photonics

MDPI

Article

Chance-Constrained Receiver–Scheduler Co-Design via Probabilistic Decodability Graphs for Reliable SIC in Overlapping Multi-Cell NOMA VLC Networks

Tingting Qin and Yang Tu *

School of Big Data and Intelligent Engineering, Chongqing College of International Business and Economics, Chongqing 401520, China; chongqingqtt@ccibe.edu.cn
* Correspondence: ytu@ccibe.edu.cn

Abstract

Overlapping optical cells create geometry-dependent inter-cell interference, while receiver-geometry and channel-estimation errors can reverse the effective non-orthogonal multiple access (NOMA) decoding order and increase successive interference cancelation (SIC) failures. This paper develops a chance-constrained receiver–scheduler co-design framework for a multi-cell NOMA visible-light communication network with an asymmetrically clipped DC-biased optical orthogonal frequency-division multiplexing physical layer. Correlated position, photodetector-orientation, and channel-estimation errors are propagated through nonlinear geometry-based scenarios. For each SIC direction, a joint three-SINR event defines a layer-, resource-, and direction-labeled probabilistic decodability graph. Candidate NOMA and orthogonal modes are screened on optimization scenarios, admitted by independent one-sided confidence bounds, and selected through resource-constrained mixed-integer linear programming. With the matching fixed, hierarchical powers are adapted under empirical conditional-value-at-risk constraints using trust-region sequential quadratic programming. Because candidate-edge certificates need not remain valid after global matching and power redistribution, the frozen complete assignment is independently recertified before held-out testing. Under the specified uncertainty generator, the proposed method maintains selected-pair outage probabilities of approximately 2.7×10^{-3}–3.3×10^{-3} over the half-power-angle sweep, compared with 0.027–0.060 for nominal-CSI allocation. Additional experiments quantify network-wide outage, model misspecification, unbalanced deployments, feasibility, and computational cost. The results support reliable slow-timescale scheduling under the adopted link and uncertainty models, without implying distribution-free, waveform-level, or real-time guarantees.

Keywords: visible light communication; non-orthogonal multiple access; successive interference cancellation; imperfect channel state information; probabilistic decodability graph; chance-constrained resource allocation

Check for updates

Received: 26 July 2026
Revised: 17 August 2026
Accepted: 19 August 2026
Published: 21 August 2026

1. Introduction

Visible light communication (VLC) uses intensity-modulated light-emitting diodes (LEDs) for illumination and downlink transmission and is a key component of networked LiFi systems [1–4]. In dense indoor deployments, adjacent optical cells overlap. A boundary user then receives both its serving signal and resource-dependent interference from neighboring LEDs. Power-domain non-orthogonal multiple access (NOMA) can improve resource reuse, but its gain depends on user grouping, power allocation, decoding direction, and reliable successive interference cancelation (SIC) [5–7]. These decisions are coupled across cells because changing one LED or layer power alters both desired reception and inter-cell interference [8–11].

In practice, the scheduler observes imperfect channel state information (CSI). Position and photodetector (PD) orientation errors affect distance, irradiance and incidence angles, and the field-of-view condition, while estimation residuals perturb the resulting gains. A single geometry error changes every LED link of the same user, so the desired and interfering channels are correlated. Imperfect CSI and residual cancelation degrade NOMA-VLC outage and error performance [12–14]; robust VLC allocation has also been studied under bounded or stochastic uncertainty [15,16]. These methods do not, however, determine whether a candidate pair, its SIC direction, and the complete three-step decoding chain remain reliable after multi-cell interference and global resource decisions change.

Complementary physical-layer studies have examined directly modulated DCO-OFDM and DM/DD-OFDM optical links under channel impairments and have investigated PAPR reduction for optical IM/DD OFDM [17–19]. They motivate practical concerns such as received power, noise, BER, clipping, and nonlinear distortion but do not address NOMA scheduling or chance-constrained SIC certification. The present work therefore treats these studies as physical-layer context rather than as competing resource-allocation methods.

The system model uses an aggregate ADO-OFDM link abstraction with fixed DC bias, layer spectral-use factors, constant layer-noise powers, and an ACO-to-DCO leakage coefficient. Rate thresholds are spectral-efficiency targets within this abstraction, not waveform-derived BER or packet-error requirements. Likewise, confidence statements are conditional on the scenario generator used for certification: nonlinear geometry mapping does not invalidate independent binomial certification, but a misspecified error law is not covered. The numerical study therefore states and justifies the fixed DC operating point, operating-point noise, equal active-block power division, and aggregate ACO-to-DCO leakage as system-level abstractions. It separately examines uncertainty–law mismatch and unbalanced deployments, while limiting the practical claims to slow-timescale centralized scheduling.

This paper develops a probabilistic-decodability-graph-based hierarchical power-adaptation (PDG-HPA) framework. Each directed graph edge identifies an ordered user pair, serving LED, ACO/DCO layer, and resource block; certified OMA edges provide an explicit resource-accounted fallback. Candidate coefficients are screened on optimization scenarios, while fresh scenarios provide one-sided confidence bounds for mode admission. A resource-constrained mixed-integer linear program (MILP) selects compatible modes. For the selected matching, inter-LED, inter-layer, and intra-pair powers are adapted through fixed-threshold empirical CVaR constraints and trust-region sequential quadratic programming (SQP).

Candidate-edge certificates are admission filters, not a network-level guarantee. After matching and power adaptation, the complete assignment, activity indicators, powers, and interference terms are frozen and jointly recertified on an independent scenario set. An untouched test set is used only for final evaluation.

The methodological contribution is not the individual use of PDG, CVaR, MILP, or SQP. It is their reliability-consistent composition: directed receiver modes are statistically admitted, globally scheduled with resource-accounted fallback, continuously adapted, and then recertified as one simultaneous assignment. Matched counterfactuals quantify the separate effects of graph certification and risk-constrained power adaptation.

The principal contributions are as follows.

https://doi.org/10.3390/photonics13080795

- A geometry-based uncertainty model preserves the common position and orientation perturbation across all LED links of one user. Both SIC directions are evaluated through a joint three-SINR event. The certificate is explicitly generator-conditional, and controlled deployment–law mismatches are tested separately.
- A layer-, block-, LED-, and direction-labeled PDG combines confidence-qualified NOMA admission with certified OMA fallback. The side-constrained MILP selects pairing, decoding direction, optical layer, and resource block; a small-network extension also permits competing LED labels for joint association–scheduling.
- Fixed-matching power adaptation uses sequential rate lower bounds, original empirical CVaR and physical-power constraints, and a trust-region SQP acceptance test. Monotonicity is claimed only for accepted inner updates under one matching; no global-optimality or outer-loop convergence claim is made.
- The evaluation separates selected-pair and network-wide outage, matched and misspecified uncertainty laws, balanced and unbalanced topologies, conditional and unconditional performance, and feasible and unsuccessful runs. Matched baselines, multi-start references, and computational audits delimit the incremental benefit and operating scale.

The remainder of this paper is organized as follows. Section 2 reviews related scheduling and robust-allocation work. Section 3 presents the model and co-design problem. Section 4 develops PDG-HPA. Section 5 reports the numerical evaluation, and Section 6 concludes the paper.

2. Reliability-Aware Scheduling and Resource Allocation in Multi-Cell NOMA VLC Networks

2.1. NOMA Grouping, Optical OFDM, and Multi-Cell Interference

Early NOMA-VLC studies established the rate and error-performance gains of optical power-domain superposition and SIC [5,6]. Subsequent work considered gain-ratio allocation, two-user optimization, and joint grouping and power allocation [8,20–24]. These methods show that NOMA performance depends on channel disparity and power coefficients, but their groups and decoding orders are usually constructed from nominal or deterministic channel metrics rather than statistically certified under coupled geometry and interference uncertainty.

ACO-OFDM and DCO-OFDM satisfy the real and nonnegative constraints of intensity modulation/direct detection with different power and spectral efficiencies; ADO-OFDM superposes their complementary components [25,26]. Nonorthogonal multicarrier access, including SEFDM-NOMA, further couples waveform- and user-domain interference [27].

Recent physical-layer studies have analyzed a directly modulated DCO-OFDM visible-light link under atmospheric turbulence, a low-cost directly modulated OFDM optical-wireless architecture, and hybrid SLM-based PAPR reduction for optical IM/DD OFDM [17–19]. Their BER, received-power, noise, PAPR, and nonlinear distortion analyses complement the aggregate ADO-OFDM model used here. They do not, however, address multi-cell NOMA pairing, SIC-direction selection, or chance-constrained mode admission.

Multi-cell VLC studies have addressed grouping, dimming-aware interference mitigation, coordinated transmission, and hybrid NOMA/OFDMA [8–11]. The authors' earlier overlap-aware ADO-OFDM NOMA design jointly adapted grouping and three power levels under nominal CSI [28]. It is therefore a relevant deterministic baseline, but it neither certifies both SIC directions nor controls the complete decoding event out of sample.

Angle-diversity receivers provide a complementary physical method for suppressing multi-cell interference [29]. The present single-PD model instead optimizes receiver modes under uncertain geometry; diversity-combined gains could be supplied to the same scenario framework but are not implemented here.

2.2. Geometry and CSI Uncertainty

Imperfect CSI and residual SIC have been analyzed in NOMA-VLC through outage, BER, and achievable-rate metrics [12–14]. These studies evaluate the consequences of channel error for a prescribed transmission structure; they do not decide whether a risky pair should be excluded or whether its SIC direction should be reversed.

Receiver orientation and channel acquisition have also been studied in spatial-modulation VLC. Neural-network channel estimation and detection have been evaluated for mobile, randomly oriented receivers [30], and estimation and error statistics have been derived for narrow-FOV receivers with random orientation [31]. These works support the orientation and estimation components of the present generator but do not consider directed NOMA admission or confidence-certified multi-cell scheduling.

Robust resource allocation has been considered for integrated visible-light positioning/communication, RSMA, and blockage-aware hybrid access [15,16,32]. Worst-case methods protect a prescribed uncertainty set but can be conservative; chance constraints instead expose a reliability-efficiency tradeoff [33]. Scenario methods provide finite-sample tools [34,35], and CVaR controls tail loss [36]. In this work, optimization-set CVaR constrains empirical three-SINR shortfall, whereas fresh certification samples provide one-sided Clopper–Pearson admission bounds. The two roles are kept separate.

Finite-sample coverage remains conditional on the adopted generator. It does not protect against biased localization, heavy-tailed or mixture laws, omitted residual correlation, incorrect variances, or calibration error. The numerical study therefore freezes the nominal design and tests these misspecifications on independent deployment samples; it does not claim distributional robustness.

2.3. Graph Scheduling and Research Gap

Matching theory represents association, channel allocation, and pairing in wireless systems [37,38]. A conventional pairing edge, however, does not distinguish opposite SIC directions, optical layers, serving LEDs, resource blocks, or the cost of OMA fallback. Adding these labels and capacity constraints yields a side-constrained binary mode-selection problem.

Attribution also requires matched baselines. Nominal-CSI grouping, overlap-aware grouping, and bounded worst-case allocation differ in several mechanisms simultaneously. The revised study therefore adds scenario-robust HPA without PDG certification and PDG scheduling with deterministic power, while holding the mode universe, OMA accounting, direction options, scenarios, and budgets fixed. It reports common-feasible and unconditional performance together with feasibility, rather than ranking methods only by feasible-run averages.

The remaining problem is thus to (i) generate and test both directed NOMA modes through a correlated three-SINR event; (ii) separate coefficient tuning, statistical admission, complete-assignment recertification, and held-out testing; (iii) schedule NOMA and OMA modes with explicit resource costs; and (iv) adapt hierarchical powers without assuming that independently certified edges compose after global interference changes. PDG-HPA addresses this gap through directed probabilistic edges, side-constrained MILP scheduling, fixed-threshold CVaR-constrained SQP, and joint recertification of the frozen assignment. Its novelty is this reliability-consistent composition and audited evaluation, not any one constituent optimization tool.

3. System Model and Reliability-Constrained Formulation

As illustrated in Figure 1, we consider an indoor overlapping multi-cell NOMA-VLC network served by a ceiling-mounted 2×2 LED array. Each LED forms an optical cell on the receiver plane. Reuse of the same ACO/DCO resources by adjacent LEDs produces geometry-dependent inter-cell interference in the overlap regions. Each user is equipped with one photodetector (PD), and its estimated position and orientation may differ from their actual values. A single geometry perturbation therefore changes the desired and interfering links of that user jointly and may also reverse the feasible SIC direction of a candidate NOMA pair.

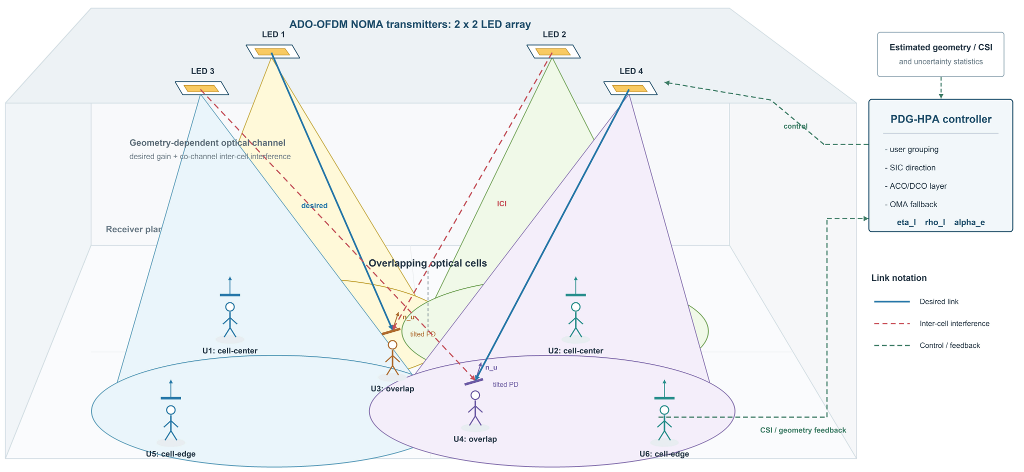

Figure 1. System model and control information flow of the considered overlapping multi-cell ADO-OFDM NOMA-VLC network. Solid blue, dashed red, and dashed green lines denote desired optical links, inter-cell interference, and control/feedback paths, respectively. Arrowheads indicate the direction of optical transmission or control/feedback signaling.

From estimated geometry, CSI, and uncertainty statistics, the controller selects the serving mode, SIC direction, ACO/DCO layer, and resource block and then adapts the hierarchical powers. Reliability is tested against fixed decoding thresholds; achieved rates enter only the scenario-average utility.

3.1. Network, Association, and ADO-OFDM Resources

Let $\mathcal{L} = \{1, \ldots, L\}$, $\mathcal{U} = \{1, \ldots, U\}$, and $\mathcal{B} = \{1, \ldots, B\}$ denote the LED, user, and resource-block sets. The two ADO-OFDM layers are indexed by $q \in \mathcal{Q} = \{\mathrm{A}, \mathrm{D}\}$, where A and D denote the ACO-OFDM and DCO-OFDM layers, respectively. The factor $\nu_q \in (0, 1]$ represents the net spectral-use efficiency of layer q after Hermitian symmetry, clipping, and reserved subcarriers [26]. All rates below are normalized spectral efficiencies in bit/s/Hz.

For the stated full-FFT normalization, Hermitian symmetry and the odd/even ADO allocation give $\nu_{\mathrm{A}} = \nu_{\mathrm{D}} = 0.25$. Cyclic-prefix, pilot, coding, and waveform-clipping losses are not modeled.

Association is determined on a slow timescale from the nominal channel:

$$\ell(u) = \arg\max_{l \in \mathcal{L}} P_l^{\mathrm{ref}} \left(R_{\mathrm{PD}} \widehat{h}_{l,u} \right)^2, \tag{1}$$

where R_{PD} is the PD responsivity and P_l^{ref} is a fixed reference power. Equal reference powers are used in the simulations. Ties are broken deterministically. A candidate NOMA mode contains users with the same serving LED. The directed mode

$$e = (s, w, l, q, b) \tag{2}$$

means that $l = \ell(s) = \ell(w)$, user w is decoded first, and user s decodes the signal of w before decoding its own signal. Both directions of a user pair are treated as distinct candidates.

The main algorithm uses (1). Association error is diagnosed from the scenario-wise strongest LED

$$\ell_\star^{(n)}(u) = \arg\max_{l\in\mathcal{L}} P_l^{\mathrm{ref}}\left(R_{\mathrm{PD}} h_{l,u}^{(n)}\right)^2,$$

with $\ell_\star^{(n)}(u) \neq \ell(u)$. A small-network reference permits competing labels $l \in \mathcal{L}_u^{\mathrm{cand}}$ and jointly selects association and scheduling.

Let p_l^q be the electrical modulation power assigned to layer q of LED l. The implemented power constraints are

$$\begin{aligned} \sum_{l\in\mathcal{L}}\sum_{q\in\mathcal{Q}} p_l^q &= P_{\mathrm{ac}}, \\ P_l^{\min} \le \sum_{q\in\mathcal{Q}} p_l^q &\le P_l^{\max}(I_l^{\mathrm{DC}}), \qquad \forall l \in \mathcal{L}, \\ p_l^q &\ge P_{\mathrm{layer}}^{\min}, \qquad \forall l \in \mathcal{L},\ q \in \mathcal{Q}. \end{aligned} \tag{3}$$

The DC bias I_l^{DC} is fixed to satisfy the illumination requirement and determines the available modulation headroom. It is not optimized in this study.

Here, $P_l^{\max}(I_l^{\mathrm{DC}})$ is the RMS modulation headroom at a fixed illumination point. It is not a sample-wise current constraint; clipping, dimming transients, and illumination uniformity are outside the link abstraction.

For binary NOMA and OMA variables y_e and $z_u^{l,q,b}$, define

$$a_{l,b}^q = \sum_{e\in\mathcal{M}_{l,b}^q} y_e + \sum_{u:(u,l,q,b)\in\mathcal{O}_{l,b}^q} z_u^{l,q,b}. \tag{4}$$

The resource constraint imposed below ensures $a_{l,b}^q \in \{0,1\}$. Consistently with the simulation code, the power of one LED/layer is divided equally among its active blocks:

$$\kappa_{l,b}^q = \begin{cases} a_{l,b}^q / N_l^q, & N_l^q > 0, \\ 0, & N_l^q = 0, \end{cases} \qquad N_l^q = \sum_{b\in\mathcal{B}} a_{l,b}^q, \qquad p_{l,b}^q = \kappa_{l,b}^q p_l^q. \tag{5}$$

Thus, inactive blocks transmit no modulation power. For a fixed mode selection, the collection $\boldsymbol{\kappa} = \{\kappa_{l,b}^q\}_{l,q,b}$ is fixed during the continuous power update and is recomputed after the graph matching changes.

Equal active-block division is used by every method to avoid a power variable for each LED–layer–block tuple; it is not claimed to be an optimal subcarrier-power law.

For a selected NOMA mode $e = (s, w, l, q, b)$, let $p_{e,w}$ and $p_{e,s}$ be the powers assigned to the first-decoded and SIC users:

$$p_{e,w} + p_{e,s} = p_{l,b}^q, \qquad p_{e,w} \ge 0, \quad p_{e,s} \ge 0. \tag{6}$$

The corresponding electrical superposition is

$$x_{l,b}^q = \sqrt{p_{e,w}}\, d_w + \sqrt{p_{e,s}}\, d_s, \tag{7}$$

where d_w and d_s have unit variance. The hierarchical coefficients reported by the controller are recovered as

$$\eta_l = \frac{p_l^{\mathrm{A}} + p_l^{\mathrm{D}}}{P_{\mathrm{ac}}}, \qquad \rho_l = \frac{p_l^{\mathrm{A}}}{p_l^{\mathrm{A}} + p_l^{\mathrm{D}}}, \qquad \alpha_e = \frac{p_{e,w}}{p_{e,w} + p_{e,s}}. \tag{8}$$

A zero ratio is defined as zero. In the implementation, strictly positive numerical lower bounds are used for selected NOMA user powers to prevent division by zero; they do not impose an additional power-ordering rule.

3.2. Correlated Geometry and CSI Uncertainty

For a line-of-sight LED–PD link, the DC optical gain is [1]

$$h_{l,u}(\mathbf{r}_u, \mathbf{n}_u) = \frac{(m_l+1)A_{\mathrm{PD}}}{2\pi d_{l,u}^2} \cos^{m_l}(\phi_{l,u}) T_s(\psi_{l,u}) g(\psi_{l,u}) \cos(\psi_{l,u}) \mathbf{1}_{\{0 \le \psi_{l,u} \le \Psi_{\mathrm{FOV}}\}}, \tag{9}$$

where A_{PD} is the PD area, $m_l = -\ln 2 / \ln(\cos \Phi_{1/2,l})$ is the Lambertian order, $d_{l,u}$ is the link distance, $\phi_{l,u}$ and $\psi_{l,u}$ are the irradiance and incidence angles, $T_s(\cdot)$ is the filter gain, $g(\cdot)$ is the concentrator gain, and Ψ_{FOV} is the PD field of view.

Let $\widehat{\mathbf{r}}_u$ and $\widehat{\mathbf{n}}_u$ denote the estimated user position and unit PD normal. The scenario generator used in the numerical study is represented by

$$\begin{aligned}
\mathbf{r}_u &= \Pi_{\mathcal{R}}(\widehat{\mathbf{r}}_u + \Delta\mathbf{r}_u), \qquad \Delta\mathbf{r}_u \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma}_{r,u}), \\
\widetilde{\mathbf{n}}_u &= \widehat{\mathbf{n}}_u + \delta_{u,1}\mathbf{t}_{u,1} + \delta_{u,2}\mathbf{t}_{u,2}, \qquad \delta_{u,k} \sim \mathcal{N}(0, \sigma_{\mathrm{ori},u}^2/2), \\
\mathbf{n}_u &= \Pi_+(\widetilde{\mathbf{n}}_u / \|\widetilde{\mathbf{n}}_u\|_2), \\
h_{l,u} &= \Big[h_{l,u}(\mathbf{r}_u, \mathbf{n}_u) + \sigma_{\mathrm{est},u}\widehat{h}_{l,u}\varepsilon_{l,u}\Big]_+, \qquad \varepsilon_{l,u} \sim \mathcal{N}(0,1).
\end{aligned} \tag{10}$$

Here, $\Pi_{\mathcal{R}}$ clips a perturbed position to the physical room, $\mathbf{\Sigma}_{r,u}$ has zero variance in the fixed receiver-plane height coordinate used by the simulations, $\mathbf{t}_{u,1}$ and $\mathbf{t}_{u,2}$ form an orthonormal tangent basis at $\widehat{\mathbf{n}}_u$, and Π_+ renormalizes the PD normal while retaining its upward-facing component. The same $(\Delta\mathbf{r}_u, \delta_{u,1}, \delta_{u,2})$ is used to generate all L links of user u. Consequently, geometry-induced correlation among the desired and interfering channels is preserved. The final additive term represents a relative channel-estimation residual. Every optimization, certification, and test scenario recomputes the nonlinear Lambertian gain and the hard FOV indicator.

Let $\widehat{\mathbb{P}}_{\widehat{\boldsymbol{\theta}}}$ denote the fitted scenario law. All nominal probability statements use this law. The binomial certificates remain valid for i.i.d. success indicators after the nonlinear geometry and FOV transformations, but not for a different deployment law $\mathbb{P}_\star$. Section 5 therefore tests bias, heavy tails, mixtures, residual correlation, variance error, and finite-data calibration. The channel is quasi-static, LOS-dominant, and block-flat.

3.3. Control Timescale, Acquisition Overhead, and CSI Aging

The controller operates on a slow scheduling timescale. Pilots and receiver sensing provide the state estimate, while a longer window fits $\widehat{\boldsymbol{\theta}}$. The chosen mode, SIC direction, resources, and quantized powers remain fixed until the next update.

Let T_{upd} denote the scheduling-update period and let T_{pil}, T_{sens}, T_{fb}, and T_{sig} denote the airtime used for pilots, acquisition and reporting of position/orientation information, uplink feedback, and downlink configuration signaling, respectively. A transparent rate-accounting factor is

$$\eta_{\mathrm{ctl}} = \left[1 - \frac{T_{\mathrm{pil}} + T_{\mathrm{sens}} + T_{\mathrm{fb}} + T_{\mathrm{sig}}}{T_{\mathrm{upd}}}\right]_+. \tag{11}$$

The net rate is $\eta_{\mathrm{ctl}} R_u^{(n)}$. With optimization time T_{opt}, the information age is $T_{\mathrm{fb}} + T_{\mathrm{opt}} + T_{\mathrm{sig}}$. The simulations assume negligible motion during one update and therefore do not validate rapidly varying CSI.

No implementation-specific control-frame timing is assumed in the numerical study. The rates reported in Section 5 are therefore pre-control-overhead spectral efficiencies; (11) provides a post-processing conversion to net rate once the pilot, sensing, feedback, signaling, and update durations of a particular implementation are specified.

3.4. Interference, Fixed Decoding Thresholds, and SIC Events

For scenario n, define

$$H_{l,u}^{(n)} = \left(R_{\mathrm{PD}} h_{l,u}^{(n)}\right)^2. \tag{12}$$

For user u served by LED l, the implemented interference-plus-noise term on (q,b) is

$$\begin{aligned} C_{u,b}^{q,(n)} &= \sigma_{u,b,q}^2 + \sum_{\substack{j\in\mathcal{L}\\ j\neq l}} a_{j,b}^q H_{j,u}^{(n)} p_{j,b}^q \\ &\quad + \mathbf{1}_{\{q=\mathrm{D}\}} \zeta_{\mathrm{A}\to\mathrm{D}} \sum_{j\in\mathcal{L}} a_{j,b}^{\mathrm{A}} H_{j,u}^{(n)} p_{j,b}^{\mathrm{A}}. \end{aligned} \tag{13}$$

The second line models residual ACO-to-DCO reconstruction leakage. Setting $\zeta_{\mathrm{A}\to\mathrm{D}} = 0$ gives ideal ADO-layer separation. No DCO-to-ACO residual term is included in the implemented model. The desired received electrical power on the selected resource is

$$A_{l,u,b}^{q,(n)} = H_{l,u}^{(n)} p_{l,b}^q. \tag{14}$$

The constant $\sigma_{u,b,q}^2$ approximates receiver noise at the fixed DC/background operating point:

$$\sigma_{u,b,q}^2 = \sigma_{\mathrm{th},u,b,q}^2 + 2q_{\mathrm{e}} R_{\mathrm{PD}} \left(P_{\mathrm{bg},u} + P_{\mathrm{dc},u}\right) B_{\mathrm{e}},$$

where the terms represent thermal and operating-point shot noise. Signal-dependent variation is omitted, and $\zeta_{\mathrm{A}\to\mathrm{D}}$ is a lumped post-separation leakage coefficient; no waveform-level robustness is claimed.

For directed NOMA mode $e = (s, w, l, q, b)$, the three receiver SINRs are

$$\begin{aligned} \gamma_{w\to w}^{e,(n)} &= \frac{H_{l,w}^{(n)} p_{e,w}}{H_{l,w}^{(n)} p_{e,s} + C_{w,b}^{q,(n)}}, \\ \gamma_{s\to w}^{e,(n)} &= \frac{H_{l,s}^{(n)} p_{e,w}}{H_{l,s}^{(n)} p_{e,s} + C_{s,b}^{q,(n)}}, \\ \gamma_{s\to s}^{e,(n)} &= \frac{H_{l,s}^{(n)} p_{e,s}}{\xi_e H_{l,s}^{(n)} p_{e,w} + C_{s,b}^{q,(n)}}. \end{aligned} \tag{15}$$

where $\xi_e \in [0,1]$ is the residual intra-pair SIC factor.

Let $\overline{R}_{\mathrm{w}}^q$, $\overline{R}_{\mathrm{s}}^q$, and $\overline{R}_{\mathrm{O}}^q$ be the fixed spectral-efficiency thresholds for first-user decoding, the SIC user's own decoding, and OMA decoding, respectively. Their SINR thresholds are

$$\overline{\Gamma}_c^q = 2^{\overline{R}_c^q/\nu_q} - 1, \qquad c \in \{\mathrm{w}, \mathrm{s}, \mathrm{O}\}. \tag{16}$$

These thresholds are configuration parameters and are not optimized. The default simulations use $\overline{R}_{\mathrm{w}}^q = 0.10$, $\overline{R}_{\mathrm{s}}^q = 0.08$, and $\overline{R}_{\mathrm{O}}^q = 0.06$ bit/s/Hz on both layers.

The unit-gap mapping makes these normalized scheduler targets, not BER or packet-error claims for a specified modulation and code. A practical link may replace $\overline{\Gamma}_c^q$ by a measured decoding threshold without changing the graph or chance constraints.

The directed three-SINR pair event is

$$\mathcal{A}_e^{(n)} = \left\{ \gamma_{w\to w}^{e,(n)} \geq \overline{\Gamma}_{\mathrm{w}}^{q},\ \gamma_{s\to w}^{e,(n)} \geq \overline{\Gamma}_{\mathrm{w}}^{q},\ \gamma_{s\to s}^{e,(n)} \geq \overline{\Gamma}_{\mathrm{s}}^{q} \right\}. \tag{17}$$

Its conditional reliability requirement is

$$\Pr\Big(\mathcal{A}_e \mid \widehat{\mathbf{r}}, \widehat{\mathbf{n}}, \widehat{\mathbf{h}} \Big) \geq 1 - \epsilon_{\mathrm{pair}}. \tag{18}$$

The ordering guard uses the current signal-to-interference-plus-noise quality

$$g_{l,u,b}^{q,(n)} = \frac{A_{l,u,b}^{q,(n)}}{C_{u,b}^{q,(n)}}. \tag{19}$$

For direction $s \to w$, define

$$\mathcal{O}_e^{(n)} = \left\{ g_{l,s,b}^{q,(n)} - g_{l,w,b}^{q,(n)} \geq \Delta_g \right\}, \qquad \Pr(\mathcal{O}_e) \geq 1 - \epsilon_{\mathrm{ord}}. \tag{20}$$

The ordering event is certified separately from (17); it is an admission guard and not a replacement for the three decoding inequalities.

For OMA mode $o = (u, l, q, b)$, the instantaneous spectral efficiency and success event are

$$\begin{aligned} R_{o,u}^{(n)} &= \nu_q \log_2\left(1 + \frac{H_{l,u}^{(n)} p_{l,b}^{q}}{C_{u,b}^{q,(n)}} \right), \\ \mathcal{A}_o^{(n)} &= \left\{ \frac{H_{l,u}^{(n)} p_{l,b}^{q}}{C_{u,b}^{q,(n)}} \geq \overline{\Gamma}_{\mathrm{O}}^{q} \right\}, \qquad \Pr(\mathcal{A}_o) \geq 1 - \epsilon_{\mathrm{O}}. \end{aligned} \tag{21}$$

3.5. Scenario-Average Utility and Co-Design Problem

The achieved NOMA spectral efficiencies in scenario n are

$$\begin{aligned} R_{e,w}^{(n)} &= \nu_q \log_2\Big[1 + \min\Big(\gamma_{w\to w}^{e,(n)}, \gamma_{s\to w}^{e,(n)} \Big) \Big], \\ R_{e,s}^{(n)} &= \nu_q \log_2\Big(1 + \gamma_{s\to s}^{e,(n)} \Big). \end{aligned} \tag{22}$$

For a complete mode selection $\chi = (\mathbf{y}, \mathbf{z})$ and continuous powers $\mathbf{p}$, let $R_u^{(n)}(\chi, \mathbf{p})$ denote the raw achieved rate assigned to user u by (22) or (21). The optimization-set average is

$$\widehat{R}_u^{\mathrm{opt}}(\chi, \mathbf{p}) = \frac{1}{S_{\mathrm{opt}}} \sum_{n \in \mathcal{S}_{\mathrm{opt}}} R_u^{(n)}(\chi, \mathbf{p}). \tag{23}$$

With weights $\omega_u > 0$, reference rate $R_0 > 0$, and a small numerical rate floor $R_{\mathrm{fl}} > 0$, the implemented acceptance utility is

$$F(\chi, \mathbf{p}) = \sum_{u \in \mathcal{U}} \omega_u \log\left(\frac{\max\{ \widehat{R}_u^{\mathrm{opt}}(\chi, \mathbf{p}), R_{\mathrm{fl}} \}}{R_0} \right). \tag{24}$$

The floor prevents the logarithm of zero and is not a QoS decision variable. Reliability is imposed by the fixed-threshold events (17), (20) and (21).

Let $\chi = (\mathbf{y}, \mathbf{z})$ denote the complete binary mode-selection vector.

$$
\begin{aligned}
\max_{\chi,\mathbf{p}} \quad & F(\chi,\mathbf{p}) \\
\text{s.t.} \quad & \sum_{e\in\mathcal{E}:u\in e} y_e + \sum_{(u,l,q,b)\in\mathcal{O}} z_u^{l,q,b} = 1, \quad \forall u \in \mathcal{U}, \\
& \sum_{e\in\mathcal{M}_{l,b}^{q}} y_e + \sum_{u:(u,l,q,b)\in\mathcal{O}} z_u^{l,q,b} \le 1, \quad \forall l,q,b, \\
& y_e = 1 \Rightarrow \Pr(\mathcal{O}_e) \ge 1-\epsilon_{\mathrm{ord}}, \quad \forall e \in \mathcal{E}, \\
& y_e = 1 \Rightarrow \Pr(\mathcal{A}_e) \ge 1-\epsilon_{\mathrm{pair}}, \quad \forall e \in \mathcal{E}, \\
& z_u^{l,q,b} = 1 \Rightarrow \Pr(\mathcal{A}_o) \ge 1-\epsilon_{\mathrm{O}}, \quad \forall (u,l,q,b) \in \mathcal{O}, \\
& \mathbf{p} \text{ satisfies (3), (5), and (6)}, \\
& y_e \in \{0,1\}, \qquad z_u^{l,q,b} \in \{0,1\}.
\end{aligned}
\tag{25}
$$

Problem (25) is a mixed-integer stochastic nonlinear program. The mode variables determine resource activity and therefore alter block-power fractions and interference, while the continuous powers change both candidate reliability and edge utility. Section 4 addresses this coupling through confidence-certified graph construction, exact resource-constrained binary matching, and fixed-threshold CVaR-constrained power adaptation using sequential rate lower bounds and SQP.

Matched variants share the same mode universe, resources, thresholds, power limits, OMA costs, and test scenarios; only the named admission or power rule changes. An instance is feasible only if that method returns a complete physically valid assignment within its solver budget and, when claimed, passes final certification and repair. Failure causes are retained when computing feasibility and unconditional performance.

The chance constraints are conditional on $\widehat{\mathbb{P}}_{\hat{\theta}}$; deployment-law performance under $\mathbb{P}_\star$ is reported separately. The main problem also retains the fixed association in (1); only the small-network reference expands the edge set over competing LED labels.

After design, held-out effective rates are zeroed only when the corresponding weak-user, SIC-user, ordering, or OMA condition fails. This evaluation rule does not alter the optimization utility in (24).

For ease of reference, Table 1 summarizes the principal notation used in the reliability formulation.

Table 1. Principal notation used in the reliability formulation.

Symbol	Meaning
$h_{l,u}$	Optical channel gain from LED l to user u
$\mathcal{E}$, $\mathcal{O}$	NOMA and OMA mode sets used in the joint formulation
$\mathcal{A}_e$	Joint three-SINR decoding-success event for NOMA mode e
$\mathcal{O}_e$	SIC-ordering event for directed NOMA mode e
p_e^{LCB}	Scalar candidate-edge reliability indicator obtained from the one-sided pair- and ordering-success lower confidence bounds
$p_{\mathrm{net}}^{\mathrm{LCB}}$	One-sided lower confidence bound on complete-assignment network success
ϵ_{pair}, ϵ_{ord}, ϵ_{O}, ϵ_{net}	Pair-outage, ordering-error, OMA-outage, and network-outage tolerances
$\mathbf{y}$, $\mathbf{z}$	Binary NOMA- and OMA-mode selection variables
η_l, ρ_l, α_e	Inter-LED, inter-layer, and intra-pair hierarchical power coefficients
$\mathcal{S}_{\mathrm{opt}}$, $\mathcal{S}_{\mathrm{cert}}$, $\mathcal{S}_{\mathrm{joint}}$, $\mathcal{S}_{\mathrm{test}}$	Optimization, candidate-certification, final joint-recertification, and held-out test scenario sets

4. Probabilistic-Decodability-Graph-Based Receiver–Scheduler Co-Design

PDG-HPA alternates certified candidate admission, MILP mode selection, and fixed-mode CVaR-constrained SQP power adaptation. Power changes trigger graph reconstruction. The frozen complete assignment is then jointly recertified and, if needed, repaired.

Independent optimization, certification, and test sets support design, statistical admission, and final evaluation, respectively. Every scenario recomputes all links jointly from (10).

Certification is conditional on $\widehat{\mathbb{P}}_{\widehat{\theta}}$; deployment laws $\mathbb{P}_\star$ appear only in held-out mismatch tests.

4.1. Optimization-Set Screening and Preliminary Mode Parameters

A directed NOMA candidate is indexed by $e = (s, w, l, q, b)$ as in (2). Both SIC directions are enumerated for every pair of users associated with the same LED. An OMA candidate $o = (u, l, q, b)$ assigns one complete LED/layer/block resource to user u.

The main method uses $\mathcal{L}_u^{\text{cand}} = \{\ell(u)\}$. The small-network reference adds competing LED-labeled NOMA and OMA modes and otherwise uses the same screening and certification rules.

Each NOMA candidate is evaluated directly on nonlinear optimization scenarios over the coefficient grid $\mathcal{A}_\alpha$. For trial α, let

$$
\begin{aligned}
\widehat{p}_{e,\text{pair}}(\alpha) &= \frac{1}{S_{\text{opt}}} \sum_{n \in \mathcal{S}_{\text{opt}}} \mathbf{1}_{\{\mathcal{A}_e^{(n)}(\alpha)\}}, \\
\widehat{p}_{e,\text{ord}}(\alpha) &= \frac{1}{S_{\text{opt}}} \sum_{n \in \mathcal{S}_{\text{opt}}} \mathbf{1}_{\{\mathcal{O}_e^{(n)}\}},
\end{aligned}
\tag{26}
$$

and let $Q_\beta[R_{e,w}(\alpha)]$ and $Q_\beta[R_{e,s}(\alpha)]$ be their empirical lower β-quantile rates. The numerical experiments use $\beta = 0.05$ and $\mathcal{A}_\alpha = \{0.50, 0.52, \ldots, 0.94\}$.

For a scenario loss vector $\mathbf{v} = (v^{(1)}, \ldots, v^{(S)})$, let $v_{[1]} \geq \cdots \geq v_{[S]}$ denote its descending order statistics, $q_\epsilon = \epsilon S$, $J_\epsilon = \lfloor q_\epsilon \rfloor$, and $\theta_\epsilon = q_\epsilon - J_\epsilon$. The exact empirical upper-tail CVaR used by the code is

$$
\widehat{\text{CVaR}}_\epsilon(\mathbf{v}) = \frac{\sum_{j=1}^{J_\epsilon} v_{[j]} + \theta_\epsilon v_{[J_\epsilon+1]}}{q_\epsilon},
\tag{27}
$$

where the fractional term is omitted for $\theta_\epsilon = 0$ and an empty sum is zero. This definition preserves the requested tail mass. Using the loss in (42), the preliminary coefficient maximizes

$$
\begin{aligned}
\widetilde{\alpha}_e \in \arg\max_{\alpha \in \mathcal{A}_\alpha} \Big\{ & \widehat{p}_{e,\text{pair}}(\alpha) + 0.2\, \widehat{p}_{e,\text{ord}}(\alpha) \\
& + 0.02\big(Q_\beta[R_{e,w}(\alpha)] + Q_\beta[R_{e,s}(\alpha)]\big) \\
& - 0.05 \Big[\widehat{\text{CVaR}}_{\epsilon_{\text{pair}}}(\mathbf{v}_e(\alpha))\Big]_+ \Big\}.
\end{aligned}
\tag{28}
$$

The coefficients are fixed screening weights, and ties use the first grid point. A candidate is discarded before certification if

$$
\widehat{p}_{e,\text{pair}}(\widetilde{\alpha}_e) < p_{\text{pre}} \quad \text{or} \quad \widehat{p}_{e,\text{ord}}(\widetilde{\alpha}_e) < p_{\text{pre}} \quad \text{or} \quad \widehat{\text{CVaR}}_{\epsilon_{\text{pair}}}(\mathbf{v}_e(\widetilde{\alpha}_e)) > \varepsilon_{\text{CVaR}},
\tag{29}
$$

where $p_{\text{pre}} = 0.20$. OMA uses the analogous success, quantile, and CVaR screen. Screening reduces the candidate set but supplies no confidence statement.

For the uncertainty-severity diagnostic only, NOMA candidates retain $\mathcal{L}_u^{\text{cand}} = \{\ell(u)\}$, while OMA candidates are expanded to all nominally visible LEDs. This diagnostic variant is not used for the main fixed-association performance claims.

4.2. Candidate-Level Confidence-Certified Probabilistic Decodability Graph

For certification claim c of candidate m, let $K_{m,c}$ be the number of successes among S_{cert} independent certification scenarios. A NOMA mode has separate claims $c \in \{\text{pair}, \text{ord}\}$, whereas an OMA mode has one fixed-threshold decoding claim. For $K_{m,c} > 0$, the one-sided Clopper–Pearson lower confidence bound is the following [39]:

$$p_{m,c}^{\text{LCB}} = I_{\delta_{\text{alloc}}}^{-1}(K_{m,c}, S_{\text{cert}} - K_{m,c} + 1), \tag{30}$$

and $p_{m,c}^{\text{LCB}} = 0$ when $K_{m,c} = 0$. Here, $I_{\delta}^{-1}(a, b)$ is the δ-quantile of the beta distribution with shape parameters a and b.

The bound applies to

$$p_{m,c} = \widehat{\mathbb{P}}_{\widehat{\theta}}(\mathcal{G}_{m,c}),$$

where $\mathcal{G}_{m,c}$ is the certified event. Coverage requires independent Bernoulli trials from the fitted generator and does not extend to a misspecified $\mathbb{P}_{\star}$.

The familywise confidence budget is preallocated between candidate admission and final-assignment certification:

$$\delta_{\text{FW}} = \delta_{\text{E}} + \delta_{\text{J}}, \tag{31}$$

where δ_{E} is reserved for candidate-level claims and δ_{J} is reserved for joint recertification of complete assignments. Let N_{cfg} be the number of prespecified reliability settings evaluated jointly, T_{cert} the number of candidate-certification rounds, and

$$M_{\max} = B|\mathcal{Q}|[U(U-1) + U] \tag{32}$$

an association-independent upper bound on the candidate count in one round. Since at most two claims are reserved for every candidate, the candidate-level per-claim budget is

$$\delta_{\text{E,alloc}} = \frac{\delta_{\text{E}}}{2N_{\text{cfg}} T_{\text{cert}} M_{\max}}. \tag{33}$$

Candidate claims use $\delta_{\text{E,alloc}}$. This allocation controls simultaneous confidence loss, not network decoding probability, which is treated in Section 4.5.

A NOMA candidate is certified only if

$$p_{e,\text{pair}}^{\text{LCB}} \geq 1 - \epsilon_{\text{pair}}, \qquad p_{e,\text{ord}}^{\text{LCB}} \geq 1 - \epsilon_{\text{ord}}. \tag{34}$$

Its scalar reliability indicator is

$$p_{e}^{\text{LCB}} = \min\left\{p_{e,\text{pair}}^{\text{LCB}}, p_{e,\text{ord}}^{\text{LCB}}\right\}. \tag{35}$$

An OMA candidate is retained only if

$$p_{o,\text{O}}^{\text{LCB}} \geq 1 - \epsilon_{\text{O}}, \tag{36}$$

where its successes are evaluated using the fixed event in (21). Thus, OMA is a certified fallback mode, not an automatically feasible dummy decision.

The resulting layer-, block-, and direction-labeled multigraph is

$$\mathcal{G}^{\text{PD}} = \left(\mathcal{U} \cup \mathcal{D}, \mathcal{E}^{\text{A}} \cup \mathcal{E}^{\text{D}} \cup \mathcal{E}^{\text{O}}\right), \qquad \mathcal{D} = \{d_u : u \in \mathcal{U}\}, \tag{37}$$

Here, d_u is a private OMA dummy vertex. Parallel edges represent different directions or resources. A topology is infeasible when no certified edge cover satisfies the resource limits.

Figure 2 illustrates certified ACO/DCO directions, rejected candidates, and resource-accounted OMA modes.

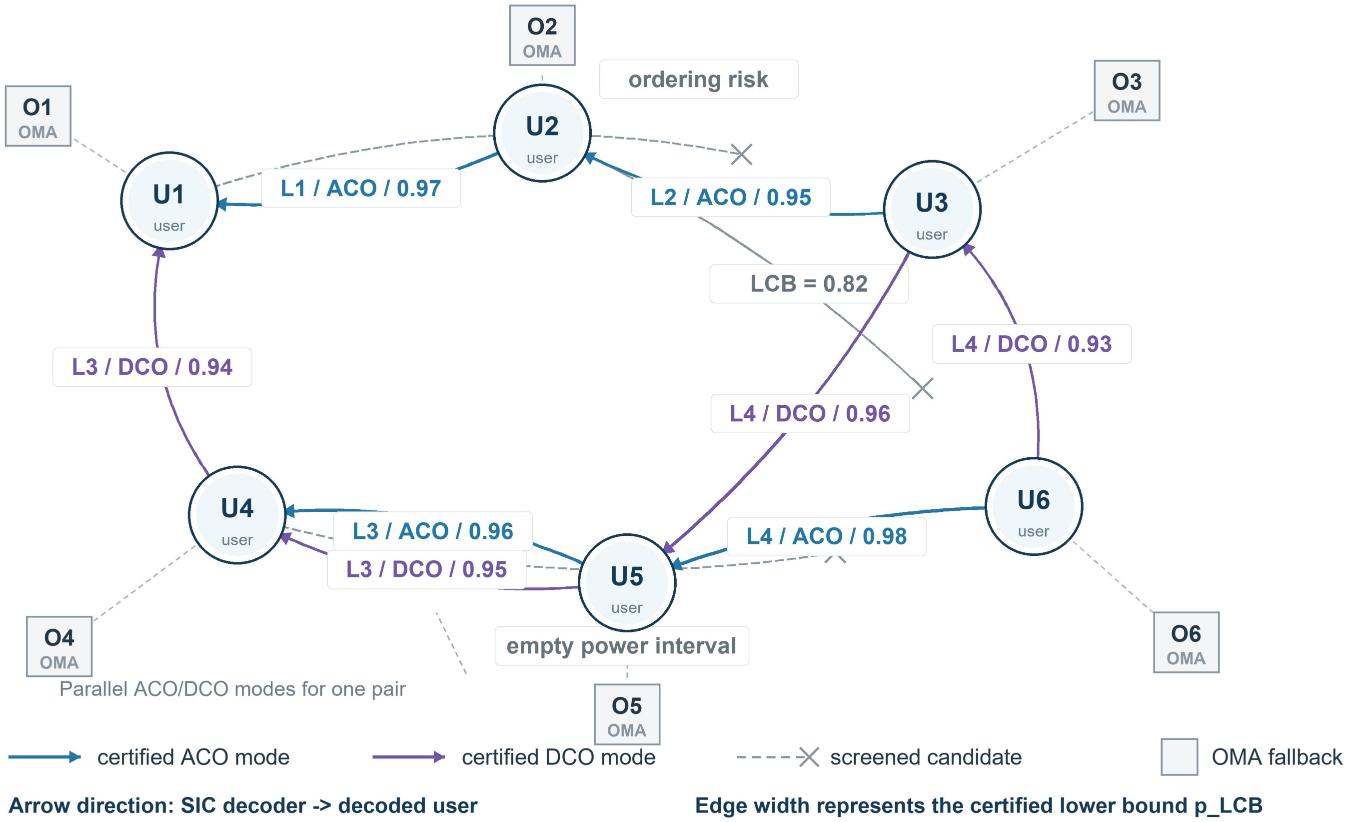

Figure 2. Illustrative six-user PDG. Circles and squares denote users and private OMA modes; blue and purple arcs denote certified ACO and DCO NOMA modes. Arrow $s \to w$ identifies the SIC user and first-decoded user, and width increases with p_e^{LCB}. Dashed crossed arcs were rejected. Probabilities are schematic.

If no failures are observed, the certification size required for a per-claim success probability of at least $1-\epsilon$ at confidence $1-\delta_{\mathrm{alloc}}$ satisfies

$$S_{\mathrm{cert}} \geq \left\lceil \frac{\ln(\delta_{\mathrm{alloc}})}{\ln(1-\epsilon)} \right\rceil. \tag{38}$$

4.3. Reliability- and Resource-Aware Mode Matching

For a certified NOMA edge, let $r_{e,w}^{\mathrm{rob}}$ and $r_{e,s}^{\mathrm{rob}}$ be the optimization-set lower-quantile rates obtained with $\widetilde{\alpha}_e$. For OMA, let $r_{o,u}^{\mathrm{rob}}$ be the corresponding lower-quantile rate. The edge weight implemented in the complete method is

$$W_e = \sum_{u\in e} \omega_u \log\left(\frac{\max\{r_{e,u}^{\mathrm{rob}}, R_{\mathrm{fl}}\}}{R_0}\right) - \lambda_r(1-p_e^{\mathrm{LCB}}) - \lambda_c - \lambda_{\mathrm{O}}\mathbf{1}_{\{e\in\mathcal{E}^{\mathrm{O}}\}}, \tag{39}$$

where the sum has two terms for NOMA and one term for OMA. The nonnegative parameters penalize residual reliability risk, resource use, and explicit OMA fallback. These edge weights guide the discrete stage; the original scenario-average utility (24) is used to evaluate continuous power updates.

Use one binary variable χ_e for every certified NOMA or OMA edge. Let $\delta(u)$ contain all edges incident to user u, and let $\mathcal{E}_{l,q,b}$ contain the edges using resource (l,q,b). The mode-selection problem is

$$
\begin{aligned}
\max_{\chi} \quad & \sum_{e\in\mathcal{E}} W_e \chi_e \\
\text{s.t.} \quad & \sum_{e\in\delta(u)} \chi_e = 1, \quad \forall u \in \mathcal{U}, \\
& \sum_{e\in\mathcal{E}_{l,q,b}} \chi_e \le 1, \quad \forall l,q,b, \\
& \chi_e \in \{0,1\}, \quad \forall e \in \mathcal{E}.
\end{aligned} \tag{40}
$$

Resource labels make (40) a 0–1 MILP; ties are resolved deterministically.

The small-network joint-association reference uses the same MILP with competing LED labels. It measures the cost of fixed association but is not presented as scalable. Natural load imbalance needs no new constraint because the MILP imposes no equal per-LED user count.

4.4. Fixed-Threshold CVaR-Constrained Power Adaptation by SQP

For a fixed matching, the activity indicators and equal active-block fractions in (5) are fixed. The continuous variables are

$$
\mathbf{x} = \Big[\{p_l^q\}_{l,q},\ \{p_{e,w}, p_{e,s}\}_{e\in\mathcal{E}_{\text{sel}}},\ \{\tau_m\}_{m\in\mathcal{M}_{\text{sel}}}\Big]^{\mathrm{T}}, \tag{41}
$$

where τ_m is the empirical CVaR threshold of selected mode m. The equality and bound constraints in (3) and (6) are imposed directly.

The implemented losses are SINR shortfalls, not cross-multiplied power margins. For NOMA mode $e = (s,w,l,q,b)$ in scenario n, define

$$
\begin{aligned}
v_{e,1}^{(n)}(\mathbf{p}) &= \overline{\Gamma}_{\mathrm{W}}^{q} - \gamma_{w\to w}^{e,(n)}(\mathbf{p}), \\
v_{e,2}^{(n)}(\mathbf{p}) &= \overline{\Gamma}_{\mathrm{W}}^{q} - \gamma_{s\to w}^{e,(n)}(\mathbf{p}), \\
v_{e,3}^{(n)}(\mathbf{p}) &= \overline{\Gamma}_{\mathrm{S}}^{q} - \gamma_{s\to s}^{e,(n)}(\mathbf{p}), \\
v_{e}^{(n)}(\mathbf{p}) &= \max_{k\in\{1,2,3\}} v_{e,k}^{(n)}(\mathbf{p}).
\end{aligned} \tag{42}
$$

For OMA mode $o = (u,l,q,b)$,

$$
v_o^{(n)}(\mathbf{p}) = \overline{\Gamma}_{\mathrm{O}}^{q} - \frac{H_{l,u}^{(n)} p_{l,b}^{q}}{C_{u,b}^{q,(n)}(\mathbf{p})}. \tag{43}
$$

For either selected mode m, the sample-average CVaR constraint is [36]

$$
\tau_m + \frac{1}{\epsilon_m S_{\text{opt}}} \sum_{n\in\mathcal{S}_{\text{opt}}} \Big[v_m^{(n)}(\mathbf{p}) - \tau_m\Big]_+ \le 0, \tag{44}
$$

where $\epsilon_m = \epsilon_{\text{pair}}$ for NOMA and $\epsilon_m = \epsilon_{\mathrm{O}}$ for OMA. CVaR controls optimization-set tail risk; independent Clopper–Pearson bounds govern statistical admission.

During one SQP update, the fitted generator, DC bias, noise, layer factors, leakage, and active-block fractions are fixed. The CVaR model does not optimize waveform or front-end effects.

To improve the continuous objective numerically, the code constructs a first-order rate lower bound around the current accepted power point $\mathbf{p}^{(i)}$, following the standard inner-approximation principle in sequential nonconvex optimization [40]. This citation motivates the local minorization step only; the complete nonlinear CVaR-constrained SQP subproblem below is not claimed to be convex. Any decoding rate can be written as

$$R^{(n)}(\mathbf{p}) = \frac{\nu_q}{\ln 2}\Big[\ln T^{(n)}(\mathbf{p}) - \ln D^{(n)}(\mathbf{p})\Big], \tag{45}$$

where $T^{(n)}$ and $D^{(n)}$ are the total and denominator powers of the corresponding SINR. Since $\ln D$ is concave,

$$-\ln D^{(n)}(\mathbf{p}) \geq -\ln D^{(n)}(\mathbf{p}^{(i)}) - \frac{D^{(n)}(\mathbf{p}) - D^{(n)}(\mathbf{p}^{(i)})}{D^{(n)}(\mathbf{p}^{(i)})}. \tag{46}$$

Substitution into (45) gives the scenario-wise local lower approximation $\underline{R}^{(n,i)}(\mathbf{p})$. For the first-decoded NOMA user, the approximation is the minimum of the two bounds corresponding to $w \to w$ and $s \to w$. The SIC-user and OMA approximations use their respective single decoding links. Consistently with the implementation, the numerical rate floor is applied to each scenario-wise approximation before scenario averaging. The resulting search surrogate is

$$\widetilde{F}^{(i)}(\mathbf{p}) = \sum_{u \in \mathcal{U}} \omega_u \log\left[\frac{\frac{1}{S_{\mathrm{opt}}}\sum_{n \in \mathcal{S}_{\mathrm{opt}}} \max\left\{\underline{R}_u^{(n,i)}(\mathbf{p}), R_{\mathrm{fl}}\right\}}{R_0}\right] - \varepsilon_{\mathrm{reg}}\|\mathbf{p}\|_2^2, \tag{47}$$

Here, $\varepsilon_{\mathrm{reg}}$ is numerical regularization. $\widetilde{F}^{(i)}$ generates a search point; because its flooring order differs from (24), it is not claimed to minorize the original utility. Acceptance always uses the original utility.

$$\begin{aligned} \max_{\mathbf{p},\boldsymbol{\tau}} \quad & \widetilde{F}^{(i)}(\mathbf{p}) \\ \text{s.t.} \quad & (3),\ (6),\ (44), \\ & |p_k - p_k^{(i)}| \leq \Delta_{\mathrm{TR}} \max\left\{|p_k^{(i)}|, p_{\mathrm{fl}}\right\}, \quad \forall k \in \mathcal{I}_p. \end{aligned} \tag{48}$$

Here, $\mathcal{I}_p$ contains the physical powers. The implementation uses $\Delta_{\mathrm{TR}} = 0.35$ and `fmincon` SQP. Because the positive-part CVaR constraints remain nonlinear, (48) is not claimed to be convex.

An SQP point is accepted only if the original CVaRs and power constraints are feasible and the original utility does not decrease, otherwise the incumbent is restored. No Young-product majorization or optimized rate variable is used.

Matched Component Baselines

All matched variants share topology and scenario seeds, the directed NOMA/OMA mode universe, resource and threshold settings, final joint recertification, and held-out tests. Scenario-robust HPA without PDG certification

removes only the candidate confidence filter but retains MILP, CVaR-SQP, OMA repair, and joint recertification.PDG scheduling with deterministic power retains certified PDG admission and matching but replaces SQP with channel-inversion LED powers and a grid-selected layer split. These variants isolate graph admission and risk-constrained power adaptation. Monolithic sample-average and distributionally robust formulations change the risk model and are not treated as matched ablations.

For the bounded worst-case HPA endpoint, let $\mathbf{z}_e$ collect the standardized Gaussian primitive position, orientation, and residual-CSI variables affecting mode e. Its per-mode uncertainty set is

$$\mathcal{U}_e^{\mathrm{wc}} = \left\{\mathbf{z}_e : \|\mathbf{z}_e\|_2 \leq \sqrt{\chi^2_{d_e, 1-\epsilon_{\mathrm{pair}}}}\right\},$$

where d_e is the primitive-vector dimension. Nonlinear geometry, FOV, and shared-user correlation are recomputed for each vector. The ellipsoid has the same fitted-Gaussian probability content as the pair target, but enforcing every point is more conservative than a chance constraint. The network version uses $1 - \epsilon_{\text{net}}$ content.

4.5. Final-Assignment Joint Recertification and Reliability Repair

Candidate certificates need not compose after matching and power redistribution. The selected assignment, resource activity, block fractions, and powers are therefore frozen after MILP–SQP, and all links are recomputed on a fresh joint-certification set.

Let $\mathcal{M}(\chi)$ denote the modes selected by the complete assignment χ. For joint-certification scenario n, define the success event of selected mode m as

$$\mathcal{G}_m^{(n)}(\chi, \mathbf{p}) = \begin{cases} \mathcal{O}_e^{(n)}(\chi, \mathbf{p}) \cap \mathcal{A}_e^{(n)}(\chi, \mathbf{p}), & m = e \text{ is NOMA}, \\ \mathcal{A}_o^{(n)}(\chi, \mathbf{p}), & m = o \text{ is OMA}. \end{cases} \tag{49}$$

Every SINR in (49) uses the frozen final activity, powers, and interference. The same-realization network event is

$$\mathcal{A}_{\text{net}}^{(n)}(\chi, \mathbf{p}) = \bigcap_{m \in \mathcal{M}(\chi)} \mathcal{G}_m^{(n)}(\chi, \mathbf{p}). \tag{50}$$

If K_{net} of S_{joint} independent scenarios satisfy (50), its one-sided Clopper–Pearson lower confidence bound is

$$p_{\text{net}}^{\text{LCB}} = I_{\delta_{\text{J,alloc}}}^{-1}\big(K_{\text{net}}, S_{\text{joint}} - K_{\text{net}} + 1\big), \tag{51}$$

with $p_{\text{net}}^{\text{LCB}} = 0$ for $K_{\text{net}} = 0$. If at most T_{joint} complete assignments are tested for each of the N_{cfg} prespecified operating configurations, the per-attempt budget is

$$\delta_{\text{J,alloc}} = \frac{\delta_{\text{J}}}{N_{\text{cfg}} T_{\text{joint}}}. \tag{52}$$

The final assignment is certified only if

$$p_{\text{net}}^{\text{LCB}} \geq 1 - \epsilon_{\text{net}}, \tag{53}$$

where ϵ_{net} is the network-outage tolerance. This is the physical probability that every selected mode succeeds in the same realization, not a per-mode target or confidence budget. The certificate applies to $\widehat{\mathbb{P}}_{\hat{\theta}}$; $\mathbb{P}_\star$ is examined only through held-out mismatch tests.

To distinguish the three probability statements used in the revised framework, let the physical outage probability of one selected mode be

$$P_m^{\text{out}}(\chi, \mathbf{p}) = \Pr\big[(\mathcal{G}_m(\chi, \mathbf{p}))^{\text{c}}\big], \qquad m \in \mathcal{M}(\chi), \tag{54}$$

where the probability is taken over a new channel realization from the specified uncertainty generator after the complete assignment and power vector have been fixed. The network-wide outage probability is instead

$$P_{\text{net}}^{\text{out}}(\chi, \mathbf{p}) = \Pr\left[\bigcup_{m \in \mathcal{M}(\chi)} (\mathcal{G}_m(\chi, \mathbf{p}))^{\text{c}}\right] = 1 - \Pr[\mathcal{A}_{\text{net}}(\chi, \mathbf{p})]. \tag{55}$$

Consequently,

$$\max_{m \in \mathcal{M}(\chi)} P_m^{\text{out}} \leq P_{\text{net}}^{\text{out}} \leq \min\left\{1, \sum_{m \in \mathcal{M}(\chi)} P_m^{\text{out}}\right\}. \tag{56}$$

Thus, certifying every selected mode at a common marginal tolerance does not by itself certify the probability that all selected modes succeed simultaneously.

$1 - \delta_{\mathrm{J,alloc}}$ is coverage over repeated certification sets, whereas $1 - \epsilon_{\mathrm{net}}$ is the required physical network-success probability. The preallocated budget covers all repair attempts; (53) controls channel outage.

If (53) fails, selected NOMA modes are ranked by their contribution to failed network scenarios. The highest-risk mode is removed, a no-good cut excludes the failed assignment, and MILP–SQP is rerun with certified OMA edges. Each repair uses a fresh joint-certification set. Exhausting T_{joint} attempts yields an unresolved computational outcome, not a certificate or proof of mathematical infeasibility. Repair does not protect against an unknown generator error.

4.6. Alternating PDG-HPA Procedure

The procedure starts from a certified all-OMA incumbent; failure is reported without relaxing QoS or reliability.

Figure 3 separates candidate admission, MILP scheduling, fixed-mode power adaptation, and final joint recertification.

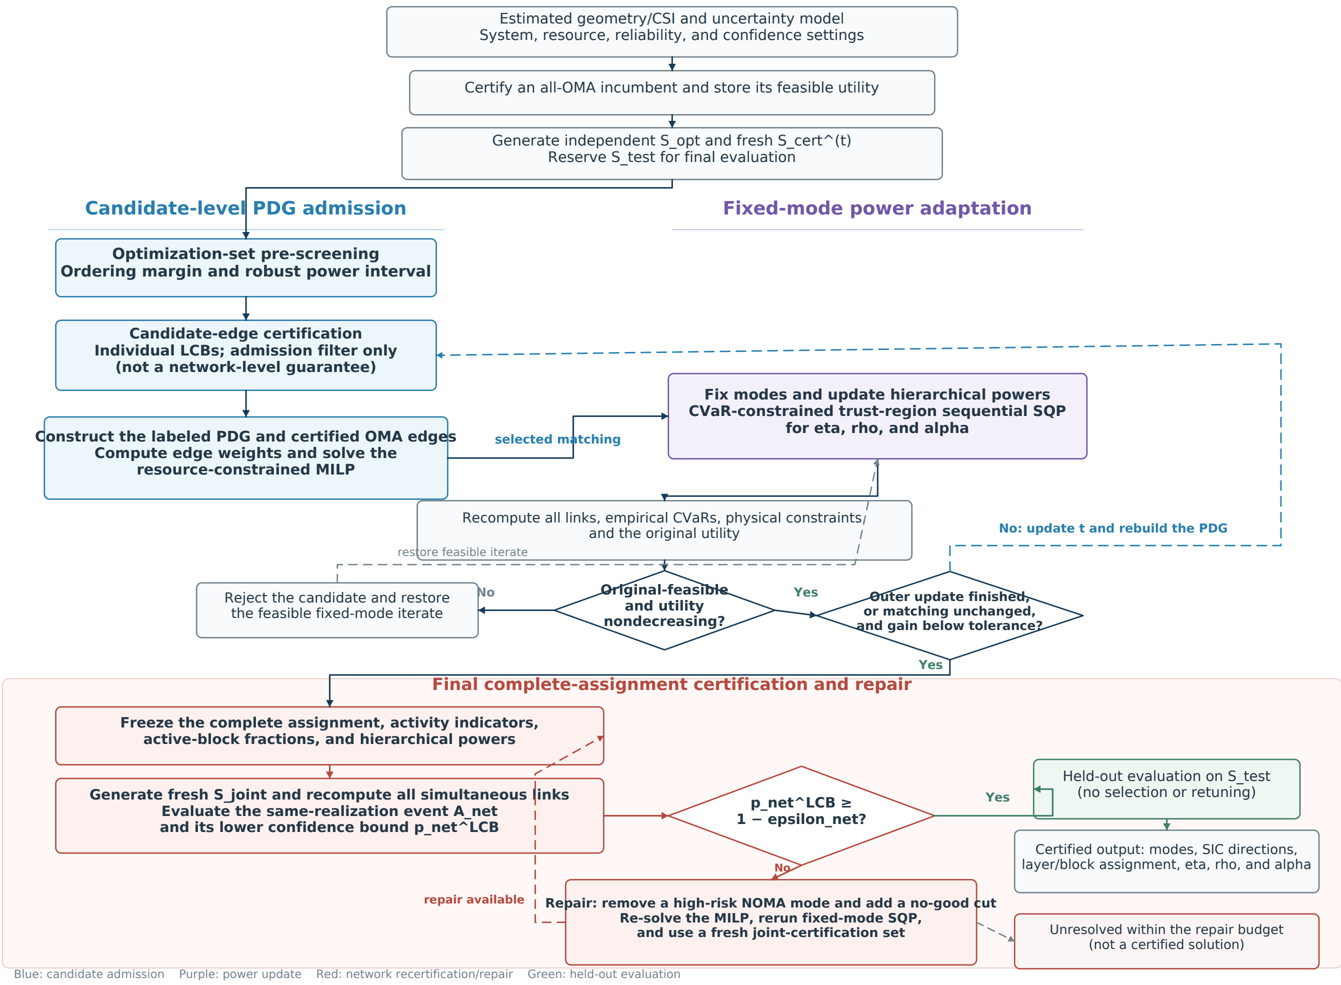

Figure 3. Alternating PDG-HPA procedure. Candidate modes are independently admitted, matched, and power-adapted. After activity, power, and interference terms are frozen, the complete assignment is jointly recertified and repaired, if necessary, before held-out testing.

Algorithm 1 states the corresponding implementation.

Algorithm 1 Alternating PDG-HPA Receiver–Scheduler Co-Design

Require: Estimated state and fitted generator; thresholds, budgets, scenario sizes, and iteration limits
Ensure: Certified NOMA/OMA modes, SIC directions, resources, and powers
Generate independent scenario families and preallocate confidence budgets; certify an all-OMA incumbent or terminate as infeasible
for $t = 1, \dots, T_{\text{out}}$ **do**
  Screen and independently certify directed candidates; build the PDG and solve the resource-constrained MILP
  if no feasible matching or fixed-mode start exists **then**
    Restore the incumbent and continue
  end if
  Freeze resource activity and run at most T_{in} trust-region CVaR–SQP updates
  Accept only original-feasible, nondecreasing updates; otherwise restore the incumbent
end for
for $j = 1, \dots, T_{\text{joint}}$ **do**
  Freeze the complete assignment and jointly recertify it on a fresh scenario set
  if $p_{\text{net}}^{\text{LCB}} \geq 1 - \epsilon_{\text{net}}$ **then**
    Accept and break
  else
    Exclude the highest-risk NOMA mode, add a no-good cut, and rerun MILP–SQP with certified OMA repair
  end if
end for
if joint recertification has not succeeded **then**
  Return an unresolved outcome
end if
Evaluate once on the held-out and prescribed mismatch sets
return modes, resources, SIC directions, and powers

Held-out scenarios affect neither design nor stopping. Matched-generator tests measure nominal generalization; mismatch sets test generator form.

4.7. Complexity and Termination Properties

Before screening, an association-independent bound on the number of directed candidates is given by (32). Fixed association uses one LED label per user; joint association can increase the candidate bound to $LM_{\max}$ and is used only for small references. Generating S correlated channel scenarios requires $\mathcal{O}(SLU)$ Lambertian channel evaluations. Preliminary NOMA grid screening requires

$$\mathcal{O}\big(M_{\text{N}}^{\max}|\mathcal{A}_\alpha|S_{\text{opt}}\big) \tag{57}$$

receiver-state evaluations, while independent certification requires $\mathcal{O}(M_{\max}S_{\text{cert}})$ evaluations. The resource-constrained matching is a 0–1 MILP and is NP-hard in general. For a fixed matching, one SQP function or constraint evaluation scales linearly with $S_{\text{opt}}|\mathcal{M}_{\text{sel}}|$; its wall-clock cost also depends on the number of SQP iterations and function evaluations. Section 5 therefore reports measured runtime rather than only an asymptotic expression.

The audit records complexity and outcome for every attempted topology. At outer round t, let $M_{\text{cand}}^{(t)}$ and $M_{\text{cert}}^{(t)}$ denote the numbers of enumerated and retained NOMA/OMA modes. The resource MILP then contains

$$N_{\text{bin}}^{(t)} = M_{\text{cert}}^{(t)}, \qquad N_{\text{eq}}^{(t)} = U, \tag{58}$$

binary variables and U cover equalities; inequalities include $L|\mathcal{Q}|B$ capacities and repair cuts. SQP evaluations, memory, termination cause, and outcome-conditioned runtime are also recorded. Scenario evaluation is parallelizable, but one MILP–SQP trajectory is serial. The implementation is a slow-timescale prototype.

4.7.1. Algorithmic Sensitivity and Small-Network Reference

The sensitivity vector is

$$\boldsymbol{\vartheta}_{\text{alg}} = \big(\Delta\alpha, \mathbf{c}_\alpha, \boldsymbol{\lambda}, p_{\text{pre}}, \varepsilon_{\text{CVaR}}, \Delta_{\text{TR}}, T_{\text{out}}, T_{\text{in}}, \mathbf{p}^{(0)}\big),$$

covering grid, screening, penalty, tolerance, trust-region, iteration, and initialization settings. One-at-a-time sweeps report their effect.

For $U \leq 8$, every certified resource-feasible discrete selection is enumerated over common power starts; for $U = 10$, the exact MILP uses 100 prespecified starts. This is a strengthened reference, not a certified global continuous optimum.

4.7.2. Monotonicity of Accepted Inner Updates

For fixed modes, acceptance requires original CVaR and power feasibility and $F^{(i+1)} \geq F^{(i)}$. Finite power, positive noise, and finite spectral-use factors bound the rates and utility; accepted fixed-mode utilities are therefore nondecreasing and bounded.

This statement does not cover PDG reconstruction, SQP stationarity, or local/global optimality of (25). Finite iteration limits ensure termination.

5. Numerical Evaluation and Discussion

5.1. Simulation Setup and Evaluation Protocol

The reference room is $5 \times 5 \times 3$ m^3, with a 2×2 ceiling-LED array and a receiver plane at 0.85 m. Unless swept, the half-power angle and PD FOV are 60° and 70°. Each realization recomputes every desired and interfering gain through (9), including FOV loss; overlap therefore follows from geometry rather than a fitted coefficient.

Besides the controlled-balanced case, we test three unfiltered families: uniform natural placement on $[0.30, 4.70]^2$ m^2; an overlap-heavy case with 70% of users within 0.35 m of a Voronoi boundary; and an overloaded case with 50% of users in one quadrant. PD tilt and azimuth are uniform on $[0°, 8°]$ and $[0, 2\pi)$, and all methods reuse the same primitive seeds.

The uncertainty-severity experiment in the uncertainty-severity experiment is designed as a diagnostic sensitivity study rather than as the unfiltered topology-stress test the unfiltered topology-stress test. In this diagnostic, NOMA candidates remain labeled by the nominal serving LED, whereas the explicit resource-accounted OMA fallback may use any nominally visible LED. This association-relaxed OMA recovery prevents local nominal load imbalance from dominating the uncertainty-severity sweep before the effects of position, orientation, and CSI uncertainty can be observed. In contrast, the main topology-stress experiment in retains the fixed nominal association of Algorithm 1 and applies no such association relaxation to the unfiltered topology families. Accordingly, the uncertainty-sensitivity should be interpreted as an uncertainty-sensitivity diagnostic, not as evidence that the fixed-association main algorithm remains feasible for arbitrary unbalanced deployments.

Let $G^{\text{ref}}_{l,u} = P^{\text{ref}}_l (R_{\text{PD}}\widehat{h}_{l,u})^2$, and let $G^{\text{ref}}_{(1),u} \geq G^{\text{ref}}_{(2),u}$ denote the two largest nominal received-power terms for user u. The explicit two-cell overlap score used in the topology-association is

$$\Omega_u = \frac{G^{\text{ref}}_{(2),u}}{G^{\text{ref}}_{(1),u} + G^{\text{ref}}_{(2),u}}, \qquad \overline{\Omega} = \frac{1}{U}\sum_{u\in\mathcal{U}} \Omega_u. \tag{59}$$

a larger $\overline{\Omega}$ means that the two strongest LEDs contribute more equally. The half-power-angle sweep changes both this score and the desired gain, so it is not interpreted as a pure-overlap experiment.

Association errors are measured by comparing the nominal serving LED with the scenario-wise strongest LED. The common physical parameters and default settings used throughout the numerical evaluation are summarized in Table 2.

Table 2. Common physical parameters and default settings.

Parameter	Value	Parameter	Value
Room size	$5 \times 5 \times 3$ m^3	Receiver-plane height	0.85 m
Number of LEDs	4 (2×2)	PD active area	10^{-4} m^2
Optical-filter gain	1	Refractive index	1.5
PD responsivity	0.53 A/W	Nominal PD tilt	0–8°
Half-power angle	60°, unless swept	PD FOV	70°
Position-error standard deviation	0.20 m	Orientation-error standard deviation	10°
Relative CSI-error standard deviation	0.06	Noise power per layer	2×10^{-14}
Residual SIC factor, ξ	0.02	ACO-to-DCO leakage, ζ	0.03
Normalized AC power budget	4	Per-LED power interval	$[0.32, 1.60]$
Minimum layer power	0.04	Initial ACO fraction	0.45
Spectral-use factors	$\nu_A = \nu_D = 0.25$	Pair-outage tolerance	$\epsilon_{\text{pair}} = 0.05$
Weak-user threshold	0.10 bit/s/Hz	SIC-user threshold	0.08 bit/s/Hz
OMA threshold	0.06 bit/s/Hz	Ordering tolerance	$\epsilon_{\text{ord}} = 0.05$
OMA-outage tolerance	$\epsilon_O = 0.05$	Confidence budget	$\delta_{\text{FW}} = 0.05$

Unless otherwise stated, the familywise confidence budget is split as $\delta_E = 0.04$ for candidate-level certification claims and $\delta_J = 0.01$ for final complete-assignment recertification, so that $\delta_{\text{FW}} = \delta_E + \delta_J = 0.05$.

Perturbations are shared across the LED links of each user, whereas optimization, candidate certification, final recertification, and testing use independent scenarios. Complete runs use two blocks per layer, two outer rounds, at most 15 SQP updates, an α grid of $\{0.50, 0.52, \ldots, 0.94\}$, and a trust-region radius of 0.35. MATLAB R2023b `fmincon` (SQP) and `intlinprog` ran on an Intel Core i5-12500H with 16 GB RAM. Sampling is listed in Table 3.

Table 3. Monte Carlo protocol for the reported experiments.

Experiment	Users	Topologies	S_{opt}	S_{cert}	S_{joint}	S_{test}
Reliability calibration (Figure 4)	12	30	500	200–2000	200–2000	2×10^4
Uncertainty-severity diagnostic (Figure 5)	16	30	300	1000	2000 per attempt	10^4
Generator mismatch (Figure 6)	12	30	300	2000	2000 per attempt	10^4
Topology/association stress (Figure 7)	16	30 per family	300	2000 per round	2000 per attempt	10^4
Joint-association reference (Figure 7d)	6, 8, 10	10 per size	300	2000 per round	2000 per attempt	10^4
Tradeoff and baselines (Figure 8)	16	30 per point	300	2000 per round	2000 per attempt	10^4
Matched components (Figure 9a)	16	30	300	2000 per round	2000 per attempt	10^4
Reference, sensitivity, and scaling (Figure 9b–f)	6–32	10 per size	100–150	500–1000	500–1000	5000/–

Here, S_{joint} denotes the number of fresh scenarios used in each final complete-assignment recertification attempt. Candidate certification uses fresh S_{cert} scenarios in each certification round, whereas each final recertification attempt uses an independent S_{joint} set.

Candidate false certification measures confidence-bound coverage. For a frozen assignment, selected-NOMA pair outage is

$$\widehat{P}_{\text{pair}}^{\text{out}} = \frac{1}{S_{\text{test}} |\mathcal{M}_{\text{N}}|} \sum_{n} \sum_{e \in \mathcal{M}_{\text{N}}} \mathbf{1}\{(\mathcal{A}_e^{(n)})^{\text{c}}\}, \tag{60}$$

whereas network-wide outage is

$$\widehat{P}_{\text{net}}^{\text{out}} = \frac{1}{S_{\text{test}}} \sum_{n} \mathbf{1}\left\{ \left[\bigcap_{m \in \mathcal{M}(\chi)} \mathcal{G}_m^{(n)} \right]^{\text{c}} \right\}. \tag{61}$$

the network metric also covers all-OMA assignments and is evaluated only after the assignment and powers are frozen.

Because feasibility differs by method, we also report the all-method common set $\mathcal{I}_\cap$, paired differences, feasibility probability, and the unconditional penalized utility

$$\widehat{F}_m^{\text{uncond}} = \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} \left[\phi_{m,i} F_{m,i}^{\text{test}} + (1 - \phi_{m,i}) F_{\text{fail}} \right], \quad F_{\text{fail}} = \sum_{u} \omega_u \log\left(\frac{R_{\text{fl}}}{R_0} \right). \tag{62}$$

every infeasible instance therefore receives the same declared penalty.

5.2. Out-of-Sample Reliability Calibration

Figure 4 tests candidate-edge admission and final joint recertification. The latter freezes activity, block fractions, powers, and interference before held-out testing.

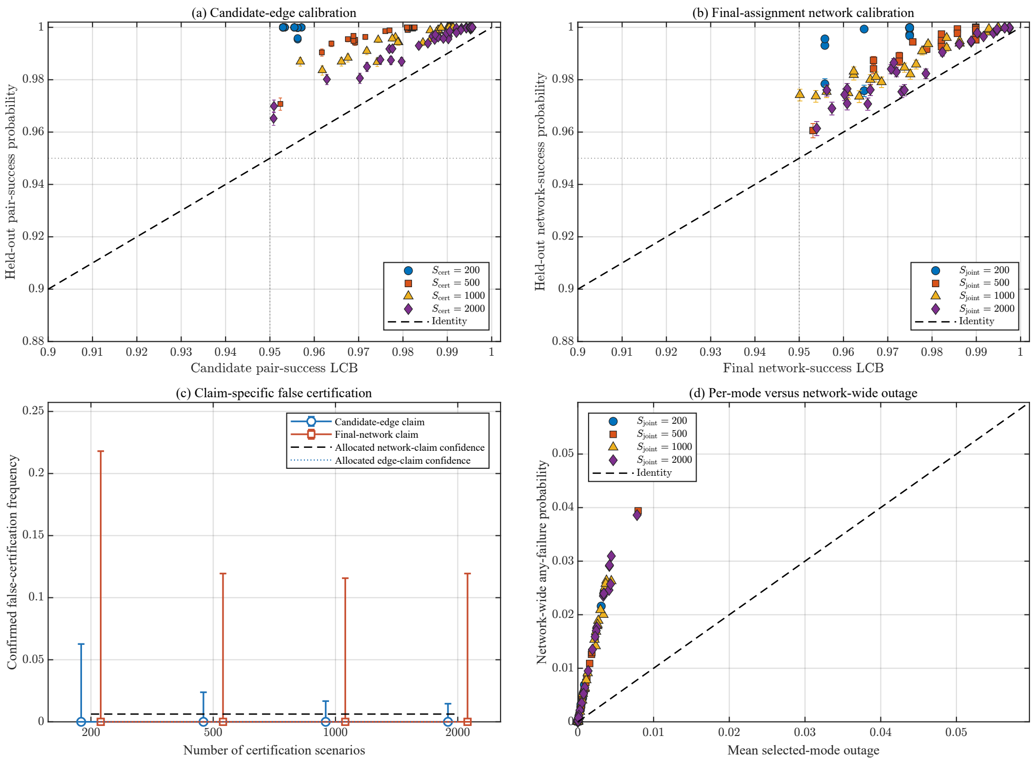

Figure 4. Out-of-sample calibration at the candidate and frozen-assignment levels. Panels show held-out success versus its LCB, false-certification frequency, and selected-mode versus same-realization network outage. Candidate certificates are admission filters; the final claim follows joint recertification.

Held-out success is consistent with the one-sided bounds, with no confirmed false certification. Panel (d) shows why selected-mode outage cannot substitute for network-wide outage.

5.3. Uncertainty Severity and Generator Misspecification

Figure 5 isolates the sensitivity of the reliability-aware design to increasing geometry and CSI uncertainty over 30 natural-placement $U = 16$ topology realizations. As specified in Section 5.1, this diagnostic retains serving-LED-labeled NOMA candidates while allowing the resource-accounted OMA fallback to use any nominally visible LED. It is therefore distinct from the unfiltered fixed-association topology-stress experiment of Figure 7. Four geometry-error profiles are considered—nominal geometry, position error only, orientation error only, and combined position–orientation error—while the relative CSI-error standard deviation σ_{est} is swept from 0 to 0.12. At every operating point, candidate admission, resource-constrained scheduling, power adaptation, frozen-assignment recertification, and held-out testing are rerun under the prescribed uncertainty profile. Unresolved outcomes remain in the unconditional statistics rather than being discarded.

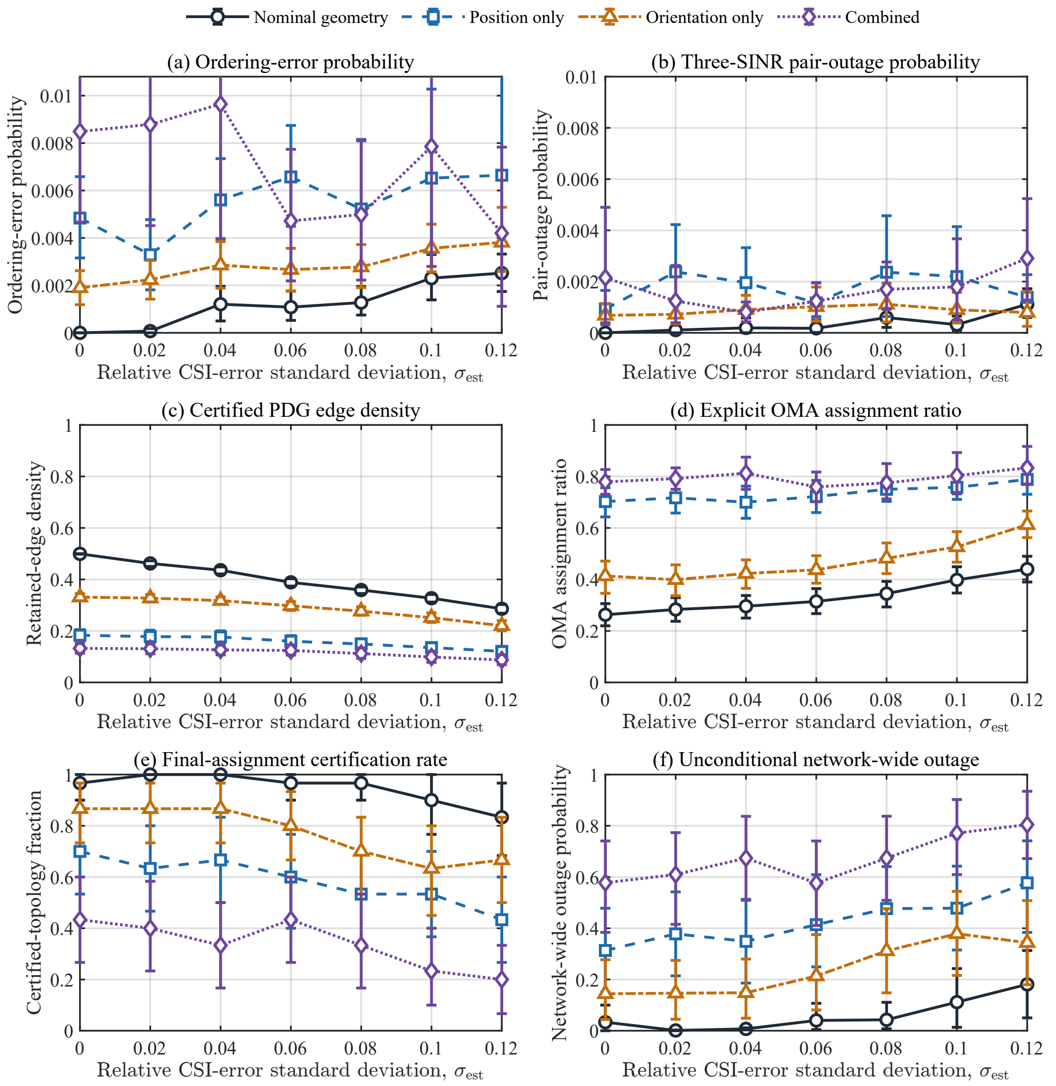

Figure 5. Association-relaxed uncertainty-severity diagnostic over 30 natural-placement $U = 16$ topology realizations. NOMA candidates remain labeled by the nominal serving LED, whereas the explicit resource-accounted OMA fallback may use any nominally visible LED. (**a**) Conditional ordering-error probability of the selected NOMA modes; (**b**) conditional three-SINR pair-outage probability; (**c**) certified PDG edge density; (**d**) explicit OMA assignment ratio; (**e**) final-assignment certification rate; and (**f**) unconditional network-wide outage probability. An unresolved instance contributes unit network-wide outage, whereas the conditional NOMA metrics in (**a**,**b**) are evaluated only for certified assignments containing selected NOMA modes. The ordering-error, pair-outage, and network-wide outage targets are 0.05. Error bars denote 95% topology-level bootstrap intervals. The unfiltered fixed-association topology stress test is reported separately in Figure 7.

The selected NOMA modes remain individually reliable even as the uncertainty level increases. Across all operating points for which selected-NOMA statistics are defined, the maximum mean ordering-error probability is 9.64×10^{-3} and the maximum mean three-SINR pair-outage probability is 2.91×10^{-3}, both well below the corresponding 0.05 targets. These small conditional outage values should not, however, be interpreted as network-wide robustness by themselves. Increasingly uncertain candidate modes are progressively removed by PDG certification or replaced by explicit OMA service, while an unresolved complete assignment remains penalized in the unconditional network metric.

This behavior is visible in Figure 5c,d. Under nominal geometry, increasing σ_{est} from 0 to 0.12 reduces the mean certified PDG edge density from 0.499 to 0.286, while the mean OMA assignment ratio among feasible assignments increases from 0.263 to 0.440. Over the same sweep, the final-assignment certification rate decreases from 0.967 to 0.833, and the unconditional network-wide outage increases from 0.033 to 0.181. Thus, even CSI uncertainty alone progressively reduces the set of certifiable NOMA opportunities and shifts the diagnostic design toward more conservative orthogonal service.

Geometry uncertainty produces a substantially stronger effect. At $\sigma_{\text{est}} = 0$, the final-certification rates are 0.700, 0.867, and 0.433 for the position-only, orientation-only, and combined profiles, respectively, with corresponding unconditional network-wide outages of 0.313, 0.144, and 0.578. At $\sigma_{\text{est}} = 0.12$, these certification rates become 0.433, 0.667, and 0.200, while the network-wide outages increase to 0.577, 0.344, and 0.804, respectively. The combined-error case is therefore the most demanding uncertainty regime in this diagnostic.

For the combined profile, the certified-edge density decreases from 0.132 at $\sigma_{\text{est}} = 0$ to 0.087 at $\sigma_{\text{est}} = 0.12$, while the OMA ratio among the remaining feasible assignments increases from 0.779 to 0.833. Local non-monotonic variations arise because candidate admission, mode selection, and power allocation are recomputed as the uncertainty setting changes; topology-level heterogeneity and finite Monte Carlo variation also contribute. The overall trend is nevertheless clear: increasing geometry and CSI uncertainty sparsifies the certified PDG, increases reliance on OMA, lowers complete-assignment certification probability, and raises unconditional network outage.

Figures 5 and 7 therefore answer complementary rather than interchangeable questions. The former isolates uncertainty severity while permitting association-relaxed OMA recovery so that local load imbalance does not by itself terminate the sensitivity experiment. The latter applies the main fixed-association procedure to unfiltered topology families and shows that topology imbalance can prevent construction of the confidence-qualified all-OMA initialization before the outer PDG–power iterations are entered. Accordingly, the favorable conditional selected-NOMA reliability in Figure 5 is not used to claim robustness of the fixed-association algorithm to arbitrary unbalanced networks.

Figure 6 examines a complementary question: sensitivity to misspecification of the uncertainty generator. Unlike the uncertainty-severity experiment in Figure 5, the design model is fixed to a fitted Gaussian uncertainty law, while only the held-out deployment generator is changed. Consequently, samples drawn from the mismatched generator cannot influence candidate selection, scheduling, power adaptation, or reliability repair.

Under the matched generator, the network-wide outage probability is 0.0088. Position bias, orientation bias, Student-$t(3)$ errors, and a Gaussian-mixture generator increase the outage to 0.0281, 0.0155, 0.0203, and 0.0317, respectively. These values remain below the nominal 0.05 target, but they show that the held-out reliability depends on the fidelity of the assumed uncertainty law.

Variance underestimation is more consequential. When the true standard deviation equals the assumed value, the network-wide outage is 0.0089, whereas increasing the true-to-assumed standard-deviation ratio to 1.5 and 2 raises it to 0.0697 and 0.1709, respectively. Thus, an uncertainty model that is structurally reasonable but underestimates dispersion can violate the nominal reliability target.

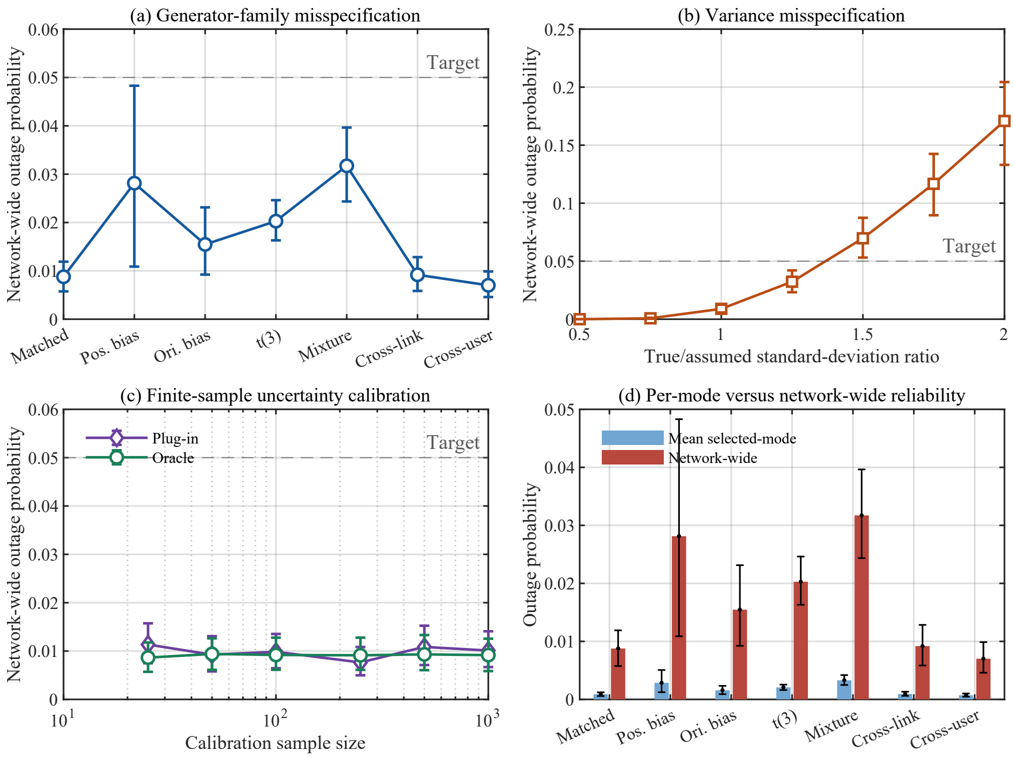

Figure 6. Generator-misspecification sensitivity over 30 unbalanced $U = 12$ topologies. The experiment evaluates alternative error laws, variance mismatch, finite-data calibration versus oracle uncertainty parameters, and selected-mode versus network-wide outage. The dashed line denotes the nominal outage target of 0.05, and error bars denote 95% topology-level bootstrap intervals.

Finite calibration data introduce an additional estimation effect. With 25 calibration samples, the plug-in design gives a network-wide outage of 0.0113, compared with 0.0086 for the corresponding oracle-parameter design. The gap narrows as the calibration sample size increases. These results indicate that the certification guarantee is conditional on the fitted uncertainty generator and its estimated parameters.

Accordingly, Figure 6 should be interpreted as a sensitivity analysis rather than evidence of distribution-free robustness. Together, Figures 5 and 6 separate two distinct effects: the former varies the severity of the modeled uncertainty, whereas the latter evaluates the consequences of using an imperfect uncertainty generator.

5.4. Topology Imbalance and Association Robustness

Figure 7 evaluates the complete algorithm under controlled-balanced, naturally unbalanced, boundary/overlap-heavy, and single-cell-overloaded deployments with $U = 16$. The experiment reports the nominal LED-load distribution, optical-overlap score, unconditional effective sum rate, feasibility probability, penalized network-wide outage, and serving-LED association error conditioned on at least one visible LED. All topology families use the same physical parameters, reliability thresholds, resource budgets, and scenario protocol. Consistently with Algorithm 1, each run first constructs a confidence-qualified all-OMA incumbent. Only runs for which this resource-constrained initialization exists enter the two prescribed PDG–power outer rounds, followed by frozen-assignment joint recertification and, when necessary, reliability repair. A separate small-network reference compares fixed nominal association with joint association–scheduling using the same paired topology–scenario instances and resource budgets.

Unlike the uncertainty-severity diagnostic in Figure 5, no association relaxation is used for the topology-family comparison in this subsection. The natural, overlap-heavy, and overloaded families are evaluated unfiltered under the fixed nominal association of the main algorithm.

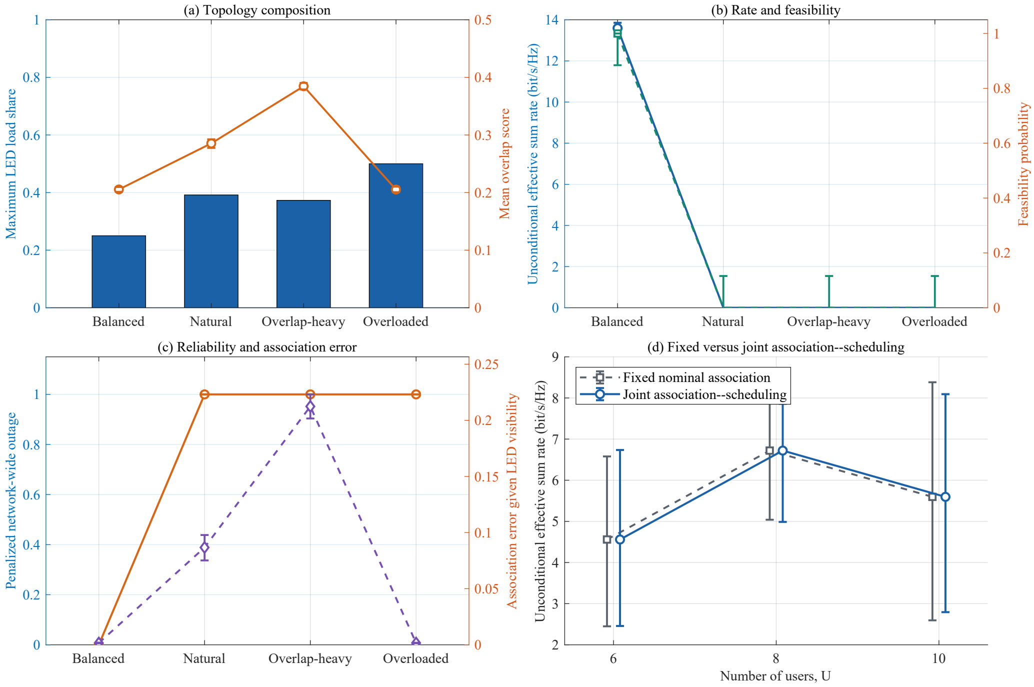

Figure 7. Topology-imbalance and association stress tests. (**a**) Maximum nominal LED-load share and mean optical-overlap score, shown by the blue bars and orange solid line, respectively; (**b**) unconditional effective sum rate and feasibility probability, shown by the blue solid and green dashed lines, respectively; (**c**) penalized network-wide outage and serving-LED association-error probability conditioned on at least one visible LED, shown by the purple dashed and orange solid lines, respectively; and (**d**) raw unconditional effective sum rates obtained independently with fixed nominal association and joint association–scheduling in the small-network reference, shown by the gray dashed and blue solid lines, respectively. The coincident curves in panel (**b**) for the unbalanced topology families reflect the corresponding zero effective rate and zero feasibility probability and are not a rendering artifact. An unresolved instance contributes zero effective rate and unit penalized network-wide outage. No fixed-incumbent replacement is applied in panel (**d**). Bootstrap intervals are shown for overlap, rate, outage, association error, and the small-network rates, whereas feasibility intervals are exact 95% binomial intervals. The maximum-load-share bars show topology-family means. Each topology family contains 30 realizations, and each small-network operating point contains 10 paired realizations.

All 30 controlled-balanced instances admit a confidence-qualified all-OMA initialization and subsequently return a jointly recertified complete assignment. Their mean unconditional effective sum rate is 13.592 bit/s/Hz, with a penalized network-wide outage of 2.0×10^{-5}. In contrast, none of the 30 naturally unbalanced, 30 boundary/overlap-heavy, or 30 overloaded-cell deployments admits a confidence-qualified all-OMA edge cover under the fixed nominal association and the prescribed per-LED resource limits. These 90 instances therefore terminate before entering the outer PDG–power iterations. Under the unconditional accounting rule, an unresolved instance contributes zero effective rate and unit penalized network-wide outage. Consequently, the latter three topology families have zero unconditional effective sum rate, zero feasibility probability, and unit penalized outage. The unit values in Figure 7b,c are therefore declared infeasibility penalties rather than measured decoding outages of feasible assignments.

The topology audit further identifies the local resource bottleneck. With two ACO/DCO layers and two resource blocks, each LED provides four layer–block slots.

The mean maximum LED-load shares are 0.250, 0.392, 0.373, and 0.500 for the controlled-balanced, naturally unbalanced, overlap-heavy, and overloaded-cell families, respectively. Relative to the four-slot per-LED capacity, the corresponding mean excess loads of the most heavily loaded LED are 0, 2.267, 1.967, and 4.000 users. The controlled-balanced deployments therefore exactly match the nominal per-LED OMA capacity, whereas the three unbalanced families introduce local load concentrations that prevent construction of the required fixed-association all-OMA initialization.

The optical and association statistics further distinguish the different sources of topology stress. The mean overlap scores are 0.205, 0.285, 0.385, and 0.205 for the controlled-balanced, naturally unbalanced, overlap-heavy, and overloaded-cell families, respectively. Their serving-LED association-error probabilities, conditioned on at least one visible LED, are 0.00160, 0.08670, 0.21237, and 0.00159, respectively. The overlap-heavy family therefore exhibits the largest uncertainty-induced strongest-LED reversal probability. By contrast, the overloaded-cell family has almost the same association-error probability as the controlled-balanced deployment despite its complete loss of feasibility. Its failure is therefore primarily attributable to nominal per-LED resource concentration rather than uncertainty-induced association reversal.

Figure 7d provides a separate small-network reference comparing fixed nominal association with joint LED association–scheduling on the same paired topology–scenario instances. Unlike the incumbent-protected comparison, the two curves report the raw unconditional outcomes of the two designs; the joint solution is not replaced by the fixed-association solution when it is worse or unresolved. For $U = 6$, both designs achieve an unconditional effective sum rate of 4.559 bit/s/Hz with feasibility 0.7. For $U = 8$, both achieve 6.721 bit/s/Hz with feasibility 0.9, whereas for $U = 10$, both achieve 5.596 bit/s/Hz with feasibility 0.6. The paired rate difference is zero for all 30 reference instances, yielding a 95% bootstrap interval of $[0, 0]$ at each tested user size.

The enlarged association candidate set therefore provides no measurable gain in this small-network reference under the adopted confidence allocation, reliability thresholds, and resource model. This negative result is important for interpreting the topology stress test: permitting additional LED labels alone does not guarantee recovery from local capacity imbalance when alternative associations must satisfy the same confidence-qualified decoding requirements. Accordingly, Figure 7 should not be interpreted as demonstrating robustness to arbitrary unbalanced deployments or uniform superiority of joint association–scheduling. Rather, it identifies an applicability boundary of the present all-OMA-initialized PDG–HPA procedure under fixed association in its main operating mode.

Certified OMA fallback provides a reliability-preserving alternative when a NOMA mode cannot be admitted reliably and orthogonal layer–block capacity remains available; however, it is not an unlimited feasibility mechanism. In the $U = 16$ topology-stress experiment, all certified feasible assignments use the resource-accounted OMA fallback. By contrast, under the strongly unbalanced topology families, a confidence-qualified all-OMA edge cover cannot be constructed under the prescribed fixed association and per-LED resource limits. Thus, OMA fallback can replace unreliable NOMA service when sufficient orthogonal resources remain available, but it cannot overcome a hard local resource-capacity bottleneck.

Accordingly, Figure 7 should be interpreted primarily as demonstrating the reliability role and the limits of resource-accounted OMA fallback under topology imbalance, association uncertainty, and per-LED resource constraints. It should not be interpreted as evidence that NOMA multiplexing gains persist under highly unbalanced deployments or that the proposed procedure guarantees universal feasibility. The principal throughput and NOMA-reliability conclusions therefore remain restricted to the balanced, pre-associated operating conditions evaluated in the main performance experiments. Developing scalable joint association–scheduling mechanisms that can enter the NOMA optimization even when a confidence-qualified all-OMA initialization is unavailable remains an important direction for future work.

5.5. Reliability–Efficiency Tradeoff and Matched Baselines

Figure 8 first varies ϵ_{pair} at $\epsilon_{\text{ord}} = 0.05$, and then compares the proposed PDG-HPA with five baselines under common topology seeds, scenario seeds, and resource budgets: fixed PA, GRPA, adapted β-MCGAD, nominal-CSI HPA, and bounded worst-case HPA. The bounded worst-case design provides an uncertainty-set-based robust benchmark. The matched-component study in the following subsection further introduces scenario-robust HPA without PDG certification and PDG scheduling with deterministic power so that the contributions of probabilistic graph admission and risk-constrained power adaptation can be isolated without changing the remaining mode universe, resource accounting, or test scenarios. Angle-dependent rate and outage comparisons use the all-method common-feasible subset, whereas feasibility is evaluated over all attempted instances.

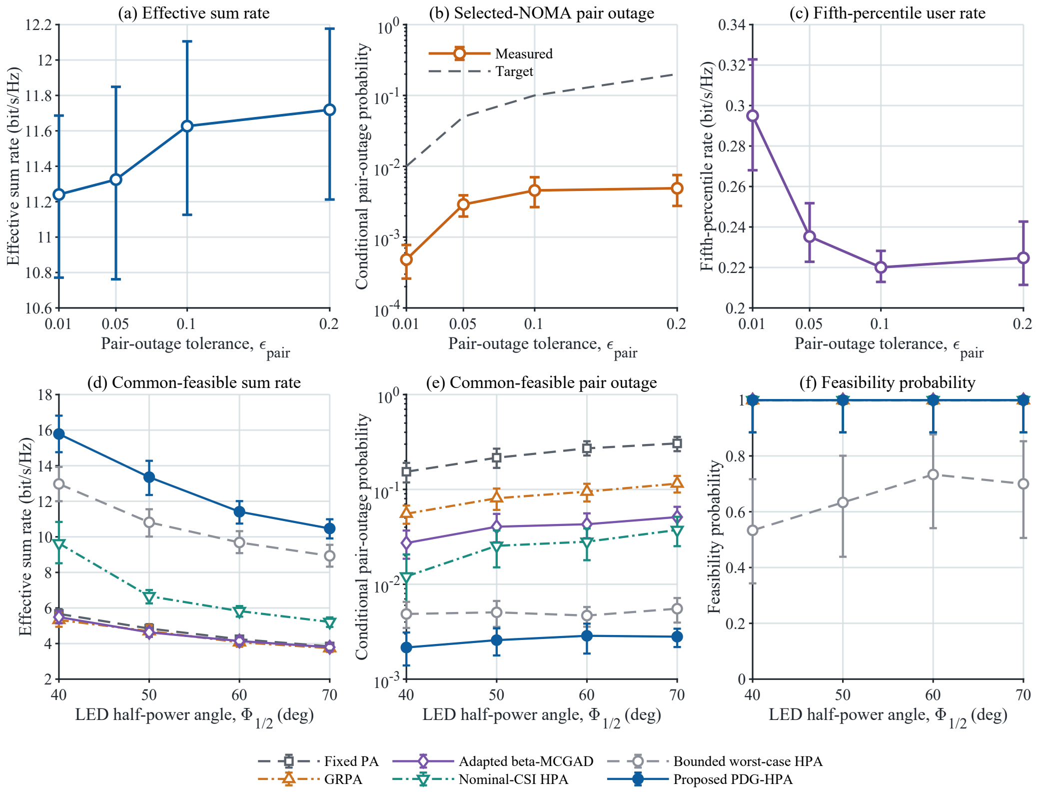

Figure 8. Reliability-efficiency tradeoff and baseline comparison. Panels (**a**–**c**) vary the pair-outage tolerance; (**d**,**e**) compare the proposed PDG-HPA with five baselines on the all-method common-feasible set over LED half-power angle; and (**f**) reports feasibility over all attempted instances. Error bars denote 95% intervals.

Because the half-power angle changes both (59) and the desired optical gain, the corresponding curves quantify their coupled physical effect rather than the effect of cell overlap alone.

Relaxing ϵ_{pair} from 0.01 to 0.20 increases the effective sum rate from 11.241 to 11.719 bit/s/Hz, while the conditional selected-NOMA pair-outage probability rises from 4.83×10^{-4} to 4.89×10^{-3} and remains below the corresponding reliability target throughout the sweep. The fifth-percentile user rate decreases from 0.295 to 0.225 bit/s/Hz, while the recorded OMA assignment ratio decreases from 0.821 to 0.538. Thus, relaxing pair admission increases aggregate resource reuse and throughput, but at the cost of lower-tail user performance.

Across the 120 topology–angle instances, PDG-HPA returns a certified feasible assignment in 118 cases (98.3%), whereas bounded worst-case HPA is feasible in 78 cases (65.0%). Because feasibility differs substantially across methods, rate and outage are compared on the all-method common-feasible subset rather than on method-specific successful runs. Feasibility over all attempts, unconditional penalized performance, and paired differences are reported separately to avoid selection bias.

5.6. Matched Components, Reference Quality, and Computational Limits

The two matched counterfactual baselines retain the same topology–scenario instances, directed NOMA/OMA mode universe, resource budgets, SIC-direction options, OMA accounting, and held-out evaluation protocol as the complete method, while changing only PDG admission or continuous power adaptation. The remaining one-component ablations alter only the specifically named modeling or scheduling mechanism. An infeasible topology–method instance contributes zero effective rate and unit penalized network outage. Figure 9 reports unconditional performance, reference-search quality, one-factor hyperparameter sensitivity, termination outcomes, problem size, and measured runtime.

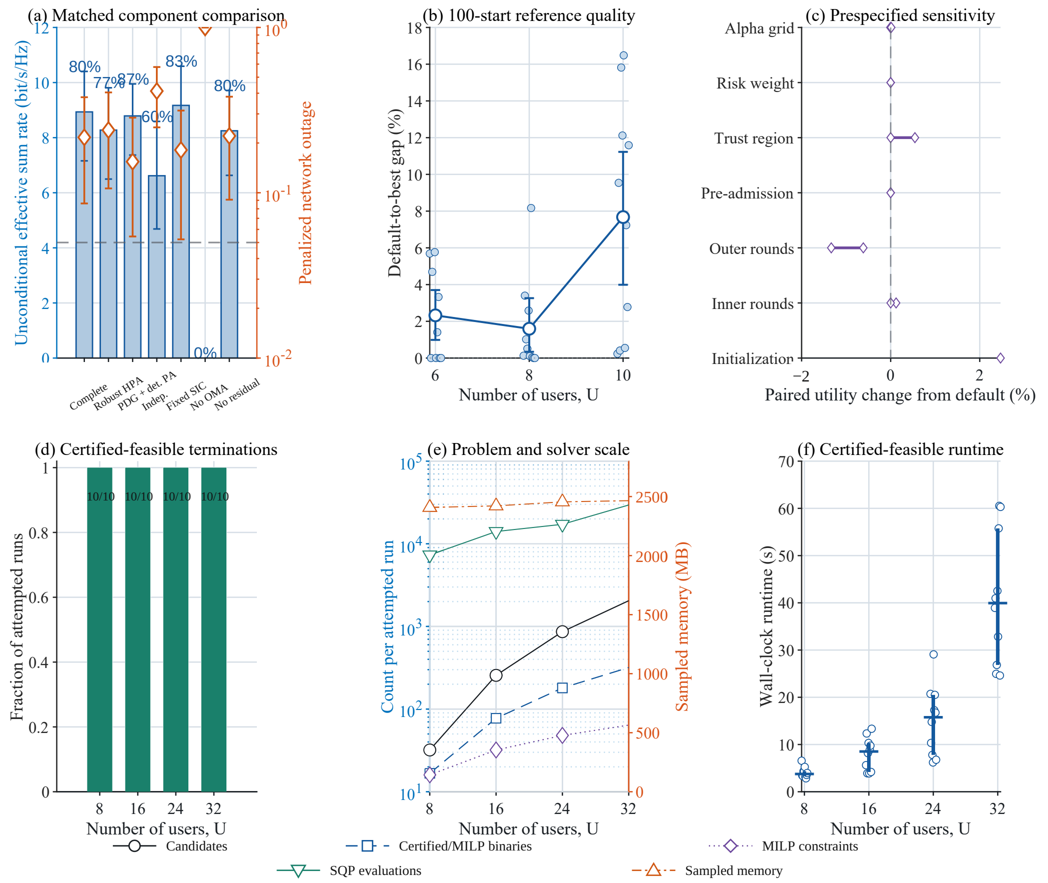

Figure 9. Matched-component, solution-quality, and computational audits. (**a**) Unconditional effective sum rate and penalized network-outage statistic, where an infeasible instance contributes zero rate and unit outage; percentage labels denote feasibility, and the dashed line marks the nominal network-outage target of 0.05. (**b**) Per-topology default-to-best gaps relative to the prespecified 100-start reference. (**c**) One-factor-at-a-time paired utility changes. (**d**) Termination outcomes. (**e**) Candidate and certified-mode/MILP-binary counts, MILP constraints, SQP evaluations, and sampled MATLAB process memory. (**f**) Wall-clock runtime distributions of certified-feasible runs. Error bars in (**a**,**b**) denote 95% topology-level bootstrap intervals.

Across the 30 common topologies in Figure 9a, PDG-HPA achieves an unconditional effective sum rate of 8.934 bit/s/Hz, compared with 8.275 bit/s/Hz for scenario-robust HPA without PDG certification and 8.789 bit/s/Hz for PDG scheduling with deterministic power. These values correspond to gains of 8.0% and 1.65%, respectively. Their feasibility probabilities are 0.800, 0.767, and 0.867, while their penalized network-outage statistics are 0.216, 0.240, and 0.154. Thus, PDG admission improves the unconditional rate relative to the matched no-PDG variant. Risk-constrained power adaptation provides a smaller rate improvement over deterministic power but does not dominate it in feasibility or penalized network outage. Because the latter statistic assigns unit outage to unresolved instances, it summarizes reliability and computational feasibility jointly and should not be interpreted as the decoding-outage probability conditioned on a feasible assignment.

The remaining one-component ablations isolate additional modeling and scheduling mechanisms. The independent-error variant reduces feasibility to 0.600 and increases the penalized outage statistic to 0.412. Fixing the nominal SIC direction gives feasibility and outage values of 0.833 and 0.181, whereas omitting residual-interference modeling gives 0.800 and 0.221. Without certified OMA fallback, none of the 30 attempted instances are feasible. These results show that the observed reliability–efficiency behavior depends jointly on probabilistic admission, risk-constrained power adaptation, uncertainty modeling, SIC-direction selection, and resource-accounted OMA fallback. In particular, the zero feasibility of the no-OMA variant shows that the certified OMA fallback is essential for obtaining complete assignments under the tested configuration.

Figure 9b gives mean default-to-best gaps of 2.32%, 1.60%, and 7.68% for $U = 6, 8, 10$, respectively; the corresponding maximum observed gaps are 5.77%, 8.17%, and 16.48%. Because the reference is the best result among 100 prespecified power starts rather than a certified global continuous optimum, these values quantify initialization sensitivity rather than global optimality. Panel (c) shows generally small one-factor utility changes, although the alternative initialization changes the mean utility by approximately 2.47%. This experiment does not establish robustness to simultaneous hyperparameter changes.

All ten formal scaling runs at each tested size in panel (d) terminate with a certified feasible assignment. No initialization, MILP/SQP solver-limit, or final-recertification failure is observed in these controlled scaling runs. This result should be distinguished from the unbalanced $U = 16$ topology stress tests in Figure 7, where the failure audit attributes all 90 unresolved instances to the absence of a confidence-qualified all-OMA edge cover under fixed association and per-LED resource limits before the outer PDG–power iterations are entered.

For $U = 8, 16, 24, 32$, the mean wall-clock runtimes are 4.08, 8.04, 15.02, and 40.82 s, while the corresponding medians are 3.76, 8.52, 15.77, and 39.94 s. Mean candidate counts increase from 32 to 2048, certified-mode/MILP-binary counts from 16.7 to 318.3, and MILP constraint counts from 16 to 64. Median SQP function-evaluation counts are 7326, 14,072, 17,206, and 29,466.5. The sampled MATLAB process memory remains between 2408 and 2465 MB; these values are process-level snapshots rather than incremental or peak algorithmic memory. Consequently, the measured growth supports centralized slow-timescale reconfiguration for the tested sizes but does not establish real-time operation, arbitrary-scale scalability, convergence to a stationary point, or global optimality.

6. Conclusions

This paper formulated reliable scheduling in overlapping multi-cell NOMA VLC as a receiver–scheduler co-design problem under correlated geometry and CSI uncertainty. PDG-HPA enumerates both SIC directions, admits NOMA and resource-accounted OMA modes through independent confidence bounds, and combines MILP scheduling with fixed-threshold, CVaR-constrained trust-region SQP power adaptation. Candidate-edge certificates are used only for admission; after the activity pattern, powers, and interference terms are fixed, the complete assignment is jointly recertified before held-out testing.

The evaluation distinguishes selected-pair from network-wide outage, matched from misspecified uncertainty laws, balanced from unbalanced topologies, and conditional from unconditional performance. On 30 matched-ablation topologies, PDG-HPA achieves an unconditional effective sum rate of 8.934 bit/s/Hz, compared with 8.275 bit/s/Hz without PDG certification and 8.789 bit/s/Hz with deterministic power adaptation. The deterministic-power PDG variant is feasible more often, whereas the complete method provides the largest unconditional rate. Reporting feasibility and network-wide outage together with throughput is therefore essential; no single method dominates every metric.

The reliability certificates are conditional on the adopted or calibrated uncertainty generator and its estimated parameters. Moreover, the reported runtime and signaling model support centralized slow-timescale reconfiguration rather than real-time per-frame scheduling. The present results should therefore be interpreted as system-level evidence for the simulated LOS-dominant setting rather than as a distribution-free, waveform-level, hardware, or real-time guarantee. The local MILP–SQP procedure also does not provide a global-optimality guarantee. Measured-channel, waveform-level, mobility, signaling-overhead, and hardware validation, together with scalable joint association–scheduling, remain directions for future work.

Author Contributions: Conceptualization, T.Q. and Y.T.; methodology, T.Q. and Y.T.; software, T.Q. and Y.T.; validation, T.Q. and Y.T.; formal analysis, T.Q. and Y.T.; investigation, T.Q. and Y.T.; resources, Y.T.; data curation, T.Q. and Y.T.; writing—original draft preparation, T.Q. and Y.T.; writing—review and editing, T.Q. and Y.T.; visualization, T.Q. and Y.T.; supervision, Y.T.; project administration, Y.T.; funding acquisition, Y.T. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the 2025–2026 Key Research Project of Chongqing College of International Business and Economics (Grant No. KYZK2025004).

Data Availability Statement: The MATLAB simulation code, processed numerical data, and audit files supporting the findings and figures of this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

ACO-OFDM	Asymmetrically Clipped Optical Orthogonal Frequency-Division Multiplexing
ADO-OFDM	Asymmetrically Clipped DC-Biased Optical Orthogonal Frequency-Division Multiplexing
CSI	Channel State Information
CVaR	Conditional Value-at-Risk
DCO-OFDM	DC-Biased Optical Orthogonal Frequency-Division Multiplexing
FOV	Field of View
HPA	Hierarchical Power Allocation
LCB	Lower Confidence Bound
MILP	Mixed-Integer Linear Programming
NOMA	Non-Orthogonal Multiple Access
OMA	Orthogonal Multiple Access
PDG	Probabilistic Decodability Graph

PDG-HPA	Probabilistic-Decodability-Graph-Based Hierarchical Power Allocation
SIC	Successive Interference Cancellation
SQP	Sequential Quadratic Programming
VLC	Visible Light Communication

References

1. Komine, T.; Nakagawa, M. Fundamental Analysis for Visible-Light Communication System Using LED Lights. *IEEE Trans. Consum. Electron.* **2004**, *50*, 100–107. [CrossRef]
2. Pathak, P.H.; Feng, X.; Hu, P.; Mohapatra, P. Visible Light Communication, Networking, and Sensing: A Survey, Potential and Challenges. *IEEE Commun. Surv. Tutor.* **2015**, *17*, 2047–2077. [CrossRef]
3. Haas, H.; Yin, L.; Wang, Y.; Chen, C. What Is LiFi? *J. Light. Technol.* **2016**, *34*, 1533–1544. [CrossRef]
4. He, C.; Chen, C. Advanced Technologies in Optical Wireless Communications. *Photonics* **2025**, *12*, 759. [CrossRef]
5. Marshoud, H.; Kapinas, V.M.; Karagiannidis, G.K.; Muhaidat, S. Non-Orthogonal Multiple Access for Visible Light Communications. *IEEE Photon. Technol. Lett.* **2016**, *28*, 51–54. [CrossRef]
6. Yin, L.; Popoola, W.O.; Wu, X.; Haas, H. Performance Evaluation of Non-Orthogonal Multiple Access in Visible Light Communication. *IEEE Trans. Commun.* **2016**, *64*, 5162–5175. [CrossRef]
7. Sadat, H.; Abaza, M.; Mansour, A.; Alfalou, A. A Survey of NOMA for VLC Systems: Research Challenges and Future Trends. *Sensors* **2022**, *22*, 1395. [CrossRef] [PubMed]
8. Zhang, X.; Gao, Q.; Gong, C.; Xu, Z. User Grouping and Power Allocation for NOMA Visible Light Communication Multi-Cell Networks. *IEEE Commun. Lett.* **2017**, *21*, 777–780. [CrossRef]
9. Eltokhey, M.W.; Khalighi, M.-A.; Ghassemlooy, Z. Dimming-Aware Interference Mitigation for NOMA-Based Multi-Cell VLC Networks. *IEEE Commun. Lett.* **2020**, *24*, 2541–2545. [CrossRef]
10. Arfaoui, M.A.; Ghrayeb, A.; Assi, C.; Qaraqe, M. CoMP-Assisted NOMA and Cooperative NOMA in Indoor VLC Cellular Systems. *IEEE Trans. Commun.* **2022**, *70*, 6020–6034. [CrossRef]
11. Guo, X.; Luo, Y. Hybrid NOMA/OFDMA Visible Light Communication System with Coordinated Multiple Point Transmission. *Opt. Express* **2022**, *30*, 47404–47420. [CrossRef] [PubMed]
12. Tran, G.N.; Kim, S. Performance Evaluation of Short Packet Communications in NOMA VLC Systems with Imperfect CSI. *IEEE Access* **2022**, *10*, 49781–49793. [CrossRef]
13. Dixit, V.; Kumar, A. An Exact BER Analysis of NOMA-VLC System with Imperfect SIC and CSI. *AEU–Int. J. Electron. Commun.* **2021**, *138*, 153864. [CrossRef]
14. Mohapatra, S.; Satapathy, G.; Dash, S.P.; Sahu, P.R. Performance Analysis of Visible Light Communication System with Imperfect CSI. *IEEE Commun. Lett.* **2020**, *24*, 2844–2848. [CrossRef]
15. Ma, S.; Yang, R.; Du, C.; Li, H.; Wu, Y.; Al-Dhahir, N.; Li, S. Robust Power Allocation for Integrated Visible Light Positioning and Communication Networks. *IEEE Trans. Commun.* **2023**, *71*, 4764–4777. [CrossRef]
16. Hu, J.; Sun, C.; Wang, J.; Gao, X.; Xia, L.; Wang, Q. Resource Efficient and Robust RSMA for Visible Light Communications. *IEEE Photon. J.* **2024**, *16*, 7302212 . [CrossRef]
17. Alhalabi, M.; Sönmezocak, T.; El-Nahal, F.I. Design and Performance Analysis of a Directly Modulated Direct Current-Biased Optical Orthogonal Frequency-Division Multiplexing Visible-Light Optical Wireless Link under Atmospheric Turbulence. *Appl. Sci.* **2026**, *16*, 6324. [CrossRef]
18. Alhalabi, M.; Taşpınar, N.; Wadi, M.; Salemdeeb, M. Performance Analysis of a Low-Cost Directly Modulated OFDM-Based Optical Wireless Communication System. *Wirel. Pers. Commun.* **2026**, *146*, 1415–1441. [CrossRef]
19. Alhalabi, M.; Taşpınar, N.; Sönmezocak, T. A Novel Hybrid Whale Optimization Algorithm-Based SLM (HWOA-SLM) for PAPR Reduction in Optical IM/DD OFDM Systems. *Appl. Sci.* **2026**, *16*, 2349. [CrossRef]
20. Tao, S.; Yu, H.; Li, Q.; Tang, Y. Performance Analysis of Gain Ratio Power Allocation Strategies for Non-Orthogonal Multiple Access in Indoor Visible Light Communication Networks. *EURASIP J. Wirel. Commun. Netw.* **2018**, *2018*, 154. [CrossRef]
21. Shen, H.; Wu, Y.; Xu, W.; Zhao, C. Optimal Power Allocation for Downlink Two-User Non-Orthogonal Multiple Access in Visible Light Communication. *J. Commun. Inf. Netw.* **2017**, *2*, 57–64. [CrossRef]
22. Liu, H.; Huang, B.; Yang, J.; Chen, Y.; Zhang, T.; Yang, S.; Chen, K. Joint User Grouping and Power Allocation in VLC-NOMA System. *Phys. Commun.* **2022**, *54*, 101841. [CrossRef]
23. Shanker, O.K.H.; Sah, A.K. Multi-User Grouping for NOMA Enabled VLC System. *IEEE Commun. Lett.* **2024**, *28*, 1072–1076. [CrossRef]
24. Tu, Y.; Li, C.; Liu, R.; Jiang, T.; Chi, J. Novel Power Allocation Grouping Scheme for Multi-User VLC NOMA Systems. In *CCF Computer Network Teaching and Computer Network Engineering Discipline Construction*; Communications in Computer and Information Science; Springer: Singapore, 2026; Volume 2809, pp. 161–171. [CrossRef]
25. Armstrong, J.; Lowery, A.J. Power Efficient Optical OFDM. *Electron. Lett.* **2006**, *42*, 370–372. [CrossRef]

26. Dissanayake, S.D.; Armstrong, J. Comparison of ACO-OFDM, DCO-OFDM and ADO-OFDM in IM/DD Systems. *J. Light. Technol.* **2013**, *31*, 1063–1072. [CrossRef]
27. He, C.; Tu, Y.; Lim, Y. The Combination of SEFDM and NOMA in a Multi-User Visible Light Communication System. In Proceedings of the 2024 IEEE Opto-Electronics and Communications Conference (OECC), Melbourne, Australia, 30 June–4 July 2024; pp. 1–3. [CrossRef]
28. Tu, Y.; Li, C.; Pham, C.V. Interference-Aware User Grouping and Power Allocation for Overlapping Multi-LED ADO-OFDM NOMA VLC Networks. *Photonics* **2026**, *13*, 241. [CrossRef]
29. Román Cañizares, M.; Palacios Játiva, P.; Azurdia-Meza, C.A.; Montejo-Sánchez, S.; Céspedes, S. Impact of Diversity Combining Schemes in a Multi-Cell VLC System with Angle Diversity Receivers. *Photon. Netw. Commun.* **2022**, *43*, 13–22. [CrossRef]
30. Palitharathna, K.W.S.; Suraweera, H.A.; Godaliyadda, R.I.; Herath, V.R.; Thompson, J.S. Neural Network-Based Channel Estimation and Detection in Spatial Modulation VLC Systems. *IEEE Commun. Lett.* **2022**, *26*, 1598–1602. [CrossRef]
31. Nagaraja, K.; Dash, S.P.; Ghose, D. Channel Estimation and Error Analysis of a Narrow FoV VLC System With Random Receiver Orientation. *IEEE Open J. Commun. Soc.* **2025**, *6*, 3920–3931. [CrossRef]
32. Qin, T.; Tu, Y. Risk-Aware Illumination-Constrained Resource Allocation for Hybrid VLC/RF Indoor Networks under Random Optical Blockage. *Photonics* **2026**, *13*, 569. [CrossRef]
33. Nemirovski, A.; Shapiro, A. Convex Approximations of Chance Constrained Programs. *SIAM J. Optim.* **2006**, *17*, 969–996. [CrossRef]
34. Calafiore, G.C.; Campi, M.C. The Scenario Approach to Robust Control Design. *IEEE Trans. Autom. Control* **2006**, *51*, 742–753. [CrossRef]
35. Campi, M.C.; Garatti, S. The Exact Feasibility of Randomized Solutions of Uncertain Convex Programs. *SIAM J. Optim.* **2008**, *19*, 1211–1230. [CrossRef]
36. Rockafellar, R.T.; Uryasev, S. Optimization of Conditional Value-at-Risk. *J. Risk* **2000**, *2*, 21–41. [CrossRef]
37. Gu, Y.; Saad, W.; Bennis, M.; Debbah, M.; Han, Z. Matching Theory for Future Wireless Networks: Fundamentals and Applications. *IEEE Commun. Mag.* **2015**, *53*, 52–59. [CrossRef]
38. Liang, W.; Ding, Z.; Li, Y.; Song, L. User Pairing for Downlink Non-Orthogonal Multiple Access Networks Using Matching Algorithm. *IEEE Trans. Commun.* **2017**, *65*, 5319–5332. [CrossRef]
39. Clopper, C.J.; Pearson, E.S. The Use of Confidence or Fiducial Limits Illustrated in the Case of the Binomial. *Biometrika* **1934**, *26*, 404–413. [CrossRef]
40. Scutari, G.; Facchinei, F.; Lampariello, L. Parallel and Distributed Methods for Constrained Nonconvex Optimization—Part I: Theory. *IEEE Trans. Signal Process.* **2017**, *65*, 1929–1944. [CrossRef]