TERRAIN-AWARE INFRASTRUCTURE PLANNING: AUTOMATED PLACEMENT OF RELIABLE GATEWAYS FOR LARGE-SCALE IOT NETWORKS

Simeon Monov, Denis Bechiragich

Abstract. Deploying large-scale Internet of Things (IoT) infrastructures in complex geographical environments presents significant challenges for maintaining reliable wireless connectivity. Terrain features such as elevation changes, obstacles, and line-of-sight limitations strongly affect signal propagation and coverage quality. This study investigates the optimal placement of LoRa gateways to ensure robust communication for distributed sensor networks operating in three-dimensional terrain. The problem is formulated as a Minimum k-Dominating Set [1] over a terrain-aware geometric graph, where edges represent feasible communication links derived from Digital Elevation Models (DEM) [2] and line-of-sight constraints. Our approach is to identify the smallest set of gateway locations providing k-redundant coverage. The resulting optimization problem is NP-hard, making exact solutions computationally challenging for large-scale instances. To address this complexity, the study compares exact optimization methods based on Integer Linear Programming with metaheuristic optimization approaches. The proposed methodology constructs a DEM-based connectivity graph and evaluates algorithm performance across different terrain scenarios, measuring execution time, scalability, and optimality gap relative to exact solutions. The results demonstrate the potential of terrain-aware optimization models for automated infrastructure planning in real-world IoT deployments.

Key Words: *IoT, infrastructure planning, terrain-aware optimization, minimum k-dominating set, Wireless Sensor Networks, gateway placement, Digital Elevation Models, LoRa.*

Introduction

Large-scale vending machine networks require reliable low-power communication for telemetry, stock reporting, status diagnostics, and remote monitoring. Operators often deploy machines across wide urban or regional areas, creating a need for scalable wireless infrastructure. Among available technologies, LoRa is a practical solution due to its long communication range, low energy requirements, and cost-effective deployment model [3]. Such connectivity enables future capabilities such as predictive maintenance and better downtime prevention strategies.

Planning gateway infrastructure for vending machine networks is a challenging task. Many simplified approaches estimate coverage using only geographic distance. In real environments, terrain elevation, hills, and line-of-sight obstruction may significantly affect wireless communication quality. As a result, distance-only planning may lead to unreliable or overly optimistic deployment strategies.

At the same time, deploying an excessive number of gateways increases infrastructure costs, maintenance effort, internet connectivity requirements, and long-term operational expenses. For this reason, operators seek reliable coverage using the smallest practical number of gateways.

This paper addresses the problem through a terrain-aware optimization methodology based on Digital Elevation Models (DEM) [2] and graph optimization techniques. Exact, heuristic, and metaheuristic approaches are evaluated across deployment scenarios of increasing scale.

Problem Formulation

The gateway placement problem is modeled as a graph $G = (V, E)$, where each node represents a vending machine location, and each edge denotes a feasible wireless communication link between two locations. Connectivity between nodes is determined using terrain-aware line-of-sight analysis based on Digital Elevation Model (DEM) data.

The exact baseline is formulated as an Integer Linear Programming (ILP) model [4]. The ILP formulation introduces a binary decision variable $x_j \in \{0,1\}$, indicating whether node j is selected as a gateway. The objective is to minimize the total number of selected gateways:

$$\min \sum_{j \in V} x_j$$

subject to the coverage constraint:

$$\sum_{j \in V} a_{ij}\, x_j \geq k, \forall i \in V$$

where $a_{ij}$is the adjacency matrix of the graph. The value $a_{ij} = 1$ indicates that a feasible communication link exists between nodes i and j, while $a_{ij} = 0$ indicates that no such link exists. Self-coverage is included by setting the diagonal elements of the adjacency matrix to 1.

This constraint ensures that each vending machine location is covered by at least k gateways, following the idea of multiple domination requirements in graph-based coverage problems [5]. A node selected as a gateway can serve itself and all neighboring nodes connected through the graph. The redundancy criterion improves network reliability in the event of gateway unavailability or communication disruptions.

Under these assumptions, the problem can be formulated as a Minimum k-Dominating Set problem [1], closely related to k-tuple domination variants in graph theory [5], where the goal is to identify the smallest subset of nodes that satisfies the required coverage constraint for the entire network.

Methodology

The proposed methodology consists of four main components: scenario generation, terrain-aware connectivity modeling, gateway placement optimization, and performance evaluation.

In the first component, multiple independent test instances are generated for deployment scenarios of increasing scale. City-scale locations are sampled uniformly within the modeled urban area, whereas regional and national locations are clustered around populated centers. The evaluated scenarios include city-scale deployments in Plovdiv, regional-scale scenarios, and national-scale deployments across Bulgaria. For the city-scale scenarios, candidate instances are retained only when every node has at least two available gateway candidates, including self-coverage, thereby ensuring feasibility of the required redundancy constraint.

In the second component, Digital Elevation Model (DEM) data is used to assign an elevation value to each generated vending machine location. The connectivity graph is then constructed by testing each candidate pair of nodes against two conditions: maximum communication distance and terrain-based line-of-sight visibility. A pair of nodes is considered a feasible communication link only if the Euclidean distance between them is within the predefined transmission radius and the terrain does not obstruct the direct line between the two antenna positions.

The line-of-sight test is implemented as a geometric ray-casting procedure over the DEM surface. For each candidate link, the straight line between the endpoints is discretized using a scenario-specific number of samples. At each sample point, the terrain elevation is extracted from the DEM and compared with the interpolated height of the direct line between the two antennas. The antenna height is added to the terrain elevation of both endpoints before interpolation. If the terrain elevation at any intermediate point exceeds the interpolated line-of-sight height, the link is marked as blocked and no edge is added to the graph.

This model intentionally focuses on terrain obstruction rather than full radio-frequency propagation. Therefore, Fresnel zone clearance, diffraction, fading, link budget, and LoRa-specific physical-layer parameters are not modeled in the current implementation. The resulting adjacency matrix represents a simplified but terrain-aware connectivity graph, where edges correspond to candidate LoRa links that satisfy both distance and DEM-based visibility constraints.

In the third component, the gateway placement problem is solved using several optimization approaches, including Integer Linear Programming (ILP) [4], Greedy heuristic [6], Simulated Annealing (SA) [7], Variable Neighborhood Descent (VND) [8], Hybrid Genetic Algorithm (GA) [9], and Ant Colony Optimization (ACO) [10].

The Greedy heuristic is used as a fast constructive method for producing a feasible gateway placement. It starts with an empty set of selected gateways and maintains the current coverage level of each node. At each iteration, the algorithm selects the non-selected node that provides the largest immediate improvement, measured as the number of still under-covered nodes that would be covered by selecting this candidate. The process is repeated until every node satisfies the required k-coverage constraint.

After a feasible solution is constructed, a redundancy removal step is applied. Selected gateways are tested one by one and removed if the k-coverage constraint remains satisfied after their removal. This post-processing step reduces unnecessary gateways while preserving feasibility.

In this study, the Greedy heuristic serves two purposes. First, it provides a fast standalone baseline for comparison with the exact ILP reference and the metaheuristic methods. Second, it provides an initial feasible solution for the metaheuristic algorithms, which allows them to refine an already valid placement instead of starting from a random configuration.

In the final component, gateway count and runtime are used to compare solution quality and execution efficiency, with results summarized by the mean, sample standard deviation, and median across independent instances.

Experimental Setup and Results

The proposed approach was evaluated across four deployment scenarios of increasing scale: a city-scale scenario in Plovdiv with 100 vending machine locations (20 instances), a denser city-scale scenario in Plovdiv with 300 locations (20 instances), a regional-scale scenario with 1000 locations (5 instances), and a national-scale scenario across Bulgaria with 3000 locations (3 instances).

ILP served as the exact optimization approach and as the reference for evaluating solution quality. For the scenarios with 100, 300, and 1000 nodes, the solver confirmed that the obtained solutions were optimal. For the scenario with 3000 nodes, the solver reached the predefined time limit, and the best feasible ILP solution found within that limit was used for comparison. All scenarios used a 1500 m transmission radius, a 2 m antenna height, fixed seeds for pseudorandom generation, and a common 120 s optimization time limit. To reflect scenario-specific coverage and computational requirements, the city scenarios used k=2 and 30 LoS discretization steps, while the regional and national scenarios used k=1 and 10 steps.

Table 1. Statistical Summary of Algorithm Performance for the 100-Node Scenario (20 Instances)

Algorithm	Gateways Mean ± SD	Median	Runtime Mean ± SD (s)
ILP	36.75 ± 4.14	37.5	0.296 ± 0.056
Greedy	37.80 ± 4.07	38.5	0.005 ± 0.001
SA	36.75 ± 4.14	37.5	1.270 ± 0.204
VND	37.15 ± 3.99	37.5	0.173 ± 0.072
Hybrid GA	37.50 ± 3.89	38.0	0.769 ± 0.089
ACO	36.75 ± 4.14	37.5	59.727 ± 6.375

Across 20 independent 100-node instances, the exact ILP reference solution averaged 36.75 gateways (SD = 4.14; median = 37.5). SA and ACO matched this reference mean, while VND, Hybrid GA, and Greedy exceeded it by only 0.40, 0.75, and 1.05 gateways, respectively. Greedy was fastest (0.005 s mean), whereas ACO required 59.727 s; all runs were feasible without timeouts.

Table 2. Statistical Summary of Algorithm Performance for the 300-Node Scenario (20 Instances)

Algorithm	Gateways Mean ± SD	Median	Runtime Mean ± SD (s)
ILP	49.25 ± 3.45	48.0	1.861 ± 0.194
Greedy	52.25 ± 3.60	52.0	0.008 ± 0.001
SA	49.85 ± 3.34	49.0	2.832 ± 0.350
VND	50.85 ± 3.45	50.0	4.002 ± 1.274
Hybrid GA	51.75 ± 3.63	51.0	1.750 ± 0.245
ACO	49.70 ± 3.47	49.0	120.000

Across 20 independent 300-node instances, the exact ILP reference solution averaged 49.25 gateways (SD = 3.45; median = 48.0). ACO and SA were closest, exceeding this reference mean by only 0.45 and 0.60 gateways, respectively. VND, Hybrid GA, and Greedy exceeded it by 1.60, 2.50, and 3.00 gateways. Greedy was fastest (0.008 s mean). Although ACO reached the 120 s limit in all 20 runs, it returned feasible solutions averaging 49.70 gateways.

Table 3. Statistical Summary of Algorithm Performance for the 1000-Node Scenario (5 Instances)

Algorithm	**Gateways Mean ± SD**	**Median**	**Runtime Mean ± SD (s)**
ILP	467.80 ± 15.85	464.0	19.963 ± 3.389
Greedy	475.00 ± 16.39	472.0	0.091 ± 0.013
SA	471.60 ± 17.78	468.0	39.113 ± 7.842
VND	475.00 ± 16.39	472.0	120.000
Hybrid GA	474.40 ± 16.41	472.0	39.491 ± 5.338
ACO	475.00 ± 16.39	472.0	120.000

Across five independent 1000-node instances, the exact ILP reference solution averaged 467.80 gateways (SD = 15.85; median = 464.0). SA produced the closest non-exact solutions, averaging 471.60 gateways, only 3.80 (0.81%) above the reference. Hybrid GA averaged 474.40 gateways, while Greedy, VND, and ACO averaged 475.00. Greedy was fastest (0.091 s mean). Within the 120 s limit, VND and ACO did not improve the gateway count of the initial Greedy solutions and reached the limit in all five runs, although all returned feasible solutions.

Table 4. Statistical Summary of Algorithm Performance for the 3000-Node Scenario (3 Instances)

Algorithm	**Gateways Mean ± SD**	**Median**	**Runtime Mean ± SD (s)**
ILP	904.33 ± 25.58	916.0	188.773 ± 2.693
Greedy	936.67 ± 24.91	949.0	0.662 ± 0.049
SA	932.33 ± 24.79	943.0	120.000
VND	936.67 ± 24.91	949.0	120.000
Hybrid GA	936.67 ± 24.91	949.0	120.000
ACO	936.67 ± 24.91	949.0	120.000

Across three independent 3000-node instances, the time-limited ILP reference returned the best feasible solutions (mean = 904.33; SD = 25.58; median = 916.0). SA was the only non-exact method to improve on Greedy, averaging 932.33 gateways, 3.10% above the ILP reference. VND, Hybrid GA, and ACO matched the Greedy mean of 936.67 gateways. Greedy was fastest (0.662 s mean), while all four metaheuristics reached the 120 s limit. The ILP end-to-end mean of 188.773 s includes model construction before the bounded solver execution. All runs returned feasible solutions.

Overall, ILP provided the reference for solution quality; for the 3000-node scenario, this reference was the best feasible ILP solution found within the configured limit. Greedy consistently achieved the shortest runtimes. SA improved upon Greedy across all tested scales and provided the strongest overall balance between solution quality and computational time among the evaluated metaheuristics. ACO was competitive at 100 and 300 nodes but reached the time limit as the scale increased. At 1000 and 3000 nodes, several methods returned the initial Greedy gateway count, indicating that larger instances require additional tuning or decomposition into coordinated subproblems.

Conclusion and Future Work

This paper presented a terrain-aware methodology for optimizing LoRa gateway placement in large-scale vending machine networks. By combining Digital Elevation Model (DEM) data, line-of-sight validation, graph-based modeling, and optimization techniques, the proposed approach enables more realistic infrastructure planning compared to distance-only methods.

Across 48 independent instances, ILP provided the solution-quality reference. At 3000 nodes, the time-limited ILP solutions were treated as feasible references rather than proven optimal solutions. Greedy consistently achieved the shortest runtimes, while SA improved upon the Greedy solution across all tested scales and provided the strongest overall balance between solution quality and computational time among the evaluated metaheuristics. ACO was competitive on smaller instances but reached the time limit as the scale increased. These findings emphasize the trade-off between solution quality and computational cost: modest gateway reductions can lower infrastructure expenses, whereas rapid Greedy solutions support time-sensitive planning.

Future work may investigate coordinated decomposition of national-scale deployments into regional or local subproblems, together with additional deployment constraints, real vending machine data, building-aware line-of-sight modeling, and adaptive optimization strategies for evolving networks.

Acknowledgments

This study is supported by the project SP25-FMI-008 “Research and implementation of modern language models and artificial neural networks for automated processing, forecasting, and structuring of data and texts in specific application domains” at the Paisii Hilendarski University of Plovdiv.

References

[1] T. Haynes, S. Hedetniemi, P. Slater, *Fundamentals of Domination in Graphs*, Marcel Dekker, Inc., New York, 1998, ISBN: 978-0824700331

[2] Z. Li, Q. Zhu, C. Gold, *Digital Terrain Modeling: Principles and Methodology*, CRC Press, Boca Raton, 2004, ISBN: 978-0415324625

[3] A. Augustin, J. Yi, T. Clausen, W. Townsley, A Study of LoRa: Long Range & Low Power Networks for the Internet of Things, *Sensors*, 2016, vol. 16 (9), Article 1466, ISSN: 1424-8220, DOI: https://doi.org/10.3390/s16091466

[4] G. Nemhauser, L. Wolsey, *Integer and Combinatorial Optimization*, Wiley, New York, 1988, ISBN: 978-0471828198

[5] C. Liao, G. Chang, K-Tuple Domination in Graphs, *Information Processing Letters*, 2003, vol. 87 (1), pp. 45–50, ISSN: 0020-0190, DOI: https://doi.org/10.1016/S0020-0190(03)00233-3

[6] V. Chvátal, A Greedy Heuristic for the Set-Covering Problem, *Mathematics of Operations Research*, 1979, vol. 4 (3), pp. 233–235, ISSN: 0364-765X, DOI: https://doi.org/10.1287/moor.4.3.233

[7] S. Kirkpatrick, C. Gelatt Jr., M. Vecchi, Optimization by Simulated Annealing, *Science*, 1983, vol. 220 (4598), pp. 671–680, ISSN: 0036-8075, DOI: https://doi.org/10.1126/science.220.4598.671

[8] N. Mladenović, P. Hansen, Variable Neighborhood Search, *Computers & Operations Research*, 1997, vol. 24 (11), pp. 1097–1100, ISSN: 0305-0548, DOI: https://doi.org/10.1016/S0305-0548(97)00031-2

[9] D. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley, Reading, Massachusetts, 1989, ISBN: 978-0201157673

[10] M. Dorigo, V. Maniezzo, A. Colorni, Ant System: Optimization by a Colony of Cooperating Agents, *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, 1996, vol. 26 (1), pp. 29–41, ISSN: 1083-4419, DOI: https://doi.org/10.1109/3477.484436

Simeon Monov[1,*], Denis Bechiragich[1]

[1] Paisii Hilendarski University of Plovdiv,

Faculty of Mathematics and Informatics,

236 Bulgaria Blvd., 4027 Plovdiv, Bulgaria

Corresponding author: smonov@uni-plovdiv.bg