

Small Cell Integration into Urban Furniture for Smart City Development

Mohammed Sh. Ali¹ , Mohammed Othman Hasan¹ , Rana H. A. Zubo^{1,*}  and Raed A. Abd-Alhameed^{2,3} 

¹Northern Technical University, Iraq

mohammed.shifaa25@ntu.edu.iq; mohammed.othman25@ntu.edu.iq; r.h.a.zubo@ntu.edu.iq

²University of Bradford, United Kingdom

r.a.a.abd@bradford.ac.uk

³Alfarqadein University, Iraq

r.a.a.abd@bradford.ac.uk

*Correspondence: r.h.a.zubo@ntu.edu.iq

Received: 9 May 2026; Accepted: 28 September 2026; Published: 1 October 2026

Abstract: Fifth-generation networks in mid-band spectrum need dense small-cell infrastructure. The binding constraint is site cost, not radio performance. Urban furniture removes part of it: street lighting columns, bus shelters and traffic signals already offer mounting height, mains power and, increasingly, fibre. This paper asks which of those assets an operator with a fixed capital budget should equip. We formulate a budget-aware greedy site-selection algorithm over a discretised service area and evaluate it against seven alternatives: random selection, a snapped lattice, k-means clustering, cost-unaware greedy maximum coverage, simulated annealing, a genetic algorithm and binary particle swarm optimisation. The coverage objective is interference-aware over a constrained, heterogeneously priced candidate set, a combination the surveyed studies do not make. Every selected cell that does not serve a test point interferes there, and propagation follows the Third Generation Partnership Project Urban Micro street-canyon model with line-of-sight and non-line-of-sight states and per-link shadow fading. Demonstrated empirically rather than assumed, the objective is then neither monotone nor submodular, so the classical constant-factor guarantee does not apply and the method is a heuristic. Across fifty paired Monte-Carlo realisations under matched expenditure it exceeds all six remaining baselines at every threshold, by margins whose confidence intervals exclude zero after correction, and it matches simulated annealing at the tightest threshold on about one thirteenth of the computation while trailing it by 0.54 and 0.32 percentage points at the looser two. Under a common budget it comes within 1.2 percentage points of the best baseline on a third to just over a half of the capital. On real urban-furniture inventories extracted from OpenStreetMap for three Swiss city centres, that matched-spend ordering holds and the advantage over cost-unaware selection is largest on the densest inventory, but under a common budget cost-unaware selection is ahead in all three cities. Runs at 28 and 39 GHz show the interference-driven stopping point disappearing as cells shrink, and a probabilistic reading of coverage matches the threshold metric within one percentage point while showing the reliably covered area to be far smaller. We release the simulator and the extracted inventories.

Keywords: *Budget-constrained optimisation; Interference modelling; Network densification; Small cell deployment; Street furniture; Submodular maximisation; Urban infrastructure planning*

1. Introduction

The fifth-generation (5G) mobile standard is organised around three service classes: enhanced mobile broadband, ultra-reliable low-latency communication and massive machine-type communication [1, 2]. Delivering the throughput expected of the first class in dense cities depends on spectrum above 3 GHz, where propagation is far less forgiving than in the sub-1 GHz bands that carried earlier generations. Attenuation rises with carrier frequency. Building fabric blocks more of the signal, and cell radii shrink from

kilometres to a few hundred metres. Capacity in such networks comes mainly from spatial reuse, and spatial reuse comes from density.

Radios are rarely the obstacle to density. A modern outdoor small cell is a compact, low-power unit that bolts onto a pole. The expense sits around it: securing rights to a location, obtaining planning consent, arranging a power feed, providing backhaul, and paying recurring site rental. Industry analysis prepared for CTIA estimated that regulatory review and site acquisition accounted for a substantial share of the total cost of small-cell deployment in the United States, and that streamlining them could release capital for additional sites¹. That same analysis put the number of small cells required in the United States alone at hundreds of thousands of units.

Urban furniture cuts into part of that cost structure. Street lighting columns, bus shelters, traffic signal masts, information kiosks and closed-circuit television poles already stand where people are, already carry mains power, already have a known owner, and in many cities are already being surveyed under smart-city asset programmes, with inventories now built automatically from street imagery [3-6]. A municipality that controls these assets can offer them under a single master agreement instead of negotiating location by location. Several European projects have demonstrated the concept at pilot scale on lighting infrastructure [7], and commercial smart-pole products that combine lighting, sensing and radio have reached the market.

Given a municipal inventory of such assets, an operator faces a concrete decision: with a fixed budget, which subset to equip. The assets are not interchangeable. A column on a wide boulevard may serve a large area; one in a narrow side street may serve almost nothing. Costs differ as well, since some poles need structural reinforcement, some need a new power feed, and some sit within reach of existing fibre where others do not. This is a budgeted selection problem over a heterogeneous, spatially structured candidate set, and it is the problem the paper addresses.

The original version of this work presented a budget-aware greedy selection algorithm, argued that the coverage objective was monotone and submodular, and invoked the classical $(1 - 1/e)$ approximation guarantee for budgeted maximum coverage. Review made clear that the argument does not survive contact with an interference-limited model, and we have rebuilt the present version around that finding. The contributions follow:

- We specify the interference-aware coverage model in full: serving-cell association by strongest received power, aggregate interference from every other selected cell, frequency reuse one, full-buffer transmission, and 3GPP Urban Micro propagation with per-link line-of-sight state and log-normal shadow fading. Section 3 states every quantity needed to reproduce the model.
- We withdraw the claim of submodularity, prove why it fails, and measure directly how often it fails. Under the interference model, adding a site reduces coverage in slightly more than half of the cases tested, and the diminishing-returns inequality is violated in a comparable fraction. We retain the $(1 - 1/e)$ guarantee only for a clearly delimited noise-limited surrogate, and otherwise present the method as a heuristic.
- The comparison set grows from two weak baselines to seven, among them k-means clustering, cost-unaware greedy maximum coverage and three metaheuristics (simulated annealing, a genetic algorithm and binary particle swarm optimisation), with the last three given a common objective-evaluation budget larger than the proposed method consumes.
- We report two budget regimes rather than one. In the first, every method receives the same aggregate budget. In the second, every method is held to the aggregate spend that the proposed method itself incurs. That second regime answers the objection that equal budgets do not imply equal expenditure, and it is the regime in which the proposed method is most fairly assessed.
- Reporting is stronger throughout: paired differences with 95% confidence intervals, Holm-corrected p-values, Wilcoxon signed-rank tests as a distribution-free check, Shapiro–Wilk assumption checks, standardised effect sizes, and per-method accounting of sites selected, capital spent and budget left unused.

¹ CTIA and Accenture Strategy, “Impact of Federal Regulatory Reviews on Small Cell Deployment”, Washington, DC, USA, 2018, Available: https://api.ctia.org/docs/default-source/default-document-library/small-cell-deployment-regulatory-review-costs_3-12-2018.pdf.

- Three robustness axes appear that the earlier version did not cover: cost-range width at all three signal quality thresholds rather than only the least demanding one, partial network load through an interferer activity factor, and spatially non-uniform traffic demand. Two further axes are added in this revision: carrier frequency at 28 and 39 GHz, and a load-balanced association rule evaluated alongside a probabilistic reading of coverage.

- We answer what the normalised cost index means in money by proving, and then verifying numerically, that the selection is invariant to any positive rescaling of costs and the budget. The index is a statement about cost ratios, so an operator can substitute their own currency figures and reproduce these results exactly. We give the decomposition they would use and report the cost multiples the method selects in practice.

- The case study now rests on three real urban-furniture inventories rather than one, drawn from OpenStreetMap for city-centre tiles in Zürich, Bern and Basel, and covering street lighting columns, bus stops and traffic-signal masts rather than almost exclusively streetlights. We use the whole inventory as the candidate set instead of a random subsample, and we describe the experiment throughout as a real-location simulation rather than as real-infrastructure validation.

- We release the simulator, the randomisation protocol, the seeds, the OpenStreetMap queries, the extracted inventories and the software environment, rather than offering them on request.

The remainder of the paper is organised as follows. Section 2 reviews small-cell technology, the convergence of telecommunications and municipal infrastructure, and the site-selection literature, closing with a comparative summary that positions this work. Section 3 specifies the system model. Section 4 formulates the optimisation problem, examines its structural properties and presents the algorithm. Section 5 describes the experimental design and the statistical procedure. Section 6 reports results. Section 7 discusses interpretation, practical implications and limitations, and Section 8 concludes.

2. Related Work

Three bodies of work bear on the problem. The first is small-cell technology and the host assets open to it. The second is the convergence of telecommunications with municipal infrastructure, together with the regulatory setting that governs it. The third is the set of methods operators use to choose where radio sites go. This section takes each in turn and closes with a comparative summary.

2.1. Small Cells and Urban-Furniture Host Assets

Small cells are low-power radio access nodes. They complement macro base stations by adding capacity where demand concentrates. The Small Cell Forum classifies them by coverage footprint and intended environment, running from femtocells that serve a single dwelling through picocells and microcells that serve public indoor and outdoor areas. Outdoor urban small cells, the class considered here, typically radiate at 20–30 dBm from a mounting height of four to eight metres. That places them below the surrounding rooftop line, inside the street canyon. The Small Cell Forum's 2024 market forecast projects continued growth in deployments through 2030, with public small cells in urban environments growing at a compound annual rate of 10%².

Several groups have examined street lighting infrastructure as a host, from both technical and commercial angles. Benseny *et al.* analysed city strategies for deploying a 5G small-cell network on light poles, treating asset ownership and municipal strategy as first-order variables alongside the radio design [7]. Work on automated site identification starts from the same premise and runs the other way: Zhang and Alshaykh used Google Street View imagery and a Mask R-CNN detector to build a utility-pole inventory suitable for 5G planning, which indicates that candidate inventories can come from imagery rather than from manual survey [6]. Power and backhaul remain the practical constraint. To contain the resulting energy cost, others propose dynamic deployment schemes that switch cells on according to demand [8].

Crowdsourced observation data offers one check on the premise that mid-band densification is under way rather than merely projected. An October 2024 snapshot of the OpenCellID database covering the

² Small Cell Forum, "Market Forecast Report 2024", Document SCF050, Dursley, United Kingdom, 3 July 2024, Available: <https://www.smallcellforum.org/docs/scf-market-forecast-report/>.

United States contains 563 885 records, of which 80.9% are Long-Term Evolution (LTE) and 16.2% are 5G New Radio, and the New Radio share differs by more than an order of magnitude between operators. Data of this kind is biased towards populated areas and towards the handsets and applications of contributing volunteers, so these proportions should be read as an indication of relative deployment maturity and not as a census. The figures appear here only to motivate the mid-band focus of the study. No part of the site-selection algorithm or its evaluation uses OpenCellID data, and the released archive carries the full breakdown by radio-access technology and by operator rather than the body of this paper.

This literature establishes two things: urban furniture makes a credible host, and inventories of it can be assembled. It leaves open which members of a given inventory an operator should choose under a capital constraint. That gap is the one addressed here.

2.2. Smart-City Infrastructure Convergence and Its Regulatory Setting

Smart-city programmes fold information and communication technologies (ICT) into urban systems, aiming to improve service delivery and resource efficiency [3]. Chourabi *et al.* proposed an integrative framework that identifies eight factors shaping smart-city initiatives: management and organisation, technology, governance, policy context, people and communities, economy, built infrastructure, and the natural environment [4]. Three of those factors bear directly on small-cell siting. Built infrastructure fixes which assets exist and where they stand; governance fixes who may authorise their use; policy context fixes how long authorisation takes and what it costs.

Salih *et al.* survey Internet-of-Things applications in urban development, examining smart-city case studies across six cities together with the challenges those deployments encountered [5]. Regulation has itself become a subject of study. Meese *et al.* examine regulatory trends and local-government responses to small-cell deployment across several jurisdictions, giving particular attention to urban amenity (the visual impact of equipment on streetscapes) as a factor shaping those responses [9]. In the United States, the National League of Cities has published guidance for municipalities negotiating such agreements³. A Government Accountability Office review found federal coordination of deployment to be incomplete⁴. Together, this work motivates the cost heterogeneity modelled in Section 3: two physically similar poles can carry sharply different acquisition costs depending on ownership, consent regime and amenity sensitivity.

2.3. Site-Selection Methods

Work on choosing radio site locations falls into four groups.

Stochastic-geometry models treat base-station positions as a random point process and derive coverage and rate distributions analytically. Dhillon *et al.* developed the canonical K-tier downlink model on this basis [10], and the framework underpins much of the ultra-dense-network literature [11, 12]. Its value lies in closed-form results that expose scaling behaviour: Kamel *et al.* and Ge *et al.* both use it to examine how the gains from densification are bounded once interference rather than noise limits performance [11, 12]. For the present problem the limitation is structural, not a matter of accuracy. A homogeneous Poisson point process assumes a site may go anywhere with equal probability and at equal cost. Urban-furniture deployment violates both assumptions. Positions are confined to an existing, spatially clustered inventory, and costs differ across it. Random models remain a reasonable approximation when the candidate set is dense relative to the intended cell radius, because the achievable geometry then approaches the unconstrained one; they become unreliable when the inventory is sparse or strongly aligned to a street grid.

Combinatorial-optimisation formulations treat the problem as covering a discretised area with a chosen subset of candidates. The budgeted maximum coverage problem is the closest classical formulation. Khuller *et al.* established that a modified greedy algorithm achieves a $(1 - 1/e)$ approximation for it, building on the analysis of monotone submodular maximisation by Nemhauser *et al.* [13, 14]. Minoux's accelerated greedy algorithm exploits the same submodularity to avoid re-evaluating every candidate at every step

³ National League of Cities, "Small Cell Wireless Technology in Cities: A Municipal Action Guide", Washington, DC, USA, 2018, Available: <https://www.nlc.org/resource/small-cell-wireless-technology-in-cities/>.

⁴ United States Government Accountability Office, "5G Deployment: FCC Needs Comprehensive Strategic Planning to Guide Its Efforts", Report GAO-20-468, Washington, DC, USA, June 2020, Available: <https://www.gao.gov/products/gao-20-468>.

[15]. These guarantees form the intellectual basis of the algorithm used here, and Section 4.2 examines with care the conditions under which they do and do not transfer to an interference-limited setting.

Metaheuristics dominate the applied placement literature. Nikam *et al.* applied a tailored simulated-annealing scheme to small-cell positioning under coverage and cost objectives [16], and Li applied simulated annealing to 5G base-station siting [17]. Venkateswararao and Swain used binary particle swarm optimisation for energy-efficient small-cell deployment in ultra-dense networks [18]. Multi-objective evolutionary approaches appear where coverage must be traded against cost or energy, including recent quantum-inspired variants for three-dimensional base-station deployment [19]. Mei and Zhang formulated joint base-station and intelligent reflecting surface deployment as a graph-based modelling and optimisation problem [20]. None of these methods assumes anything structural about the objective, which is precisely why they remain applicable when submodularity fails. The price is computational, and the comparison in Section 6 quantifies it.

Data-driven and hybrid approaches have grown fast. Aloupogianni *et al.* present an artificial-intelligence-driven approach to optimising small-cell deployment for beyond-5G networks [21]. Ferreira *et al.* combine computational modelling with real measured data to position small cells in 5G networks, the closest published approach to the validation this study does not attempt [22]. Mahbub and Barua likewise base positioning decisions on measured propagation characteristics [23]. Two further strands sit close to this work without overlapping it. Network digital twins aim to keep a live model of the radio access network for planning and what-if analysis [24], and federated learning moves model training to the network edge so that operators can learn from data they cannot pool [25]. Both target the operation of a deployed network rather than the choice of which assets to equip, so neither supplies a baseline for the problem studied here. The present study is explicitly simulation-based and claims no measurement validation; Section 7.3 states this as a limitation rather than presenting the real-location experiments as field validation.

Two further bodies of work are relevant but distinct. Operations-research facility-location theory supplies the exact formulations (maximum covering location [26], p-median, set covering) of which the present problem is a budgeted variant. Church and ReVelle's maximal covering formulation maximises demand covered by a fixed number of facilities; the budgeted form used here replaces that cardinality constraint with a spend constraint over heterogeneously priced sites, which is the variant Khuller *et al.* analyse [14]. Geographic information system (GIS) methods combined with multi-criteria decision analysis (MCDA) are used in planning practice to rank candidate sites against weighted criteria such as land use, visibility and access. We note that family for completeness rather than review it, because it produces a ranking rather than an optimised subset and does not model interference, so it complements the approach taken here rather than competes with it.

2.4. Comparative Summary and Position of This Work

Table 1 sets representative studies against the dimensions that matter for the present problem: whether the candidate set is constrained to pre-existing assets, whether per-site costs are heterogeneous, whether the study models interference, and how it validates the method.

In Table 1, "Real-asset candidates" marks whether siting is restricted to a pre-existing asset inventory, "Per-site cost varies" whether per-site costs differ across candidates, and "Interference modelled" whether the coverage criterion accounts for inter-cell interference. The classification follows the scope each study states for itself and should be read as a positioning aid rather than as an assessment of those works. Works on digital twins [24] and federated learning [25] are discussed in Section 2.3 but are absent from the table, because they address the operation of a network already deployed rather than the choice of which assets to equip.

Table 1. Representative small-cell and base-station placement studies

Study	Method	Real-asset candidates	Per-site cost varies	Interference modelled	Validation
Dhillon <i>et al.</i> [10]	Stochastic geometry	No	No	Yes	Analysis
Kamel <i>et al.</i> [11]	Survey, stoch. geometry	No	No	Yes	Survey
Ge <i>et al.</i> [12]	Analytical limits	No	No	Yes	Analysis
Benseny <i>et al.</i> [7]	Techno-economic analysis	Yes (light poles)	Partly	No	City case studies

Study	Method	Real-asset candidates	Per-site cost varies	Interference modelled	Validation
Zhang and Alshaykh [6]	Mask R-CNN pole inventory	Yes (utility poles)	No	No	Detection accuracy
Nikam <i>et al.</i> [16]	Simulated annealing	No	In objective	No (signal only)	Simulation
Li [17]	Simulated annealing	No	No	No (signal only)	Simulation
Venkateswararao and Swain [18]	Binary PSO, energy	No	Energy only	Yes	Simulation
Mei and Zhang [20]	Graph-based joint BS/IRS	Discrete set	No	Yes	Simulation
Chou <i>et al.</i> [19]	Quantum-inspired MO	No	Yes	Yes	Simulation
Aloupogianni <i>et al.</i> [21]	Machine-learning optimisation	Partly	Partly	Yes	Simulation
Ferreira <i>et al.</i> [22]	Hybrid model plus measurement	Yes	No	Yes	Field measurement
Mahbub and Barua [23]	Measurement-driven	No	No	No (signal only)	Measurement
This work	Budget-aware greedy vs. seven baselines	Yes (urban furniture)	Yes (aggregate budget)	Yes (full SINR)	Paired Monte-Carlo

Two observations follow. First, studies that model interference with care tend to treat position as continuous or unconstrained; studies that constrain positions to real assets tend to use a signal-strength criterion rather than a signal-to-interference one. Second, heterogeneous per-site cost is uncommon, and where it does appear it usually takes the form of a fixed count constraint rather than an aggregate budget. The combination examined here (a constrained real-asset candidate set, heterogeneous per-site cost, an aggregate budget, and an interference-limited coverage criterion) is not well covered.

The contribution is framed to match. Budget-aware greedy selection is no new algorithm, and this paper claims none; Section 4 attributes it plainly. The claim is narrower. The combination above defines a problem whose structure differs from the one the classical guarantee addresses, the difference is measurable and consequential, and a systematic comparison under matched expenditure shows where a simple cost-aware heuristic stands relative to general-purpose metaheuristics on this problem.

3. System Model

3.1. Service Area, Test Points and Candidate Sites

The service area is a square of side 1 km, discretised into a uniform lattice of test points spaced 10 m apart. That gives $G = 10\,000$ points, one at the centre of each $10\text{ m} \times 10\text{ m}$ cell. Test points sit at height $h_{UT} = 1.5\text{ m}$. Coverage is evaluated at these points only. The model makes no attempt at indoor penetration, so all results should be read as outdoor street-level coverage.

The candidate set L contains M urban-furniture locations. In the synthetic study, we draw candidate coordinates independently from a uniform distribution over the service area and redraw them at every Monte-Carlo iteration; we state that distribution explicitly because the earlier version of this work left it unspecified. In the real-location study of Section 6.11, the candidate set is the complete OpenStreetMap inventory of street lighting columns, bus stops and traffic-signal masts within the tile, with no subsampling. Each candidate carries a mounting height $h_{BS} = 6\text{ m}$, representative of a lighting column.

Each candidate l has a deployment cost c_l drawn independently from the discrete uniform distribution on $\{1, \dots, c_{\max}\}$. The cost is a normalised index, not a monetary figure. It is intended to capture the relative burden of equipping one asset rather than another (structural reinforcement, a new power feed, backhaul provision, consent and rental), and Section 7.3 discusses at length what can and cannot be inferred from it. The aggregate budget is set as a fixed fraction of the total cost of the candidate set, $B = \varrho * \sum_l c_l$ with $\varrho = 0.20$, so that the budget scales with problem size and the three cost ranges tested remain comparable.

3.2. What the Cost Index Means, and How to Put Money on It

A normalised cost index invites the question of what it corresponds to in real economic terms. The index is the ratio of a site's total cost of ownership to the cheapest site in the same inventory. For candidate l an operator would build it as $c_l = (E + S_l + P_l + K_l + A_l + T R_l) / \min$ over the inventory of the same quantity, where E is the radio unit, S_l structural reinforcement, P_l the power connection, K_l backhaul provision, A_l consent and permitting, R_l annual site rental and T the planning horizon in years. Nothing in this paper estimates those quantities. They differ by city, by asset owner and by consent regime, and inventing figures for them would put a number in the paper that no reader could check.

What can be established is that the choice of monetary unit does not matter. The selection rule ranks candidates by marginal coverage gain per unit cost and admits a candidate when the running spend plus its cost stays inside the budget. Multiplying every cost and the budget by the same positive factor leaves both the ranking and the feasibility test unchanged, so the selected set is identical. An operator can therefore substitute their own currency figures and reproduce these results exactly, provided the ratios between sites match the ones used here. Table 2 verifies the claim numerically rather than leaving it as an assertion: across 30 instances at each threshold, scaling by one half, by two, by 137 and by 12 500 reproduced the same selection every time.

Each cell of Table 2 is the percentage of instances in which scaling every site cost and the budget by the factor λ produced exactly the same selected set as the unscaled problem. The two large factors stand in for a change of currency unit.

Table 2. Cost scale invariance: share of instances with an unchanged selected set (%)

γ_{th} (dB)	Instances	$\lambda = 0.5$	$\lambda = 2.0$	$\lambda = 137.0$	$\lambda = 12500.0$
0	30	100	100	100	100
3	30	100	100	100	100
6	30	100	100	100	100

The second half of the answer is what the method buys. Over 448 selected sites at 3 dB, the median selected site cost twice the cheapest site in its inventory, the mean was 3.09 times, and the ninetieth percentile was 7 times. Just over half the selections, 56.7%, cost no more than twice the cheapest, and 18.3% cost more than five times. The candidate costs are drawn uniformly on the integers 1 to 10, with mean and median 5.5, so a median selected-site ratio of 2.0 shows the rule is strongly biased towards cheap assets. It is not exclusively so: 18.3% of selected sites cost more than five times the cheapest. It buys expensive ones where the coverage justifies the ratio, and an operator holding real figures can read those multiples directly against their own cost spread.

3.3. Propagation Model

Path loss follows the Urban Micro (UMi) street-canyon model of 3GPP TR 38.901⁵ at a default carrier frequency $f_c = 3.5$ GHz; Section 6.9 repeats the study at 28 and 39 GHz. For a link of three-dimensional length d_{3D} and two-dimensional length d_{2D} , the line-of-sight (LOS) path loss is

$$PL_{LOS} = 32.4 + 21 \log_{10}(d_{3D}) + 20 \log_{10}(f_c), \quad 10 \text{ m} \leq d_{2D} \leq d'_{BP} \quad (1)$$

and, beyond the breakpoint,

$$PL_{LOS} = 32.4 + 40 \log_{10}(d_{3D}) + 20 \log_{10}(f_c) - 9.5 \log_{10}[(d'_{BP})^2 + (h_{BS} - h_{UT})^2], \quad d'_{BP} \leq d_{2D} \leq 5 \text{ km} \quad (2)$$

In equations (1) and (2), f_c is in gigahertz and distances are in metres. The breakpoint distance is $d'_{BP} = 4 (h_{BS} - 1)(h_{UT} - 1) f_c / c$, with c the speed of light, which gives $d'_{BP} = 116.7$ m for the heights used here. The non-line-of-sight (NLOS) path loss is

$$PL_{NLOS} = \max\{PL_{LOS}, 35.3 \log_{10}(d_{3D}) + 22.4 + 21.3 \log_{10}(f_c) - 0.3 (h_{UT} - 1.5)\} \quad (3)$$

and the probability that a link is in line of sight, used to draw the LOS state in equation (4), is

$$Pr[LOS] = \min(18 / d_{2D}, 1) + \exp(-d_{2D} / 36) \cdot [1 - \min(18 / d_{2D}, 1)] \quad (4)$$

⁵ 3rd Generation Partnership Project (3GPP), "Study on Channel Model for Frequencies from 0.5 to 100 GHz", Technical Report TR 38.901, Version 16.1.0, December 2019, Sophia Antipolis, France, Published by 3GPP, Available: https://www.3gpp.org/ftp/Specs/archive/38_series/38.901/.

which equals one for $d_{2D} \leq 18$ m, as the model requires. Equation (3) takes the maximum with the line-of-sight value so that the non-line-of-sight loss never falls below it. We add shadow fading as a zero-mean Gaussian term in decibels with standard deviation 4 dB in LOS and 7.82 dB in NLOS conditions, following the same specification. The model leaves out building penetration loss, and it leaves out small-scale fading; Section 7.3 discusses the consequences.

We draw both the LOS state and the shadow-fading sample once per (candidate, test point) link at the start of a Monte-Carlo iteration, then hold them fixed for the whole of that iteration. This is deliberate. Every placement method within an iteration then meets exactly the same channel realisation, which is what makes the paired statistical comparisons of Section 5.4 valid, and it removes the possibility that a method looks better merely because it was given a more favourable set of random draws.

3.4. Link Budget, Cell Association and Interference

Each small cell transmits at $P_t = 24$ dBm with an antenna gain of $G_t = 5$ dBi, omnidirectional in azimuth. The model includes no down-tilt, no sectorisation and no beamforming pattern. The receive antenna gain is $G_r = 0$ dBi. The received power at test point g from candidate l is

$$P_{rx}(l, g) \text{ [dBm]} = P_t + G_t + G_r - PL(l, g) - X_\sigma(l, g) \quad (5)$$

where PL is the path loss of equations (1)-(3) under the drawn LOS state of equation (4), and X_σ is the shadow-fading sample.

Thermal noise power is $N = N_0 + 10 \log_{10}(W) + F$ with noise power spectral density $N_0 = -174$ dBm/Hz, bandwidth $W = 20$ MHz and receiver noise figure $F = 7$ dB, giving $N = -93.99$ dBm. We state the spectral density explicitly because the earlier version gave the noise figure and the bandwidth but left it out.

Given a selected set S and the received powers of equation (5), the serving cell of test point g is the member of S with the strongest received power at that point, $s^*(g) = \arg \max_{l \in S} P_{rx}(l, g)$. Every other member of S acts as an interferer on the same bandwidth: frequency reuse is one, all cells transmit continuously, and the default model applies no scheduling, no power control, no coordinated muting and no load balancing. Section 6.10 evaluates a load-aware association variant of this rule. The signal-to-interference-plus-noise ratio (SINR) at g is, therefore, in linear units:

$$SINR_g(S) = P_{rx}(s^*(g), g) / (\alpha \cdot \sum_{l \in S, l \neq s^*(g)} P_{rx}(l, g) + N) \quad (6)$$

where α in $(0, 1]$ is an interferer activity factor: the fraction of time a non-serving cell occupies the shared resource. Setting $\alpha = 1$ corresponds to a fully loaded network and is the default, the most pessimistic interference assumption available within this model. Section 6.6 reports how the results move as α is reduced. Modelling load in this aggregate way is a simplification. It captures the average effect of partial occupancy, but not the spatial correlation of demand, user mobility, or scheduler behaviour.

3.5. Coverage Objective

A test point counts as covered when its SINR meets a threshold γ_{th} . Writing $w_g \geq 0$ for the traffic-demand weight of test point g , the coverage objective is

$$C(S) = (\sum_g w_g \cdot 1[SINR_g(S) \geq \gamma_{th}]) / (\sum_g w_g) \quad (7)$$

with $1(\cdot)$ the indicator function. In the uniform-demand case, the default, $w_g = 1$ for all g and $C(S)$ is the fraction of the area meeting the threshold. Section 6.7 reports a non-uniform case: there w_g is the sum of a unit background and four Gaussian hotspots of peak weight four and standard deviation 120 m, placed at random within the area at each iteration.

We evaluate three thresholds: $\gamma_{th} = 0$ dB for bare connectivity, 3 dB for moderate quality, and 6 dB for a demanding service level. Defining coverage by a hard threshold on a single realisation is a simplification. Since shadow fading is drawn per link and held fixed, each Monte-Carlo iteration represents one realisation of the shadowing field rather than an outage probability averaged over it; the distribution across iterations, reported throughout as a standard deviation and a confidence interval, is what carries the uncertainty. Small-scale fading would randomise the instantaneous SINR further, and the model leaves it out. Section 6.10 and Table 17 report the same selections read as a coverage probability, obtained by redrawing the shadowing field on fixed geometry, so that the threshold metric can be compared against a probabilistic one.

Table 3. Simulation, propagation, algorithm and computing parameters

Parameter	Value	Unit
Simulation environment		
Service area	1 × 1	km ²
Test-point spacing	10	m
Number of test points, G	10 000	-
Candidate locations, M	100	-
Candidate distribution (synthetic)	Uniform over the service area	-
Monte-Carlo iterations per scenario	50	-
Per-site cost distribution	Discrete uniform on [1, c_max]	normalised units
Cost range c_max (default; sensitivity)	10 (30, 50)	normalised units
Budget ratio ϱ	0.20	-
Propagation and signal model		
Carrier frequency, f_c (default; sensitivity)	3.5 (28, 39)	GHz
Channel bandwidth, W	20	MHz
Path-loss model	3GPP TR 38.901 UMi street canyon	-
Small-cell height, hBS	6	m
Test-point height, hUT	1.5	m
Breakpoint distance, d'BP	116.7	m
Transmit power, P_t	24	dBm
Transmit antenna gain, G_t (omnidirectional)	5	dBi
Receive antenna gain, G_r	0	dBi
Shadow fading std. dev., LOS / NLOS	4.00 / 7.82	dB
Noise power spectral density, N_0	-174	dBm/Hz
Receiver noise figure, F	7	dB
Thermal noise power, N	-93.99	dBm
Cell association (default; variant)	Strongest received power (load-aware)	-
Frequency reuse	1	-
Interferer activity factor, α (default; sensitivity)	1.00 (0.50, 0.25)	-
SINR thresholds, γ_{th}	0, 3, 6	dB
Algorithm configuration		
Metaheuristic objective-evaluation budget	3 000	evaluations
Genetic algorithm population / mutation rate	50 / 2M	-
PSO particles / inertia / c1 / c2	50 / 0.7 / 1.5 / 1.5	-
Simulated annealing temperature schedule	Geometric, 0.05 → 10 ⁻⁴	-
Computing environment		
Language and libraries	Python 3.11.15, NumPy 2.4.4, SciPy 1.17.1, pandas 3.0.2, Matplotlib 3.10.9	-
Platform	Linux-6.18.44-fc-v24-x86_64-with-glibc2.39	-
Logical processors	2	-
Master seed	20260820	-

4. Problem Formulation and Algorithm

4.1. Budgeted Maximum Coverage Formulation

The candidate set L carries costs $\{c_l\}$, the aggregate budget is B , and the coverage objective C comes from equation (7). Selection then reads

$$\text{maximise } C(S) \text{ over } S \subseteq L, \text{ subject to } \sum_{l \in S} c_l \leq B \quad (8)$$

Equation (8) is a budgeted maximum coverage problem. The classical form takes the objective as a union of fixed coverage sets, and in that form the problem is NP-hard. Khuller *et al.* showed that a modified greedy algorithm reaches a $(1 - 1/e)$ approximation [14]: run the cost-benefit greedy rule, separately consider the single best affordable candidate, and return the better of the two. The analysis rests on the objective being monotone and submodular, properties that Nemhauser *et al.* established for coverage functions [13].

Equation (7) does not fit that form. A candidate's coverage set is not fixed: whether a test point is covered depends on which other candidates joined the selection, through the interference term in equation (6). The next subsection makes the consequence precise.

4.2. Structural Properties of the Interference-Aware Objective

Proposition 1. The objective C of equation (7) is in general neither monotone nor submodular in S .

Proof. One counterexample per property suffices, and both come from the interference term. Take monotonicity first. Consider a single test point g covered by a set S with SINR exactly at the threshold, and let e be a candidate whose received power at g is strictly positive but strictly below that of the serving cell $s^*(g)$. Adding e leaves the serving cell unchanged and raises the denominator of equation (6) by $\alpha * \text{Prx}(e, g) > 0$, so $\text{SINR}_g(S + e) < \text{SINR}_g(S)$ and the point falls below the threshold. Then $C(S + e) < C(S)$, and C is not monotone. Submodularity requires that the marginal gain of adding e is non-increasing in the set to which it is added, that is $C(A + e) - C(A) \geq C(B + e) - C(B)$ for all A subset of B . Let every member of A lie far from a region R , and set $B = A + f$, where f reaches R strongly enough to keep its points below the threshold as an interferer but not strongly enough to serve them at the threshold itself. Adding f therefore removes R from the covered set, which is the non-monotonicity established above. Now let e reach R more strongly than f does. Adding e to A yields a modest gain, because the distant members of A already cover part of the region. Adding the same e to B yields a larger gain, because e becomes the serving cell in R and demotes f to one interferer among several, restoring the points that f had removed. The marginal gain therefore increases with the set, violating the inequality. QED.

Those two arguments are existence arguments. They show that the guarantee cannot be invoked, and they say nothing about how often violations occur in practice. Measurement answers that question. The test drew random selected sets over 30 independent channel realisations; for each realisation, and at each of the three thresholds, it checked 20 random additions for a coverage decrease and 20 nested pairs (A subset of B , where B contains ten members that A does not) for a submodularity violation. Table 4 reports the outcome. Coverage falls on slightly more than half of all additions, and the diminishing-returns inequality fails on between 56% and 60% of nested pairs. Violations of both kinds are the typical behaviour of the objective at the densities considered, not rare corner cases.

In Table 4, the coverage-decreasing share is the fraction of random single-site additions that reduced coverage, and the submodularity-violation share is the fraction of nested pairs whose marginal gain grew with the size of the set. Thirty independent channel realisations were used per threshold.

Table 4. Empirical monotonicity and submodularity of the interference-aware objective

γ_{th} (dB)	Additions tested	Coverage-decreasing (%)	Largest decrease (pp)	Nested pairs tested	Submodularity violations (%)	Largest violation (pp)
0	600	52.5	2.01	600	56.5	1.62
3	600	53.3	2.41	600	59.8	2.13
6	600	52.2	2.83	600	58.2	1.93

The earlier version of this paper claimed that C is monotone submodular and that the $(1 - 1/e)$ guarantee therefore applies. That claim is withdrawn. A restricted statement survives.

Proposition 2. Define the noise-limited surrogate objective C_0 by replacing the denominator of equation (6) with N , which amounts to ignoring interference. Then C_0 is monotone and submodular, and the modified greedy algorithm of Khuller *et al.* applied to C_0 returns a set S with $C_0(S) \geq (1 - 1/e) C_0(S^*)$, where S^* is the optimal feasible set for C_0 .

Proof. Under C_0 , the coverage indicator of a test point g turns only on whether some selected candidate exceeds the threshold receive power at g , so $C_0(S)$ is a weighted coverage function of the fixed sets $U_1 = \{g : \text{Prx}(l, g) \geq \gamma_{\text{th}} * N\}$. Weighted coverage functions are monotone and submodular [13], and the cost function is additive by construction. The result is then exactly Theorem 3 of Khuller *et al.* [14]. QED.

Proposition 2 describes a different objective from the one behind the reported results. It stays for two reasons. First, it explains why a greedy rule is a sensible starting point at all: in the noise-limited regime, which is the regime of a sparse deployment at the first few sites, the algorithm inherits a genuine guarantee. Second, it marks exactly where that reasoning stops. Every coverage figure reported here comes from the full interference-aware objective of equation (7), and this paper treats the algorithm throughout as a heuristic with no approximation guarantee.

4.3. The Proposed Algorithm

Algorithm 1 states the method. At each step the algorithm evaluates the marginal gain of every affordable unselected candidate against the current selection, under the full interference-aware objective, and adds the candidate with the largest gain per unit cost. It stops once no affordable candidate produces a positive gain, which under the interference model can happen well before the budget runs out. That behaviour is a property of the problem rather than a defect of the implementation, and Section 6.3 reports how much budget is left unused as a result.

Algorithm 1. Budget-aware greedy site selection under interference

```

Input: candidates  $L$ , costs  $\{c_l\}$ , budget  $B$ , power matrix  $\text{Prx}$ , threshold  $\gamma_{\text{th}}$ , weights  $\{w_g\}$ 
Output: selected set  $S$  with total cost at most  $B$ 
1:  $S \leftarrow$  empty set; spend  $\leftarrow 0$ ; best  $\leftarrow 0$ 
2: repeat
3:   cand  $\leftarrow$  null; best_ratio  $\leftarrow 0$ 
4:   for each unselected  $l$  in  $L$  with spend +  $c_l \leq B$  do
5:     gain  $\leftarrow C(S + l) - C(S)$  // full interference-aware evaluation
6:     if gain /  $c_l >$  best_ratio then cand  $\leftarrow l$ ; best_ratio  $\leftarrow$  gain /  $c_l$ 
7:   end for
8:   if cand = null then break // no affordable candidate helps
9:    $S \leftarrow S + \text{cand}$ ; spend  $\leftarrow$  spend +  $c_{\text{cand}}$ 
10: until false
11:  $l_{\text{star}} \leftarrow \arg \max_{\{l \text{ in } L, c_l \leq B\}} C(\{l\})$  // single-site check
12: if  $C(\{l_{\text{star}}\}) > C(S)$  then  $S \leftarrow \{l_{\text{star}}\}$ 
13: return  $S$ 

```

Line 5 carries the substantive change from the earlier version. The earlier marginal gain came from stand-alone coverage sets, which do not depend on the current selection. Under interference those sets misstate the true marginal contribution, so the pseudocode and the implementation both sat at odds with the stated objective. Both now evaluate $C(S + l)$ directly. The evaluation runs incrementally: for every test point the algorithm keeps the running sum of received powers from the selected set and the current strongest received power, so testing a candidate costs $O(G)$ rather than $O(|S|G)$.

Lines 11-13 hold the single-site comparison of Khuller *et al.*, which returns the better of the greedy set and the best affordable single site. It is what the $(1 - 1/e)$ guarantee of the modified greedy algorithm rests on, so the implementation runs it even though Proposition 1 removes the guarantee here, and it guards against the pathological case in which the ratio rule fills the budget with cheap, individually weak sites. A single cell has no interferers, so the coverage of every singleton can be evaluated in one pass; the check costs $M + 1$ objective evaluations and, across every instance in this study, never changed the returned set. We report it because a reader comparing the released code against Algorithm 1 should find the same procedure in both.

4.4. Lazy Evaluation as an Accelerator

Minoux's accelerated greedy algorithm keeps candidates in a max-heap keyed by their last computed gain-to-cost ratio and re-evaluates only the top entry at each step [15]. With a submodular objective, a stale ratio is an upper bound on the current one. An entry that stays on top after re-evaluation is then provably the true maximiser, and the selection sequence matches that of the exhaustive rule. The earlier version of this paper asserted that identity. Under the interference-aware objective the premise fails, because a stale ratio is no longer an upper bound, so the assertion is incorrect.

This version keeps the accelerator, reports it as a separate variant and measures what using it costs. Over 30 independent instances at each threshold the lazy variant reproduced the exhaustive selection sequence in only 6.7% of cases. Its coverage ran on average 0.60, 0.89 and 1.51 percentage points lower at 0, 3 and 6 dB respectively, with a worst case of 4.42 percentage points. In exchange it spent roughly a quarter of the objective evaluations. Table 15 reports the trade-off in full. Every result labelled "proposed" in this paper uses the exhaustive rule of Algorithm 1, and the lazy variant appears beside it so that a reader who needs the speed can see the price.

4.5. Complexity

Building the received-power matrix costs $O(MG)$. Inside the selection loop, each step evaluates at most M candidates at $O(G)$ each, and at most M steps occur, which gives $O(M^2G)$ in the worst case for the exhaustive rule. For the lazy variant the number of re-evaluations per step depends on the data and admits no bound better than M in the worst case. Rather than assert a tighter bound, Table 15 reports the empirical evaluation counts directly. With $M = 100$ and $G = 10\,000$, the exhaustive rule needed a mean of 1 414 to 1 738 objective evaluations depending on the threshold, and the lazy variant 369 to 426.

Single-threaded runtimes measured on the reference platform (Table 3) came to 0.081-0.101 s for the exhaustive rule and 0.019-0.023 s for the lazy variant, against 1.77-2.19 s for each of the three metaheuristics at an evaluation budget of 3 000. The released run metadata records the full platform specification and the software versions.

5. Experimental Design

5.1. Methods Compared

Nine methods enter the comparison: the proposed algorithm, its lazy variant and seven baselines. Each one works under the same aggregate budget, and each is scored with the same interference-aware objective.

- Random selection. The method visits candidates in random order and accepts each one it can afford.
- Uniform grid (snapped). The method builds an ideal square lattice of target points, visits the lattice points in random order, snaps each one to the nearest not-yet-used candidate and accepts that candidate if it is affordable.
- k-means clustering. Test points are clustered by Lloyd's algorithm, weighted by traffic demand where a demand field is in use, and each centroid is snapped to the nearest candidate the remaining budget affords. The number of clusters is set from the budget and the mean site cost, matching the lattice baseline. Clustering of this kind is a common alternative for site selection.
- Cost-unaware greedy maximum coverage. Identical to Algorithm 1, except that it ranks candidates by absolute marginal gain rather than gain per unit cost. This baseline is the one that separates the contribution of cost-awareness from the contribution of greedy selection itself.
- Simulated annealing. Binary encoding over candidates, geometric cooling from $T = 0.05$ to $1e-4$, and a move set of swap, add, and drop-then-refill. A random repair operator restores feasibility after every move.
- Genetic algorithm. Population 50 with binary tournament selection, uniform crossover, bit-flip mutation at rate $2/M$ and elitism of one, plus the same repair and refill operators.
- Binary particle swarm optimisation. Fifty particles, inertia 0.7, cognitive and social coefficients 1.5, velocity clamped to ± 4 , a sigmoid transfer function, and the same repair and refill operators.
- Proposed budget-aware greedy (Algorithm 1), with the lazy variant reported separately.

Each of the three metaheuristics receives a budget of 3 000 objective evaluations. The genetic algorithm consumes 2 942 of them, because it evaluates an initial population of 50, then carries one elite individual per generation forward unevaluated across 59 generations, and scores the returned set once at the end, and the other two consume 3 001. All three exceed the 1 414-1 738 evaluations the proposed method uses. Starving the metaheuristics of search effort therefore does not bias the comparison towards the proposed method on the synthetic instances. On the larger real inventories of Section 6.11 the same budget does become binding, and the wins it produces are reported there as evidence about the metaheuristics at this budget rather than about the problem.

5.2. Two Budget Regimes

Equal nominal budgets do not mean equal expenditure, because a method that stops early leaves capital unspent. Two regimes therefore carry the results.

In the equal-budget regime, every method receives the same aggregate budget $B = 0.20 * \sum_l c_l$. That matches the question an operator with a fixed capital allocation would ask. It also flatters methods that spend everything, since they may buy coverage the proposed method declines to buy.

The matched-spend regime runs the proposed method first, records what it spent, and then gives every other method exactly that amount as its budget. Equalising expenditure gives the fairer basis for attributing a coverage difference to the selection rule rather than to the level of investment. Both regimes are reported in full, with per-method site counts, expenditure, unused budget and the spread of selected-site costs.

5.3. Monte-Carlo Protocol

Each scenario runs 50 Monte-Carlo iterations. Within an iteration the simulator draws the candidate positions, the per-site costs, the aggregate budget, the LOS state of every link and the shadow-fading sample of every link once, then hands the same draw to every method, so all methods face identical instances. Method-internal randomness comes from a separate stream, seeded so that changing one method cannot perturb the instances another method sees. The released code records the master seed, and the per-iteration seed is a deterministic function of it, so any single instance can be regenerated on its own.

Unless stated otherwise, $M = 100$ candidates, $c_{\max} = 10$, $\rho = 0.20$ and $\alpha = 1$. Scenarios that vary $c_{\max}$, α and the demand weighting are described where they are reported.

5.4. Statistical Analysis

All methods see the same instances, so the comparisons are paired. For each scenario, threshold and competing method, the analysis takes the paired difference in coverage between the proposed method and the competitor across the 50 iterations, then reports the mean difference, its 95% confidence interval from the t-distribution, the paired t-statistic and the standardised effect size (Cohen's d_z). Two robustness measures accompany every comparison: a Wilcoxon signed-rank test, which makes no normality assumption about the differences, and a Shapiro-Wilk test of that assumption. Within each scenario and threshold, the Holm procedure [27] corrects the family of comparisons against the proposed method for multiple testing, and the Holm-corrected p-values are the ones reported for the t-tests. The Wilcoxon p-values beside them are uncorrected and are given as a distribution-free check on the direction and rough magnitude of each result, not as a second significance test. The Shapiro-Wilk statistics are in the released results rather than in a table. Exact p-values appear where they exceed $1e-3$, and scientific notation covers the rest; no p-value is reported as zero.

5.5. Reproducibility

An archive accompanying this paper releases the simulator, the experiment driver, the statistical code and the figure scripts. It carries a metadata file recording the master seed, the number of iterations, the exact package versions, the platform string, the processor count and the wall-clock time of the synthetic run, with the real-location and diagnostic runs recorded separately in the archive. The same archive holds the Overpass queries, the bounding boxes, the extraction date and the extracted coordinates for all three real-location inventories, so the candidate sets can be rebuilt exactly rather than re-queried against a changed map. Reproducing the full result set takes eight commands, listed in order in the archive README, and no manual configuration.

6. Results

Except for Figure 1, which plots individual link samples from the diagnostic run, every figure in this section is a mean over 50 paired Monte-Carlo iterations. The interval after the plus-or-minus sign gives the half-width of the 95% confidence interval of the mean. @@Ts@@ 7 and 10 report differences between methods as paired differences, with confidence intervals and Holm-corrected p-values.

6.1. Consistency Check of the Propagation Generator

A comparison of placement methods is worth making only if the simulator injects shadow fading correctly on top of the analytical path-loss curves. The check generated twenty thousand links at distances log-uniformly distributed between 10 m and 500 m and examined the residual between each simulated sample and the corresponding 3GPP mean. Table 5 reports the outcome; Figure 1 plots the samples against the analytical curves.

In Table 5, the chi-square statistic tests the residual variance against the specified shadow-fading variance; a p-value far from zero indicates consistency.

Table 5. Consistency of the path-loss generator with the 3GPP UMi specification at 3.5 GHz

Condition	Samples	Specified σ (dB)	Residual RMSE (dB)	Bias (dB)	Residual s.d. (dB)	χ^2	p
LOS	9339	4.00	3.99	-0.004	3.99	9274.6	0.65
NLOS	10661	7.82	7.80	+0.040	7.80	10594.7	0.66

Figure 1 plots line-of-sight and non-line-of-sight samples separately. The plot is an internal consistency check of the generator, not an independent validation of the 3GPP model against measurement.




Figure 1. Simulated path-loss samples against the 3GPP UMi mean curves at 3.5 GHz

Residual standard deviations match the specified values to within 0.03 dB, and the biases sit within 0.04 dB of zero. A chi-square test of the residual variance against the specified variance returns $p = 0.65$ for line-of-sight and $p = 0.66$ for non-line-of-sight links, so there is no evidence of a discrepancy. The earlier version quoted root-mean-square errors of 3.5 and 6.5 dB and called them "within sampling variability" without running a test; the present figures of 3.99 and 7.80 dB, against specified standard deviations of 4.00 and 7.82 dB, come with the test that justifies the claim. The scope is narrow. The simulated samples are compared against analytical curves produced by the same code, so the test confirms that shadow fading is injected with the specified distribution on top of those curves; it cannot detect an error in the path-loss formulas themselves, which were instead checked term by term against Table 7.4.1-1 of the specification. The exercise checks the implementation against the specification, not the specification against measurement.

6.2. Comparison under a Common Aggregate Budget

Under the equal-budget regime every method receives the same aggregate budget, 110.5 normalised cost units on average. Table 6 reports coverage for the proposed method, its lazy variant and all seven baselines; Figure 2 shows the same data with confidence intervals.

Entries in Table 6 are means over 50 paired Monte-Carlo iterations with the half-width of the 95% confidence interval of the mean, for $M = 100$ candidates, cost range $[1, 10]$ and budget ratio 0.20.

Table 6. Coverage efficiency (%) under a common aggregate budget

Method	0 dB	3 dB	6 dB
Proposed (BA-greedy)	84.10 $\pm$ 0.14	68.33 $\pm$ 0.20	56.11 $\pm$ 0.28
Proposed (lazy)	83.63 $\pm$ 0.17	67.26 $\pm$ 0.25	54.25 $\pm$ 0.25
Simulated annealing	84.74 $\pm$ 0.11	69.52 $\pm$ 0.14	57.24 $\pm$ 0.16
Genetic algorithm	83.01 $\pm$ 0.14	67.11 $\pm$ 0.18	54.42 $\pm$ 0.17
Binary PSO	82.26 $\pm$ 0.11	66.14 $\pm$ 0.13	53.35 $\pm$ 0.15
Greedy (cost-unaware)	84.42 $\pm$ 0.12	68.55 $\pm$ 0.20	55.85 $\pm$ 0.23
k-means clustering	83.29 $\pm$ 0.19	67.62 $\pm$ 0.25	55.10 $\pm$ 0.30
Uniform grid	82.42 $\pm$ 0.30	66.41 $\pm$ 0.40	53.74 $\pm$ 0.46
Random	77.10 $\pm$ 0.53	59.44 $\pm$ 0.65	45.93 $\pm$ 0.71

In Figure 2, error bars give 95% confidence intervals of the mean over 50 paired iterations ($M = 100$ candidates, cost range $[1, 10]$).

Three features stand out. Simulated annealing takes the highest coverage at every threshold. The proposed method, cost-unaware greedy selection and simulated annealing sit in a tight group spanning at most 1.39 percentage points, k-means clustering follows 0.7 to 0.8 points behind that group, and the genetic algorithm, particle swarm optimisation, lattice placement and random selection trail further. Any structured method pulls further ahead of random selection as the threshold tightens; for the proposed method the gap grows from 7.01 percentage points at 0 dB to 10.18 at 6 dB.




Figure 2. Coverage efficiency against SINR threshold under a common aggregate budget

Table 7 quantifies these differences through paired comparisons. The proposed method also leads k-means clustering by 0.81, 0.71 and 1.00 percentage points at 0, 3 and 6 dB, each significant after correction. Against random selection, the snapped lattice, particle swarm optimisation and the genetic algorithm, the proposed method holds a significant advantage at every threshold after Holm correction, with effect sizes between 1.23 and 4.13. It trails simulated annealing by 0.64, 1.19 and 1.13 percentage points at 0, 3 and 6 dB respectively, all significant. Against cost-unaware greedy selection the picture is mixed: 0.32 percentage points behind at 0 dB, statistically indistinguishable at 3 and 6 dB. The Wilcoxon tests agree with the t-tests in every case, and where the Shapiro-Wilk test indicates non-normal differences the two procedures still reach the same conclusion.

Table 7 reports the paired differences in percentage points between the proposed method and each competitor; positive values favour the proposed method. Intervals are 95% t-intervals on the paired difference, p-values are Holm-corrected within each threshold, and d_z is Cohen's standardised effect size for paired samples.

Table 7. Paired differences in coverage efficiency under a common aggregate budget

γ_{th} (dB)	Compared with	Difference (pp)	95% CI	p (Holm)	p (Wilcoxon)	d_z
0	Proposed (lazy)	+0.47	[+0.34, +0.60]	1.3e-08	7.6e-09	+1.01
0	Simulated annealing	-0.64	[-0.79, -0.49]	1.5e-10	3.4e-08	-1.20
0	Genetic algorithm	+1.10	[+0.92, +1.27]	3.2e-16	1.2e-09	+1.78
0	Binary PSO	+1.85	[+1.69, +2.00]	4.0e-28	7.6e-10	+3.43
0	Greedy (cost-unaware)	-0.32	[-0.52, -0.12]	0.002	0.009	-0.46
0	k-means clustering	+0.81	[+0.57, +1.05]	2.9e-08	1.8e-07	+0.96
0	Uniform grid	+1.68	[+1.38, +1.98]	2.1e-14	8.0e-10	+1.58
0	Random	+7.01	[+6.51, +7.50]	3.4e-31	7.6e-10	+4.02
3	Proposed (lazy)	+1.07	[+0.76, +1.37]	1.7e-08	9.6e-08	+1.00
3	Simulated annealing	-1.19	[-1.40, -0.98]	2.3e-14	2.1e-09	-1.59
3	Genetic algorithm	+1.23	[+0.98, +1.47]	7.4e-13	5.1e-09	+1.43
3	Binary PSO	+2.20	[+2.00, +2.40]	3.0e-26	7.5e-10	+3.12
3	Greedy (cost-unaware)	-0.22	[-0.49, +0.05]	0.106	0.141	-0.23
3	k-means clustering	+0.71	[+0.36, +1.06]	3.4e-04	5.8e-04	+0.58
3	Uniform grid	+1.92	[+1.48, +2.37]	7.0e-11	4.1e-09	+1.23
3	Random	+8.90	[+8.29, +9.51]	9.2e-32	7.6e-10	+4.13
6	Proposed (lazy)	+1.85	[+1.47, +2.24]	2.3e-12	9.8e-14	+1.38
6	Simulated annealing	-1.13	[-1.43, -0.83]	3.0e-09	3.2e-08	-1.07

γ_{th} (dB)	Compared with	Difference (pp)	95% CI	p (Holm)	p (Wilcoxon)	d _z
6	Genetic algorithm	+1.69	[+1.40, +1.97]	2.8e-15	2.8e-09	+1.68
6	Binary PSO	+2.76	[+2.44, +3.07]	6.7e-22	7.6e-10	+2.48
6	Greedy (cost-unaware)	+0.26	[-0.07, +0.59]	0.115	0.104	+0.23
6	k-means clustering	+1.00	[+0.58, +1.43]	3.4e-05	6.0e-05	+0.67
6	Uniform grid	+2.37	[+1.87, +2.86]	3.0e-12	9.6e-10	+1.36
6	Random	+10.18	[+9.46, +10.91]	4.3e-31	1.8e-15	+4.00

6.3. Budget Utilisation and Cost Efficiency

The equal-budget comparison understates the position, because the methods do not spend comparable amounts. Table 8 reports, for every method and threshold, the number of sites selected, the capital spent, the budget left unused, the range of selected-site costs, the objective evaluations consumed and the measured runtime.

In Table 8, "Utilisation" is the fraction of the aggregate budget spent. Costs are in normalised units, and runtimes are single-threaded means on the reference platform of Table 3.

Table 8. Resource accounting under a common aggregate budget

Method	γ_{th}	Sites	Cost spent	Unused	Utilisation (%)	Cost range	Evaluations	Runtime (s)
Proposed (BA-greedy)	0	13.1	36.0	74.5	32.6	1.0–7.7	1414	0.081
Proposed (BA-greedy)	3	13.7	39.6	70.9	35.9	1.0–7.4	1453	0.086
Proposed (BA-greedy)	6	17.3	58.9	51.6	53.3	1.0–8.3	1738	0.101
Proposed (lazy)	0	10.0	24.1	86.4	21.8	1.0–7.2	426	0.023
Proposed (lazy)	3	9.0	21.8	88.7	19.7	1.0–6.8	386	0.021
Proposed (lazy)	6	8.9	22.3	88.2	20.2	1.0–6.7	369	0.019
Simulated annealing	0	24.7	108.4	2.1	98.1	1.0–9.7	3001	1.895
Simulated annealing	3	24.2	107.7	2.8	97.4	1.0–9.8	3001	1.906
Simulated annealing	6	24.2	108.0	2.5	97.7	1.0–9.6	3001	1.878
Genetic algorithm	0	27.0	110.0	0.4	99.6	1.0–9.7	2942	2.152
Genetic algorithm	3	26.6	110.0	0.5	99.6	1.0–9.8	2942	2.190
Genetic algorithm	6	26.1	110.0	0.4	99.6	1.0–9.7	2942	2.130
Binary PSO	0	20.6	110.0	0.4	99.6	1.2–10.0	3001	1.787
Binary PSO	3	21.1	110.0	0.4	99.6	1.1–10.0	3001	1.789
Binary PSO	6	20.8	110.0	0.4	99.6	1.1–9.9	3001	1.767
Greedy (cost-unaware)	0	10.3	56.4	54.1	51.0	1.4–9.3	1071	0.065
Greedy (cost-unaware)	3	9.3	50.7	59.8	45.9	1.7–9.4	976	0.060
Greedy (cost-unaware)	6	8.9	49.5	60.9	44.8	1.8–9.2	942	0.058
k-means clustering	0	19.1	104.3	6.2	94.4	1.2–9.8	1	0.925
k-means clustering	3	19.1	104.3	6.2	94.4	1.2–9.8	1	0.931
k-means clustering	6	19.1	104.3	6.2	94.4	1.2–9.8	1	0.934
Uniform grid	0	20.4	109.1	1.4	98.7	1.1–9.9	1	0.001
Uniform grid	3	20.4	109.1	1.4	98.7	1.1–9.9	1	0.001
Uniform grid	6	20.4	109.1	1.4	98.7	1.1–9.9	1	0.001
Random	0	21.2	110.0	0.4	99.6	1.0–9.8	1	0.000
Random	3	21.2	110.0	0.4	99.6	1.0–9.8	1	0.000
Random	6	21.2	110.0	0.4	99.6	1.0–9.8	1	0.000




Figure 3. Capital spent and number of sites selected under a common aggregate budget

Figure 3 shows the capital spent (left) and the number of sites selected (right) by each method, with 95% confidence intervals of the mean over 50 paired iterations ($M = 100$ candidates, cost range $[1, 10]$). The dashed line marks the aggregate budget. Exhausting the budget does not buy proportionally more coverage.

The proposed method spends between 36.0 and 58.9 normalised units of the 110.5 available, that is between 32.6% and 53.3% of the budget. Random selection, the genetic algorithm and particle swarm optimisation spend 99.6%, the snapped lattice 98.7% and simulated annealing 97.4-98.1%. The algorithm makes the reason plain: selection stops when no affordable candidate improves coverage, and under interference that point arrives well before the money runs out. The candidates left unbought offer no positive marginal gain at the point the algorithm stops, which under interference means the coverage they would add is at least cancelled by the coverage they would remove elsewhere.

Read against Table 6, the proposed method reaches 99.2% of simulated annealing's coverage at 0 dB on a third of the capital and one twenty-third of the computation, and 98.0% of it at 6 dB on 53% of the available budget. Figure 4 plots the resulting position in the coverage-cost plane. The proposed method and its lazy variant lie on the efficient frontier at every threshold; the methods that exhaust the budget cluster at high cost without a matching coverage advantage.

Figure 4 plots each method at each SINR threshold; bars give 95% confidence intervals of the mean on both axes over 50 paired iterations ($M = 100$ candidates, cost range $[1, 10]$). Points towards the upper left mark the more cost-efficient methods.




Figure 4. Mean coverage efficiency against mean capital spent under a common aggregate budget

6.4. Comparison under Matched Expenditure

A second regime separates the effect of the selection rule from the effect of the level of investment: every method runs again with its budget set to the amount the proposed method spent on the same instance. Table 9 reports coverage, Table 10 the paired comparisons; Figure 5 shows the data.

In Table 9, every method receives exactly the capital the proposed method spent on the same instance; entries are means over 50 paired iterations with the half-width of the 95% confidence interval.

Table 9. Coverage efficiency (%) under matched aggregate expenditure

Method	0 dB	3 dB	6 dB
Proposed (BA-greedy)	83.99 ± 0.15	68.49 ± 0.24	56.32 ± 0.31
Proposed (lazy)	83.37 ± 0.15	66.95 ± 0.25	53.96 ± 0.28
Simulated annealing	84.53 ± 0.15	68.80 ± 0.18	56.36 ± 0.24
Genetic algorithm	83.07 ± 0.12	67.03 ± 0.17	53.98 ± 0.22
Binary PSO	82.69 ± 0.17	66.37 ± 0.18	53.20 ± 0.20
Greedy (cost-unaware)	81.68 ± 1.33	66.45 ± 1.09	54.81 ± 0.66
k-means clustering	79.26 ± 2.65	65.10 ± 1.20	52.86 ± 1.35
Uniform grid	79.12 ± 1.42	64.55 ± 0.96	52.47 ± 1.09
Random	73.51 ± 1.84	56.98 ± 1.06	44.00 ± 0.93

In Figure 5, error bars give 95% confidence intervals of the mean over 50 paired iterations.

Table 10 reports the paired differences in percentage points between the proposed method and each competitor; positive values favour the proposed method.




Figure 5. Coverage efficiency against SINR threshold under matched aggregate expenditure

Table 10. Paired differences in coverage efficiency under matched aggregate expenditure

γ_{th} (dB)	Compared with	Difference (pp)	95% CI	p (Holm)	p (Wilcoxon)	d _z
0	Proposed (lazy)	+0.62	[+0.45, +0.79]	7.2e-09	8.7e-09	+1.05
0	Simulated annealing	-0.54	[-0.71, -0.37]	1.0e-07	5.0e-07	-0.93
0	Genetic algorithm	+0.92	[+0.74, +1.10]	5.3e-13	1.1e-09	+1.45
0	Binary PSO	+1.29	[+1.11, +1.48]	7.4e-18	9.6e-10	+1.98
0	Greedy (cost-unaware)	+2.30	[+1.00, +3.60]	0.001	6.6e-05	+0.50
0	k-means clustering	+4.73	[+2.12, +7.35]	0.001	2.9e-07	+0.51
0	Uniform grid	+4.87	[+3.50, +6.24]	1.5e-08	2.2e-09	+1.01
0	Random	+10.48	[+8.69, +12.27]	4.9e-15	1.8e-15	+1.66
3	Proposed (lazy)	+1.54	[+1.18, +1.90]	1.1e-10	2.0e-09	+1.21
3	Simulated annealing	-0.32	[-0.53, -0.10]	0.005	0.010	-0.42
3	Genetic algorithm	+1.46	[+1.21, +1.72]	1.1e-14	3.0e-09	+1.62
3	Binary PSO	+2.12	[+1.80, +2.43]	2.4e-17	1.2e-09	+1.92
3	Greedy (cost-unaware)	+2.04	[+0.99, +3.08]	5.8e-04	6.3e-07	+0.55
3	k-means clustering	+3.39	[+2.27, +4.50]	5.0e-07	9.1e-10	+0.86
3	Uniform grid	+3.94	[+3.09, +4.79]	1.1e-11	5.3e-15	+1.31
3	Random	+11.51	[+10.51, +12.51]	4.4e-27	7.6e-10	+3.26
6	Proposed (lazy)	+2.37	[+1.94, +2.80]	3.3e-14	1.5e-09	+1.56
6	Simulated annealing	-0.04	[-0.25, +0.17]	0.706	0.546	-0.05
6	Genetic algorithm	+2.34	[+2.10, +2.59]	1.1e-23	8.0e-10	+2.73
6	Binary PSO	+3.12	[+2.76, +3.48]	4.9e-22	8.9e-15	+2.49
6	Greedy (cost-unaware)	+1.51	[+0.98, +2.03]	1.1e-06	4.5e-08	+0.81
6	k-means clustering	+3.47	[+2.30, +4.63]	7.3e-07	7.6e-10	+0.85
6	Uniform grid	+3.85	[+2.90, +4.79]	3.9e-10	1.2e-09	+1.16
6	Random	+12.32	[+11.46, +13.18]	1.6e-31	1.8e-15	+4.08

Matched expenditure changes the ordering. The margin over k-means clustering widens sharply, to 4.73, 3.39 and 3.47 percentage points at 0, 3 and 6 dB, because clustering places its centroids by geography and consults cost only to check that the nearest candidate is affordable, so it never trades coverage against price. The proposed method now exceeds cost-unaware greedy selection by 2.30, 2.04 and 1.51 percentage points at 0, 3 and 6 dB, all significant after correction. It beats the genetic algorithm by 0.92 to 2.34 points, particle swarm optimisation by 1.29 to 3.12 points, the snapped lattice by 3.85 to 4.87 points and random selection by 10.48 to 12.32 points. Simulated annealing still leads, by 0.54 and 0.32 percentage points at 0 and 3 dB; at 6 dB the difference of 0.04 points is not significant after correction (Holm-corrected $p = 0.706$, Wilcoxon $p = 0.546$).

Cost-awareness, not greediness alone, is what the proposed method contributes. Given the same money, ranking candidates by gain per unit cost beats ranking them by gain. Against a well-tuned

metaheuristic the method trades a fraction of a percentage point of coverage for a runtime roughly thirteen times shorter, and at the most demanding threshold it gives up nothing measurable.

Table 11 gives the matched-spend counterpart of Table 8, so that the two regimes can be compared on the same terms. Under matched expenditure every method is handed the same aggregate sum, so the interesting columns are how many sites that sum buys and how much of it each method commits. The proposed method and the metaheuristics spend all of it; the lazy variant commits between a third and just over a half, and cost-unaware greedy selection and k-means clustering leave between 11% and 39% unspent, because both stop when their own rule runs out of candidates rather than when the money does.

The columns of Table 11 match Table 8. Every method receives the sum the proposed method spends at that threshold, so utilisation below 100% means the method stopped before the budget was exhausted rather than that it was denied capital.

Table 11. Resource accounting under matched aggregate expenditure

Method	γ th	Sites	Cost spent	Unused	Utilisation (%)	Cost range	Evaluations	Runtime (s)
Proposed (BA-greedy)	0	14.2	42.1	0.0	100.0	1.0–7.6	1362	0.082
Proposed (BA-greedy)	3	15.6	49.6	0.0	100.0	1.0–8.0	1473	0.089
Proposed (BA-greedy)	6	18.5	64.7	0.0	100.0	1.0–8.7	1719	0.109
Proposed (lazy)	0	10.1	23.2	18.9	55.1	1.0–6.6	419	0.025
Proposed (lazy)	3	8.8	20.2	29.4	40.8	1.0–6.1	382	0.022
Proposed (lazy)	6	8.9	22.1	42.6	34.1	1.0–7.0	369	0.020
Simulated annealing	0	10.7	39.8	2.4	94.4	1.1–7.2	3001	1.097
Simulated annealing	3	12.4	47.7	1.9	96.2	1.1–8.4	3001	1.202
Simulated annealing	6	15.6	63.2	1.5	97.7	1.1–8.5	3001	1.370
Genetic algorithm	0	11.5	42.1	0.0	100.0	1.1–8.1	2942	1.348
Genetic algorithm	3	13.9	49.6	0.0	100.0	1.0–8.2	2942	1.449
Genetic algorithm	6	17.1	64.7	0.0	100.0	1.0–9.0	2942	1.638
Binary PSO	0	9.4	42.1	0.0	100.0	1.1–8.3	3001	1.085
Binary PSO	3	10.1	49.6	0.0	100.0	1.3–8.9	3001	1.174
Binary PSO	6	13.4	64.7	0.0	100.0	1.1–9.2	3001	1.346
Greedy (cost-unaware)	0	7.4	36.7	5.4	87.1	1.7–8.8	655	0.042
Greedy (cost-unaware)	3	6.8	35.2	14.4	70.9	1.7–8.8	648	0.042
Greedy (cost-unaware)	6	7.8	39.3	25.4	60.7	1.7–8.7	770	0.048
k-means clustering	0	6.9	35.9	6.2	85.2	2.1–8.5	1	0.367
k-means clustering	3	8.2	43.5	6.1	87.8	1.7–8.8	1	0.431
k-means clustering	6	10.8	57.8	6.9	89.3	1.9–9.3	1	0.556
Uniform grid	0	7.5	39.8	2.3	94.4	2.1–8.5	1	0.000
Uniform grid	3	9.2	47.3	2.3	95.3	1.8–9.1	1	0.000
Uniform grid	6	12.2	62.3	2.4	96.3	1.6–9.2	1	0.000
Random	0	8.8	42.1	0.0	100.0	1.4–9.0	1	0.000
Random	3	10.3	49.6	0.0	100.0	1.3–9.1	1	0.000
Random	6	12.7	64.7	0.0	100.0	1.5–9.4	1	0.000

Four of the baselines turn markedly more variable in this regime. k-means clustering is the most variable method in the study, with standard deviations of 4.2 to 9.3 percentage points; random selection, the snapped lattice and cost-unaware greedy selection show 2.3 to 6.5, against 0.5 to 1.1 for the proposed method. Under a tight budget a rule that ignores cost sometimes commits its capital to a handful of expensive sites, and occasionally that fails badly. The cost-aware rule wins on robustness across instances as well as on the average.

6.5. Sensitivity to the Cost Range

The earlier version tested cost-range sensitivity only at 0 dB, where coverage sits close to saturation and the test says little. Table 12 repeats it at all three thresholds, with confidence intervals and with the resulting site counts and expenditure.

Table 12. Sensitivity of the proposed method to the width of the per-site cost range

Cost range	0 dB	Sites	Spend	3 dB	Sites	Spend	6 dB	Sites	Spend
[1, 10]	84.10 $\pm$ 0.14	13.1	36.0	68.33 $\pm$ 0.20	13.7	39.6	56.11 $\pm$ 0.28	17.3	58.9
[1, 30]	84.12 $\pm$ 0.15	14.6	110.4	68.61 $\pm$ 0.25	16.1	138.8	56.63 $\pm$ 0.28	19.7	188.4
[1, 50]	84.14 $\pm$ 0.15	15.0	176.9	68.71 $\pm$ 0.22	16.7	223.1	56.23 $\pm$ 0.33	18.5	257.6

Table 12 covers all three SINR thresholds. Coverage is a percentage with the half-width of the 95% confidence interval; sites and cost are means over 50 iterations.

Figure 6 shows each SINR threshold separately; error bars give 95% confidence intervals of the mean.




Figure 6. Coverage efficiency of the proposed method against the width of the cost range

Coverage varies by at most 0.04 percentage points across the three cost ranges at 0 dB, 0.38 at 3 dB and 0.52 at 6 dB. Each of those spreads is smaller than the width of the confidence intervals at the individual points, so there is no evidence that the width of the cost distribution affects the coverage the method achieves. What the method buys does change. As the cost range widens the algorithm generally selects slightly more sites, though at 6 dB the count falls back from 19.7 to 18.5 between the two widest ranges and spends substantially more in absolute normalised units, because the budget is proportional to the total cost of the candidate set and therefore grows with c_{max} . The insensitivity here is an empirical result, not a consequence of the scale invariance of Section 3.2. Widening the range from [1, 10] to [1, 50] is not a rescaling; it changes the relative spread of the costs as well as their unit, which is why the site counts and the absolute expenditure move while the coverage does not. It should not be read as a general claim of robustness.

6.6. Sensitivity to Offered Load

The default model has every non-serving cell transmitting continuously, the worst case for interference. Table 13 reports coverage as the interferer activity factor drops to 0.50 and 0.25, standing in for networks at partial load.

Table 13 reports the proposed method at each SINR threshold; $\alpha = 1$ is the fully loaded default.

Table 13. Coverage efficiency (%) against the interferer activity factor α

Activity factor α	0 dB	3 dB	6 dB
0.25	99.81 ± 0.02	95.85 ± 0.08	84.00 ± 0.17
0.50	96.18 ± 0.06	84.02 ± 0.17	68.88 ± 0.22
1.00	84.10 ± 0.14	68.33 ± 0.20	56.11 ± 0.28

In Figure 7, error bars give 95% confidence intervals of the mean over 50 paired iterations ($M = 100$ candidates, cost range [1, 10]). Lower activity corresponds to a more lightly loaded network and less aggregate interference.




Figure 7. Coverage efficiency of the proposed method against the interferer activity factor

Coverage climbs sharply as load falls. At 0 dB it goes from 84.10% at full load to 96.18% at half load and 99.81% at quarter load; at 6 dB, from 56.11% to 68.88% and 84.00%. The near-total coverage reported in the earlier version of this paper reappears only in the lightly loaded, low-threshold corner of this space. Two conclusions follow. First, the headline coverage figures in this paper are conservative by construction, and a real network carrying partial load would sit between the $\alpha = 1$ and $\alpha = 0.5$ rows. Second, a coverage claim for an interference-limited deployment means nothing unless the load assumption comes with it, since the same placement spans a 28-percentage-point range at 6 dB depending on that assumption.

6.7. Non-Uniform Traffic Demand

Table 14 reports a scenario that weights test points by a demand field, a unit background plus four Gaussian hotspots, so that coverage of busy locations counts for more.

The traffic field behind Table 14 is a unit background plus four Gaussian hotspots of peak weight 4 and standard deviation 120 m, repositioned at random each iteration.

Table 14. Demand-weighted coverage efficiency (%) with non-uniform traffic demand

Method	0 dB	3 dB	6 dB
Proposed (BA-greedy)	85.04 $\pm$ 0.24	69.57 $\pm$ 0.34	56.91 $\pm$ 0.41
Proposed (lazy)	84.79 $\pm$ 0.24	68.93 $\pm$ 0.37	56.15 $\pm$ 0.45
Greedy (cost-unaware)	85.87 $\pm$ 0.21	70.62 $\pm$ 0.32	58.27 $\pm$ 0.44
Uniform grid	81.56 $\pm$ 0.44	65.48 $\pm$ 0.52	52.90 $\pm$ 0.55
Random	77.31 $\pm$ 0.63	59.56 $\pm$ 0.72	45.94 $\pm$ 0.74

The ordering against random selection and the snapped lattice holds, and the margins widen slightly, to between 7.74 and 10.97 percentage points over random selection. Against cost-unaware greedy selection the proposed method now trails by 0.83 to 1.36 percentage points under the equal-budget regime shown, a larger and more consistent deficit than the equal-budget result of Section 6.2, where the proposed method trails only at 0 dB and is marginally ahead at 6 dB. When demand concentrates, the highest-weight regions are worth covering even at high cost, and a rule that discounts by cost is slightly reluctant to buy them. That is a genuine limitation of gain-per-cost ranking under skewed demand. It points towards a demand-aware cost discount as a worthwhile extension rather than towards a defence of the present rule.

6.8. Cost of the Lazy Accelerator

Table 15 puts numbers on the trade-off introduced in Section 4.4.

In Table 15, "Identical sequence" is the percentage of instances in which the lazy variant selected exactly the same sites in the same order; the coverage gap is how far the lazy variant falls short, in percentage points.

Table 15. Lazy accelerator against exhaustive re-evaluation

γ th (dB)	Identical sequence (%)	Mean gap (pp)	Max gap (pp)	Evaluations exact	Evaluations lazy	Runtime exact (s)	Runtime lazy (s)
0	6.7	0.60	2.72	1414	426	0.081	0.023
3	6.7	0.89	4.08	1453	386	0.086	0.021
6	6.7	1.51	4.42	1738	369	0.101	0.019

The lazy variant reproduces the exhaustive selection in 6.7% of instances. It loses 0.60 to 1.51 percentage points of coverage on average, up to 4.42 points in the worst case, and in exchange uses roughly a quarter of the objective evaluations and slightly less than a quarter of the runtime. At the scale studied here the exhaustive rule already runs in about a tenth of a second, so the accelerator is not worth its cost. It earns its place only for candidate sets one or two orders of magnitude larger, where the quadratic term in Section 4.5 begins to bite.

Tables 13, 16 and 17 report the proposed method alone, because their purpose is to show how it responds to a changed assumption rather than to re-rank the methods; Table 14 reports five methods. The deterministic baselines and cost-unaware greedy selection were run in every sensitivity scenario and the paired comparisons are in the released results, but the three metaheuristics were not, because each costs about 3 000 objective evaluations per instance and the sensitivity sweeps multiply the instance count. Two comparisons in the released results go against the proposed method and are worth naming here: at 0 dB, cost-unaware greedy selection is ahead by 0.33 and 0.31 percentage points in the two wider cost ranges and

by 0.32 points under the load-aware association scenario, each significant after correction. This is the same 0 dB deficit that Section 6.2 reports, and it persists under those changes of assumption.

6.9. Sensitivity to Carrier Frequency

All results so far sit at 3.5 GHz. Because the case for urban furniture is strongest where cells are smallest, Table 16 repeats the equal-budget experiment at 28 GHz and 39 GHz, keeping every other parameter fixed except the breakpoint distance, which is proportional to the carrier frequency and therefore moves with it and using the same 3GPP UMi model, which is specified up to 100 GHz.

Table 16 gives means over 50 paired iterations with the half-width of the 95% confidence interval; all other parameters are held at the values in Table 3.

Table 16. Coverage efficiency (%) and sites bought as the carrier frequency rises

Carrier (GHz)	0 dB	Sites	3 dB	Sites	6 dB	Sites
3.5	84.10 ± 0.14	13.1	68.33 ± 0.20	13.7	56.11 ± 0.28	17.3
28	80.27 ± 0.13	35.7	63.19 ± 0.15	33.2	49.58 ± 0.15	31.6
39	77.43 ± 0.14	36.7	60.55 ± 0.15	34.5	47.31 ± 0.16	32.9

Coverage falls with frequency, by 3.83 percentage points at 0 dB and 6.53 points at 6 dB moving from 3.5 GHz to 28 GHz, and by a further 2.84 and 2.27 points at 39 GHz. The more informative change is what the algorithm buys. At 3.5 GHz it selects 13 to 17 sites and spends 36 to 59 normalised units. At 28 GHz it selects 31.6 to 35.7 sites on average and spends about 107, which is almost the whole budget. The plausible mechanism is that higher frequency shrinks each cell, so far fewer candidates become redundant and the interference-driven stopping point that leaves between 47% and 67% of the budget unspent at 3.5 GHz largely disappears.

This matters for how the paper should be read. The headline finding that a cost-aware rule reaches good coverage while declining to spend most of the budget is a mid-band result. In millimetre-wave deployment the same rule still works but the economics reverse: the constraint becomes how many sites the budget affords rather than how many the interference environment can tolerate.

6.10. Association Rule and Probabilistic Coverage

Two modelling simplifications drew comment. The default rule attaches each test point to the strongest received signal, which ignores load balancing, and coverage is scored as a hard threshold on one shadowing realisation rather than as a probability. Table 17 addresses both on the same selections.

In Table 17, "Load-aware" attaches each test point to the cell maximising received power scaled by a bias in the inverse square root of that cell's load, clipped to [0.25, 4] and recomputed over five fixed-point rounds. "Load CV" is the coefficient of variation of the per-cell load under that rule. The probability columns hold the chosen sites and their mean path loss fixed, draw one line-of-sight field, and redraw shadow fading 40 times on it.

Table 17. Association rule and probabilistic coverage for the proposed method

γ_{th} (dB)	Coverage, max-power	Coverage, load-aware	Load CV	Mean coverage probability	Area with probability ≥ 0.95
0	84.12 ± 0.17	84.01 ± 0.18	0.183	83.34 ± 0.16	40.05 ± 0.78
3	68.65 ± 0.26	68.65 ± 0.26	0.194	67.69 ± 0.19	28.31 ± 0.95
6	56.06 ± 0.38	56.06 ± 0.38	0.190	55.23 ± 0.28	24.18 ± 0.86

Load-aware association changes coverage by at most 0.11 percentage points, and at 3 and 6 dB it changes nothing at all, while spreading load across cells with a coefficient of variation near 0.19. The reason is structural rather than a modelling artefact, though not because the interference is unaffected. Every selected cell transmits at full power whatever the association, so the total received power at a test point is fixed; moving that point to a less loaded neighbour moves power from the numerator of equation (6) to the denominator, and the signal-to-interference-plus-noise ratio can only fall. The change is small because the bias rarely overturns a large power difference, and it is one-directional: at 0 dB the load-aware rule gives lower coverage in all fifty iterations. Load balancing therefore redistributes served traffic at a slight cost in coverage, and coverage is the quantity this paper optimises. A model in which cells could sleep or reduce power under light load would behave differently, and Section 7.3 lists that as a limitation.

The probabilistic view is less comfortable. Averaging over redrawn shadow fading gives mean coverage probabilities of 83.34%, 67.69% and 55.23%, within one percentage point of the single-realisation figures, so the headline numbers are not inflated. The area covered reliably is another matter. Only 40.05% of the area meets the 0 dB threshold with probability at least 0.95, falling to 28.31% at 3 dB and 24.18% at 6 dB. A threshold metric computed on one realisation describes the average outcome well and the dependable outcome poorly. Any planning target drawn from figures of this kind should state which of the two it means.

6.11. Real-Location Study: Three Swiss City Centres

The experiments reported so far place candidates uniformly at random. This section swaps them for the coordinates of real urban furniture, to test whether the spatial structure of a genuine street network (clustering along carriageways, gaps across blocks, concentration at junctions) changes the standing of the methods.

The candidate inventories come from OpenStreetMap, one 1 km² tile in the centre of Zürich, Bern and Basel, covering street lighting columns, bus stops and traffic-signal masts. Nodes closer together than 2 m collapse into one, since two cells cannot sit at effectively the same position. Table 18 reports the resulting inventories, which differ markedly in size and in composition: signal masts dominate Zürich, bus stops and lamps dominate Bern, and Basel is the sparsest of the three at about 37% of Zürich's candidate density. That spread is useful. It lets the same comparison run across a 2.7-fold range of candidate density.

The inventories in Table 18 were extracted on 20 August 2026 from the overpass.osm.ch mirror (data © OpenStreetMap contributors, ODbL). The released archive holds the Overpass queries and the extracted coordinates.

Table 18. OpenStreetMap urban-furniture inventories for the three 1 km² study tiles

City	Bounding box (S, W, N, E)	Candidates	Street lamps	Bus stops	Traffic signals
Zürich	47.3700, 8.5300, 47.3790, 8.5433	154	57	10	87
Bern	46.9440, 7.4380, 46.9530, 7.4511	98	40	45	13
Basel	47.5500, 7.5800, 47.5590, 7.5933	57	7	11	39

Two departures from the earlier version of this study deserve mention. First, the candidate set is the whole inventory rather than a random sample of 100 positions drawn afresh each iteration. Random subsampling corresponds to nothing an operator does: the inventory is what it is, and the planning problem is to choose within it. Variation across Monte-Carlo iterations therefore comes from the per-site costs and the channel realisation, which is what the paired comparison requires, and not from an artificial resampling of the asset base. Second, the candidate set is no longer almost entirely streetlights; each city carries the three asset classes in materially different proportions.

The three tiles are the central squares of the three largest German-speaking Swiss cities, chosen because the Swiss mirror was the only reachable Overpass endpoint and because OpenStreetMap tagging of street furniture is dense in all three. We should be plain about the limits of that choice. It was not made to sample a population of cities, and no argument is offered that these tiles are representative of anything wider. What three inventories buy over one is a spread in density and in asset mix wide enough to ask whether the results depend on either, which the single Berlin tile of the previous version could not support. Section 7.3 states the geographical limitation.

This remains a real-location simulation and not a validation against real infrastructure. The coordinates are real. The propagation is simulated, the per-site costs come from the same normalised distribution used elsewhere in the paper, and nothing is known about structural capacity, power availability, backhaul reach or planning consent at any of these positions. The coverage figures below are not offered as figures that would be measured in the field. The earlier version called the corresponding experiment a real-infrastructure case study, which overstated what such data can support.

Figure 8 also circles the sites chosen by the proposed algorithm in one representative realisation at $\gamma_{\text{th}} = 3$ dB. Clustering along carriageways and at junctions is the structural feature the uniformly random synthetic study lacks.




Figure 8. Candidate inventories for the three study tiles, coloured by asset class

Table 19 reports coverage under both budget regimes; Figure 9 shows the matched-spend results with confidence intervals.

Entries in Table 19 are means over 50 paired Monte-Carlo iterations with the half-width of the 95% confidence interval of the mean.

Table 19. Coverage efficiency (%) on real urban-furniture inventories

Method	Equal budget, 0 dB	3 dB	6 dB	Matched spend, 0 dB	3 dB	6 dB
Zürich						
Proposed (BA-greedy)	83.00 ± 0.13	66.87 ± 0.18	53.91 ± 0.18	82.94 ± 0.15	66.86 ± 0.19	54.01 ± 0.17
Proposed (lazy)	82.74 ± 0.16	66.18 ± 0.21	53.12 ± 0.26	82.55 ± 0.17	66.06 ± 0.25	53.15 ± 0.25
Simulated annealing	81.27 ± 0.13	64.63 ± 0.19	51.55 ± 0.23	83.52 ± 0.13	67.34 ± 0.18	54.21 ± 0.26
Genetic algorithm	78.54 ± 0.11	60.70 ± 0.14	46.72 ± 0.22	81.87 ± 0.11	65.23 ± 0.19	51.83 ± 0.19
Binary PSO	77.98 ± 0.12	59.85 ± 0.16	46.02 ± 0.19	81.05 ± 0.18	64.53 ± 0.17	51.36 ± 0.18
Greedy (cost-unaware)	83.51 ± 0.13	67.71 ± 0.15	55.00 ± 0.14	78.81 ± 1.07	62.80 ± 0.89	50.81 ± 0.68
k-means clustering	78.83 ± 0.13	61.03 ± 0.17	47.05 ± 0.19	75.27 ± 1.60	59.55 ± 1.74	46.91 ± 1.25
Uniform grid	77.58 ± 0.23	59.20 ± 0.31	44.87 ± 0.38	76.58 ± 1.56	60.77 ± 1.45	48.33 ± 1.23
Random	71.12 ± 0.66	51.28 ± 0.81	36.80 ± 0.83	70.07 ± 1.79	52.27 ± 1.58	39.92 ± 1.16
Bern						
Proposed (BA-greedy)	82.20 ± 0.16	65.75 ± 0.22	52.39 ± 0.27	82.28 ± 0.16	65.81 ± 0.21	52.52 ± 0.23
Proposed (lazy)	81.99 ± 0.20	65.48 ± 0.21	52.18 ± 0.28	81.94 ± 0.21	65.41 ± 0.27	51.95 ± 0.26
Simulated annealing	81.61 ± 0.16	64.88 ± 0.17	51.45 ± 0.23	83.43 ± 0.14	66.87 ± 0.29	53.70 ± 0.22
Genetic algorithm	78.77 ± 0.15	61.12 ± 0.16	47.69 ± 0.22	82.01 ± 0.17	65.35 ± 0.20	51.89 ± 0.18
Binary PSO	78.35 ± 0.12	60.35 ± 0.20	46.87 ± 0.16	81.41 ± 0.17	64.74 ± 0.26	51.28 ± 0.22
Greedy (cost-unaware)	83.29 ± 0.14	67.29 ± 0.11	54.27 ± 0.16	80.51 ± 0.99	64.22 ± 1.14	50.78 ± 0.94
k-means clustering	78.94 ± 0.17	61.43 ± 0.24	47.87 ± 0.29	77.43 ± 1.85	59.12 ± 3.45	46.05 ± 2.25
Uniform grid	77.65 ± 0.29	59.39 ± 0.37	45.46 ± 0.43	76.88 ± 1.84	58.36 ± 3.38	47.22 ± 1.61
Random	67.74 ± 1.16	48.50 ± 1.25	34.86 ± 1.20	66.72 ± 1.56	47.83 ± 1.75	34.66 ± 1.41
Basel						
Proposed (BA-greedy)	79.61 ± 0.19	63.32 ± 0.17	50.54 ± 0.16	79.59 ± 0.17	63.29 ± 0.20	50.35 ± 0.18
Proposed (lazy)	79.50 ± 0.22	63.04 ± 0.19	50.08 ± 0.23	79.43 ± 0.19	62.98 ± 0.25	49.70 ± 0.22
Simulated annealing	81.02 ± 0.11	64.54 ± 0.15	51.57 ± 0.10	81.04 ± 0.15	64.67 ± 0.17	51.65 ± 0.17
Genetic algorithm	78.73 ± 0.15	61.68 ± 0.20	48.58 ± 0.18	79.68 ± 0.16	63.01 ± 0.21	49.82 ± 0.20
Binary PSO	78.92 ± 0.17	61.70 ± 0.17	48.56 ± 0.19	79.42 ± 0.18	62.54 ± 0.16	49.48 ± 0.19
Greedy (cost-unaware)	80.38 ± 0.11	63.87 ± 0.13	50.85 ± 0.14	78.40 ± 0.77	61.65 ± 0.87	49.23 ± 0.61
k-means clustering	78.20 ± 0.21	60.25 ± 0.27	46.51 ± 0.30	75.82 ± 1.57	58.77 ± 1.12	46.04 ± 1.05
Uniform grid	76.21 ± 0.57	58.07 ± 0.75	44.34 ± 0.82	75.74 ± 0.94	59.04 ± 1.00	46.72 ± 0.57
Random	68.93 ± 0.91	49.46 ± 1.01	35.66 ± 1.05	68.58 ± 1.72	50.28 ± 1.75	36.86 ± 1.59

In Figure 9, error bars give 95% confidence intervals of the mean over 50 paired iterations.

Table 20 reports the differences in percentage points between the proposed method and each competitor, under both budget regimes and by city. Positive values favour the proposed method and negative values the competitor; p-values are Holm-corrected within each city, regime and threshold.




Figure 9. Coverage efficiency against SINR threshold by city under matched expenditure

Table 20. Paired differences in coverage efficiency on the real inventories

γ_{th} (dB)	Compared with	Difference (pp)	95% CI	p (Holm)	d _z
Zürich, common aggregate budget					
0	Proposed (lazy)	+0.26	[+0.12, +0.40]	6.2e-04	+0.52
0	Simulated annealing	+1.73	[+1.56, +1.91]	4.7e-25	+2.88
0	Genetic algorithm	+4.46	[+4.30, +4.62]	2.4e-45	+7.96
0	Binary PSO	+5.02	[+4.85, +5.20]	7.4e-46	+8.18
0	Greedy (cost-unaware)	-0.51	[-0.69, -0.32]	3.0e-06	-0.77
0	k-means clustering	+4.18	[+4.00, +4.35]	4.3e-42	+6.79
0	Uniform grid	+5.42	[+5.18, +5.66]	9.4e-41	+6.34
0	Random	+11.89	[+11.24, +12.53]	7.8e-37	+5.22
3	Proposed (lazy)	+0.69	[+0.47, +0.91]	5.2e-08	+0.91
3	Simulated annealing	+2.24	[+1.99, +2.49]	1.2e-22	+2.53
3	Genetic algorithm	+6.17	[+5.95, +6.39]	2.8e-45	+7.93
3	Binary PSO	+7.02	[+6.80, +7.23]	1.1e-48	+9.35
3	Greedy (cost-unaware)	-0.84	[-1.06, -0.63]	3.8e-10	-1.13
3	k-means clustering	+5.84	[+5.61, +6.07]	4.0e-43	+7.13
3	Uniform grid	+7.67	[+7.30, +8.04]	2.7e-39	+5.91
3	Random	+15.59	[+14.78, +16.41]	1.1e-37	+5.45
6	Proposed (lazy)	+0.80	[+0.57, +1.02]	5.7e-09	+1.00
6	Simulated annealing	+2.37	[+2.05, +2.68]	8.9e-20	+2.16
6	Genetic algorithm	+7.19	[+6.91, +7.47]	1.7e-43	+7.27
6	Binary PSO	+7.90	[+7.65, +8.15]	5.5e-48	+9.05
6	Greedy (cost-unaware)	-1.09	[-1.31, -0.86]	9.6e-13	-1.38
6	k-means clustering	+6.87	[+6.61, +7.12]	2.6e-44	+7.57
6	Uniform grid	+9.05	[+8.63, +9.46]	2.9e-40	+6.20
6	Random	+17.12	[+16.28, +17.95]	3.8e-39	+5.84
Bern, common aggregate budget					
0	Proposed (lazy)	+0.21	[+0.10, +0.32]	5.1e-04	+0.53
0	Simulated annealing	+0.59	[+0.33, +0.84]	4.6e-05	+0.66
0	Genetic algorithm	+3.43	[+3.21, +3.64]	1.7e-33	+4.49
0	Binary PSO	+3.85	[+3.65, +4.05]	1.3e-37	+5.51
0	Greedy (cost-unaware)	-1.09	[-1.30, -0.88]	6.8e-14	-1.51
0	k-means clustering	+3.26	[+3.04, +3.47]	1.6e-32	+4.27
0	Uniform grid	+4.55	[+4.22, +4.87]	2.6e-31	+4.00
0	Random	+14.45	[+13.29, +15.61]	6.1e-29	+3.53
3	Proposed (lazy)	+0.27	[+0.12, +0.41]	5.9e-04	+0.52
3	Simulated annealing	+0.87	[+0.59, +1.14]	1.7e-07	+0.89
3	Genetic algorithm	+4.63	[+4.33, +4.93]	4.5e-33	+4.38
3	Binary PSO	+5.40	[+5.08, +5.73]	1.8e-34	+4.71
3	Greedy (cost-unaware)	-1.54	[-1.78, -1.30]	8.5e-17	-1.81
3	k-means clustering	+4.32	[+3.98, +4.67]	4.7e-29	+3.56
3	Uniform grid	+6.36	[+5.99, +6.72]	2.6e-35	+4.92
3	Random	+17.25	[+16.04, +18.45]	1.3e-31	+4.06
6	Proposed (lazy)	+0.22	[+0.04, +0.39]	0.019	+0.34

γ th (dB)	Compared with	Difference (pp)	95% CI	p (Holm)	d _z
6	Simulated annealing	+0.94	[+0.56, +1.33]	2.0e-05	+0.70
6	Genetic algorithm	+4.70	[+4.34, +5.06]	7.0e-30	+3.72
6	Binary PSO	+5.53	[+5.19, +5.86]	2.0e-34	+4.71
6	Greedy (cost-unaware)	-1.88	[-2.20, -1.56]	2.0e-15	-1.67
6	k-means clustering	+4.52	[+4.11, +4.93]	2.1e-26	+3.11
6	Uniform grid	+6.93	[+6.47, +7.40]	2.2e-32	+4.25
6	Random	+17.54	[+16.35, +18.72]	3.2e-32	+4.20
Basel, common aggregate budget					
0	Proposed (lazy)	+0.11	[+0.01, +0.21]	0.033	+0.31
0	Simulated annealing	-1.41	[-1.63, -1.19]	8.0e-17	-1.86
0	Genetic algorithm	+0.88	[+0.63, +1.13]	1.1e-08	+1.01
0	Binary PSO	+0.69	[+0.44, +0.94]	2.8e-06	+0.78
0	Greedy (cost-unaware)	-0.76	[-0.97, -0.56]	8.2e-09	-1.04
0	k-means clustering	+1.42	[+1.12, +1.71]	3.2e-12	+1.36
0	Uniform grid	+3.40	[+2.83, +3.97]	1.7e-15	+1.71
0	Random	+10.68	[+9.72, +11.65]	2.1e-26	+3.15
3	Proposed (lazy)	+0.28	[+0.13, +0.43]	4.3e-04	+0.53
3	Simulated annealing	-1.22	[-1.48, -0.97]	1.9e-12	-1.37
3	Genetic algorithm	+1.63	[+1.38, +1.89]	6.1e-17	+1.85
3	Binary PSO	+1.62	[+1.37, +1.88]	1.5e-16	+1.80
3	Greedy (cost-unaware)	-0.55	[-0.77, -0.34]	6.9e-06	-0.74
3	k-means clustering	+3.06	[+2.72, +3.41]	2.9e-22	+2.53
3	Uniform grid	+5.25	[+4.45, +6.05]	6.0e-17	+1.86
3	Random	+13.86	[+12.83, +14.89]	3.9e-30	+3.81
6	Proposed (lazy)	+0.46	[+0.24, +0.68]	2.6e-04	+0.59
6	Simulated annealing	-1.02	[-1.20, -0.84]	5.1e-15	-1.62
6	Genetic algorithm	+1.96	[+1.72, +2.21]	3.5e-20	+2.25
6	Binary PSO	+1.98	[+1.71, +2.25]	5.9e-19	+2.09
6	Greedy (cost-unaware)	-0.31	[-0.49, -0.12]	0.002	-0.47
6	k-means clustering	+4.03	[+3.69, +4.37]	1.4e-27	+3.34
6	Uniform grid	+6.20	[+5.32, +7.08]	2.0e-18	+2.01
6	Random	+14.88	[+13.80, +15.97]	1.2e-30	+3.91
Zürich, matched expenditure					
0	Proposed (lazy)	+0.39	[+0.27, +0.51]	5.8e-08	+0.93
0	Simulated annealing	-0.58	[-0.76, -0.40]	5.8e-08	-0.93
0	Genetic algorithm	+1.07	[+0.89, +1.25]	2.2e-15	+1.69
0	Binary PSO	+1.89	[+1.64, +2.14]	6.0e-19	+2.11
0	Greedy (cost-unaware)	+4.12	[+3.04, +5.21]	2.1e-09	+1.08
0	k-means clustering	+7.67	[+6.02, +9.32]	9.0e-12	+1.32
0	Uniform grid	+6.36	[+4.77, +7.95]	6.7e-10	+1.14
0	Random	+12.87	[+11.04, +14.69]	4.2e-18	+2.00
3	Proposed (lazy)	+0.80	[+0.58, +1.03]	8.1e-09	+1.01
3	Simulated annealing	-0.48	[-0.74, -0.22]	5.2e-04	-0.53
3	Genetic algorithm	+1.63	[+1.40, +1.85]	1.3e-18	+2.05
3	Binary PSO	+2.33	[+2.06, +2.60]	5.8e-22	+2.49
3	Greedy (cost-unaware)	+4.06	[+3.18, +4.94]	1.1e-11	+1.31
3	k-means clustering	+7.31	[+5.54, +9.08]	2.0e-10	+1.17
3	Uniform grid	+6.09	[+4.63, +7.54]	1.8e-10	+1.19
3	Random	+14.59	[+12.95, +16.23]	3.0e-22	+2.53
6	Proposed (lazy)	+0.86	[+0.66, +1.07]	7.8e-11	+1.20
6	Simulated annealing	-0.21	[-0.45, +0.03]	0.089	-0.25
6	Genetic algorithm	+2.17	[+1.96, +2.39]	6.3e-25	+2.90
6	Binary PSO	+2.64	[+2.43, +2.85]	8.5e-29	+3.56
6	Greedy (cost-unaware)	+3.19	[+2.51, +3.88]	5.2e-12	+1.32
6	k-means clustering	+7.09	[+5.87, +8.31]	4.3e-15	+1.65
6	Uniform grid	+5.67	[+4.47, +6.87]	4.4e-12	+1.34
6	Random	+14.08	[+12.91, +15.26]	5.6e-28	+3.41
Bern, matched expenditure					
0	Proposed (lazy)	+0.34	[+0.22, +0.46]	1.9e-06	+0.82
0	Simulated annealing	-1.15	[-1.34, -0.95]	6.2e-15	-1.65
0	Genetic algorithm	+0.27	[+0.07, +0.47]	0.009	+0.38
0	Binary PSO	+0.87	[+0.66, +1.08]	2.2e-10	+1.20

γ_{th} (dB)	Compared with	Difference (pp)	95% CI	p (Holm)	d _z
0	Greedy (cost-unaware)	+1.77	[+0.74, +2.81]	0.002	+0.49
0	k-means clustering	+4.85	[+3.00, +6.71]	9.7e-06	+0.74
0	Uniform grid	+5.40	[+3.55, +7.24]	1.8e-06	+0.83
0	Random	+15.56	[+13.97, +17.15]	5.7e-24	+2.78
3	Proposed (lazy)	+0.40	[+0.26, +0.55]	2.9e-06	+0.81
3	Simulated annealing	-1.06	[-1.37, -0.75]	5.8e-08	-0.98
3	Genetic algorithm	+0.46	[+0.23, +0.69]	5.0e-04	+0.58
3	Binary PSO	+1.07	[+0.76, +1.38]	5.6e-08	+0.98
3	Greedy (cost-unaware)	+1.59	[+0.43, +2.74]	0.008	+0.39
3	k-means clustering	+6.69	[+3.27, +10.11]	5.3e-04	+0.56
3	Uniform grid	+7.45	[+4.10, +10.80]	1.8e-04	+0.63
3	Random	+17.98	[+16.21, +19.76]	1.1e-24	+2.88
6	Proposed (lazy)	+0.57	[+0.40, +0.74]	6.2e-08	+0.96
6	Simulated annealing	-1.18	[-1.42, -0.95]	1.3e-12	-1.42
6	Genetic algorithm	+0.63	[+0.40, +0.86]	3.6e-06	+0.77
6	Binary PSO	+1.24	[+0.95, +1.54]	2.3e-10	+1.20
6	Greedy (cost-unaware)	+1.74	[+0.83, +2.64]	3.4e-04	+0.54
6	k-means clustering	+6.47	[+4.27, +8.67]	9.7e-07	+0.84
6	Uniform grid	+5.30	[+3.72, +6.88]	6.8e-08	+0.95
6	Random	+17.86	[+16.47, +19.24]	2.2e-29	+3.67
Basel, matched expenditure					
0	Proposed (lazy)	+0.16	[+0.10, +0.23]	7.7e-05	+0.68
0	Simulated annealing	-1.45	[-1.68, -1.22]	5.7e-16	-1.76
0	Genetic algorithm	-0.09	[-0.32, +0.14]	0.423	-0.11
0	Binary PSO	+0.17	[-0.06, +0.40]	0.276	+0.21
0	Greedy (cost-unaware)	+1.19	[+0.36, +2.01]	0.017	+0.41
0	k-means clustering	+3.77	[+2.18, +5.37]	7.7e-05	+0.67
0	Uniform grid	+3.85	[+2.85, +4.85]	2.7e-09	+1.10
0	Random	+11.01	[+9.29, +12.73]	2.1e-16	+1.82
3	Proposed (lazy)	+0.31	[+0.18, +0.45]	6.9e-05	+0.66
3	Simulated annealing	-1.38	[-1.64, -1.12]	1.6e-13	-1.51
3	Genetic algorithm	+0.28	[+0.04, +0.52]	0.024	+0.33
3	Binary PSO	+0.75	[+0.51, +0.99]	4.2e-07	+0.88
3	Greedy (cost-unaware)	+1.64	[+0.74, +2.53]	0.001	+0.52
3	k-means clustering	+4.52	[+3.34, +5.69]	2.5e-09	+1.09
3	Uniform grid	+4.24	[+3.20, +5.29]	6.6e-10	+1.15
3	Random	+13.01	[+11.25, +14.77]	6.9e-19	+2.10
6	Proposed (lazy)	+0.65	[+0.50, +0.80]	8.7e-11	+1.23
6	Simulated annealing	-1.30	[-1.51, -1.09]	4.2e-16	-1.77
6	Genetic algorithm	+0.53	[+0.27, +0.79]	2.9e-04	+0.58
6	Binary PSO	+0.87	[+0.63, +1.12]	1.2e-08	+1.01
6	Greedy (cost-unaware)	+1.12	[+0.50, +1.74]	6.3e-04	+0.52
6	k-means clustering	+4.30	[+3.25, +5.36]	3.5e-10	+1.16
6	Uniform grid	+3.63	[+3.07, +4.19]	9.9e-17	+1.85
6	Random	+13.49	[+11.85, +15.12]	8.2e-21	+2.34

All three cities reproduce the pattern of the synthetic study under matched expenditure. The proposed method exceeds cost-unaware greedy selection by between 1.12 and 4.12 percentage points, the snapped lattice by between 3.63 and 7.45, and random selection by between 11.01 and 17.98, in every case with confidence intervals excluding zero after correction. Against the genetic algorithm and particle swarm optimisation the advantage is smaller and depends on the city. It reaches 2.17 and 2.64 percentage points respectively in Zürich, the densest inventory, and falls to values indistinguishable from zero in Basel at the 0 dB threshold, where the genetic algorithm is in fact 0.09 points ahead and particle swarm optimisation 0.17 points behind, neither significantly. Simulated annealing stays ahead of the proposed method by between 0.21 and 1.45 percentage points, though the 0.21-point margin in Zürich at 6 dB is not significant after correction.

The unfavourable change is larger and runs the other way. Under a common aggregate budget the proposed method trails cost-unaware greedy selection in all three cities at all three thresholds, by 0.51, 0.84 and 1.09 percentage points in Zürich, by 1.09, 1.54 and 1.88 in Bern, and by 0.76, 0.55 and 0.31 in Basel. Every one of those differences survives Holm correction, and the 1.88-point deficit in Bern at 6 dB is the largest

margin against the proposed method anywhere in this paper. The synthetic study gave a milder version of the same result, a deficit at 0 dB only, so the real inventories sharpen it rather than contradict it. The reading we take from the two regimes together is that cost-aware selection earns its advantage when capital is the binding constraint and loses it when the budget is loose enough that buying more sites is simply better. Table 20 reports both regimes so that the comparison can be made directly.

The gradient across cities is the most informative feature of this experiment. Cost-aware selection gains most over its cost-unaware counterpart in Zürich, at 3.19 to 4.12 percentage points, and less in Bern and Basel, at 1.59 to 1.77 and 1.12 to 1.64 respectively; averaged over the three thresholds the margins are 3.79, 1.70 and 1.31 percentage points. The gradient is clear between Zürich and the other two, but not between Bern and Basel, whose ranges overlap and whose order reverses at 3 dB. The plausible mechanism is that the value of discriminating between cheap and expensive assets grows with the number of assets available to discriminate among, so a sparse inventory gives the ranking rule less room to express a preference.

Under a common aggregate budget the ordering changes in two ways, one favourable to the proposed method and one not. The favourable one is that in Zürich and Bern the proposed method now exceeds simulated annealing, by 1.73 to 2.37 and 0.59 to 0.94 percentage points respectively, although in Basel it remains behind by 1.02 to 1.41. The candidate count alone does not explain this: Bern holds 98 candidates, fewer than the 100 of the synthetic study where the annealer wins, so the difference lies in the spatial structure of a real inventory rather than in its size, and a fixed budget of 3 000 objective evaluations buys less on a clustered candidate set than on a uniform one. This is a statement about the metaheuristics at this evaluation budget rather than about the problem, and Section 7.3 lists it as a limitation.

Budget utilisation spans a far wider range on the real inventories than on the synthetic candidate sets, where the proposed method spent 40.6% of its budget. Table 21 reports the accounting. In Zürich the proposed method spends 28.7 of an available 168.9 normalised units, that is 17.0% of the budget, and selects 12.4 sites where random selection takes 31.9; the corresponding figures are 24.7% in Bern and 52.1% in Basel. The denser the inventory, the larger the share of the budget the algorithm declines to spend: 17.0%, 24.7% and 52.1% spent, against inventories of 154, 98 and 57 candidates. This is a statement about budget share and not about site counts, which do not follow the same order (the densest tile buys the most sites, 12.4 against 10.1 in the sparsest). Runtimes ran between 0.04 and 0.14 s for the proposed method against 1.21 to 2.84 s for the three metaheuristics.

Table 21 is broken down by city and SINR threshold, and its columns match Table 8. Entries are means over 50 paired iterations; costs are in normalised units, and runtimes are single-threaded means on the reference platform of Table 3.

Table 21. Resource accounting on the real inventories under a common aggregate budget

Method	Sites	Spent	Unused	Util. (%)	Cost range	Evals	Runtime (s)
<i>Zürich, 0 dB (budget 168.9)</i>							
Proposed (BA-greedy)	12.2	26.6	142.3	15.7	1.0–6.6	2110	0.132
Proposed (lazy)	10.8	25.9	143.0	15.3	1.0–7.6	645	0.036
Simulated annealing	27.8	160.2	8.7	94.9	1.1–10.0	3001	2.432
Genetic algorithm	31.8	168.5	0.4	99.8	1.0–10.0	2942	2.843
Binary PSO	28.5	168.5	0.4	99.8	1.1–10.0	3001	2.582
Greedy (cost-unaware)	11.8	64.9	104.0	38.4	1.4–9.5	1896	0.124
k-means clustering	29.2	160.7	8.2	95.2	1.0–9.9	1	1.359
Uniform grid	30.9	167.4	1.5	99.1	1.1–10.0	1	0.001
Random	31.9	168.5	0.4	99.8	1.0–9.9	1	0.000
<i>Zürich, 3 dB (budget 168.9)</i>							
Proposed (BA-greedy)	12.2	27.8	141.1	16.5	1.0–6.4	2107	0.135
Proposed (lazy)	9.9	25.4	143.5	15.0	1.0–7.4	583	0.033
Simulated annealing	27.1	156.2	12.7	92.5	1.2–10.0	3001	2.401
Genetic algorithm	30.1	168.5	0.4	99.8	1.0–10.0	2942	2.799
Binary PSO	28.0	168.5	0.4	99.8	1.2–10.0	3001	2.584
Greedy (cost-unaware)	12.1	63.3	105.6	37.5	1.3–9.5	1936	0.127
k-means clustering	29.2	160.7	8.2	95.2	1.0–9.9	1	1.371
Uniform grid	30.9	167.4	1.5	99.1	1.1–10.0	1	0.001
Random	31.9	168.5	0.4	99.8	1.0–9.9	1	0.000
<i>Zürich, 6 dB (budget 168.9)</i>							
Proposed (BA-greedy)	12.8	31.6	137.3	18.7	1.0–7.3	2188	0.136

Method	Sites	Spent	Unused	Util. (%)	Cost range	Evals	Runtime (s)
Proposed (lazy)	9.7	24.6	144.3	14.6	1.0–7.1	555	0.031
Simulated annealing	27.3	154.5	14.4	91.5	1.1–9.9	3001	2.344
Genetic algorithm	29.7	168.5	0.4	99.8	1.0–10.0	2942	2.752
Binary PSO	27.7	168.5	0.4	99.8	1.1–10.0	3001	2.527
Greedy (cost-unaware)	12.5	69.0	99.9	40.9	1.2–9.7	1992	0.130
k-means clustering	29.2	160.7	8.2	95.2	1.0–9.9	1	1.374
Uniform grid	30.9	167.4	1.5	99.1	1.1–10.0	1	0.001
Random	31.9	168.5	0.4	99.8	1.0–9.9	1	0.000
Bern, 0 dB (budget 110.0)							
Proposed (BA-greedy)	10.0	30.2	79.8	27.5	1.0–7.9	1124	0.065
Proposed (lazy)	8.6	24.1	86.0	21.9	1.0–7.7	406	0.022
Simulated annealing	16.1	96.3	13.8	87.5	1.3–9.9	3001	1.695
Genetic algorithm	19.5	109.7	0.3	99.7	1.0–10.0	2942	2.039
Binary PSO	17.3	109.7	0.3	99.7	1.3–10.0	3001	1.741
Greedy (cost-unaware)	8.7	51.9	58.1	47.2	1.8–9.6	910	0.057
k-means clustering	18.2	102.6	7.4	93.2	1.1–9.8	1	0.911
Uniform grid	19.4	108.6	1.5	98.7	1.2–9.9	1	0.001
Random	20.8	109.7	0.3	99.7	1.1–9.9	1	0.000
Bern, 3 dB (budget 110.0)							
Proposed (BA-greedy)	8.9	26.4	83.6	24.0	1.0–7.9	1024	0.061
Proposed (lazy)	7.6	20.5	89.6	18.6	1.0–7.4	368	0.020
Simulated annealing	16.2	95.7	14.4	86.9	1.3–9.8	3001	1.669
Genetic algorithm	19.5	109.7	0.3	99.7	1.1–9.9	2942	2.036
Binary PSO	17.9	109.7	0.3	99.7	1.4–9.9	3001	1.764
Greedy (cost-unaware)	8.5	50.0	60.1	45.4	1.7–9.4	895	0.058
k-means clustering	18.2	102.6	7.4	93.2	1.1–9.8	1	0.902
Uniform grid	19.4	108.6	1.5	98.7	1.2–9.9	1	0.001
Random	20.8	109.7	0.3	99.7	1.1–9.9	1	0.000
Bern, 6 dB (budget 110.0)							
Proposed (BA-greedy)	8.6	25.0	85.0	22.7	1.0–7.6	995	0.057
Proposed (lazy)	7.4	21.0	89.1	19.0	1.0–7.3	352	0.019
Simulated annealing	17.0	99.3	10.7	90.2	1.4–9.9	3001	1.698
Genetic algorithm	19.5	109.7	0.3	99.7	1.0–10.0	2942	2.036
Binary PSO	18.0	109.7	0.3	99.7	1.4–10.0	3001	1.757
Greedy (cost-unaware)	8.8	51.5	58.6	46.8	1.7–9.5	918	0.058
k-means clustering	18.2	102.6	7.4	93.2	1.1–9.8	1	0.915
Uniform grid	19.4	108.6	1.5	98.7	1.2–9.9	1	0.001
Random	20.8	109.7	0.3	99.7	1.1–9.9	1	0.000
Basel, 0 dB (budget 64.2)							
Proposed (BA-greedy)	10.1	32.7	31.5	50.9	1.0–8.4	636	0.039
Proposed (lazy)	9.3	28.7	35.5	44.7	1.0–8.1	248	0.015
Simulated annealing	9.7	55.0	9.2	85.7	1.6–9.5	3001	1.275
Genetic algorithm	11.4	63.7	0.5	99.2	1.3–9.7	2942	1.578
Binary PSO	10.3	63.7	0.5	99.2	1.8–9.9	3001	1.230
Greedy (cost-unaware)	7.8	44.8	19.4	69.8	2.0–9.3	463	0.033
k-means clustering	10.7	58.4	5.8	91.0	1.4–9.5	1	0.595
Uniform grid	11.8	62.6	1.6	97.6	1.2–9.6	1	0.000
Random	12.6	63.7	0.5	99.2	1.2–9.6	1	0.000
Basel, 3 dB (budget 64.2)							
Proposed (BA-greedy)	10.0	33.6	30.6	52.4	1.0–8.4	628	0.038
Proposed (lazy)	8.7	29.6	34.6	46.0	1.0–8.3	224	0.013
Simulated annealing	9.7	56.1	8.1	87.4	1.7–9.6	3001	1.209
Genetic algorithm	11.6	63.7	0.5	99.2	1.1–9.8	2942	1.534
Binary PSO	10.5	63.7	0.5	99.2	1.6–9.8	3001	1.232
Greedy (cost-unaware)	8.4	49.4	14.8	77.0	1.7–9.6	488	0.032
k-means clustering	10.7	58.4	5.8	91.0	1.4–9.5	1	0.566
Uniform grid	11.8	62.6	1.6	97.6	1.2–9.6	1	0.000
Random	12.6	63.7	0.5	99.2	1.2–9.6	1	0.000
Basel, 6 dB (budget 64.2)							
Proposed (BA-greedy)	10.1	34.0	30.2	53.0	1.0–8.5	634	0.041
Proposed (lazy)	8.6	30.0	34.2	46.8	1.0–8.4	211	0.013
Simulated annealing	10.2	56.2	8.0	87.5	1.4–9.5	3001	1.239

Method	Sites	Spent	Unused	Util. (%)	Cost range	Evals	Runtime (s)
Genetic algorithm	11.9	63.7	0.5	99.2	1.1–9.8	2942	1.589
Binary PSO	10.5	63.7	0.5	99.2	1.6–9.9	3001	1.288
Greedy (cost-unaware)	8.1	47.2	17.0	73.6	1.9–9.5	470	0.031
k-means clustering	10.7	58.4	5.8	91.0	1.4–9.5	1	0.569
Uniform grid	11.8	62.6	1.6	97.6	1.2–9.6	1	0.000
Random	12.6	63.7	0.5	99.2	1.2–9.6	1	0.000

The three cities support the conclusions drawn from the synthetic study under matched expenditure, and suggest one the synthetic study could not reach: the benefit of cost-aware selection appears to grow with the density of the candidate inventory. Three cities cannot establish a scaling law, and the ordering between the two sparser inventories is not clean, so this is offered as a direction for a larger study rather than as a result.

7. Discussion

7.1. Interpretation

Four findings emerge. The first concerns the structure of the problem rather than any algorithm. Once interference enters the model, coverage stops being a monotone submodular function of the selected set, and it stops by a wide margin rather than marginally. Roughly half of all site additions reduce coverage somewhere. Any treatment of small-cell placement that carries over the classical coverage guarantee without examining this stands on unsound ground, and the earlier version of this paper was among them.

The second is that losing the guarantee costs less in practice than in theory. The greedy rule keeps working well, and not because the guarantee secretly survives. Table 4 shows the opposite of an approximately submodular objective: coverage falls on about half of all site additions and the diminishing-returns inequality is violated on well over half of nested pairs. What saves the greedy rule is that it never has to rely on those properties. Line 5 of Algorithm 1 evaluates the true marginal gain of every candidate at every step, so an addition that would reduce coverage is simply not made, and the violations that Table 4 counts are ones the algorithm declines to walk into. A guarantee is a worst-case statement, and this problem is not adversarial.

The third is that cost-awareness rather than greediness is the operative ingredient. Under matched expenditure the proposed method beats cost-unaware greedy selection by 1.5 to 2.3 percentage points, and with a fraction of the variance. Under an equal nominal budget the two are statistically indistinguishable at 3 and 6 dB, and the cost-unaware rule sits 0.32 percentage points ahead at 0 dB. That is what one would expect if the cost-aware rule buys comparable coverage more cheaply and then declines to spend the remainder.

The fourth concerns metaheuristics. Simulated annealing is the strongest method tested on the synthetic instances, ahead at every threshold under both regimes, if slightly. It is not uniformly strongest: on the two denser real inventories the proposed method overtakes it under a common budget, which Section 6.11 attributes to the fixed evaluation budget rather than to the method. It is also about twenty-three times slower under a common budget and thirteen times slower under matched expenditure, and its advantage shrinks as the SINR threshold tightens until, at 6 dB under matched expenditure, it disappears into the noise. A planner who must evaluate many candidate budgets and many city districts is well served by a rule that runs in about a tenth of a second and lands within 0.55 percentage points of a reasonably configured annealer under matched expenditure. For a single final design, the annealer is worth running.

7.2. Practical Implications

A small number of concrete recommendations follow, each tied to a result rather than to general reasoning.

- Budget exhaustion is not a design objective. On the synthetic instances at 3.5 GHz the cost-aware rule stopped spending at between a third and just over a half of the budget, and the methods that spent the remainder gained at most 1.2 percentage points of coverage for it. The margin is not universal: on the real inventories the best budget-exhausting method gains up to 1.88 points, and at 28 and 39 GHz, or under light

load, the rule spends almost everything. A deployment plan that consumes its entire allocation deserves a check for sites contributing more interference than signal.

- Record cost heterogeneity per asset rather than averaging it. Under matched expenditure the only difference between the cost-aware rule and its cost-unaware counterpart is that the first divides marginal gain by price; an inventory that carries a single average cost per pole cannot support it.

- Load assumptions must accompany any coverage figure. Coverage at 6 dB was 56% under full load and 84% at quarter load. The two are not the same deployment: at quarter load the algorithm buys 29.9 sites for 104.3 units against 17.3 sites for 58.9 units at full load, so the load assumption changes both what is bought and what it achieves. Planning targets quoted without a load assumption do not compare across studies or across vendors.

- Apply a gain-per-cost rule with care in demand-skewed areas. Section 6.7 shows it slightly underperforms an absolute-gain rule when demand is concentrated, because it is reluctant to buy expensive coverage of high-value locations.

The cost-reduction figures of 20-30% sometimes quoted for urban-furniture deployment come from industry analyses of site acquisition and permitting^{6,7}, not from the simulations presented here. This paper models relative cost, not money, and no monetary saving can be inferred from its results.

7.3. Limitations

The limitations below bound what may be concluded. They are stated as constraints on the claims rather than as future work, because extending the present study would not remove several of them.

- Cost model. Per-site costs are normalised indices drawn from a uniform distribution, with no empirical basis. Section 3.2 shows the selection to be invariant to any rescaling of costs and budget, so the index is a statement about cost ratios and an operator can substitute their own currency figures. What the paper does not supply is the ratios themselves for any real inventory, and a real cost spread would not be uniform. Nothing here supports a claim about absolute deployment economics.

- No physical feasibility data. The model covers none of structural capacity, power availability, backhaul reach, antenna mounting constraints or planning consent, including for the three real inventories of Section 6.11. A deployable candidate set would be smaller than the ones used here, and non-uniformly so.

- Simplified radio model. Omnidirectional antennas, no down-tilt, no beamforming, no power control, no scheduling, no coordinated muting, no macro-cell tier, no building penetration and no small-scale fading. Each of these would change absolute coverage; several would change it substantially.

- Static demand and fixed transmit power. The model contains no users; demand enters only as a static weight on test points, and mobility, temporal variation and admission control are absent. Section 6.10 shows load-aware association leaving coverage essentially unchanged, which follows from every selected cell transmitting at full power whatever the association. A model with cell sleep or power control under light load would behave differently.

- Frequency coverage. Results run at 3.5, 28 and 39 GHz. The higher bands use the same 3GPP UMi model extrapolated to those frequencies, without the blockage, beamforming and antenna-array behaviour that a millimetre-wave design would need. Section 6.9 should be read as a scaling trend, not as a millimetre-wave system study.

- Simulation only. No measurement campaign was conducted, and the real-location study runs simulated propagation over real coordinates. The consistency check of Section 6.1 validates the implementation against the specification, not the specification against reality.

- Geographical scope. The three real inventories are all central tiles of Swiss cities, and share a European street pattern, a similar building stock and one national mapping culture. OpenStreetMap completeness also varies between them, so a difference in inventory density is partly a difference in

⁶ CTIA and Accenture Strategy, "Impact of Federal Regulatory Reviews on Small Cell Deployment", Washington, DC, USA, 2018, Available: https://api.ctia.org/docs/default-source/default-document-library/small-cell-deployment-regulatory-review-costs_3-12-2018.pdf.

⁷ National League of Cities, "Small Cell Wireless Technology in Cities: A Municipal Action Guide", Washington, DC, USA, 2018, Available: <https://www.nlc.org/resource/small-cell-wireless-technology-in-cities/>.

mapping effort rather than in physical asset density. Cities with different morphology, and regions with different mapping coverage, are untested.

- Candidate-set scale. Between 57 and 154 candidates over 1 km². The quadratic term in the complexity analysis has not been stressed, and behaviour at metropolitan scale is untested.
- Fixed evaluation budget for the metaheuristics. Section 6.11 shows that 3 000 objective evaluations become insufficient for simulated annealing on the real inventories, including Bern's 98 candidates, which is fewer than the synthetic study's 100. What binds is the structure of the candidate set rather than its size, and the boundary was not located. Where the proposed method exceeds a metaheuristic on the larger inventories, that is evidence about the metaheuristic at this budget and not about the ceiling of the metaheuristic.
- Metaheuristic tuning. The three metaheuristics use standard parameter settings and a fixed evaluation budget. A more heavily tuned implementation, or a larger budget, would likely improve them; the comparison should be read as against reasonable rather than optimal configurations of those methods.

8. Conclusions

This paper examined how to select urban-furniture assets as hosts for 5G small cells under a fixed capital budget, using a fully specified interference-aware coverage model and a comparison against seven alternative placement methods.

The central technical finding is negative, and the earlier version of this work did not anticipate it: with interference in the model, the coverage objective is neither monotone nor submodular, and the classical $(1 - 1/e)$ approximation guarantee does not apply. The violation is common rather than pathological, turning up on roughly half of all site additions tested. The guarantee survives here only for an explicitly noise-limited surrogate objective, and the proposed algorithm is presented as a heuristic.

The central empirical finding is that a budget-aware greedy rule remains a strong practical choice despite the loss of the guarantee. Under matched expenditure it exceeded cost-unaware greedy selection, a genetic algorithm, binary particle swarm optimisation, lattice placement and random placement at every SINR threshold tested, with confidence intervals excluding zero after correction for multiple comparisons; at the most demanding threshold it was statistically indistinguishable from simulated annealing on about one thirteenth of the computation. Under a common aggregate budget it came within 1.19 percentage points of the best method on a third to a half of the available capital. Real urban-furniture inventories for three Swiss city centres reproduced that ordering under matched expenditure, with one exception at the least demanding threshold on the sparsest inventory. Under a common aggregate budget they did not: cost-unaware greedy selection was ahead in all three cities at all three thresholds. Those experiments also added a finding the synthetic study could not reach: under matched expenditure, the advantage of cost-aware selection over cost-unaware selection grows with the density of the candidate inventory, from 1.1 percentage points on the sparsest tile to 4.1 on the densest.

Two further results qualify how these figures should be read. Repeating the study at 28 and 39 GHz shows the interference-driven stopping point disappearing as cells shrink: at 3.5 GHz the algorithm halts with between 47% and 67% of the budget unspent, whereas at the higher bands it spends about 107 of 110.5 normalised units. Reading coverage as a probability over redrawn shadow fading returns mean values within one percentage point of the single-realisation figures, but the area meeting the threshold with probability at least 0.95 is far smaller, 40.05% at 0 dB and 24.18% at 6 dB, so a threshold metric on one realisation describes the average outcome well and the dependable outcome poorly.

The limitations set out in Section 7.3 bound these conclusions. They concern a simulated, single-tier, omnidirectional-antenna model evaluated at three carrier frequencies with normalised costs, evaluated over 1 km² tiles in three Swiss city centres and on synthetic candidate sets. They do not establish deployment cost savings, they do not constitute field validation, and they should not be read as general claims about small-cell economics.

Three extensions follow directly. The first is measurement: the inventories of Section 6.11 are real coordinates, so even a modest drive-test campaign over one of them would establish how far the simulated ordering of methods survives contact with real propagation. The second is a cost model grounded in operator or municipal data, which would turn the relative comparisons here into statements about money.

The third is demand-aware selection, since Section 6.7 shows the gain-per-cost rule is weakest exactly where demand is most concentrated, at the locations that matter most.

CRediT Author Contribution Statement

Mohammed Sh. Ali: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Visualization, Writing – original draft; Mohammed Othman Hasan: Validation, Investigation, Writing – review & editing; Rana H. A. Zubo: Conceptualization, Supervision, Validation, Project administration, Writing – review & editing; Raed A. Abd-Alhameed: Methodology, Supervision, Writing – review & editing.

Acknowledgement

This work is partially supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/Y035135/1, and HORIZON-MSCA-2022-SE-01-01-ID: 101131501, Marie Skłodowska-Curie, Research and Innovation Staff Exchange (RISE), titled: 6G Terahertz Communications for Future Heterogeneous Wireless Network (6G-TERAFIT).

Data Availability Statement

An archive accompanying this article releases the simulator, experiment driver, statistical analysis code, figure scripts, OpenStreetMap extraction pipeline, master seed and full run metadata. OpenCellID data are openly available⁸. OpenStreetMap data are (c) OpenStreetMap contributors and available under the Open Database Licence⁹.

Conflict of Interest

The authors declare no conflict of interest.

References

- [1] Mamta Agiwal, Abhishek Roy and Navrati Saxena, "Next Generation 5G Wireless Networks: A Comprehensive Survey", *IEEE Communications Surveys & Tutorials*, Online ISSN: 1553-877X, 19 February 2016, Vol. 18, No. 3, pp. 1617-1655, Published by IEEE, DOI: 10.1109/COMST.2016.2532458, Available: <https://ieeexplore.ieee.org/document/7414384>.
- [2] Jeffrey G. Andrews, Stefano Buzzi, Wan Choi, Stephen V. Hanly, Angel Lozano *et al.*, "What Will 5G Be?", *IEEE Journal on Selected Areas in Communications*, Print ISSN: 0733-8716, Online ISSN: 1558-0008, June 2014, Vol. 32, No. 6, pp. 1065-1082, Published by IEEE, DOI: 10.1109/JSAC.2014.2328098, Available: <https://ieeexplore.ieee.org/document/6824752>.
- [3] Andrea Zanella, Nicola Bui, Angelo Castellani, Lorenzo Vangelista and Michele Zorzi, "Internet of Things for Smart Cities", *IEEE Internet of Things Journal*, Online ISSN: 2327-4662, February 2014, Vol. 1, No. 1, pp. 22-32, Published by IEEE, DOI: 10.1109/JIOT.2014.2306328, Available: <https://ieeexplore.ieee.org/document/6740844>.
- [4] Hafedh Chourabi, Taewoo Nam, Shawn Walker, J. Ramon Gil-Garcia, Sehl Mellouli *et al.*, "Understanding Smart Cities: An Integrative Framework", in *Proceedings of the 45th Hawaii International Conference on System Sciences (HICSS)*, 4-7 January 2012, Maui, Hawaii, USA, Print ISBN: 978-1-4577-1925-7, Online ISBN: 978-0-7695-4525-7, pp. 2289-2297, Published by IEEE, DOI: 10.1109/HICSS.2012.615, Available: <https://ieeexplore.ieee.org/document/6149291>.
- [5] Sayeed Salih, Abdelzahir Abdelmaboud, Omayma Husain, Abdelwahed Motwakel, Hashim Elshafie *et al.*, "IoT in Urban Development: Insight into Smart City Applications, Case Studies, Challenges, and Future Prospects", *PeerJ Computer Science*, Online ISSN: 2376-5992, 29 April 2025, Vol. 11, Article No. e2816, Published by PeerJ, DOI: 10.7717/peerj-cs.2816, Available: <https://peerj.com/articles/cs-2816/>.
- [6] Yanyu Zhang and Osama Alshaykh, "5G Utility Pole Planner Using Google Street View and Mask R-CNN", in *Proceedings of the 2020 IEEE International Conference on Electro Information Technology (EIT)*, 31 July-1 August 2020, Chicago, Illinois, USA, Print ISBN: 978-1-7281-5318-6, Online ISBN: 978-1-7281-5317-9, pp. 309-312, Published by IEEE, DOI: 10.1109/EIT48999.2020.9208333, Available: <https://ieeexplore.ieee.org/document/9208333>.

⁸ OpenCellID: <https://opencellid.org>.

⁹ OpenStreetMap: <https://www.openstreetmap.org>.

- [7] Jaume Benseny, Jaspreet Singh Walia, Heikki Hämmäinen and Juha Salmelin, "City Strategies for a 5G Small Cell Network on Light Poles", in *Proceedings of the 2019 CTTE-FITCE: Smart Cities & Information and Communication Technology (CTTE-FITCE)*, 25-27 September 2019, Ghent, Belgium, Print ISBN: 978-1-7281-2122-2, Online ISBN: 978-1-7281-2121-5, pp. 1-6, Published by IEEE, DOI: 10.1109/CTTE-FITCE.2019.8894825, Available: <https://ieeexplore.ieee.org/document/8894825>.
- [8] Arijeet Ghosh and Iti Saha Misra, "A Method of Dynamic Small Cell Deployment in 5G Two-Tier HetNet for Improved Area Spectral Efficiency", in *Proceedings of the 2023 8th International Conference on Computers and Devices for Communication (CODEC)*, 14-16 December 2023, Kolkata, India, Print ISBN: 979-8-3503-1718-3, Online ISBN: 979-8-3503-1717-6, pp. 1-2, Published by IEEE, DOI: 10.1109/CODEC60112.2023.10466153, Available: <https://ieeexplore.ieee.org/document/10466153>.
- [9] James Meese, Kieran Hegarty, Rowan Wilken, Fan Yang and Catherine Middleton, "5G and Urban Amenity: Regulatory Trends and Local Government Responses around Small Cell Deployment", *Digital Policy, Regulation and Governance*, Print ISSN: 2398-5038, Online ISSN: 2398-5046, 2 July 2024, Vol. 26, No. 6, pp. 601-616, Published by Emerald, DOI: 10.1108/DPRG-10-2023-0150, Available: <https://www.emerald.com/dprg/article/26/6/601-616/1229983>.
- [10] Harpreet S. Dhillon, Radha Krishna Ganti, Francois Baccelli and Jeffrey G. Andrews, "Modeling and Analysis of K-Tier Downlink Heterogeneous Cellular Networks", *IEEE Journal on Selected Areas in Communications*, Print ISSN: 0733-8716, Online ISSN: 1558-0008, April 2012, Vol. 30, No. 3, pp. 550-560, Published by IEEE, DOI: 10.1109/JSAC.2012.120405, Available: <https://ieeexplore.ieee.org/document/6171996>.
- [11] Mahmoud Kamel, Walaa Hamouda and Amr Youssef, "Ultra-Dense Networks: A Survey", *IEEE Communications Surveys & Tutorials*, Online ISSN: 1553-877X, 23 May 2016, Vol. 18, No. 4, pp. 2522-2545, Published by IEEE, DOI: 10.1109/COMST.2016.2571730, Available: <https://ieeexplore.ieee.org/document/7476821>.
- [12] Xiaohu Ge, Song Tu, Guoqiang Mao, Cheng-Xiang Wang and Tao Han, "5G Ultra-Dense Cellular Networks", *IEEE Wireless Communications*, Print ISSN: 1536-1284, Online ISSN: 1558-0687, February 2016, Vol. 23, No. 1, pp. 72-79, Published by IEEE, DOI: 10.1109/MWC.2016.7422408, Available: <https://ieeexplore.ieee.org/document/7422408>.
- [13] George L. Nemhauser, Laurence A. Wolsey and Marshall L. Fisher, "An Analysis of Approximations for Maximizing Submodular Set Functions - I", *Mathematical Programming*, Print ISSN: 0025-5610, Online ISSN: 1436-4646, December 1978, Vol. 14, No. 1, pp. 265-294, Published by Springer, DOI: 10.1007/BF01588971, Available: <https://link.springer.com/article/10.1007/BF01588971>.
- [14] Samir Khuller, Anna Moss and Joseph (Seffi) Naor, "The Budgeted Maximum Coverage Problem", *Information Processing Letters*, Print ISSN: 0020-0190, April 1999, Vol. 70, No. 1, pp. 39-45, Published by Elsevier, DOI: 10.1016/S0020-0190(99)00031-9, Available: <https://www.sciencedirect.com/science/article/abs/pii/S0020019099000319>.
- [15] Michel Minoux, "Accelerated Greedy Algorithms for Maximizing Submodular Set Functions", in *Optimization Techniques II: Proceedings of the 8th IFIP Conference on Optimization Techniques, Lecture Notes in Control and Information Sciences (LNCIS)*, 5-9 September 1977, Würzburg, Germany, Print ISBN: 978-3-540-08708-3, Online ISBN: 978-3-540-35890-9, Series Print ISSN: 0170-8643, Vol. 7, pp. 234-243, Published by Springer, DOI: 10.1007/BFb0006528, Available: <https://link.springer.com/chapter/10.1007/BFb0006528>.
- [16] V. B. Nikam, Anuj Arora, Deeplaxmi Lambture, Jash Zaveri, Prathamesh Shinde *et al.*, "Optimal Positioning of Small Cells for Coverage and Cost Efficient 5G Network Deployment: A Smart Simulated Annealing Approach", in *Proceedings of the 2020 IEEE 3rd 5G World Forum (5GWF)*, 10-12 September 2020, Bangalore, India, Print ISBN: 978-1-7281-7300-9, Online ISBN: 978-1-7281-7299-6, pp. 454-459, Published by IEEE, DOI: 10.1109/5GWF49715.2020.9221257, Available: <https://ieeexplore.ieee.org/document/9221257>.
- [17] Zhaomin Li, "Research on 5G Base Station Location Problem Based on Simulated Annealing Algorithm", in *Proceedings of the 7th International Conference on Intelligent Information Processing (ICIIP '22)*, 29-30 September 2022, Bucharest, Romania, pp. 1-5, Published by ACM, DOI: 10.1145/3570236.3570246, Available: <https://dl.acm.org/doi/10.1145/3570236.3570246>.
- [18] Kuna Venkateswararao and Pravati Swain, "Binary-PSO-Based Energy-Efficient Small Cell Deployment in 5G Ultra-Dense Network", *The Journal of Supercomputing*, Print ISSN: 0920-8542, Online ISSN: 1573-0484, 4 June 2021, Vol. 78, No. 1, pp. 1071-1092, Published by Springer, DOI: 10.1007/s11227-021-03910-5, Available: <https://link.springer.com/article/10.1007/s11227-021-03910-5>.
- [19] Yao-Hsin Chou, Cheng-Yen Hua, Ru-Wei Tseng and Shu-Yu Kuo, "Quantum-Inspired Optimization Algorithm for 3D Multi-Objective Base-Station Deployment in Next-Generation 5G/6G Wireless Network", *Computers, Materials & Continua*, Online ISSN: 1546-2226, 12 March 2026, Vol. 87, No. 2, pp. 1-10, Published by Tech Science Press, DOI: 10.32604/cmc.2025.075705, Available: <https://www.techscience.com/cmc/v87n2/66632>.
- [20] Weidong Mei and Rui Zhang, "Joint Base Station and IRS Deployment for Enhancing Network Coverage: A Graph-Based Modeling and Optimization Approach", *IEEE Transactions on Wireless Communications*, Print ISSN: 1536-1276, Online ISSN: 1558-2248, November 2023, Vol. 22, No. 11, pp. 8200-8213, Published by IEEE, DOI: 10.1109/TWC.2023.3260805, Available: <https://ieeexplore.ieee.org/document/10086045>.

- [21] Eleni Aloupogianni, Charalampos Karyotis, Tomasz Maniak, Rahat Iqbal, Nikos Passas *et al.*, “AI-Driven Optimization of Small Cell Deployment for Beyond 5G Networks”, *Procedia Computer Science*, Online ISSN: 1877-0509, 8 July 2024, Vol. 238, pp. 908-913, Published by Elsevier, DOI: 10.1016/j.procs.2024.06.111, Available: <https://www.sciencedirect.com/science/article/pii/S1877050924013474>.
- [22] Flávio Henry Ferreira, Fabrício José Brito Barros, Miércio Cardoso de Alcântara Neto, Evelin Cardoso, Carlos Renato Lisboa Francês *et al.*, “Hybrid Computational and Real Data-Based Positioning of Small Cells in 5G Networks”, *PeerJ Computer Science*, Online ISSN: 2376-5992, 26 June 2023, Vol. 9, Article No. e1412, Published by PeerJ, DOI: 10.7717/peerj-cs.1412, Available: <https://peerj.com/articles/cs-1412/>.
- [23] Mobasshir Mahbub and Bobby Barua, “The Measurements for Optimal Small Cell Base Station Positioning Considering 5G Cellular Network”, in *Proceedings of the 2021 International Conference on Information Technology (ICIT)*, 14-15 July 2021, Amman, Jordan, Print ISBN: 978-1-6654-2871-2, Online ISBN: 978-1-6654-2870-5, pp. 280-283, Published by IEEE, DOI: 10.1109/ICIT52682.2021.9491710, Available: <https://ieeexplore.ieee.org/document/9491710>.
- [24] Irene Vilà, Oriol Sallent and Jordi Pérez-Romero, “On the Design of a Network Digital Twin for the Radio Access Network in 5G and Beyond”, *Sensors*, Online ISSN: 1424-8220, 20 January 2023, Vol. 23, No. 3, Article No. 1197, Published by MDPI, DOI: 10.3390/s23031197, Available: <https://www.mdpi.com/1424-8220/23/3/1197>.
- [25] Fotis Foukalas, “Federated-Learning-Driven Radio Access Networks”, *IEEE Wireless Communications*, Print ISSN: 1536-1284, Online ISSN: 1558-0687, August 2022, Vol. 29, No. 4, pp. 48-54, Published by IEEE, DOI: 10.1109/MWC.102.2100113, Available: <https://ieeexplore.ieee.org/document/9770092>.
- [26] Richard Church and Charles ReVelle, “The Maximal Covering Location Problem”, *Papers of the Regional Science Association*, Print ISSN: 1056-8190, Online ISSN: 1435-5957, December 1974, Vol. 32, No. 1, pp. 101-118, Published by Springer, DOI: 10.1007/BF01942293, Available: <https://link.springer.com/article/10.1007/BF01942293>.
- [27] Sture Holm, “A Simple Sequentially Rejective Multiple Test Procedure”, *Scandinavian Journal of Statistics*, Print ISSN: 0303-6898, 1979, Vol. 6, No. 2, pp. 65-70, Published by Wiley, Available: <https://www.jstor.org/stable/4615733>.



© 2026 by the author(s). Published by Annals of Emerging Technologies in Computing (AETiC), under the terms and conditions of the Creative Commons Attribution (CC BY) license which can be accessed at <http://creativecommons.org/licenses/by/4.0>.