Research Square

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

Strategic Urban Air Mobility Network Design with Disaggregated Demand and Scalable Hybrid Optimization

Hossein Zamani Saghazadeh
University of Central Florida
Soheil Sabri
soheil.sabri@ucf.edu

University of Central Florida
Parham Kebria
North Carolina Agricultural and Technical State University

Research Article

Keywords: Urban air mobility, Hub location, Disaggregated travel demand, Matheuristic optimization, Local branching, Fix-and-optimize

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11684063/v1

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

Additional Declarations: The authors declare no competing interests.

Strategic Urban Air Mobility Network Design with Disaggregated Demand and Scalable Hybrid Optimization*

Hossein Zamani Saghazadeh[a], Soheil Sabri[a,*], Parham Kebria[b]

[a]*Urban Digital Twin Lab, School of Modeling, Simulation, and Training, University of Central Florida, 3100 Technology Parkway, Orlando, 32826, Florida, United States*
[b]*Department of Electrical and Computer Engineering, North Carolina Agricultural and Technical State University, 1601 East Market Street, Greensboro, 27411, North Carolina, United States*

Abstract

Strategic urban air mobility (UAM) network design must preserve trip-specific accessibility while supporting repeated optimization and subsequent operational planning. This study develops a six-stage framework that (1) constructs census-block-resolution potential trips and candidate stations, (2) generates complete ground–air–ground paths with multiple endpoint station alternatives, (3) screens paths by absolute and relative no-overhead time advantages and constructs retained-demand scenarios, (4) optimizes hub-location problems with minimum-utilization and functional-substitution requirements, (5) selects a representative network through infrastructure-scale and cross-scenario comparisons, and (6) derives a nested deployment sequence and operational allowances. Two objectives maximize potential total passenger travel-time saving or lexicographically prioritize assigned demand before travel-time saving. A hybrid local-branching and fix-and-optimize (LBFO) method enables repeated optimization across four retained-demand scenarios and 12 station budgets. At representative budgets, runtime summed across both objectives within each scenario decreases by 83.5–99.5% relative to matched complete-domain Gurobi benchmarks, with primary-objective certified-bound shortfalls below 1% in all eight tasks. In Greater Los Angeles, the selected network has 40 stations and 160 physical corridors, assigns 30,634 potential trips, and yields 1.89 million passenger-minutes of no-overhead travel-time saving. The demand-weighted median operational allowance is 25.7 min, and 95% of assigned demand has an allowance of at least 19.0 min while retaining a 40% relative-saving requirement. These outputs connect strategic network selection to phased deployment and subsequent scheduling, fleet-sizing, and vertiport-capacity analysis.

Keywords: Urban air mobility, Hub location, Disaggregated travel demand, Matheuristic optimization, Local branching, Fix-and-optimize

1. Introduction

Urban Air Mobility (UAM), the urban and metropolitan component of Advanced Air Mobility (AAM), has emerged as a potential complement to surface transportation through advances in electric vertical take-off and landing (eVTOL) aircraft, automation, and airspace-management technologies (Cohen et al., 2021; Garrow et al., 2021). Its passenger-transportation value, however, depends on the complete ground–air–ground journey. Station locations, access and egress conditions, network connectivity, and operational overheads, including waiting, processing, charging, and coordination delays, jointly determine whether a UAM alternative provides a meaningful advantage over direct ground travel.

Strategic UAM planning must translate spatially distributed potential demand into a connected infrastructure network. Existing research has advanced UAM demand estimation, vertiport siting, network design, and operational modeling, but individual formulations generally integrate only subsets of these elements (Wu and Zhang, 2021; Willey and Salmon, 2021; Escribano Macias et al., 2023). Their integration becomes particularly difficult when individual

*Preprint version prior to peer review.
*Corresponding author: {soheil.sabri@ucf.edu}

trips retain multiple station alternatives and network design must be repeated across demand scopes, station budgets, and planning objectives. Moreover, strategic decisions precede observations of realized UAM adoption and operational overheads. The planning representation must therefore distinguish constructed potential demand from predicted ridership and avoid prescribing one operational overhead before scheduling, fleet, and vertiport-capacity decisions are known.

This study develops a six-stage strategic demand-to-network framework to address this planning need (Figure 2). Building on the demand-informed candidate-generation procedure in a preceding study (Zamani Saghazadeh et al., 2026), the framework constructs census-block-resolution potential trips and retains multiple eligible stations at each trip end. Complete door-to-door alternatives are screened using absolute and relative no-overhead travel-time advantages, and nested retained-demand scenarios define alternative planning scopes rather than forecast market shares. A hub-location problem (HLP) then jointly selects stations, directed links, and trip assignments under minimum-utilization and functional-substitution requirements. Two objectives are considered: (1) potential total passenger travel-time saving (PTTS), which emphasizes aggregate passenger benefit, and (2) demand-coverage lexicographic optimization (DCLEX), which prioritizes assigned potential demand before travel-time benefit.

The study makes three contributions. First, it develops a hybrid local-branching and fix-and-optimize (LBFO) method for repeated HLP optimization across retained-demand scenarios, station budgets, and planning objectives, supported by matched Gurobi benchmarks (Gurobi Optimization, LLC, 2026) and controlled component comparisons. Second, it connects disaggregated potential demand and endpoint-specific station alternatives to a complete door-to-door network representation without imposing a universal catchment or predicted market share. Third, it converts the optimized network family into a representative infrastructure scale, a nested deployment sequence, and trip- and station-level operational allowances for subsequent operational planning.

The remainder of the paper is organized as follows. Section 2 reviews related infrastructure, demand, operational, and computational research. Section 3 presents the strategic planning framework, HLP formulation, LBFO method, and evaluation procedures. Section 4 describes the Greater Los Angeles application and computational protocol with results presented in Section 5. Section 6 discusses the implications and limitations, followed by conclusions in Section 7.

2. Background and Literature Review

2.1. UAM Infrastructure Siting and Network Design

UAM infrastructure planning has been addressed through covering, facility-location, clustering, hub-location, and network-design models. Covering formulations select vertiports within predefined service ranges, including binary and gradual coverage specifications (Hess et al., 2024). Clustering-based methods derive locations from spatial demand patterns (Rajendran and Zack, 2019; Sinha and Rajendran, 2023), while capacitated facility-location and maximal-covering models examine tradeoffs among coverage, accessibility, capacity, service radius, and infrastructure count (Volakakis and Mahmassani, 2025). These approaches provide valuable siting guidance, although demand-to-facility eligibility is generally established through their clustering or coverage rules.

Hub-location formulations more directly connect facility selection, origin–destination flows, and network structure (Campbell and O'Kelly, 2012). Willey and Salmon (2021) demonstrate how aircraft speed and range constrain structured UAM networks. Related hub-location and network-design studies address multiple allocation and incomplete connectivity (Kitthamkesorn and Chen, 2024), demand uncertainty and disruption redundancy (Jin et al., 2024; Wei et al., 2024), and coordinated land–air and operational planning (Zhang et al., 2025; Wang et al., 2025). Collectively, these studies show that station selection and connectivity depend on demand representation, network structure, capacity, uncertainty, redundancy, and operational coupling.

2.2. Demand, Accessibility, and Operational Representation

In the absence of mature revealed UAM demand, studies have constructed or behaviorally estimated potential demand. Haan et al. (2021) combine mobile-device data, census attributes, and stated-preference mode choice, whereas Wu and Zhang (2021) integrate simulated disaggregated demand with vertiport location, access and egress modes, and generalized cost. Other network models represent traveler response through elastic-demand, random-utility, or

binary-logit formulations (Rath and Chow, 2022; Kitthamkesorn and Chen, 2024; Kim et al., 2025). These studies demonstrate that demand representations serve different purposes and must be interpreted consistently with their underlying data and behavioral assumptions.

Accessibility assumptions shape the station alternatives available to each trip. Venkatesh et al. (2020) use a common catchment around home and work blocks, whereas Jin et al. (2024) represent zones by centroids and impose a common maximum access or egress time. The present framework retains endpoint-specific station alternatives and evaluates their road-network access and egress times within complete ground–air–ground paths, leaving station-pair selection to the HLP.

Operational time also remains simplified in several strategic siting and network-design models. Jin et al. (2024) define UAM travel time as the sum of access, flight, and egress times, without explicitly representing waiting, processing, charging, or coordination delays. Kim et al. (2025) assume sufficient aircraft availability and no passenger waiting in their skyport-location model. Venkatesh et al. (2020) add a fixed 3 min for boarding, deboarding, and taxiing, whereas Willey and Salmon (2021) use a fixed 5 min for boarding and deboarding on each flight leg. Such assumptions support tractable strategic models, but operational studies show that throughput and delay depend on vertiport layout, pad and stand resources, service times, and directional demand imbalance (Vascik and Hansman, 2019; Rimjha and Trani, 2021). Integrated queueing and network-simulation studies further link passenger waiting to resource sizing, charging, dispatch, repositioning, and congestion (Escribano Macias et al., 2023; Onat et al., 2024; Li and Zhang, 2026). These findings show why one prescribed operational overhead cannot convey the station- and demand-specific conditions needed for subsequent operational planning. Because these outcomes are unavailable at the strategic-planning stage, the present framework does not prescribe one realized operational overhead for the entire network. Instead, it derives the maximum overhead that each selected trip can absorb while retaining its required advantage over direct ground travel. The resulting trip- and station-level allowance distributions provide inputs and feasibility benchmarks for subsequent operational models.

2.3. Computational Scalability and Research Positioning

The computational burden of UAM network design increases when disaggregated trips can retain multiple feasible station pairs. Previous studies address complex network-design formulations using problem-specific methods, including structured-network search (Willey and Salmon, 2021), cross-entropy optimization (Kim et al., 2025), and mixed-integer Bayesian optimization (Zhang et al., 2025). These studies demonstrate the value of problem-specific search strategies, although their procedures are designed for their respective formulations.

Motivated by this need for computational scalability, the present study combines local branching, which defines mixed-integer neighborhoods around an incumbent (Fischetti and Lodi, 2003), with fix-and-optimize, which reoptimizes selected decisions while temporarily fixing the remainder (Helber and Sahling, 2010). This design follows the broader matheuristic principle of integrating exact optimization with heuristic search (Puchinger and Raidl, 2005). Variable neighborhood search (VNS) is retained as a controlled comparator (Hansen et al., 2010).

Taken together, the reviewed literature has separately advanced demand representation, infrastructure siting, network design, operational analysis, and computational solution methods. The research gap addressed here is their limited integration into a scalable strategic workflow that retains multiple endpoint-specific station alternatives for disaggregated trips, evaluates complete ground–air–ground paths, and enforces network-wide utilization and functional-substitution requirements. The workflow must also support repeated optimization across planning scopes and station budgets and translate the resulting network family into deployment and operational-allowance inputs. The proposed framework addresses this integrated planning need without interpreting constructed potential demand as predicted adoption or assuming one realized operational overhead for all trips.

3. Strategic Demand-to-Network Planning Framework

This section translates spatially disaggregated potential demand and feasible candidate stations into a strategic UAM network-design problem. The framework constructs trip-specific door-to-door UAM alternatives, screens them into retained-demand scenarios, and optimizes the resulting HLP domains under two planning objectives. A hybrid LBFO procedure supports repeated solves across scenarios and station budgets, after which the selected networks are evaluated for representative infrastructure scale, phased deployment, and trip-specific operational allowances.

3.1. Demand, Candidate Stations, and Door-to-Door Paths

The demand and candidate-station inputs follow the improved Visual Assessment of cluster Tendency (iVAT)-guided fuzzy c-means framework developed in a preceding work (Zamani Saghazadeh et al., 2026). That framework produces the constructed OD trip set R and physical candidate-station set H. Demand is represented at census-block resolution rather than as aggregate zonal OD flows, so each constructed OD record is evaluated separately. Each trip $r \in R$ has an origin o_r, destination d_r, and demand weight q_r, preserving its direct-ground travel time, eligible station alternatives, feasible paths, and travel-time advantages. These trips represent constructed potential UAM demand rather than predicted ridership, adoption, or mode share. Section S1 of the Supplementary Material summarizes the associated demand allocation, spatial processing, and candidate-construction procedures.

The upstream spatial structure also determines the station alternatives at each trip end. The sets $H_r^O \subseteq H$ and $H_r^D \subseteq H$ contain the eligible origin- and destination-side stations for trip r, respectively. Fuzzy memberships and overlapping local clustering zones allow different trip ends to retain different alternatives. This keeps station-pair selection within the HLP rather than assigning each trip to one pair before optimization. Eligibility therefore follows the spatial demand structure rather than a common catchment radius, fixed access-time threshold, or predetermined number of alternatives. Because each eligible origin–destination station combination produces a candidate path, this flexibility also enlarges the HLP domain.

For $a \in H_r^O$, let α_{ra} denote the road-network access time from o_r to station a. For $b \in H_r^D$, let β_{rb} denote the corresponding egress time from station b to d_r. A candidate path $k = (r(k), a(k), b(k))$ therefore represents the complete itinerary $o_{r(k)} \to a(k) \to b(k) \to d_{r(k)}$, as shown in Figure 1.

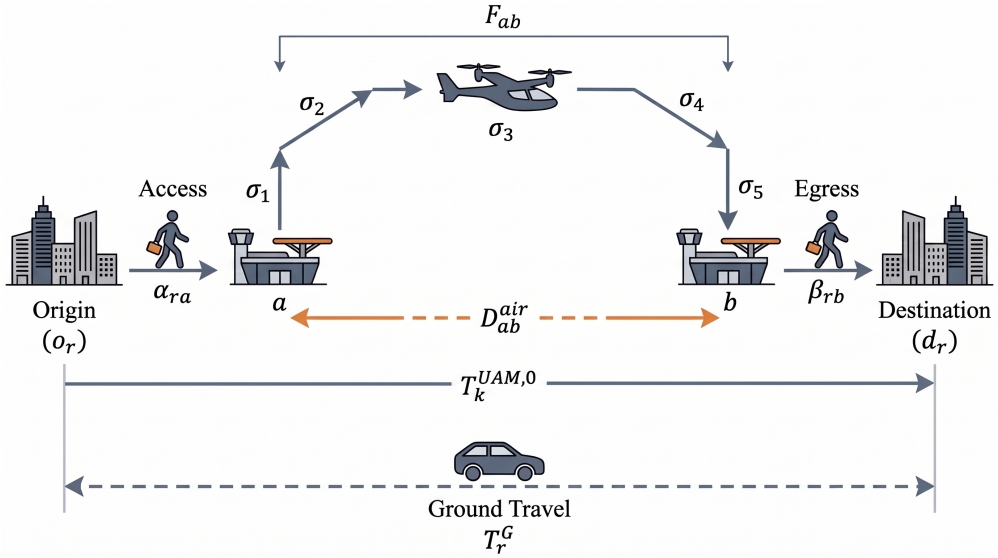


Figure 1: Candidate door-to-door UAM path $k = (r, a, b)$; ground access α_{ra}, station-to-station flight F_{ab}, and ground egress β_{rb}, with direct ground time T_r^G as the baseline. Flight time comprises vertical climb (σ_1), forward climb (σ_2), cruise (σ_3), forward descent (σ_4), and vertical descent (σ_5)

Road-network travel times for the direct-ground, access, and egress legs are obtained from OpenRouteService (openrouteservice, 2026). Each ordered station pair (a, b) is characterized by its air distance D_{ab}^{air} and precomputed flight time F_{ab}. The phase-based flight-time model and its aircraft-performance parameters are reported in a previous work (Saghazadeh et al., 2026); the present framework uses F_{ab} as an input to the strategic HLP.

The complete candidate path set P and its induced directed station-pair set L are defined as

$$P = \left\{k = (r, a, b) : r \in R,\ a \in H_r^O,\ b \in H_r^D\right\}$$

and $L = \{(a(k), b(k)) : k \in P\}$. Thus, L contains only the directed station pairs represented by at least one candidate path; the HLP determines their activation. For each path, the no-overhead door-to-door UAM time $T_k^{\text{UAM},0}$, absolute

travel-time advantage ΔT_k^{adv}, and relative saving over direct ground travel ρ_k^0 are defined in Eq. (1).

$$\begin{aligned} T_k^{\text{UAM},0} &= \alpha_{r(k),a(k)} + F_{a(k)b(k)} + \beta_{r(k),b(k)}, \\ \Delta T_k^{\text{adv}} &= T_{r(k)}^G - T_k^{\text{UAM},0}, \qquad \rho_k^0 = \frac{\Delta T_k^{\text{adv}}}{T_{r(k)}^G}, \qquad k \in P,\ T_{r(k)}^G > 0 \end{aligned} \tag{1}$$

These measures are evaluated before operational overhead is known. Because the same absolute time advantage represents a larger proportional benefit for a short trip than for a long one, ρ_k^0 provides a duration-normalized route-quality screen, while ΔT_k^{adv} supports scenario construction based on absolute time advantage. Together, these measures exclude paths whose benefit over direct ground travel is marginal in either proportional or absolute terms. The screening and scenario-specific domain construction are presented in Section 3.2; trip-level operational allowances are evaluated only after the network and trip assignments have been selected.

3.2. Retained-Demand Scenarios and HLP-Ready Domains

Before constructing the retained-demand scenarios, paths that are physically infeasible or insufficiently competitive are removed. The initial eligible path set is defined in Eq. (2).

$$P^0 = \left\{ k \in P : \begin{array}{l} a(k) \neq b(k),\ D_{\min}^{\text{air}} \leq D_{a(k)b(k)}^{\text{air}} \leq D_{\max}^{\text{air}}, \\ T_{r(k)}^G > 0,\ \rho_k^0 \geq \rho_{\min}^{\text{pre}} \end{array} \right\} \tag{2}$$

Here, $D_{\min}^{\text{air}}$ and $D_{\max}^{\text{air}}$ are eVTOL-informed flight-distance requirements that delimit the station-to-station flight legs evaluated using the model reported by Saghazadeh et al. (2026). The threshold $\rho_{\min}^{\text{pre}}$ is a planning requirement that ensures every retained path provides a minimum proportional improvement over direct ground travel.

The set P^0 may still contain paths whose endpoint stations or directed links cannot attract meaningful demand even under the most favorable assignment from the remaining alternatives. The lower bounds $Q_{\min}^{\text{st}}$ and $Q_{\min}^{\text{link}}$ are minimum-utilization requirements; they prevent the HLP from opening a station or activating a directed link to serve only negligible demand. The fixed-point operator $\Phi_{\text{HLP}}^{\star}(\cdot)$ removes paths whose required infrastructure cannot satisfy these requirements under any assignment from the remaining domain. The objective-independent base set is therefore $P^{\text{base}} = \Phi_{\text{HLP}}^{\star}(P^0)$. The complete operator is defined in Section S2.1 of the Supplementary Material.

Retained-demand scenarios are defined by the share of constructed potential demand having at least one path with a specified absolute no-overhead advantage. For any threshold $\Gamma \geq 0$, the retained-demand share and the largest threshold preserving at least target share ξ are denoted by $\xi(\Gamma)$ and Γ_ξ, respectively. The two quantities are presented in Eq. (3).

$$\begin{aligned} \xi(\Gamma) &= \frac{\sum_{r \in R} q_r \mathbf{1}\left\{\exists\, k \in P^{\text{base}} : r(k) = r,\ \Delta T_k^{\text{adv}} \geq \Gamma\right\}}{\sum_{r \in R} q_r}, \\ \Gamma_\xi &= \sup\{\Gamma \geq 0 : \xi(\Gamma) \geq \xi\} \end{aligned} \tag{3}$$

The retained share is measured relative to the complete constructed potential-demand set R and defines the demand scope considered in strategic planning. Since $\xi(\Gamma)$ is non-increasing, a smaller target share produces a weakly larger Γ_ξ and therefore a more selective planning domain.

At the derived threshold, $P(\Gamma_\xi)$ contains the paths in P^{base} satisfying $\Delta T_k^{\text{adv}} \geq \Gamma_\xi$, and $R(\Gamma_\xi)$ contains the trips having at least one such path. Because $P^{\text{base}} \subseteq P^0$, these paths already satisfy the proportional requirement $\rho_k^0 \geq \rho_{\min}^{\text{pre}}$.

Thresholding can reduce the demand support available to some stations and directed links. HLP-safe pruning is therefore repeated to obtain the final scenario-specific domains in Eq. (4).

$$\begin{aligned} P_\xi^F &= \Phi_{\text{HLP}}^{\star}\left(P(\Gamma_\xi)\right), \qquad R_\xi^F = \left\{r(k) : k \in P_\xi^F\right\}, \\ H_\xi^F &= \bigcup_{k \in P_\xi^F} \{a(k), b(k)\}, \qquad L_\xi^F = \left\{(a(k), b(k)) : k \in P_\xi^F\right\} \end{aligned} \tag{4}$$

Because thresholding is followed by HLP-safe pruning, the final retained share can be smaller than the nominal target ξ. Trips in $R \setminus R_\xi^F$ have no retained UAM alternative in scenario ξ and are structurally ground-only before optimization.

Distinct candidate stations can provide nearly equivalent access- or egress-side service to a substantial common demand base. Without substitution control, differences of only seconds or a few minutes could cause the optimizer to open multiple stations that are mathematically distinct but strategically redundant. The base substitution set C^{sub} is therefore constructed using shared demand support and requiring the relevant 90th-percentile difference in access or egress time to be no greater than the substitution tolerance δ^{sub}. Neither station is removed during preprocessing; instead, the HLP prevents their simultaneous opening while retaining the option to select either station. For scenario ξ, the applicable set is $C_\xi^{\text{sub}} = \{(i,j) \in C^{\text{sub}} : i, j \in H_\xi^F\}$. The complete construction is provided in Section S2.2 of the Supplementary Material.

The resulting P_ξ^F, R_ξ^F, H_ξ^F, L_ξ^F, and C_ξ^{sub} constitute the scenario-specific domains passed to the strategic HLP. Figure 2 summarizes the complete sequence from the upstream demand and candidate-station inputs to network optimization and strategic evaluation.

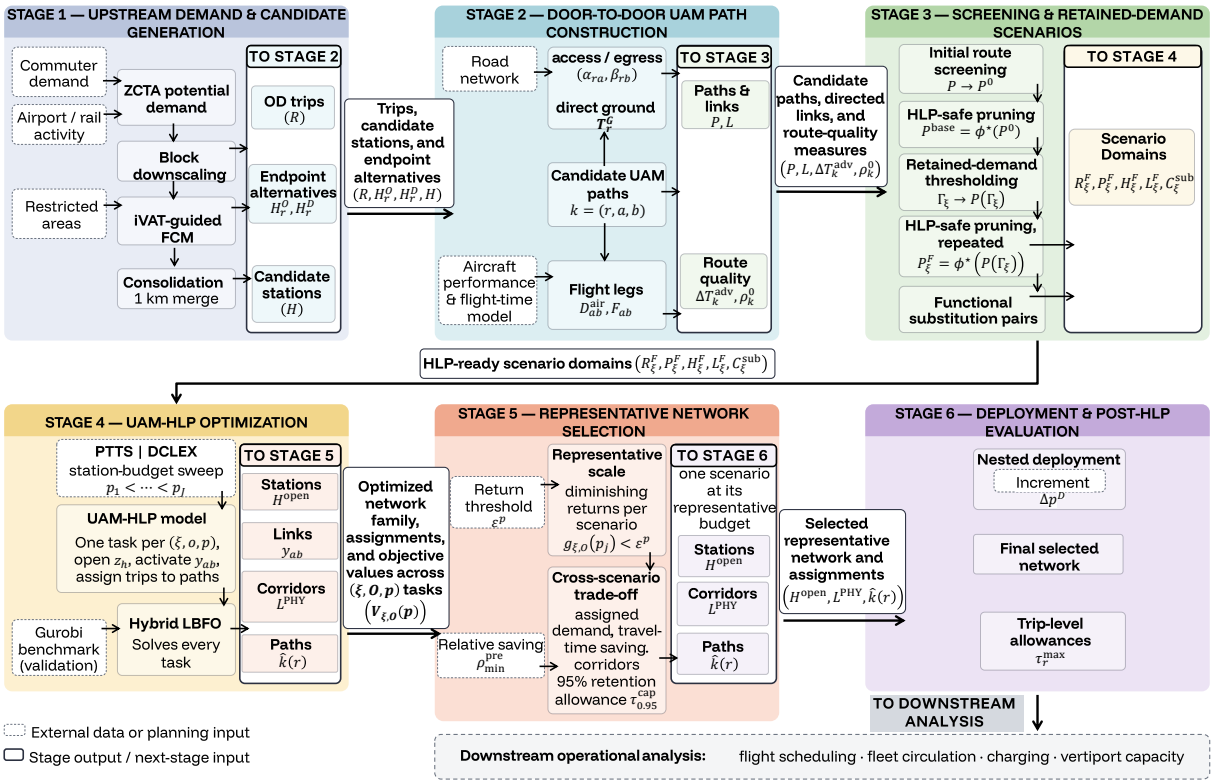


Figure 2: Six-stage demand-to-network framework and interstage information flow. Stage 1 converts commuter, airport and rail, and spatial data into potential OD trips, candidate stations, and trip-end alternatives. Stage 2 combines these outputs with road-network and aircraft inputs to construct complete ground–air–ground paths, directed links, and route-quality measures. Stage 3 screens and prunes these paths to form scenario-specific HLP domains. Stage 4 solves the PTTS and DCLEX station-budget tasks using hybrid LBFO, with Gurobi providing benchmark validation. Stage 5 selects a representative network based on diminishing returns and cross-scenario trade-offs. Stage 6 derives nested deployment and trip-level allowance distributions for downstream operational analysis. Dashed boxes denote external inputs; outlined boxes denote outputs transferred between stages.

3.3. Strategic HLP Formulation and Planning Objectives

The scenario-specific domains constructed in Section 3.2 define the inputs to the strategic HLP. For a fixed scenario ξ, the station budget p specifies the maximum number of stations that may be opened. For $k \in P_\xi^F$, the binary variable x_k equals one if path k is selected for its associated trip. Similarly, the binary variables z_h and y_{ab} indicate whether

station $h \in H_\xi^F$ is opened and directed link $(a, b) \in L_\xi^F$ is activated, respectively. The model jointly determines these network and assignment decisions subject to infrastructure, utilization, and substitution requirements. PTTS and DCLEX use the same feasible set but represent different planning objectives.

The parameter $Q_{\min}^{\text{st}}$ is the minimum assigned demand required at an opened station, and $Q_{\min}^{\text{link}}$ is the corresponding requirement for an activated directed link. Because at most one path can be selected for each retained trip, the scenario-specific domains also provide valid upper bounds for the aggregated activation constraints. Specifically, $U_{\xi h}$ bounds the number of selected trip paths that can use station h, whereas $\overline{Q}_{\xi,ab}^{\text{link}}$ bounds the assigned demand that can use directed link (a, b). These bounds are computed as:

$$\begin{aligned} U_{\xi h} &= \sum_{r \in R_\xi^F} \mathbf{1}\left\{\exists\, k \in P_\xi^F : r(k) = r,\ h \in \{a(k), b(k)\}\right\}, && h \in H_\xi^F \\ \overline{Q}_{\xi,ab}^{\text{link}} &= \sum_{r \in R_\xi^F} q_r \mathbf{1}\left\{\exists\, k \in P_\xi^F : r(k) = r,\ a(k) = a,\ b(k) = b\right\}, && (a, b) \in L_\xi^F \end{aligned} \tag{5}$$

Interpretation of framework bounds. The formulation distinguishes planning requirements, data-derived valid bounds, and eVTOL-informed flight-distance limits. The station budget p, relative-saving requirement $\rho_{\min}^{\text{pre}}$, minimum-utilization requirements $Q_{\min}^{\text{st}}$ and $Q_{\min}^{\text{link}}$, and functional-substitution tolerance δ^{sub} specify planning requirements. The limits $D_{\min}^{\text{air}}$ and $D_{\max}^{\text{air}}$ delimit the station-to-station flight distances considered based on eVTOL flight requirements. For each target share ξ, the threshold Γ_ξ is derived from the demand and travel-time data. The quantities $U_{\xi h}$ and $\overline{Q}_{\xi,ab}^{\text{link}}$ are instead valid activation bounds computed from the corresponding scenario domains. None of these quantities represents station, corridor, fleet, or airspace throughput capacity; such capacities depend on subsequent scheduling, charging, fleet-operation, and vertiport-design decisions.

$$\sum_{\substack{k \in P_\xi^F \\ r(k) = r}} x_k \le 1, \qquad r \in R_\xi^F, \tag{6a}$$

$$\sum_{h \in H_\xi^F} z_h \le p, \tag{6b}$$

$$\sum_{\substack{k \in P_\xi^F \\ h \in \{a(k), b(k)\}}} x_k \le U_{\xi h} z_h, \qquad h \in H_\xi^F, \tag{6c}$$

$$\sum_{\substack{k \in P_\xi^F \\ h \in \{a(k), b(k)\}}} q_{r(k)} x_k \ge Q_{\min}^{\text{st}} z_h, \qquad h \in H_\xi^F, \tag{6d}$$

$$Q_{\min}^{\text{link}} y_{ab} \le \sum_{\substack{k \in P_\xi^F \\ a(k) = a,\ b(k) = b}} q_{r(k)} x_k \le \overline{Q}_{\xi,ab}^{\text{link}} y_{ab}, \qquad (a, b) \in L_\xi^F, \tag{6e}$$

Constraints (6a)–(6e) define the assignment, infrastructure, and utilization requirements shared by PTTS and DCLEX. Constraint (6a) permits at most one UAM path per retained trip, while Constraint (6b) imposes the station budget. Constraints (6c) and (6d) require every selected path to use opened stations and every opened station to satisfy its minimum-utilization requirement. Constraint (6e) provides the corresponding activation and utilization conditions for directed links. Functionally substitutable stations cannot be opened simultaneously, so $z_i + z_j \le 1$ for every $(i, j) \in C_\xi^{\text{sub}}$. The station budget is expressed as an upper bound because opening additional stations may be infeasible under the utilization or substitution requirements.

The two objectives share this feasible set but prioritize different network outcomes. Assigned potential demand

and potential total passenger travel-time saving are defined as

$$Q^{\text{assigned}} = Q^{\dagger} = \sum_{k \in P^F_{\xi}} q_{r(k)} x_k, \qquad PTTS^{\text{total}} = \sum_{k \in P^F_{\xi}} q_{r(k)} \Delta T^{\text{adv}}_k x_k. \tag{7}$$

Because ΔT^{adv}_k is measured in minutes, $PTTS^{\text{total}}$ is expressed in passenger-minutes of no-overhead travel-time advantage relative to direct ground travel. Let $Q^{\dagger *}$ denote the optimal assigned demand obtained in the first stage of DCLEX. The objective structures are

$$\begin{aligned} &\text{PTTS:} \quad \max \ PTTS^{\text{total}}, \\ &\text{DCLEX:} \quad \max \ Q^{\dagger}, \qquad \max \ PTTS^{\text{total}} \quad \text{s.t.} \quad Q^{\dagger} = Q^{\dagger *}. \end{aligned} \tag{8}$$

PTTS may favor assignments with larger time advantages, whereas DCLEX gives strict priority to assigning the greatest potential demand and uses PTTS only to distinguish among maximum-coverage solutions.

3.4. Hybrid LBFO Solution Method

Census-block-resolved demand and endpoint-specific station alternatives are the two principal sources of model scale. Evaluating each OD record separately preserves differences in access, egress, and travel-time advantage, while retaining multiple endpoint alternatives leaves station-pair selection within the HLP. These choices support trip-level assignment and demand-weighted post-optimization analysis but generate several hundred thousand candidate paths, making repeated complete-domain optimization computationally prohibitive. Each combination of retained-demand scenario, planning objective, and station budget defines one HLP optimization task. The proposed hybrid LBFO method preserves the formulation in Section 3.3 while restricting each solve to a promising region of the station-opening space.

Fix-and-optimize iteratively reoptimizes selected subsets of binary decisions while keeping the remaining binary decisions fixed (Helber and Sahling, 2010). Local branching instead restricts the search to a neighborhood of an incumbent integer solution through an explicit Hamming-distance constraint (Fischetti and Lodi, 2003). LBFO therefore follows an integrative matheuristic structure in which restricted mixed-integer linear programming (MILP) optimization is embedded within a problem-specific neighborhood search (Puchinger and Raidl, 2005). Throughout, the incumbent denotes the best validated solution obtained for the current task.

LBFO proceeds through four phases. INIT ranks the candidate stations, constructs the initial restricted station pool, and produces a validated incumbent. S1 and S2 rebuild the pool around that incumbent and explore progressively wider neighborhoods, while ELITE performs a final search with greater emphasis on accumulated solution memory. In each subsequent phase, all incumbent-open stations remain in the pool, while selected incumbent-closed stations are admitted as alternatives using the phase-specific ranking. For phase ϑ, station variables within the resulting pool $S_{\vartheta} \subseteq H^F_{\xi}$ remain free, whereas candidate stations outside S_{ϑ} are temporarily fixed closed. Paths requiring a fixed-closed station are consequently excluded, together with their induced directed links. Reoptimizing the HLP over the remaining free decisions constitutes the fix-and-optimize component of LBFO. S1, S2, and ELITE additionally impose a local-branching constraint that limits the change from the incumbent station configuration. The objective and constraints of Section 3.3 otherwise remain unchanged.

Three profile families control the restricted search. The weighting profile W determines the relative importance of the components used to rank candidate stations. The pool-breadth profile PB determines how many stations are retained in each restricted domain, while the neighborhood profile N specifies how far the station configuration may move from the incumbent during S1, S2, and ELITE. Together, these profiles determine which station decisions are released and how far the restricted solution may move from the incumbent; they do not modify the underlying complete-domain HLP. Table 1 summarizes the score groups used to construct the restricted station pools.

For every score group, each raw component is independently min–max normalized over the stations evaluated in the corresponding ranking. If a component is constant over those stations, its normalized values are set to zero and the component does not affect their ordering. For a score group with m ordered components, let $\widehat{v}_{hc} \in [0, 1]$ denote the normalized value of component c for station h, and let r^W_c be its relative weight under profile W. The resulting station

score is

$$\mathrm{Score}_h^W = \sum_{c=1}^{m} \frac{r_c^W}{\sum_{j=1}^{m} r_j^W} \widehat{v}_{hc}. \tag{9}$$

Table 1: LBFO uses the listed score groups to construct restricted station pools after independently min–max normalizing each component to $[0, 1]$; the listed order determines how profile W assigns component weights.

Score group	Use	Components in profile-weight order
PTTS base	Objective-aware ranking	Station PTTS potential; incident directed-link PTTS potential; station demand potential; station degree.
DCLEX base	Objective-aware ranking	Station demand potential; incident directed-link demand potential; station degree; secondary station PTTS potential.
INIT	Initial restricted pool	Base score; memory support; visit deficit.
S1	First neighborhood pool	Incumbent opportunity; memory support; base score; pair opportunity; visit deficit.
S2	Second neighborhood pool	Incumbent opportunity; memory support; pair opportunity; visit deficit; base score.
ELITE	Final neighborhood pool	Memory support; incumbent opportunity; pair opportunity; base score; visit deficit.

Station PTTS potential aggregates the path-level PTTS contributions involving a station, whereas station demand potential counts each trip at most once for that station. Incident-link PTTS and demand potential aggregate the corresponding quantities over the station's incoming and outgoing candidate directed links. Station degree is the number of such links incident to the station. The objective-specific base score combines these quantities to emphasize passenger travel-time saving under PTTS and supported demand under DCLEX, with station PTTS used as a secondary DCLEX criterion.

The phase-specific scores incorporate information generated during the search. Incumbent opportunity measures the potential improvement offered by a closed station through paths connecting it to the current open-station network. Pair opportunity represents the objective potential of directed links connecting pairs of closed stations. Memory support gives greater priority to stations that repeatedly occur in strong validated solutions from earlier phases or previously solved budgets, while visit deficit promotes stations that have appeared less frequently in earlier restricted pools. Memory provides soft guidance only; remembered stations are not forced open, and none of these scoring components modifies the HLP objective or constraints.

The tested weighting profiles are specified as follows: $W1 = (1, 1, \ldots, 1)$ assigns equal importance to all components, $W2 = (2, 1, \ldots, 1)$ emphasizes the first component, $W3 = (2, 2, 1, \ldots, 1)$ emphasizes the first two, and $W4 = (5, 3, 2, 1, 1)$ imposes a steeper descending hierarchy. For a score group with fewer than five components, $W4$ uses the first m entries of its stated sequence. Pool breadth is scaled to both the station budget and the number of candidates in the scenario domain. Under pool profile PB, the target pool size for phase ϑ is defined in Eq. (10).

$$n_{\vartheta,PB}^{\mathrm{tar}} = \min\left\{|H_\xi^F|,\ p + \max\left\{e_{PB},\ \lceil \mu_{PB} p \rceil,\ \left\lceil \nu_{\vartheta,PB} |H_\xi^F| \right\rceil\right\}\right\} \tag{10}$$

Here, e_{PB} maintains a minimum number of alternatives beyond the station budget, μ_{PB} scales the additional breadth with p, and $\nu_{\vartheta,PB}$ scales exploration with the scenario-domain size. Table 2 reports the three jointly varied pool-breadth profiles.

Once INIT produces a validated incumbent with station-opening vector $\bar{z}$, the remaining phases restrict the Hamming distance from that incumbent using Eq. (11).

$$\sum_{h \in S_\vartheta:\, \bar{z}_h = 1} (1 - z_h) + \sum_{h \in S_\vartheta:\, \bar{z}_h = 0} z_h \le \delta_\vartheta, \qquad \vartheta \in \{\mathrm{S1}, \mathrm{S2}, \mathrm{ELITE}\} \tag{11}$$

The first term in Eq. (11) counts incumbent-open stations that are closed, while the second counts incumbent-closed stations that are opened. The constraint bounds the total number of changes in station opening status. When the incumbent and candidate configurations contain the same number of open stations, a radius of $2s$ permits at most s station exchanges. The tested neighborhood profiles are specified in Eq. (12), where the three entries apply to S1,

Table 2: LBFO pool-breadth profiles

Profile	Interpretation	e_{PB}	μ_{PB}	$\nu_{\text{INIT},PB}$	$\nu_{\text{S1},PB}$	$\nu_{\text{S2},PB}$	$\nu_{\text{ELITE},PB}$
$PB1$	Lean	5	0.15	0.06	0.03	0.05	0.08
$PB2$	Medium	10	0.25	0.10	0.05	0.08	0.10
$PB3$	Broad	15	0.40	0.15	0.08	0.12	0.15

S2, and ELITE, respectively.

$$N1 = [2, 4, 6], \qquad N2 = [2, 6, 10], \qquad N3 = [2, 8, 12] \tag{12}$$

Each radius is attempted once in phase order, and an improving validated candidate becomes the incumbent for the following phase.

A hybrid treatment of the path-assignment variables reduces the computational cost of restricted solves. PTTS and both DCLEX stages are first attempted with binary station and link decisions but continuous path variables, $0 \le x_k \le 1$. If the resulting assignment is integral within numerical tolerance, it proceeds directly to validation. Otherwise, the selected station network is fixed, x_k is restored to binary, and the link and assignment variables are reoptimized. If this recovery fails, the station decisions are released and the restricted HLP is solved with all variable groups binary. The relaxation is therefore only an intermediate computational device; every accepted solution is integral and validated against the complete scenario-specific HLP. The DCLEX recovery procedure also preserves its lexicographic ordering: assigned demand is optimized first, and PTTS is improved only while that value remains fixed. Each validated candidate updates solution memory, but it replaces the incumbent only if it improves the applicable objective ordering. After ELITE, LBFO returns the best validated integral incumbent obtained across all four phases.

Detailed implementation settings and computational specifications are reported in Section S3 of the Supplementary Material. The MATLAB implementation, configuration files, and reproduction instructions used for the reported experiments are publicly available.

3.5. Strategic Evaluation Procedures

Each HLP solve produces one network for a given retained-demand scenario, planning objective, and station budget. Three subsequent procedures translate these task-level solutions into strategic planning information: (1) representative-scale selection, (2) nested deployment, and (3) trip-specific operational-allowance evaluation.

For subsequent comparisons, Eq. (13) defines the opened-station H^{open} and physical-corridor sets L^{phy}. A physical corridor is counted once when its associated directed link is activated in either direction.

$$\begin{aligned} H^{\text{open}} &= \left\{ h \in H^F_\xi : z_h = 1 \right\}, \\ L^{\text{phy}} &= \left\{ \{a, b\} : (a, b) \in L^F_\xi,\ y_{ab} = 1 \right\} \end{aligned} \tag{13}$$

The cardinalities $|H^{\text{open}}|$ and $|L^{\text{phy}}|$ give the numbers of opened stations and selected physical corridors, respectively.

Representative infrastructure scale. The station budget p in Constraint (6b) is a planning input rather than a known or cost-optimal network size. Because reliable capital and operating costs for metropolitan UAM infrastructure are not yet available, a representative budget is identified from diminishing returns in network benefit.

The HLP is evaluated over an ordered sequence of budgets $p_1 < p_2 < \cdots < p_J$. For scenario ξ and objective $o \in \{\text{PTTS}, \text{DCLEX}\}$, let $V_{\xi,o}(p)$ denote the primary objective value at budget p, equal to $PTTS^{\text{total}}$ under PTTS and $Q^{\dagger}$ under DCLEX. The percentage gain from moving to the next tested budget is given by Eq. (14).

$$g_{\xi,o}(p_j) = 100\, \frac{V_{\xi,o}(p_{j+1}) - V_{\xi,o}(p_j)}{V_{\xi,o}(p_j)}, \qquad j = 1, \ldots, J-1 \tag{14}$$

Given a diminishing-return threshold ϵ^p, the transition budget for each objective is the first tested budget whose next increment produces a gain below this threshold. The representative budget for scenario ξ is the larger of the

PTTS and DCLEX transition budgets, so both planning priorities have reached diminishing returns. To assess the infrastructure productivity of expanding the retained-demand scope, cross-scenario changes in assigned demand and total PTTS are also normalized by the corresponding net increase in physical corridors.

Nested deployment sequence. Once a terminal network is selected from the representative solutions, a nested procedure identifies how it can be constructed in stages. Let H^* denote the opened stations in the terminal network. The largest deployment stage is the selected network itself; the smaller candidate stages, separated by the deployment increment Δp^{D}, are solved sequentially downward from it and satisfy $p_1^{\mathrm{D}} < \cdots < |H^*|$. At every stage, the path domain is restricted to retained paths whose trip is served in the terminal solution and whose endpoint stations and directed link belong to that solution. This prevents intermediate stages from relying on infrastructure or markets outside the final system.

The station-budget constraint is replaced by the exact station count and nesting requirements in Eq. (15).

$$\sum_{h \in H^*} z_h^{(p_j^{\mathrm{D}})} = p_j^{\mathrm{D}}, \qquad z_h^{(p_j^{\mathrm{D}})} \le \bar{z}_h^{(p_{j+1}^{\mathrm{D}})}, \quad h \in H^* \tag{15}$$

Here, $\bar{z}_h^{(p_{j+1}^{\mathrm{D}})}$ is the station-opening decision fixed at the next larger stage, so each stage opens a subset of the stations available at the stage above it. Only station openings are constrained to be nested. Path assignments and directed-link activations are reoptimized at each stage subject to the original feasibility requirements; consequently, the served trips and physical corridors need not be nested.

Trip-specific operational allowances. Operational overhead is not known when the strategic network is selected. Its effects are therefore evaluated after optimization without changing the selected infrastructure or trip assignments. Let R^{UAM} denote the trips assigned to UAM and let $\hat{k}(r)$ be the selected path of $r \in R^{\mathrm{UAM}}$.

For each assigned trip, $\tau_r^{\max}$ denotes the largest additional operational time that its selected path can absorb while retaining the relative-saving requirement $\rho_{\min}^{\mathrm{pre}}$. The scenario threshold Γ_ξ remains a no-overhead domain-construction threshold and is not reapplied as an operational service requirement. These trip-specific allowances induce the demand-retention function $S^Q(\tau)$, which measures the share of assigned demand whose allowance is at least threshold τ, and the corresponding retention benchmark τ_c^{cap} for target proportion c, as defined in Eq. (16).

$$\begin{aligned} \tau_r^{\max} &= \Delta T_{\hat{k}(r)}^{\mathrm{adv}} - \rho_{\min}^{\mathrm{pre}} T_r^G, \qquad r \in R^{\mathrm{UAM}} \\ S^Q(\tau) &= 100\, \frac{\sum_{r \in R^{\mathrm{UAM}}} q_r\, \mathbf{1}\{\tau_r^{\max} \ge \tau\}}{Q^\dagger} \\ \tau_c^{\mathrm{cap}} &= \sup\left\{\tau \ge 0 : S^Q(\tau) \ge 100c\right\}, \qquad 0 < c \le 1 \end{aligned} \tag{16}$$

Because every selected path satisfies the pre-optimization relative saving requirement, $\tau_r^{\max}$ is nonnegative. The function $S^Q(\tau)$ reports the percentage of assigned demand whose individual allowance equals or exceeds τ. Accordingly, $\tau_{0.95}^{\mathrm{cap}}$ is the allowance equaled or exceeded by at least 95% of assigned demand, while $\tau_{0.50}^{\mathrm{cap}}$ is the demand-weighted median allowance.

These quantities summarize trip-specific time-saving margins rather than imposing a common operational delay or predicting realized delays or passenger behavioral tolerance. Their distributions are also summarized by station using the demand weights of assigned trips whose selected paths use that station as either endpoint, without distinguishing access from egress. Determining the heterogeneous overheads realized in service requires subsequent scheduling, route-planning, and fleet-sizing analysis.

4. Greater Los Angeles Case Study and Computational Setup

This section translates the proposed demand-to-network framework into a reproducible Greater Los Angeles case study and establishes the computational design used to evaluate LBFO. Section 4.1 describes the regional demand and candidate-station inputs, constructs the retained-demand scenario domains, and specifies the planning parameters.

Section 4.2 then defines the LBFO calibration, transfer, controlled-comparison, and Gurobi benchmarking protocols, together with the solution-quality measures and computing environment. These domains and procedures form the basis for the computational and planning results in Section 5.

4.1. Case-Study Data, Scenario Domains, and Parameterization

The case study covers Greater Los Angeles, a polycentric metropolitan region with dispersed residential and employment centers and three major transportation anchors. The constructed potential-demand set described in Section 3.1 combines calibrated home-to-work movements (U.S. Census Bureau, LEHD Program, 2019; U.S. Census Bureau, 2019b,a) with passenger activity associated with Los Angeles International Airport (LAX) and Hollywood Burbank Airport (Bureau of Transportation Statistics, 2019b,a). Rail-associated demand is represented through Los Angeles Union Station (Amtrak, 2020). Figure 3 maps the constructed demand and candidate stations together with the allowed and restricted census blocks used in candidate generation (Zamani Saghazadeh et al., 2026).

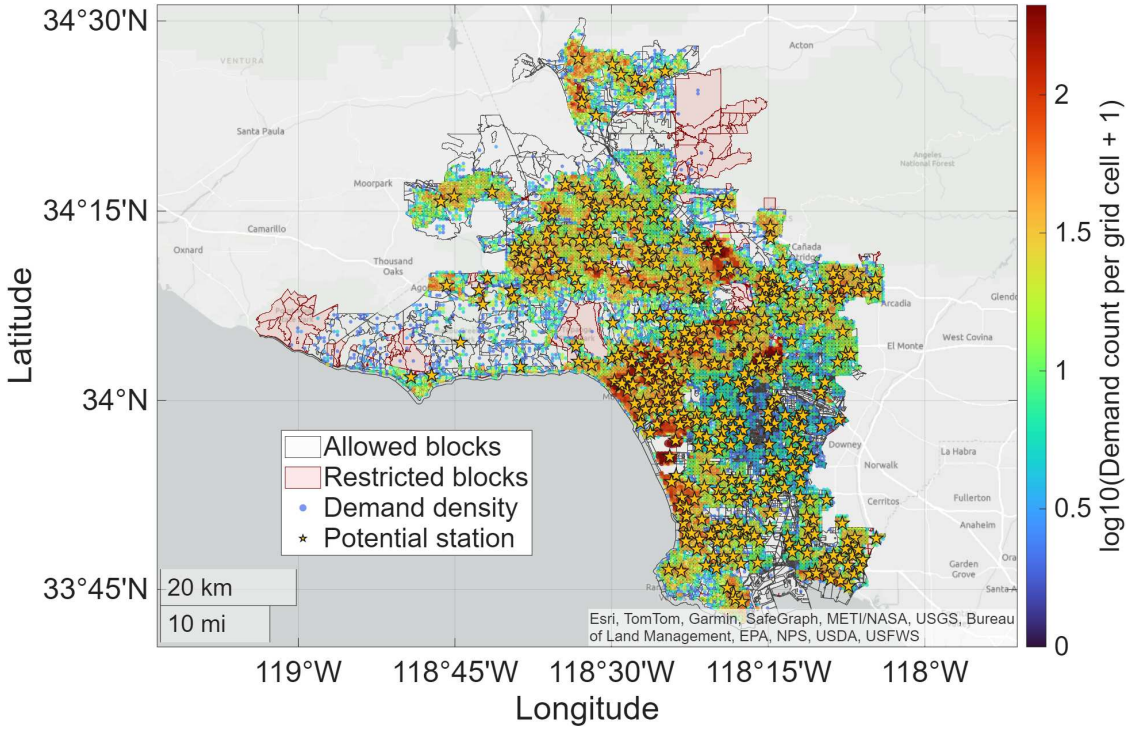


Figure 3: Constructed potential UAM demand, 430 physical candidate stations, and candidate-generation restrictions in Greater Los Angeles

Table 3 reports the principal domain sizes and planning parameters. The value $\rho^{\text{pre}}_{\min} = 0.40$ is adopted as a planning assumption, consistent with the travel-time-saving eligibility threshold used in earlier UAM network-planning and vertiport-location studies (Rajendran and Zack, 2019; Escribano Macias et al., 2023).

The retained-demand construction in Section 3.2 is evaluated at target shares of 10%, 30%, 60%, and 90%, denoted Xi10, Xi30, Xi60, and Xi90, respectively. Among trips represented in P^{base}, approximately 85% have a best no-overhead travel-time advantage between 25 and 65 min (Figure 4(a)). This concentration produces the steep middle portion of the retained-demand curve in Figure 4(b). Lowering Γ_ξ expands the planning scope but admits trips with smaller no-overhead time advantages.

Table 4 reports the resulting thresholds and final HLP domains. Route screening reduces the 1,286,305 complete candidate paths to 969,407 initially eligible paths, and base HLP-safe pruning retains 862,198 paths. We report the intermediate reductions and scenario-specific pruning audit in Section S2 and Table S1 of the Supplementary Material.

Increasing the target share enlarges both the HLP domain and the number of station-pair alternatives per retained trip. The four scenarios therefore provide a controlled range of instance sizes for evaluating LBFO transfer and

Table 3: The Greater Los Angeles case study uses the listed domain sizes and HLP settings. Demand weights and minimum-utilization thresholds are applied separately to 07:00–12:00 demand and approximate reverse movements during 15:00–20:00; the thresholds are activity requirements rather than capacity limits.

Parameter	Symbol	Value
Constructed OD trip records	$\|R\|$	156,215
Physical candidate stations	$\|H\|$	430
Study-area ZCTAs	$\|Z\|$	201
Trip demand weight	q_r	1
Tested station budgets	p	$\{5, 10, \ldots, 60\}$
Representative-budget threshold	ϵ^p	5%
Deployment-stage increment	Δp^{D}	5 stations
Minimum station-pair air distance	$D^{\mathrm{air}}_{\mathrm{min}}$	15 km
Maximum station-pair air distance	$D^{\mathrm{air}}_{\mathrm{max}}$	128 km
Minimum relative no-overhead saving	$\rho^{\mathrm{pre}}_{\mathrm{min}}$	0.40
Substitution time-difference tolerance	δ^{sub}	5 min
Minimum station load	$Q^{\mathrm{st}}_{\mathrm{min}}$	100 trip ends
Minimum directed-link flow	$Q^{\mathrm{link}}_{\mathrm{min}}$	20 potential passenger trips

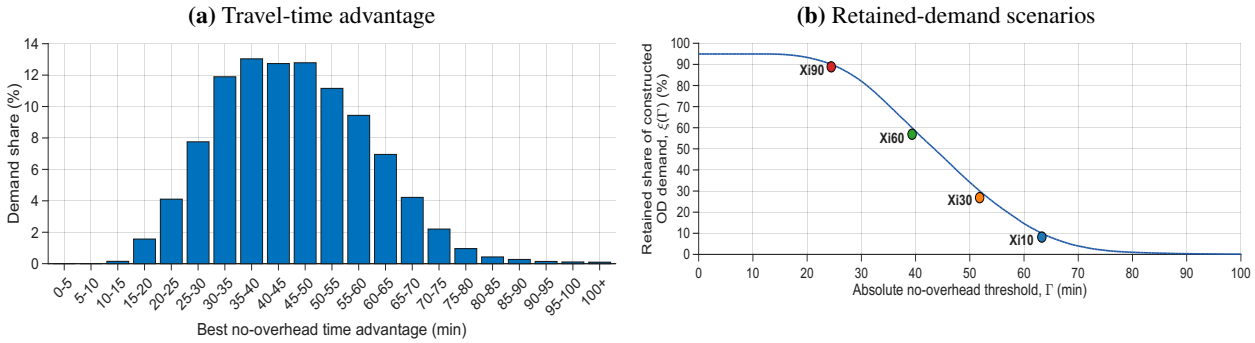


Figure 4: Route-quality and retained-demand characterization: (a) best no-overhead travel-time advantage among trips in P^{base}; and (b) retained share versus Γ, including the final shares after HLP-safe pruning.

scalability while preserving the trip-specific representation.

Table 4: Retained-demand scenarios and final HLP domains

Scenario	Target ξ (%)	Γ_ξ (min)	Retained share (%)	$\|P^F_\xi\|$	$\|R^F_\xi\|$	$\|H^F_\xi\|$	$\|L^F_\xi\|$	Paths per retained trip	$\|C^{\mathrm{sub}}_\xi\|$
Xi10	10	63.29	8.21	31,133	12,819	72	316	2.43	25
Xi30	30	51.82	26.87	130,125	41,975	196	1,261	3.10	78
Xi60	60	39.35	56.92	374,291	88,923	315	3,440	4.21	133
Xi90	90	24.45	88.85	763,339	138,795	376	6,841	5.50	153

Note: Shares are relative to all constructed trips and may fall below their targets after scenario-specific HLP-safe pruning. Paths per retained trip is $|P^F_\xi|/|R^F_\xi|$.

4.2. Computational Protocol and Benchmarking

LBFO calibration on Xi10 evaluates all $4 \times 3 \times 3 = 36$ combinations of the W, N, and PB profiles under both objectives and all 12 station budgets (total of 864 optimization tasks). Five shortlisted configurations are transferred unchanged to Xi30. Controlled comparisons isolate the contributions of solution memory and the hybrid assignment treatment, while pure VNS provides an external heuristic baseline. The selected LBFO configuration is then fixed without further retuning for the production analysis. This analysis comprises 96 HLP tasks spanning four retained-demand scenarios, 12 station budgets, and two objectives. A complete Gurobi benchmark sweep over all 12 station

budgets is performed only for Xi10, for which the scenario domain remains computationally tractable. For Xi30–Xi90, Gurobi is instead evaluated at each scenario's representative budget because repeating the complete budget sweep over these substantially larger domains would be computationally prohibitive.

Solution quality is evaluated against complete-domain Gurobi results. For primary objective value V, the percentage shortfalls from the Gurobi incumbent and its strongest valid upper bound are defined as:

$$\delta^{\mathrm{inc}} = 100\,\frac{\left[V^{\mathrm{G,inc}} - V^{\mathrm{LBFO}}\right]_+}{V^{\mathrm{G,inc}}} \quad , \quad \delta^{\mathrm{bnd}} = 100\,\frac{\left[V^{\mathrm{G,bnd}} - V^{\mathrm{LBFO}}\right]_+}{V^{\mathrm{G,bnd}}} \tag{17}$$

where $[u]_+ = \max\{u, 0\}$, $V = PTTS^{\mathrm{total}}$ for PTTS, and $V = Q^{\dagger}$ for the lexicographically primary DCLEX stage. Only integral LBFO solutions validated against the complete scenario-specific HLP are retained. Gurobi incumbents and bounds are used solely for external evaluation and do not inform LBFO scoring, memory, pool construction, or incumbent generation.

All experiments are implemented in MATLAB R2025b with Gurobi 13.0.2 and executed on a 64-bit Windows 11 workstation with an Intel Core Ultra 7 265K processor and 32 GB of RAM; Gurobi is limited to 12 threads. All reported complete-domain Gurobi benchmarks use the same hybrid assignment treatment as the production LBFO method. Here, `MIPGap` denotes Gurobi's relative MIP optimality-gap tolerance. It is set to zero for the complete-domain Xi10 benchmarks. For larger complete-domain benchmarks and restricted LBFO solves, it is set to 0.005 (0.5%) for PTTS and DCLEX Stage 2 and to zero for DCLEX Stage 1. Reported LBFO runtimes cover the complete solution and validation procedure but exclude reporting exports.

5. Computational Results and Planning Interpretation

This section first evaluates LBFO calibration, transferability, and controlled computational performance. The validated production solutions are then interpreted in terms of representative network scale, selected-network structure, phased deployment, and trip-level operational allowances.

5.1. LBFO Calibration, Transfer, and Benchmark Performance

Figure 5 summarizes the complete Xi10 calibration of the 36 W–N–PB configurations. For each configuration, the plotted mean shortfalls are the unweighted arithmetic averages of the task-level values defined in Eq. (17) across the 24 objective–budget tasks, and total runtime is summed over the same tasks.

Configurations requiring the unrestricted station domain in at least one task are unsuitable for scalable restricted-domain optimization and are therefore excluded. A configuration is also excluded if its maximum certified-bound shortfall exceeds 1% in any calibration task. This maximum shortfall criterion serves as a robustness screen rather than the sole ranking measure; the configurations that pass it are compared using their mean shortfalls and total runtime. The remaining configurations exhibit different quality–runtime trade-offs; $W4$–$N2$–$PB1$ is the fastest, whereas configurations using $PB2$ generally provide smaller mean shortfalls at greater computational cost. Because no finalist dominates across quality and runtime, all five are transferred without retuning to the larger Xi30 domain.

The transfer results in Panel A of Table 5 separate the finalists more clearly. The $W1$–$N1$–$PB2$ configuration produces the smallest mean incumbent and certified-bound shortfalls for both objectives and is the only finalist whose maximum certified-bound shortfall remains below 1% for both PTTS and DCLEX. It is therefore selected for the controlled comparisons and subsequent production analysis.

With the profiles fixed, Figure 6 evaluates the remaining architectural choices over the complete Xi10 budget sweep. The production method is compared with its binary-assignment and no-memory variants, pure VNS, and complete-domain Gurobi.

The memory-enabled binary LBFO variant achieves mean certified-bound shortfalls of 0.0016% for PTTS and 0.0092% for DCLEX, with corresponding maxima of 0.0060% and 0.0991%, but requires 23.66 min across the two objective sweeps, comparable to the 23.53 min required by the hybrid-formulation Gurobi benchmarks. The production hybrid treatment reduces this runtime to 9.92 min, a reduction of approximately 58%, while returning validated integral solutions with mean certified-bound shortfalls of 0.1028% for PTTS and 0.0675% for DCLEX and corresponding maxima below 0.47%. This result supports using continuous path variables during the initial restricted solves, followed by integral recovery whenever required.

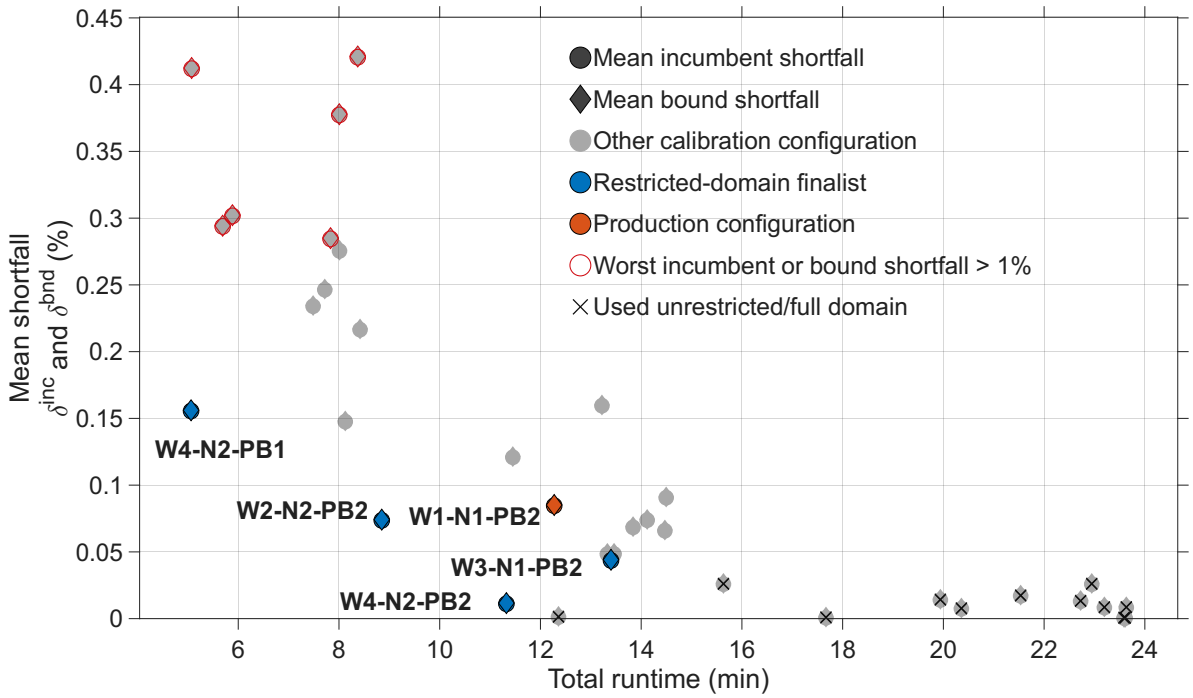


Figure 5: Quality–runtime trade-offs for 36 LBFO configurations across 24 Xi10 tasks. Circles and diamonds show mean incumbent and certified-bound shortfalls. Blue marks the five restricted-domain finalists; orange marks the selected configuration. Red outlines indicate a maximum certified-bound shortfall above 1%, and crosses indicate use of the unrestricted station domain

The smaller Xi10 domain does not provide conclusive evidence for solution memory; removing memory slightly reduces both the mean and maximum certified-bound shortfalls for PTTS but increases both measures for DCLEX. The matched Xi30 comparison in Panel B of Table 5 provides a clearer distinction. Memory reduces the mean incumbent shortfall by more than three quarters for both objectives. It also reduces the mean certified-bound shortfall from 2.3317% to 0.9429% for PTTS and from 1.9861% to 0.3866% for DCLEX, while reducing the corresponding maxima from 3.1830% to 0.9603% and from 2.6356% to 0.6009%, respectively.

Pure VNS satisfies the same HLP constraints but explores station configurations without the objective-aware scores and accumulated solution memory used by LBFO. Despite a comparable total runtime, its mean certified-bound shortfalls increase to 6.67% for PTTS and 7.44% for DCLEX, with corresponding maxima of 14.15% and 15.75%. The combined calibration, transfer, and controlled comparisons therefore support the memory-enabled hybrid $W1$–$N1$–$PB2$ configuration. This configuration is fixed without scenario-specific retuning for the 96 production tasks used in the remaining analysis.

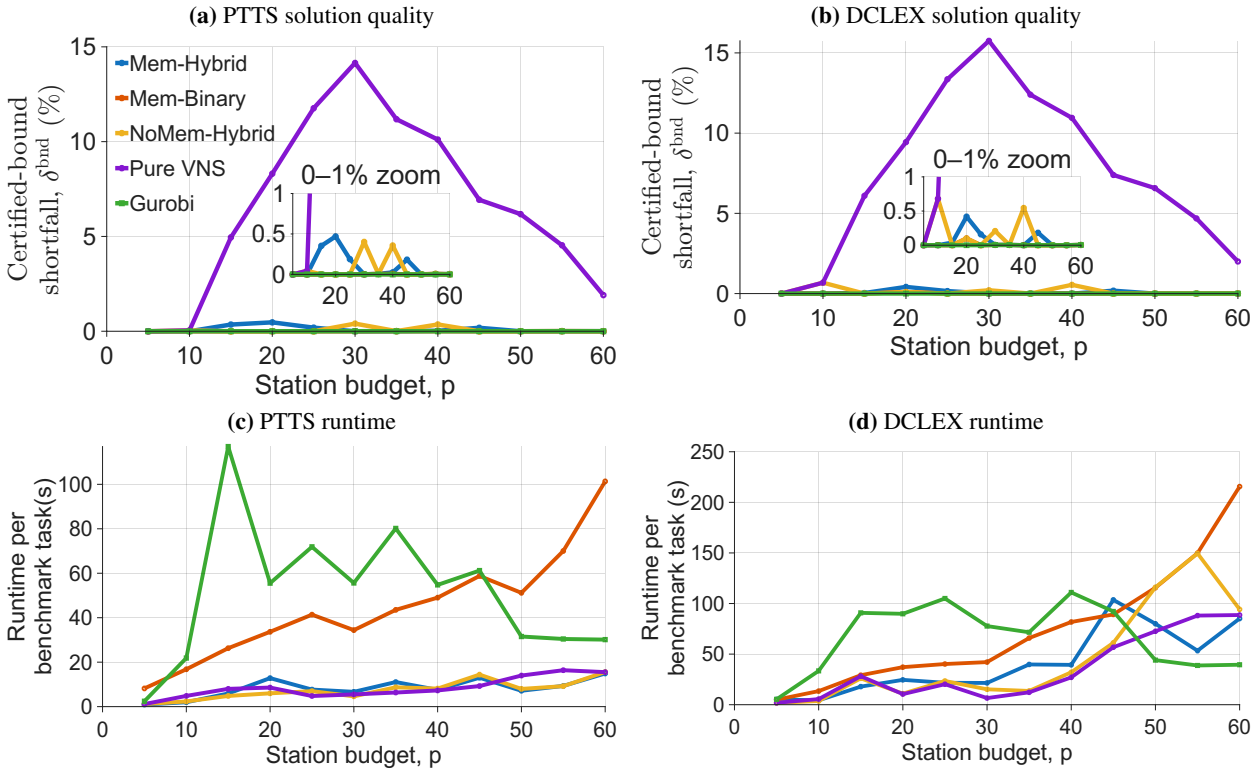


Figure 6: Controlled comparison over the complete Xi10 budget sweep: certified-bound shortfall for (a) PTTS and (b) DCLEX, and task runtime for (c) PTTS and (d) DCLEX. Insets magnify the 0–1% range; the legend in panel (b) applies throughout.

Table 5: Xi30 transfer performance and solution-memory comparison

Panel A. Shortlisted-configuration transfer

Configuration	Objective	Runtime (min)	Mean δ^{inc} (%)	Mean δ^{bnd} (%)	Maximum δ^{bnd} (%)
*W*4–*N*2–*PB*1	PTTS	1.34	1.7382	2.2270	2.5745
*W*4–*N*2–*PB*1	DCLEX	4.28	1.4599	1.4603	1.6046
*W*2–*N*2–*PB*2	PTTS	5.76	0.5272	1.0220	1.1511
*W*2–*N*2–*PB*2	DCLEX	7.19	0.4326	0.4330	0.6403
*W*4–*N*2–*PB*2	PTTS	11.36	0.5119	1.0068	1.1055
*W*4–*N*2–*PB*2	DCLEX	13.46	0.5727	0.5732	0.6043
*W*1–*N*1–*PB*2	**PTTS**	**5.81**	**0.4477**	**0.9429**	**0.9603**
*W*1–*N*1–*PB*2	**DCLEX**	**7.17**	**0.3861**	**0.3866**	**0.6009**
*W*3–*N*1–*PB*2	PTTS	5.14	0.5119	1.0068	1.1055
*W*3–*N*1–*PB*2	DCLEX	7.60	0.5737	0.5742	0.6043

Panel B. Matched solution-memory comparison

Method	Objective	Runtime (min)	Mean δ^{inc} (%)	Mean δ^{bnd} (%)	Maximum δ^{bnd} (%)
Memory-enabled hybrid	**PTTS**	**9.50**	**0.4477**	**0.9429**	**0.9603**
Memory-enabled hybrid	**DCLEX**	**8.46**	**0.3861**	**0.3866**	**0.6009**
No-memory hybrid	PTTS	4.29	1.8435	2.3317	3.1830
No-memory hybrid	DCLEX	7.85	1.9856	1.9861	2.6356

Note: Matched tasks use $p \in \{30, 35, 40\}$ for PTTS and $p \in \{35, 40, 45\}$ for DCLEX. Runtimes are totals over the matched tasks; mean shortfalls are their unweighted averages, and maxima are taskwise. Panels A and B are separate experiments. Bold marks the production configuration and method.

5.2. Network Scale and Representative-Budget Behavior

Figure 7 reports the next-step primary-objective gains from the production sweeps. Applying the 5% diminishing-return rule in Eq. (14) gives representative budgets of 35, 40, 45, and 45 stations for Xi10 through Xi90, respectively.

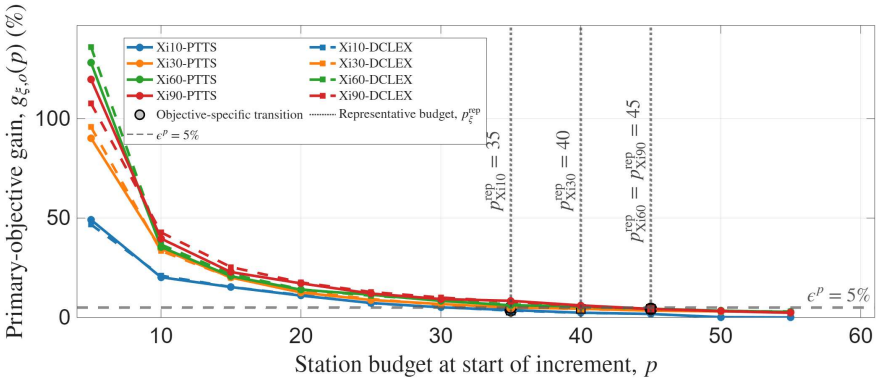


Figure 7: Next-step primary-objective gains and representative station budgets. Outlined markers show the first gain below $\epsilon^p = 5\%$; dotted lines mark the resulting p_ξ^{rep} values

Table 6 summarizes the representative network outcomes and matched Gurobi benchmarks solved over the complete scenario-specific HLP domain at the same station budgets. As the planning scope expands from Xi10 to Xi90, the complete model grows from 31,879 decision variables and 14,337+1 constraints to 770,610 decision variables and 153,491+1 constraints, depending on the objective stage. The corresponding two-objective Gurobi runtime at the representative budget increases sharply from 2.53 min for Xi10 to 11,763.85 min for Xi90, illustrating why complete Gurobi budget sweeps become impractical for the broader scenarios. Across all eight representative-budget comparisons, the largest incumbent and certified-bound shortfalls are 0.4622% and 0.9573%, respectively. Summed across the two objectives within each scenario, LBFO reduces runtime by 83.5–99.5% relative to the corresponding Gurobi benchmarks while maintaining a certified-bound shortfall below 1% in every task. For Xi90, the two LBFO tasks require 110.70 min in total, compared with 11,763.85 min for the corresponding Gurobi benchmarks.

Moving from Xi10 to Xi30 increases assigned demand from 11,338 to 30,634 trips and total PTTS from 0.80 to 1.89 million passenger-minutes, while adding 60 physical corridors. Broader scenarios increase aggregate coverage but require substantially denser networks. Incremental returns per net added physical corridor are approximately 322 assigned trips and 18.2×10^3 passenger-minutes of PTTS from Xi10 to Xi30. They decline to 236–279 trips and $10.4\text{–}11.9 \times 10^3$ passenger-minutes from Xi30 to Xi60, before increasing to 610–633 trips and $18.4\text{–}19.4 \times 10^3$ passenger-minutes from Xi60 to Xi90. This nonmonotonic pattern shows that corridor productivity alone cannot determine the representative planning scope.

Figure 8 provides a complementary comparison based on the trip-specific operational allowances defined in Eq. (16). At each threshold, the curves report the share of assigned demand whose trip-level allowance equals or exceeds that value. The thresholds exceeded by 95% of assigned demand are approximately 24.3 min for Xi10, 19.0 min for Xi30, 10.8–11.6 min for Xi60, and 3.9–4.1 min for Xi90.

The decline in operational allowance is consistent with the retained-demand construction. Increasing ξ lowers Γ_ξ and admits trips with smaller no-overhead absolute advantages. As Eq. (16) shows, the remaining allowance also depends on direct-ground travel time; empirically, the broader scenarios contain more assigned trips closer to the 40% relative-saving boundary. They therefore trade greater market coverage for denser infrastructure and smaller trip-level operational allowances.

Taken together, the results identify Xi30 at $p = 40$ as a balanced representative planning case. Compared with Xi10, it substantially expands coverage and passenger benefit, while 95% of its assigned demand has a trip-level operational allowance of at least approximately 19 min. The broader Xi60 and Xi90 scenarios provide greater coverage but require considerably more physical corridors and have much smaller 95%-retention allowances. PTTS and DCLEX also produce the same representative Xi30 network and aggregate outcomes. This network is therefore carried into Section 5.3 for structural and phased-deployment analysis. Its selection represents a case-study planning compromise rather than an economic optimum or a universally preferred demand scenario.

Table 6: Representative-network outcomes and matched Gurobi benchmarks

Panel A. Network outcomes at p_ξ^{rep}

Scenario	Objective	p_ξ^{rep}	$\lvert L^{\text{phy}}\rvert$	$Q^\dagger$ (coverage, %)	$PTTS^{\text{total}}$ (10^6 passenger-min)
Xi10	PTTS	35	100	11,338 (7.26)	0.80
Xi10	DCLEX	35	100	11,338 (7.26)	0.80
Xi30	PTTS	40	160	30,634 (19.61)	1.89
Xi30	DCLEX	40	160	30,634 (19.61)	1.89
Xi60	PTTS	45	298	63,262 (40.50)	3.32
Xi60	DCLEX	45	278	63,504 (40.65)	3.30
Xi90	PTTS	45	367	106,953 (68.47)	4.66
Xi90	DCLEX	45	350	107,446 (68.78)	4.62

Panel B. Task-level benchmark performance

Scenario	Objective	LBFO runtime (min)	Gurobi runtime (min)	δ^{inc} (%)	δ^{bnd} (%)
Xi10	PTTS	0.13	1.34	0.0051	0.0060
Xi10	DCLEX	0.28	1.19	0.0000	0.0037
Xi30	PTTS	1.45	29.29	0.4622	0.9573
Xi30	DCLEX	1.59	576.84	0.4323	0.4323
Xi60	PTTS	30.06	722.56	0.1692	0.5393
Xi60	DCLEX	49.10	988.71	0.2889	0.2889
Xi90	PTTS	53.70	5,760.03	0.1121	0.7701
Xi90	DCLEX	57.00	6,003.82	0.0595	0.0595

Note: Coverage is relative to all constructed OD trips. Runtimes are for the listed objective at its representative budget. DCLEX shortfalls concern Stage 1 assigned demand. Gurobi time excludes the LBFO run used to generate its MIP start.

5.3. Selected Network Structure and Phased Deployment

The selected 40-station Xi30 contains 160 physical corridors. Its graph is connected and contains no bridge corridor or articulation station. Eight selected-station pairs lie within 5 km by road, forming six proximity groups involving 14 stations, yet none of these stations can be removed without reducing assigned demand. Thus, geographic proximity does not imply functional redundancy in the selected network. Complete graph and replacement diagnostics are reported in Section S4 of the Supplementary Material.

Across the eight representative scenario–objective networks in Table 6, 11 stations appear in every solution and form a persistent cross-scenario backbone. They span major transportation anchors near LAX, Los Angeles Union Station, and Hollywood Burbank Airport; urban-core locations in Downtown Los Angeles and Century City; and regional gateways in Valencia, Simi Valley, the western San Fernando Valley, Palos Verdes, Agoura Hills, and Long Beach. Their repeated selection under both objectives and all four retained-demand scenarios indicates their strategic importance within the case study.

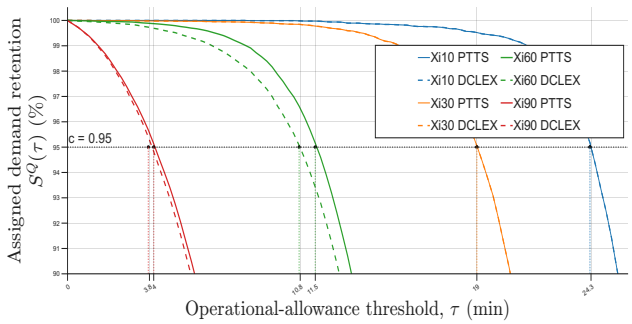


Figure 8: Assigned-demand retention versus operational-allowance threshold τ for the representative scenario networks. Markers and dotted lines identify $\tau_{0.95}^{\text{cap}}$

Using the terminal network as the fixed deployment domain and the nesting requirements in Eq. (15), the deployment sequence is evaluated over $p^{\mathrm{D}} \in \{5, 10, \ldots, 40\}$. Figure 9 shows four spatial milestones, and Table 7 summarizes their performance.

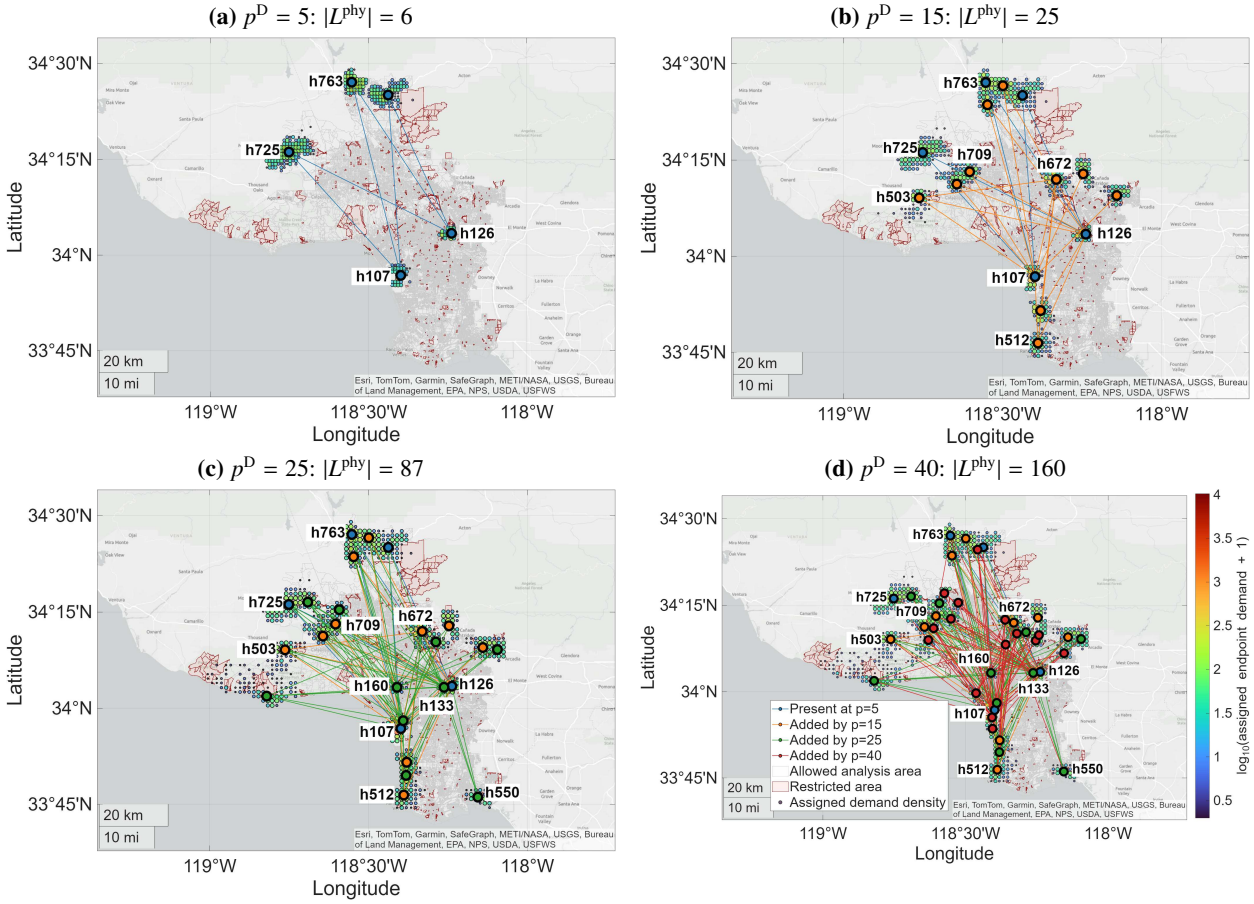


Figure 9: Nested deployment of the selected Xi30 network. Colors show the first displayed milestone containing each station and physical corridor. Panel (d) provides the shared legend and demand-density scale, labels the 11 backbone stations, and outlines the six road-proximity groups

The five-station core contains four backbone stations and captures 23.02% of terminal assigned demand and 23.71% of terminal PTTS using six physical corridors. Both capture rates reach approximately 60% at $p^{\mathrm{D}} = 15$ and 81% at $p^{\mathrm{D}} = 25$, by which point all 11 backbone stations are present. Later stages extend and densify the airport–urban-core–regional-gateway structure. Expanding from 25 to 40 stations adds 5,774 assigned trips and approximately 345,437 passenger-minutes of PTTS while requiring 73 additional physical corridors. The early capture rates demonstrate the benefits of incremental deployment but do not justify stopping before reaching the selected terminal network. Complete deployment results and the station-opening sequence are provided in Section S4.2 of the Supplementary Material. The next subsection examines trip- and station-level operational allowances.

Table 7: Selected-network deployment milestones

p^{D}	$\lvert L^{\mathrm{phy}}\rvert$	$Q^{\dagger}$	Demand captured (%)	$PTTS^{\mathrm{total}}$ (10^6 passenger-min)	PTTS captured (%)
5	6	7,053	23.02	0.448	23.71
15	25	18,496	60.38	1.136	60.10
25	87	24,860	81.15	1.544	81.72
40	160	30,634	100.00	1.890	100.00

Note: Capture is relative to the terminal $p^{\mathrm{D}} = 40$ network.

5.4. Trip- and Station-Level Operational Allowances

The operational-allowance measures defined in Eq. (16) quantify the additional time that each selected path can absorb while retaining the 40% relative saving requirement. They are evaluated after network optimization using the selected paths in the representative Xi30 network at $p = 40$. These measures characterize the remaining time-saving margin of the strategic assignments; they do not predict realized waiting, processing, charging, or coordination delays.

Figure 10 shows that the trip-level allowances are concentrated mainly between 20 and 30 min, with a wider right tail. The demand-weighted median is 25.7 min, while $\tau^{\mathrm{cap}}_{0.95} = 19.0$ min indicates that 95% of assigned demand has an individual allowance of at least 19.0 min. Each $\tau^{\max}_r$ provides the downstream operational model with a trip-specific service threshold; an assigned trip retains the required 40% advantage over direct ground travel only if its realized operational overhead does not exceed this allowance. The allowances can therefore be used to determine how much assigned demand is served within the strategic travel-time requirement after scheduling, fleet circulation, charging, and vertiport operations are modeled.

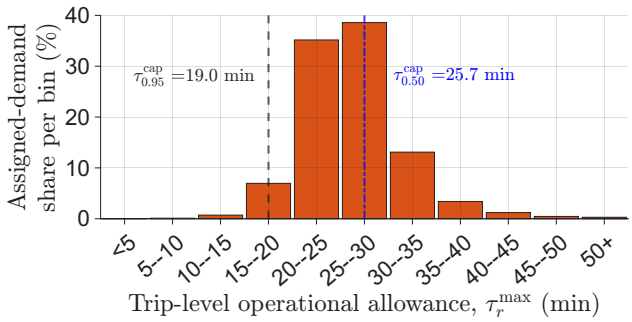


Figure 10: Assigned-demand-weighted distribution of trip-level operational allowances in the selected Xi30 network. Vertical lines identify the demand-weighted median and the 95%-retention allowance $\tau^{\mathrm{cap}}_{0.95}$.

Figure 11 provides two complementary station-level inputs for downstream operational modeling. Panel (a) groups the trip-specific allowances by the backbone station used, showing that their demand-weighted distributions differ despite the persistent selection of these stations. The station-specific 95%-retention allowances, represented by the lower endpoints of the thin intervals, range from 17.2 min at $h503$ to 21.0 min at $h725$. Panel (b) reports the corresponding outgoing and incoming assigned trip-end demand during 07:00–12:00 and reveals substantial directional imbalance across the stations. The complete station inventory is reported in Table A.2 in Appendix A.

Together, the two panels show how the selected network can inform subsequent operational planning. The station-level allowance distributions provide benchmarks for time-sensitive decisions, including passenger consolidation and service scheduling, whereas the directional-demand profiles indicate where aircraft imbalances and asymmetric infrastructure use may arise. These outputs help prioritize stations for fleet-positioning, scheduling, and capacity analysis, while realized delays, throughput, and service performance remain to be determined by downstream operational models. Additional relative-saving distributions and operational-allowance retention benchmarks are provided in Section S5 of the Supplementary Material.

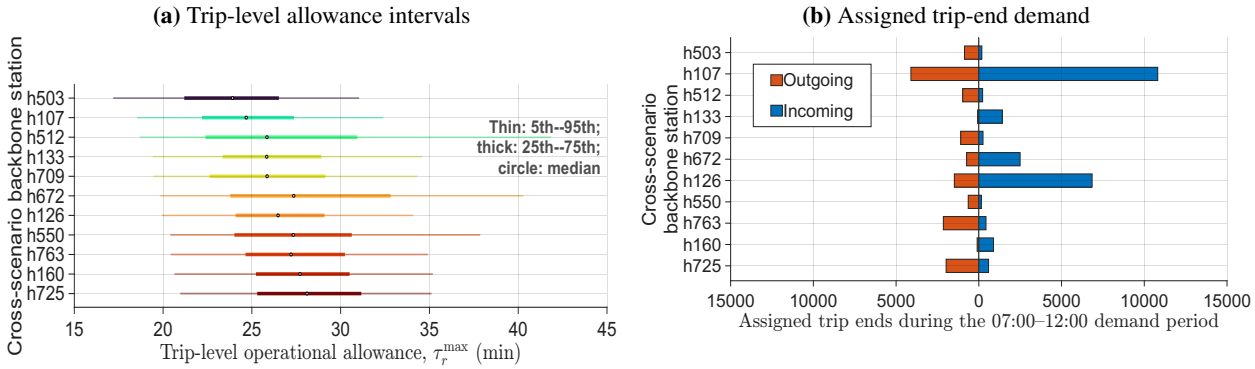


Figure 11: Operational characteristics of the 11 cross-scenario backbone stations: (a) demand-weighted distributions of trip-level allowances by station, with thin lines showing the 5th–95th percentiles, thick lines the interquartile range, circles the median, and the lower thin-line endpoint the station-specific 95%-retention allowance; and (b) outgoing (left) and incoming (right) assigned trip-end demand during 07:00–12:00

6. Discussion

6.1. Computational Scalability and Contribution

Capturing trip-level differences in UAM accessibility comes at a computational cost. Spatially disaggregated demand and multiple station alternatives preserve differences in complete door-to-door journeys, but enlarge the path-assignment domain. The challenge extends beyond solving one network: strategic planning requires repeated optimization across demand scenarios, station budgets, and competing objectives. In this study, the largest formulation contains 770,610 variables and approximately 153,492 constraints, while the production analysis comprises 96 HLP tasks.

Earlier UAM planning models have limited the size and complexity of the optimization domain through spatial aggregation, clustering-based site preselection, or preprocessed single-allocation structures (Escribano Macias et al., 2023; Sinha and Rajendran, 2023; Rath and Chow, 2022). Complementing these modeling choices, tailored algorithms address the remaining combinatorial search. Willey and Salmon (2021) develop structured-network heuristics, evaluated against exact solutions on 12-candidate test networks and a 38-candidate application. For choice-based skyport location, Kim et al. (2025) report cross-entropy algorithms that are 4–95 times faster than CPLEX for one formulation and 2–35 times faster for another. The present study advances this computational emphasis by retaining trip-specific station-pair decisions within complete door-to-door paths and making their repeated optimization practical across demand scenarios, station budgets, and objectives. Matched complete-domain benchmarks quantify the resulting runtime savings and certified-bound shortfalls.

At representative budgets, LBFO reduces runtime summed across both objectives within each scenario by 83.5–99.5% relative to matched complete-domain Gurobi benchmarks, while maintaining primary-objective certified-bound shortfalls below 1% in all eight benchmarked tasks. Both approaches use the same hybrid assignment treatment, so the improvement extends beyond relaxing the path variables. Controlled comparisons further support the roles of solution memory and hybrid assignment, while LBFO achieves smaller mean shortfalls than the implemented VNS comparator at comparable runtime. Following calibration and transfer-based selection, the configuration remains fixed throughout production. The resulting computational savings enable systematic comparisons across demand scopes and infrastructure budgets while retaining trip-specific station alternatives. Gurobi benchmarking covers the complete Xi10 budget sweep and representative budgets in the larger scenarios.

6.2. Strategic and Operational Planning Implications

Repeated optimization across retained-demand scenarios reveals that broader coverage is not an unqualified improvement. Larger domains assign more trips and generate greater aggregate PTTS, but require denser corridor networks and include trips with smaller advantages and operational allowances. The selected Xi30 network therefore represents a case-study balance rather than a forecast or economic optimum. Agreement between PTTS and DCLEX

strengthens its selection, although their divergence in broader scenarios confirms that coverage and passenger benefit need not produce the same design. The persistent backbone and nested deployment sequence further identify priorities for staged infrastructure development.

The operational allowances make the resulting assignments testable in subsequent operational analysis. Earlier strategic models use fixed boarding, deboarding, or taxiing additions (Venkatesh et al., 2020; Willey and Salmon, 2021), while others assume no passenger waiting (Kim et al., 2025). Such assumptions can support scenario analysis, but do not establish whether each assignment remains competitive under realized operations. Here, $\tau_r^{\max}$ gives the trip-specific threshold illustrated in Figure 10; exceeding it violates the required 40% travel-time saving. Assuming too little operational overhead can make some trips appear to meet the required travel-time saving when they do not. On the other hand, assuming too much can exclude trips that could meet the requirement. Station-level allowance distributions and directional demand (Figure 11) complement these thresholds by identifying where limited time margins and asymmetric aircraft flows may coincide, guiding subsequent scheduling, fleet-positioning, and vertiport-capacity analysis.

6.3. Scope and Research Extensions

The framework can incorporate additional planning information through behavioral demand weights, station and corridor costs, and robust or stochastic demand scenarios. These additions would extend the present potential-demand and infrastructure-scale comparisons toward adoption and investment decisions. Operational feedback can likewise return simulated delays or throughput restrictions as revised path attributes or constraints, following sequential infrastructure–operations approaches (Escribano Macias et al., 2023) and external optimization interfaces such as AAMSim (Li and Zhang, 2026). The complete workflow can be applied to other regions using local demand, candidate stations, and transportation data. The present computational results provide a benchmark for evaluating these regional applications and future model extensions.

7. Conclusions

This study developed a hybrid local-branching and fix-and-optimize (LBFO) method for repeated UAM hub-location optimization with disaggregated potential demand and multiple station alternatives. Embedded in a demand-to-network framework, the method supports comparisons across retained-demand scenarios, station budgets, and complementary passenger-benefit and demand-coverage objectives.

At representative budgets, LBFO reduced runtime summed across both objectives within each scenario by 83.5–99.5% relative to matched complete-domain Gurobi benchmarks, with primary-objective certified-bound shortfalls below 1% in all eight tasks. Benchmarking covered the complete Xi10 budget sweep and representative budgets in the larger scenarios. These results demonstrate the computational practicality of repeated network comparisons while retaining trip-specific station alternatives.

The selected Greater Los Angeles network contains 40 stations and 160 physical corridors, assigns 30,634 potential trips, and generates 1.89×10^6 passenger-minutes of no-overhead travel-time saving. Its nested deployment sequence identifies infrastructure priorities, while trip-specific operational allowances provide service thresholds for subsequent operational models. For 95% of assigned demand, these allowances are at least 19.0 min while retaining the 40% travel-time-saving requirement. The workflow can be applied to other regions using local demand, candidate-station, and transportation inputs.

CRediT authorship contribution statement

Hossein Zamani Saghazadeh: Conceptualization, methodology, writing—original draft preparation. Soheil Sabri: Conceptualization, validation, supervision, review and editing. Parham Kebria: Validation, review and editing.

Declaration of competing interest

The authors declare no conflict of interest.

Funding

This research received no external funding.

Appendix A. Selected-Station Inventory

Tables A.1 and A.2 report the spatial, deployment, assignment, network, and operational-allowance characteristics of the 40 stations in the selected Xi30 network. Bold identifiers denote the 11 cross-scenario backbone stations.

Table A.1: Selected-station locations and deployment stages.

Station	ZCTA	Latitude	Longitude	First deployment stage, p^{D}
h107	90045	33.94706	-118.39793	5
h126	90012	34.05723	-118.23767	5
h645	91387	34.41726	-118.43772	5
h725	93065	34.26944	-118.75120	5
h763	91354	34.45081	-118.55264	5
h493	90278	33.85914	-118.38006	10
h512	90275	33.77408	-118.38894	10
h672	91504	34.19999	-118.33086	10
h709	91304	34.21915	-118.60374	10
h756	91355	34.39305	-118.54729	10
h503	91301	34.15202	-118.76279	15
h656	91104	34.15823	-118.14019	15
h679	91214	34.21401	-118.24622	15
h719	91307	34.18792	-118.64368	15
h762	91350	34.44270	-118.49895	15
h160	90067	34.05448	-118.41052	20
h289	91201	34.17214	-118.28774	20
h494	90265	34.03132	-118.82049	20
h550	90814	33.76896	-118.15541	20
h713	93063	34.27539	-118.69104	20
h133	90017	34.05483	-118.26260	25
h341	90045	33.96777	-118.39065	25
h486	90277	33.82483	-118.38144	25
h659	91107	34.15284	-118.09437	25
h728	91311	34.25631	-118.59184	25
h258	91030	34.11147	-118.15470	30
h309	91206	34.14810	-118.25384	30
h488	90266	33.89285	-118.40473	30
h499	91302	34.14950	-118.62975	30
h715	91335	34.21109	-118.55135	30
h283	90291	33.99626	-118.46324	35
h334	90245	33.92603	-118.40727	35
h358	91608	34.13664	-118.35886	35
h701	91367	34.18378	-118.61166	35
h775	91326	34.28456	-118.57386	35
h238	91506	34.16881	-118.31959	40
h419	91352	34.20820	-118.36139	40
h676	91207	34.16402	-118.24248	40
h758	91387	34.41113	-118.45758	40
h768	91325	34.25753	-118.52554	40

Note: First stage is the smallest nested network containing the station.

Table A.2: Selected-station directional demand, physical topology, and 95%-retention allowances.

Station	Outgoing trip-end demand	Incoming trip-end demand	Total trip-end demand	Physical degree	95% allowance (min)
h107	4,107	10,814	14,921	28	18.53
h126	1,480	6,848	8,328	22	19.92
h645	550	118	668	5	13.41
h725	1,977	589	2,566	7	20.96
h763	2,144	435	2,579	8	20.42
h493	1,371	137	1,508	6	21.90
h512	972	230	1,202	8	18.69
h672	737	2,497	3,234	12	19.84
h709	1,091	254	1,345	6	19.45
h756	1,110	257	1,367	6	17.76
h503	865	184	1,049	7	17.19
h656	752	303	1,055	4	20.79
h679	880	234	1,114	4	18.61
h719	849	209	1,058	6	19.48
h762	1,049	225	1,274	6	18.50
h160	103	887	990	22	20.64
h289	457	174	631	3	21.71
h494	698	106	804	8	11.72
h550	641	150	791	4	20.41
h713	909	276	1,185	6	19.40
h133	70	1,421	1,491	21	19.42
h341	314	880	1,194	20	18.41
h486	959	118	1,077	7	23.74
h659	505	254	759	4	21.94
h728	597	163	760	5	19.63
h258	375	124	499	3	20.41
h309	338	147	485	3	21.56
h488	718	55	773	6	22.90
h499	369	75	444	4	16.09
h715	414	96	510	3	19.95
h283	485	20	505	4	21.18
h334	133	743	876	19	20.00
h358	0	344	344	8	20.75
h701	250	111	361	6	20.51
h775	329	83	412	3	21.70
h238	141	306	447	6	20.09
h419	202	330	532	8	20.86
h676	334	136	470	3	19.54
h758	1,112	237	1,349	6	19.59
h768	247	64	311	3	20.68

Note: Demand covers 07:00–12:00. Physical degree counts incident corridors, and bold identifies the 11 cross-scenario backbone stations. The reported allowance is equaled or exceeded by 95% of assigned trip-end demand.

Data and code availability

The public data sources used in this study are identified and cited in the manuscript. The processed datasets, MATLAB implementation, configuration files, and reproduction instructions needed to reproduce the reported experiments are available at https://github.com/UrbanDigitalTwin/Urban-and-Advanced-Air-Mobility.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used OpenAI's ChatGPT to support manuscript drafting and language refinement. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Supplementary material

The accompanying supplementary document provides demand and candidate-station construction details, HLP-safe pruning and functional-substitution procedures, LBFO robustness and implementation controls, complete network and deployment audits, and additional operational-allowance benchmarks.

References

Amtrak, 2020. Amtrak Fact Sheet: Fiscal Year 2019, State of California. Technical Report. Amtrak Government Affairs. URL: https://www.amtrak.com/content/dam/projects/dotcom/english/public/documents/corporate/statefactsheets/CALIFORNIA19.pdf. accessed: 2026-11-06.

Bureau of Transportation Statistics, 2019a. Reporting Carrier On-Time Performance (1987–Present): 2019 Monthly Files. URL: https://transtats.bts.gov/PREZIP/. january–December 2019 reporting-carrier files, including scheduled departure and arrival times. Accessed: 2026-11-06.

Bureau of Transportation Statistics, 2019b. T-100 Domestic Segment (U.S. Carriers): 2019. URL: https://www.bts.gov/browse-statistical-products-and-data/bts-publications/data-bank-28ds-t-100-domestic-segment-data-us. monthly passenger and departure data; scheduled passenger service, class F. Accessed: 2026-11-06.

Campbell, J.F., O'Kelly, M.E., 2012. Twenty-five years of hub location research. Transportation Science 46, 153–169. doi:10.1287/trsc.1120.0410.

Cohen, A.P., Shaheen, S.A., Farrar, E.M., 2021. Urban air mobility: History, ecosystem, market potential, and challenges. IEEE Transactions on Intelligent Transportation Systems 22, 6074–6087. doi:10.1109/TITS.2021.3082767.

Escribano Macias, J.J., Khalife, C., Slim, J., Angeloudis, P., 2023. An integrated vertiport placement model considering vehicle sizing and queuing: A case study in london. Journal of Air Transport Management 113, 102486. doi:10.1016/j.jairtraman.2023.102486.

Fischetti, M., Lodi, A., 2003. Local branching. Mathematical Programming 98, 23–47. doi:10.1007/s10107-003-0395-5.

Garrow, L.A., German, B.J., Leonard, C.E., 2021. Urban air mobility: A comprehensive review and comparative analysis with autonomous and electric ground transportation for informing future research. Transportation Research Part C: Emerging Technologies 132, 103377. doi:10.1016/j.trc.2021.103377.

Gurobi Optimization, LLC, 2026. Gurobi Optimizer Reference Manual. URL: https://www.gurobi.com. software version 13.0.2.

Haan, J., Garrow, L.A., Marzuoli, A., Roy, S., Bierlaire, M., 2021. Are commuter air taxis coming to your city? a ranking of 40 cities in the united states. Transportation Research Part C: Emerging Technologies 132, 103392. doi:10.1016/j.trc.2021.103392.

Hansen, P., Mladenović, N., Moreno Pérez, J.A., 2010. Variable neighbourhood search: Methods and applications. Annals of Operations Research 175, 367–407. doi:10.1007/s10479-009-0657-6.

Helber, S., Sahling, F., 2010. A fix-and-optimize approach for the multi-level capacitated lot sizing problem. International Journal of Production Economics 123, 247–256. doi:10.1016/j.ijpe.2009.08.022.

Hess, T., Karatas, M., Eriskin, L., 2024. A maximal covering approach for vertiport location problem, in: 2024 International Conference on Decision Aid Sciences and Applications (DASA), IEEE. pp. 1–6. doi:10.1109/DASA63652.2024.10836496.

Jin, Z., Ng, K.K.H., Zhang, C., 2024. Robust optimisation for vertiport location problem considering travel mode choice behaviour in urban air mobility systems. Journal of the Air Transport Research Society 2, 100006. doi:10.1016/j.jatrs.2024.100006.

Kim, G., Jeong, Y., Shin, Y., 2025. An efficient algorithm for multi-modal skyport location problem with choice-based demand. Socio-Economic Planning Sciences 100, 102227. doi:10.1016/j.seps.2025.102227.

Kitthamkesorn, S., Chen, A., 2024. Maximum capture problem for urban air mobility network design. Transportation Research Part E: Logistics and Transportation Review 187, 103569. doi:10.1016/j.tre.2024.103569.

Li, B., Zhang, Y., 2026. AAMSim: A comprehensive simulator for advanced air mobility operations and planning. Transportation Research Part C: Emerging Technologies 192, 105858. doi:10.1016/j.trc.2026.105858.

Onat, E.B., Bulusu, V., Chakrabarty, A., Hansen, M., Sengupta, R., Sridhar, B., 2024. Evaluating eVTOL network performance and fleet dynamics through simulation-based analysis, in: AIAA Scitech 2024 Forum, American Institute of Aeronautics and Astronautics. p. 0336. doi:10.2514/6.2024-0336.

openrouteservice, 2026. openrouteservice. [Online]. Available: https://openrouteservice.org/. Accessed: 2026-05-28.

Puchinger, J., Raidl, G.R., 2005. Combining metaheuristics and exact algorithms in combinatorial optimization: A survey and classification, in: Artificial Intelligence and Knowledge Engineering Applications: A Bioinspired Approach, Springer, Berlin, Heidelberg. pp. 41–53. doi:10.1007/11499305_5.

Rajendran, S., Zack, J., 2019. Insights on strategic air taxi network infrastructure locations using an iterative constrained clustering approach. Transportation Research Part E: Logistics and Transportation Review 128, 470–505. doi:10.1016/j.tre.2019.06.003.

Rath, S., Chow, J.Y.J., 2022. Air taxi skyport location problem with single-allocation choice-constrained elastic demand for airport access. Journal of Air Transport Management 105, 102294. URL: https://www.sciencedirect.com/science/article/pii/S0969699722001132, doi:10.1016/j.jairtraman.2022.102294.

Rimjha, M., Trani, A., 2021. Urban air mobility: Factors affecting vertiport capacity, in: 2021 Integrated Communications Navigation and Surveillance Conference (ICNS), IEEE. pp. 1–14. doi:10.1109/ICNS52807.2021.9441631.

Saghazadeh, H.Z., Ayalew, Y., Ahmari, R., Kebria, P., Homaifar, A., 2026. Demand-driven vertiport siting and discrete-event fleet simulation for on-demand urban air mobility network design. arXiv preprint arXiv:2608.14974 URL: https://arxiv.org/abs/2608.14974, doi:10.48550/arXiv.2608.14974, arXiv:2608.14974.

Sinha, A.A., Rajendran, S., 2023. Study on facility location of air taxi skyports using a prescriptive analytics approach. Transportation Research Interdisciplinary Perspectives 18, 100761. URL: https://www.sciencedirect.com/science/article/pii/S2590198223000088, doi:10.1016/j.trip.2023.100761.

U.S. Census Bureau, 2019a. 2015–2019 American Community Survey Five-Year Public Use Microdata Sample: California Person Records. URL: https://www2.census.gov/programs-surveys/acs/data/pums/2019/5-Year/. california person file csv_pca.zip; departure-for-work variable JWDP and person weight PWGTP. Accessed: 2026-11-06.

U.S. Census Bureau, 2019b. American Community Survey (ACS) 2019 5-Year Estimates. [Online]. Available: https://data.census.gov. Accessed: 2026-01-18.

U.S. Census Bureau, LEHD Program, 2019. Longitudinal Employer–Household Dynamics (LEHD) Origin–Destination Employment Statistics (LODES), 2019. [Online]. Available: https://onthemap.ces.census.gov. Accessed: 2026-01-18.

Vascik, P.D., Hansman, R.J., 2019. Development of vertiport capacity envelopes and analysis of their sensitivity to topological and operational factors, in: AIAA Scitech 2019 Forum, American Institute of Aeronautics and Astronautics. p. 0526. doi:10.2514/6.2019-0526.

Venkatesh, N., Payan, A.P., Justin, C.Y., Kee, E., Mavris, D., 2020. Optimal siting of sub-urban air mobility (sUAM) ground architectures using network flow formulation, in: AIAA AVIATION 2020 Forum, p. 2921. doi:10.2514/6.2020-2921. aIAA Paper 2020.

Volakakis, V., Mahmassani, H.S., 2025. Strategic vertiport placement for airport access: Utilizing urban air mobility for accelerated and reliable transportation. Infrastructures 10, 242. URL: https://www.mdpi.com/2412-3811/10/9/242, doi:10.3390/infrastructures10090242.

Wang, Z., Lv, D., Jia, S., Wang, K., Qu, X., 2025. Urban air mobility network design and operations strategy in an urban agglomeration. Transportation Research Part E: Logistics and Transportation Review 203, 104316. doi:10.1016/j.tre.2025.104316.

Wei, Q., Gao, Z., Clarke, J.P., Topcu, U., 2024. Risk-aware urban air mobility network design with overflow redundancy. Transportation Research Part B: Methodological 185, 102967. doi:10.1016/j.trb.2024.102967.

Willey, L.C., Salmon, J.L., 2021. A method for urban air mobility network design using hub location and subgraph isomorphism. Transportation Research Part C: Emerging Technologies 125, 102997. doi:10.1016/j.trc.2021.102997.

Wu, Z., Zhang, Y., 2021. Integrated network design and demand forecast for on-demand urban air mobility. Engineering 7, 473–487. doi:10.1016/j.eng.2020.11.007.

Zamani Saghazadeh, H., Esmaeili Ghayoumabadi, S., Sabri, S., 2026. Robust vertiport candidate generation under demand uncertainty using iVAT-guided fuzzy clustering. Conference paper presented at the Annual Modeling and Simulation Conference (ANNSIM). URL: https://scs.org/wp-content/uploads/2026/04/69dc69088c6c0_Preliminary-2026-ANNSIM-Program-2.pdf. may 2026.

Zhang, H., Huo, J., Chen, C., Liu, Z., 2025. A composite transportation network design problem with land-air coordinated operations. Transportation Research Part C: Emerging Technologies 171, 104967. doi:10.1016/j.trc.2024.104967.