be exploited as found. Under PASS, pinch placement makes that disparity a design variable. User grouping and antenna assignment therefore stop being separable subproblems, and the semantic weights decide which disparities are worth creating in the first place. It is this coupling rather than the co-existence of PASS, NOMA, and learning, that the present work addresses.

Building on this observation, we formulate the coupled problem over user pairing, PA assignment, and intra-group power allocation under a priority-weighted utility with practical SIC constraints, and we characterize how that problem decomposes, leading to two structural results. The power split separates across groups once pairing and assignment are fixed, which lets a bounded α-scan sit inside the reward of a learning agent at no cost in optimality. Pairing and assignment do not separate, because the weak-strong labeling within a pair is itself decided by the assignment, so any pairing rule evaluated from a static channel statistic optimizes a surrogate. We then develop a Q-learning agent that co-optimizes assignment and power, evaluate the channel disparity it engineers against proximity-based, random, and fixed-antenna assignment, and report the cost that resolving pairing separately still incurs.

II. System Model and Problem Formulation

This section presents the system model and semantic-aware problem formulation of the proposed NOMA-PASS-based NGMA framework.

A. NOMA-PASS System Model

We consider a downlink PASS system, where a base station (BS) feeds one or multiple dielectric waveguides and creates radiating points by activating pinching antennas (PA) at selected locations along the waveguides [5], [6], [9]. Let $\mathcal{K} = \{1, \ldots, K\}$ denote the set of users and $\mathcal{M} = \{1, \ldots, M\}$ denote the set of pinch antenna locations. The location of user k is denoted by ψ_k and that of PA m by ψ_m^{Pin}.

$$d_{k,m} \triangleq \left\| \boldsymbol{\psi}_k - \boldsymbol{\psi}_m^{\text{Pin}} \right\| \tag{1}$$

The channel from PA m to user k is modeled based on [9], as

$$h_{k,m} = \frac{\sqrt{\eta}}{d_{k,m}} \exp\left(-j\frac{2\pi}{\lambda} d_{k,m}\right) \exp(-j\theta_m) \tag{2}$$

where $\eta = \left(\frac{\lambda}{4\pi}\right)^2$ is the free-space propagation constant, λ is the wavelength, and θ_m captures a PA-dependent phase offset due to waveguide and activation effects.

B. Joint User Grouping, PA Assignment, and Power Allocation

In the single waveguide case, we adopt a standard NOMA power allocation abstraction, where power is divided evenly across candidate PAs. This yields the per-PA power budget for all $m \in \mathcal{M}$.

$$P_m = \frac{P_{\text{tx}}}{M} \tag{3}$$

where P_{tx} denotes the total BS transmit power. For multiple waveguides, the total transmit power is first divided equally across W waveguides, i.e., each waveguide is allocated P_{tx}/W. The allocated power is then divided equally among the active PAs on each waveguide.

We consider downlink NOMA transmission with two-user groups under SIC. Let $P = K/2$ denote the number of groups (assuming K is even), and let $\mathcal{P} = \{1, \ldots, P\}$ index these groups. A pairing is represented by a mapping $\pi : \mathcal{K} \to \mathcal{P}$ such that each group contains exactly two users. Each pinching antenna is assigned either to one group or left inactive. Let $a_m \in \{0\} \cup \mathcal{P}$ denote the assignment label, where $a_m = 0$ indicates that PA m is inactive and $a_m = p$ indicates that PA m serves group p. The set of PAs assigned to group p is then given by

$$\mathcal{M}_p \triangleq \{m \in \mathcal{M} : a_m = p\} \tag{4}$$

Given $\mathcal{M}_p$, the effective channel of user k in group p is defined as

$$\tilde{h}_{k,p} \triangleq \sum_{m \in \mathcal{M}_p} h_{k,m} \sqrt{P_m} \tag{5}$$

which aggregates the contributions of all PAs assigned to that group.

C. Downlink NOMA SINR and Rate Analysis with SIC

Based on the effective channel defined in (5), consider a two-user group p containing users $\{u, v\}$. To guarantee successful SIC, the weak user's stream must be decodable at both users. The weak–strong ordering based on effective channel power is defined as

$$w(p) = \arg \min_{k \in \{u,v\}} |\tilde{h}_{k,p}|^2, \quad s(p) = \arg \max_{k \in \{u,v\}} |\tilde{h}_{k,p}|^2 \tag{6}$$

The quantity SIC operates on is the ratio of the two, which we call the channel disparity of group p.

$$\Delta_p \triangleq 10 \log_{10} \left(\frac{|\tilde{h}_{s(p),p}|^2}{|\tilde{h}_{w(p),p}|^2} \right) \tag{7}$$

Under fixed antennas Δ_p is set by user geometry. Under PASS it is a function of $\mathcal{M}_p$, and therefore a design variable rather than an inherited one.

Let $\alpha_{p,w}$ and $\alpha_{p,s}$ denote the intra-group NOMA power fractions assigned to the weak and strong users that satisfy

$$\alpha_{p,w} + \alpha_{p,s} = 1, \quad \alpha_{p,w} \geq \alpha_{p,s} > 0 \tag{8}$$

Evaluating (5) at user k with another group's antenna set gives the inter-group leakage that user sees,

$$I_k \triangleq \sum_{q \neq \pi(k)} \left| \tilde{h}_{k,q} \right|^2 \tag{9}$$

with noise power σ^2, the SINRs for decoding the weak user's message at the weak and strong users are given by

$$\gamma_{w\to w}^{(p)} = \frac{|\tilde{h}_{w(p),p}|^2 \alpha_{p,w}}{|\tilde{h}_{w(p),p}|^2 \alpha_{p,s} + I_{w(p)} + \sigma^2}, \tag{10}$$

$$\gamma_{w\to s}^{(p)} = \frac{|\tilde{h}_{s(p),p}|^2 \alpha_{p,w}}{|\tilde{h}_{s(p),p}|^2 \alpha_{p,s} + I_{s(p)} + \sigma^2} \tag{11}$$

To ensure decodability under SIC, the weak user's effective SINR is defined as the minimum SINR across both users.

$$\gamma_{p,w} \triangleq \min\{\gamma_{w\to w}^{(p)},\ \gamma_{w\to s}^{(p)}\} \tag{12}$$

The strong user, after canceling the weak user's signal, decodes its own message with SINR as shown below

$$\gamma_{p,s} = \frac{|\tilde{h}_{s(p),p}|^2 \alpha_{p,s}}{I_{s(p)} + \sigma^2} \tag{13}$$

The results reported here solve the $I_k = 0$ instance of this model. That instance is exact when the PA sets serving distinct groups are spatially separated relative to the users they serve, and is an approximation otherwise, since proximity to a serving PA does not by itself imply distance from another group's. The objective in (17) does not penalize configurations that violate this, so we state it as a scope of the present model rather than as a property of PASS. Enforcing separation acts through the same assignment variable and is a direct extension. The achievable rates (in bits/s/Hz) are given by

$$R_{p,w} = \log_2(1 + \gamma_{p,w}), \quad R_{p,s} = \log_2(1 + \gamma_{p,s}) \tag{14}$$

D. Priority-Weighted Semantic Utility

Building on the achievable rates defined in (14), each user k is assigned a semantic importance weight $w_k \geq 0$. We model the semantic fidelity $\xi(\gamma) \in [a_1, a_2]$ as a logistic function of the SINR γ:

$$\xi(\gamma) = a_1 + \frac{a_2 - a_1}{1 + \exp(-(c_1\gamma + c_2))} \tag{15}$$

where (a_1, a_2, c_1, c_2) are shaping parameters that control the fidelity response to SINR. The weights w_k carry the application-level priority and are the component that materially changes which allocation is selected, while $\xi(\gamma)$ contributes a saturating response that leaves the ordering of candidate allocations largely intact. Section IV confirms this, showing the selected allocations insensitive to c_1 across nearly two orders of magnitude. For group p, the corresponding semantic fidelities are $\xi_{p,w} = \xi(\gamma_{p,w})$ and $\xi_{p,s} = \xi(\gamma_{p,s})$. The semantic-aware contribution of each user is then given by

$$G_{p,w} = w_{w(p)}\, \xi_{p,w}\, R_{p,w}, \quad G_{p,s} = w_{s(p)}\, \xi_{p,s}\, R_{p,s} \tag{16}$$

E. Problem Formulation

We formulate a joint optimization problem that captures user grouping, PA assignment, and NOMA power allocation. Specifically, our goal is to jointly optimize the user pairing π, PA assignment $\mathbf{a} = [a_1, \ldots, a_M]$, and NOMA power splits $\boldsymbol{\alpha} = \{\alpha_{p,w}, \alpha_{p,s}\}_{p=1}^{P}$ to maximize the total semantic utility:

$$\begin{aligned} \max_{\pi,\mathbf{a},\boldsymbol{\alpha}} \quad & \sum_{p=1}^{P} (G_{p,w} + G_{p,s}) && (17) \\ \text{s.t.} \quad & a_m \in \{0\} \cup \mathcal{P}, \quad \forall m \in \mathcal{M}, \quad \alpha_{p,w} + \alpha_{p,s} = 1, \\ & \alpha_{p,w} \geq \alpha_{p,s} > 0, \quad \forall p, \quad \sum_{m=1}^{M} P_m \leq P_{\text{tx}}, \\ & \Delta_p \geq \Delta_{\min}, \quad \forall p \end{aligned}$$

The last constraint sets a minimum SIC margin through (7). It is a design knob rather than a modeling requirement, and Section IV reports the unconstrained instance $\Delta_{\min} = 0$, so the disparity observed there is induced by (17) rather than imposed on it. The resulting problem is combinatorial and non-convex, and it decomposes in one direction but not the other.

For a fixed pairing and assignment it does separate. Group p enters (17) only through $G_{p,w} + G_{p,s}$, which by (5)–(16) depends on $\mathcal{M}_p$, the two users of p, and $\alpha_{p,w}$ alone, while the remaining constraints are either fixed by $\mathbf{a}$ or act within a group. A sum of terms in disjoint free variables is maximized termwise, so

$$\max_{\pi,\mathbf{a},\boldsymbol{\alpha}} \sum_{p=1}^{P} (G_{p,w} + G_{p,s}) = \max_{\pi,\mathbf{a}} \sum_{p=1}^{P} \max_{\alpha_{p,w}} (G_{p,w} + G_{p,s})\,. \tag{18}$$

The power split can therefore be removed from the search and recovered exactly once pairing and assignment are fixed, at no cost in optimality.

Pairing and assignment do not separate, and this governs the rest of the paper because the weak-strong labeling is itself a function of the assignment. Place PAs at $\pm x_0$ along the waveguide and users u and v at $\mp x_0$, so that $d_{u,1} < d_{u,2}$ and $d_{v,2} < d_{v,1}$. By (2) and (5), assigning only the first PA gives $|\tilde{h}_{u,p}|^2 > |\tilde{h}_{v,p}|^2$ and assigning only the second reverses it, so u is the strong user under one assignment and the weak user under the other. Since (16) applies $w_{w(p)}$ and $w_{s(p)}$ by that labeling, a candidate pair has no value independent of the assignment, and any pairing rule evaluated beforehand optimizes a surrogate.

This is a property of PASS-NOMA rather than of either component, since under fixed antennas the effective channel does not depend on the assignment and the ordering is immutable. It is the reconfigurability introduced by PASS that removes this separability. The two observations delimit the solver. The power split may leave the search without loss, while pairing and assignment must be searched against one another, which shapes the approach of Section III. Section IV quantifies what is given up when pairing is resolved separately.

III. Learning-Based Joint Optimization

The structure of the problem determines the shape of the solver. Since the power split separates once pairing and assignment are fixed, $\boldsymbol{\alpha}$ is removed from the search and optimized inside the reward for every candidate assignment, which by (18) is exact rather than approximate. Assignment and power allocation are therefore co-optimized. The pairing, by contrast, cannot be separated without loss, yet a flat action space over $(\pi, \mathbf{a}, \boldsymbol{\alpha})$ has $(K-1)!!\,(P+1)^M|\mathcal{A}|^P$ configurations, which no tabular representation covers at the deployment sizes of interest. We therefore resolve π as an outer stage and treat this as a bounded approximation rather than a free decomposition. At the sizes evaluated here the joint problem is small enough to enumerate, since $(K-1)!! = 3$ for $K = 4$, so we solve it exactly and use that optimum as the reference against which the staged scheme is priced in Section IV. What staging buys is a search that survives growth in K, and what is costs is reported rather than absorbed. Algorithm 1 summarizes the procedure.

A. MDP Formulation: State, Action, and Reward Design

We formulate the joint optimization problem as a Markov Decision Process (MDP), where the BS acts as an agent that sequentially assigns PAs to user groups.

At step $t \in \{1, \ldots, M\}$ the state captures the current PA index together with the partial assignment made so far,

$$s_t = (t, a_1, \ldots, a_{t-1}) \tag{19}$$

and the agent responds by selecting the assignment label for PA t,

$$a_t \in \{0\} \cup \mathcal{P} \tag{20}$$

where $a_t = 0$ leaves the PA inactive and $a_t = p$ assigns it to group p.

The state carries no channel information by construction. A Q-table is instantiated per channel realization and the channel enters only through the terminal reward, so the agent searches the assignment tree of the instance it is given rather than learning a polciy amortized across instances. Encoding the channel in the state is required only for the amortized case, which is the parameterized value function discussed below.

Directly including $\boldsymbol{\alpha}$ in the action space would cause a combinatorial explosion, so power allocation is folded into the reward instead. For any candidate grouping or assignment we select the best NOMA power split by scanning the discrete set

$$\mathcal{A} = \{0.50, 0.55, \ldots, 0.95\} \tag{21}$$

with $\alpha_{p,w} \in \mathcal{A}$ and $\alpha_{p,s} = 1 - \alpha_{p,w}$. For a fixed pair p and channels $(\tilde{h}_{w(p),p}, \tilde{h}_{s(p),p})$ we evaluate (14)–(16) at each $\alpha_{p,w} \in \mathcal{A}$ and keep the value maximizing $G_{p,w} + G_{p,s}$, which yields the best achievable semantic utility for that grouping without expanding the action space. To stabilize learning and avoid repeated expensive evaluations we adopt the sparse terminal reward

$$r_t = \begin{cases} 0 & t < M \\ \sum_{p=1}^{P} \max_{\alpha_{p,w} \in \mathcal{A}} \left(G_{p,w} + G_{p,s}\right) & t = M \end{cases} \tag{22}$$

where $G_{p,w}$ and $G_{p,s}$ are computed from the finalized assignment.

Algorithm 1 Joint Optimization for PASS-NOMA

1: **Input:** User locations $\{\boldsymbol{\psi}_k\}$, PA locations $\{\boldsymbol{\psi}_m^{\mathrm{Pin}}\}$, channel params (η, λ), semantic params $(w_k, a_1, a_2, c_1, c_2)$, $\mathcal{A}$.
2: Compute per-PA channels $h_{k,m}$ via (2).
3: **Pair users:** obtain π via heuristic or pairing Q-learning.
4: Initialize Q-table for PA assignment.
5: **for** episode $= 1$ to E **do**
6:     Sequentially assign each PA m a label $a_m \in \{0\} \cup \mathcal{P}$ using ϵ-greedy.
7:     After assigning all M PAs, compute the terminal reward using (22) (including α-scan).
8:     Update Q-values with tabular Q-learning.
9: **end for**
10: Output the better of the greedy policy and the best assignment encountered, its per-pair α, and total utility.

B. Q-Learning-Based Optimization

We model the BS as an RL agent that sequentially assigns each PA m to either a group label $p \in \mathcal{P}$ or remains inactive (0). We adopt per-episode tabular Q-learning with ϵ-greedy exploration, which is well-suited for PASS layouts with a moderate number of PAs. The Q-values are updated using the Bellman equation

$$Q(s_t, a_t) \leftarrow (1-\beta) Q(s_t, a_t) + \beta \left(r_t + \delta \max_{a'} Q(s_{t+1}, a') \right) \tag{23}$$

with learning rate β and discount factor δ. We decay ϵ from ϵ_{start} to ϵ_{end}.

The outer pairing stage sorts users by $|\sum_m h_{k,m}|^2$ and pairs the weakest with the strongest, which is precisely the surrogate the formulation predicts will be lossy.

The assignment MDP has $\left((P+1)^M - 1\right)/P$ distinct states and each terminal evaluation costs $O(P|\mathcal{A}|)$, so the table is exponential in the number of pinching antennas while the α-scan is independent of M. This is acceptable where M is set by the physical waveguide deployment and stays moderate, but it is not a general-purpose solution. For dense deployments the tabular representation must give way to a parameterized value function, which leaves both the reward-embedded α-scan and the staged decomposition intact while removing the dependence on enumerating states.

C. Convergence Behavior

The sparse terminal reward in (22) provides sufficient learning signal. Best observed utility increases monotonically with experience and per-episode rewards become less volatile as ϵ decays, with the returned assignment settling well within the episode budget used throughout. Its proximity to exhaustive search is reported in Section IV.

IV. Simulation Results and Discussion

We evaluate the framework first in MATLAB under controlled system parameters, then in NVIDIA Sionna under standardized CDL channel models. Unless stated otherwise, $K = 4$, $M = 4$, $P_{\text{tx}} = 10$ W, $\sigma^2 = 10^{-5}$ W, users are drawn uniformly over a 200×200 m area, the disparity evaluation places the M PAs on a circle of radius 70 m at height 3 m, and results are averaged over 200 trials. Q-learning runs 250 episodes per realization with $\beta = 0.35$, $\delta = 0.92$, and ϵ decayed from 0.90 to 0.05. The shaping parameters are $(a_1, a_2, c_1, c_2) = (0.30, 0.98, 0.25, -0.8)$, which put the floor and ceiling of ξ at 0.30 and 0.98 and its knee at $\gamma \approx 3.2$, near the operating SINR of the pairs. The weights w_k are drawn i.i.d. uniform on $[0, 1]$ and redrawn each trial, so no application profile is assumed and the reported gains average over priority profiles rather than being tuned to one.

A. MATLAB-Based Evaluation

We consider three PASS deployment configurations. The circular layout distributes PAs along the perimeter so most users have a nearby antenna, which stabilizes the effective channel in (5). The multiple waveguide configuration keeps $d_{k,m}$ small in more directions and achieves the highest utility, since each group can be assigned its strongest waveguide with minimal contention.

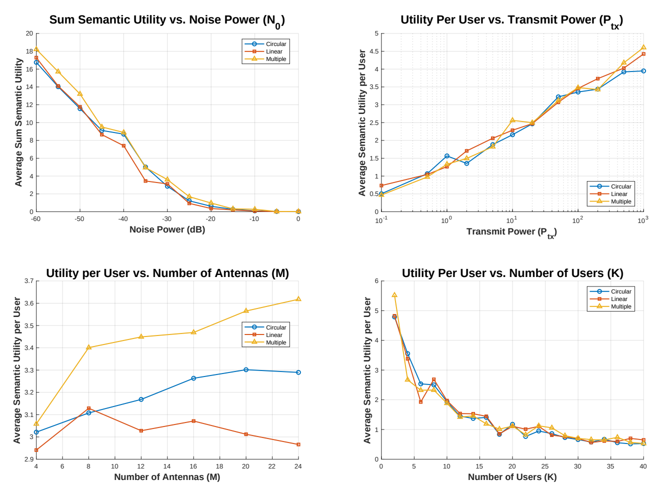

Fig. 2: Impact of key system parameters on average semantic utility per user: (a) noise power, (b) transmit power, (c) number of pinching antennas, and (d) number of users.

Fig. 2 summarizes the impact of key system parameters. Increasing noise power σ^2 in (a) steadily degrades utility since both $\gamma_{p,w}$ and $\gamma_{p,s}$ fall, reducing rates (14) and fidelity (15), with the weak user more sensitive because of intra-group interference. Increasing P_{tx} in (b) improves utility rapidly then saturates as $\xi(\gamma)$ approaches a_2. Power splitting ($P_m = P_{\text{tx}}/M$) competes with improved coverage, so circular and multiple waveguide layouts gain while the linear layout does not. In (d) the average utility per user falls as K grows, since fixed antenna resources are spread across more groups.

B. Engineered Channel Disparity

We measure the channel disparity Δ_p of (7) that each method produces, comparing the proposed agent against exhaustive search, assignment of each PA to the nearest pair centroid, uniformly random labels, and a fixed-antenna reference in which every group is served by all M radiating points with no selection freedom, at equal total radiated power.

Fig. 3(a) shows a median disparity of 8.0 dB for the proposed design, against 5.9 dB for the proximity heuristic, 6.0 dB for the fixed-antenna reference, and 5.2 dB for random assignment. Two features of this order matter. Random assignment produces the least disparity, confirming that disparity comes from choosing which antennas serve which pair rather than from dividing antennas among groups. More directly to the premise, the fixed-antenna reference is indistinguishable from the naive proximity rule, so removing the freedom to choose the assignment costs as much as never using it well. Neither reaches within 2 dB of the proposed design.

Fig. 3(b) translates disparity into performance and the two orderings agree. The proposed policy attains 98.9% of the mean utility of exhaustive search while separating clearly from all three alternatives, so the gain is attributable to the assignment decision rather than to the utility model used to score it.

Against the joint optimum over $(\pi, \mathbf{a}, \boldsymbol{\alpha})$, obtained by enumerating all three pairings each with its own exhaustive assignment, resolving pairing as an outer stage costs 4.97% of semantic utility and is strictly suboptimal in 65.5% of trials. The mechanism is the one identified in the formulation. The assignment inverts the weak-strong ordering the pairing metric assumed in 36.2% of groups. We report this rather than absorb it. Separately, varying c_1 from 0.05 to 4.0 leaves the selected allocations stable, so the conclusions do not depend on the specific values of (a_1, a_2, c_1, c_2).

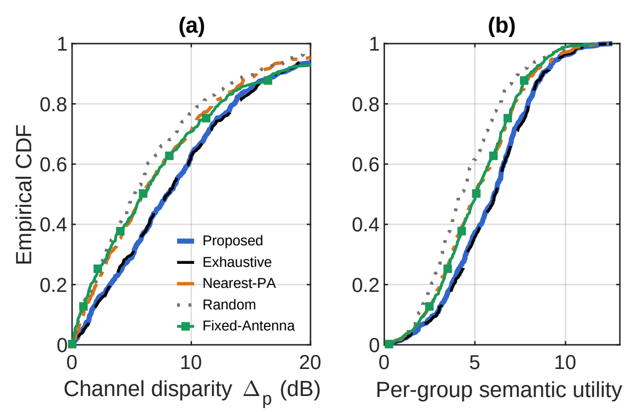

Fig. 3: (a) CDF of engineered channel disparity Δ_p and (b) CDF of per-group semantic utility, comparing the proposed policy with exhaustive search, proximity-based assignment, random assignment, and a fixed-antenna reference at equal radiated power.

C. Validation Under Realistic CDL-C Channels via Sionna

To assess robustness under realistic propagation conditions, we extend the evaluation using the CDL-C NLoS channel model [13] implemented via the Sionna library [12]. Unlike free-space models, CDL-C captures multipath propagation through 24 clusters with frequency-selective fading, representative of urban NLoS environments at 28 GHz with a delay spread of 100 ns. The effective channel is modeled as

$$h_{k,m}^{\mathrm{CDL}} = \frac{\sqrt{\eta}}{d_{k,m}} \sum_{\ell} a_{\ell}^{(k,m)} \tag{24}$$

where $\{a_{\ell}^{(k,m)}\}$ denote the complex path gains of the multipath components from the CDL-C model, normalized to unit total power and scaled by distance-dependent path loss. All other system parameters and the Q-learning policy remain identical to the free-space case to ensure a fair comparison.

CDL-C broadens the distribution of normalized channel power relative to free space and adds heavier tails, so it exercises channel uncertainty that a free-space model cannot represent. Despite this variability, Fig. 4 compares the empirical cumulative distribution functions (CDFs) of total semantic utility across trials for each method and channel model. The results show that the semantic utility distributions under free-space and CDL-C closely track each other, with Q-learning remaining near-optimal in both cases. This behavior stems from two factors. First, NOMA with SIC is inherently designed to exploit channel asymmetry. Second, the agent is re-solved per realization, so the agent tracks the channel it is given. The framework therefore generalizes beyond idealized assumptions to realistic urban NLoS environments.

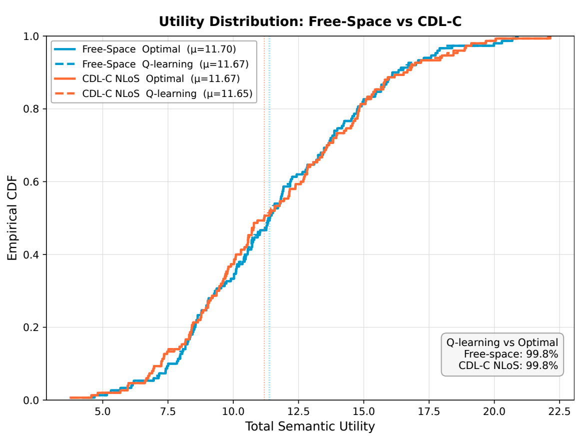

Fig. 4: CDF of total semantic utility under free-space and CDL-C NLoS channel models, comparing Q-learning with baseline methods.

V. Conclusions and Future Work

This paper treated channel disparity as the quantity a PASS-NOMA system designs rather than inherits. We established that the intra-group power split separates from the remaining decisions and can be folded into a learning agent's reward at no cost in optimality, and that user pairing and PA assignment do not separate, because the assignment determines the weak-strong ordering that any pairing rule must presuppose. The latter is a consequence of PASS reconfigurability and has no counterpart in fixed-antenna NOMA, which is why decoupled treatments are natural there and lossy here. Guided by these results, the proposed tabular Q-learning agent co-optimizes assignment and power allocation and engineers measurably greater channel disparity than proximity-based, random, and fixed-antenna assignment at equal radiated power. The result holds under Sionna CDL-C NLoS channels with only marginal degradation from free space.

Several directions remain. The most immediate is to carry the joint treatment of pairing and assignment beyond the sizes at which it can be enumerated, by admitting the pairing into the same episode under a graph-based policy such as a Graph Neural Network (GNN) that captures structured relationships between users, groups, and antennas. Policy gradient methods such as Proximal Policy Optimization (PPO) would extend the framework to continuous power allocation, and relaxing perfect SIC to model imperfect cancellation [2] would better reflect practical receiver limitations.

References

[1] 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.

[2] Y. Xu *et al.*, "QoS-Aware NOMA Design for Downlink Pinching-Antenna Systems," *IEEE Trans. Commun.*, early access, 2025, doi: 10.1109/TCOMM.2025.3605466.

[3] Z. Zhou *et al.*, "Sum-Rate Maximization for NOMA-Assisted Pinching-Antenna Systems," *IEEE Wireless Commun. Lett.*, vol. 14, no. 9, pp. 2728–2732, Sept. 2025.

[4] Z. Yang *et al.*, "Pinching Antennas: Principles, Applications and Challenges," arXiv:2501.10527, 2025.

[5] Y. Liu *et al.*, "Pinching-Antenna Systems (PASS): A Tutorial," *IEEE Trans. Commun.*, early access, 2026, doi: 10.1109/TCOMM.2026.3658289.

[6] Y. Liu *et al.*, "Pinching-Antenna Systems: Architecture Designs, Opportunities, and Outlook," *IEEE Commun. Mag.*, early access, 2025, doi: 10.1109/MCOM.001.2500037.

[7] H. Ding and K.-C. Leung, "Resource Allocation for Low-Latency NOMA-V2X Networks Using Reinforcement Learning," in *Proc. IEEE INFOCOM Wkshps.*, 2021.

[8] Y. Cao *et al.*, "Joint User Grouping and Power Optimization for Secure mmWave-NOMA Systems," *IEEE Trans. Wireless Commun.*, vol. 21, no. 5, May 2022.

[9] Z. Ding, R. Schober, and H. V. Poor, "Flexible-Antenna Systems: A Pinching-Antenna Perspective," *IEEE Trans. Commun.*, vol. 73, no. 10, Oct. 2025.

[10] Y. Fu *et al.*, "Power Minimization for NOMA-assisted Pinching Antenna Systems With Multiple Waveguides," preprint, 2025.

[11] Z. Qin *et al.*, "Semantic Communications: Principles and Challenges," arXiv:2212.00783, 2022.

[12] J. Hoydis *et al.*, "Sionna: An Open-Source Library for Next-Generation Physical Layer Research," arXiv:2203.11854, 2022.

[13] 3GPP, "Study on Channel Model for Frequencies from 0.5 to 100 GHz," TR 38.901, v17.0.0, 2022.