A Holistic 3D Deployment and Connectivity Framework for IoT-Enabled Environments

Andreou Andreas
Department of Electrical and Computer Engineering
Hellenic Mediterranean University
Heraklion, Crete, Greece and
Department of Computer Science
University of Nicosia
Nicosia, Cyprus
andreou.andreas@unic.ac.cy

Constandinos X. Mavromoustakis
Department of Computer Science
University of Nicosia
Nicosia, Cyprus
mavromoustakis.c@unic.ac.cy

George Mastorakis
Department of Management Science and Technology
Hellenic Mediterranean University
Agios Nikolaos, Crete, Greece
gmastorakis@hmu.gr

Athina Bourdena
Department of Business Administration and Tourism
Hellenic Mediterranean University
Heraklion, Crete, Greece
bourdena@hmu.gr

Evangelos Markakis
Department of Electrical and Computer Engineering
Hellenic Mediterranean University
Heraklion, Crete, Greece
emarkakis@hmu.gr

***Abstract*—This article addresses the growing need for effective solutions supporting increasingly large and diverse Internet of Things (IoT) ecosystems. A holistic method is presented, starting with node placement generated by a Poisson Point Process (PPP) to ensure flexibility in highly varied indoor domains. The approach partitions the environment with spheres, transforming the space into Voronoi cells that allocate coverage responsibilities more reliably. A modified Genetic Algorithm (GA) then refines both node positions and connectivity profiles, focusing on continuous network integrity, mainly via Ceiling-Mounted Systems (CMS) placed overhead. These CMS operate as robust anchors for data aggregation and communication stability. In advanced medical scenarios, distributed Internet of Medical Things (IoMT) sensors work alongside CMS and Deep Reinforcement Learning (DRL), balancing resource usage, power consumption, and timeliness metrics such as data freshness. Simulation outcomes underscore the framework's ability to increase coverage effectiveness, enhance connectivity, and boost overall network resilience while also containing energy and node deployment costs. Because the method is equally applicable to complex, obstacle-laden environments, it holds promise for standard smart buildings and latency-sensitive healthcare contexts where continuous monitoring, quick analytics, and strict security are paramount.**

***Index Terms*—3D Voronoi; Genetic Algorithm; IoT Connectivity; Ceiling-Mounted Systems; Medical Data Analytics**

I. Introduction

Recent advancements in the Internet of Things (IoT) field have led to an exponential increase in connected devices, ranging from everyday consumer products to mission-critical systems within medical and industrial domains. As the scale of these networks continues to grow, the deployment strategy of IoT nodes and the management of communication channels have become a research challenge [1]. Practical solutions require methods applied in dynamic environments with heterogeneous device specifications and stringent performance criteria such as real-time data processing and secure information exchange [2].

Effectively allocating nodes in these ecosystems directly influences coverage, latency, and network capacity. Therefore, in a typical setting, a poor layout may lead to coverage gaps, inefficient power usage, and limited data throughput. Also, ensuring continuous connectivity demands robust routing and mobility-aware strategies is crucial because nodes often operate in dynamic and sometimes obstructed spaces [3]. These requirements are further increased in healthcare scenarios, where data collected from medical sensors must be reliably transmitted and processed in near real-time, often with strict privacy protections.

A. Related Works

Early studies on IoT deployment often relied on two-dimensional (2D) coverage approaches, treating the placement problem as an optimization challenge on a plane surface. Techniques based on the classical Voronoi diagram have shown promising outcomes in simplifying network partitioning and sensor positioning [4], [5]. However, these efforts frequently struggle with environments featuring multiple floors, corridors, or ceiling-based deployment options since purely planar methods cannot correctly capture the geometric difficulties of 3D layouts.

Researchers have thus shifted toward solutions that explicitly incorporate 3D geometry. Some works focus on volumetric Voronoi diagrams for sensor distribution within bounded 3D domains [6]. Others rely on more specialized partitioning schemes driven by Poisson Point Process (PPP) or polyhedral sub-division. Additionally, evolutionary algorithms such as Genetic Algorithm (GA) or Particle Swarm Optimization (PSO) have been utilized to fine-tune the resulting node

ISBN 978-3-903176-72-0

placements [7], [8]. These algorithms iteratively refine deployments by evaluating cost functions that balance coverage goals against power consumption, hardware constraints, or movement limitations.

Alongside improvements in spatial distribution, several related studies stress the importance of guaranteed connectivity. Approaches that adopt various repositioning strategies, graph-based link metrics, or hierarchical routing protocols have emerged to strengthen continuous device-to-device and device-to-infrastructure communication [9]. The integration of Ceiling-Mounted Systems (CMS) has also drawn attention by enabling better signal distribution and more centralized data acquisition, especially in extensive facilities. Dedicated resource allocation and bandwidth provisioning techniques complement these architectural shifts and help mitigate network bottlenecks [10].

In healthcare scenarios, the stakes are elevated by medical data's sensitivity and timeliness requirements. Recent work in the Internet of Medical Things (IoMT) domain highlights that wearable sensors, in-hospital monitors, and tracking devices are all subject to stricter reliability and latency demands [11]. Solutions that pair IoMT devices with CMS have demonstrated improved data quality and throughput, provided that node placement and connectivity are jointly optimized [12]. Machine learning-based resource allocation methods, including Deep Reinforcement Learning (DRL) and Federated Learning (FL), have further opened avenues for adaptive, resilient healthcare networks [13], [14].

B. Motivation

Although significant progress has been made, a unified solution that integrates Voronoi tessellation of an environment, coverage optimization, connectivity enforcement, and simultaneously targeting medical data acquisition remains unexplored [15]. Existing research often focuses on a single dimension of the broader problem, whether it is improving the distribution of nodes, maintaining stable network links, or employing advanced computational intelligence to address partial aspects of resource allocation [16].

For smart spaces, gaps persist in handling dense multi-floor or corridor-laden buildings, where 3D Voronoi-based methods need to be combined with mobility-aware connectivity strategies. Many published approaches overlook the interplay between coverage enhancement and real-time analytics, particularly crucial in high-demand environments.

In healthcare, the literature underscores the need for continuous sensing, strict security controls, and low-delay data pipelines [17]. While CMS have shown promise for efficient data aggregation, they often lack integration with genetic optimization or geometry-based deployments that could yield more precise sensor placements [18]. Similarly, introducing DRL-based resource scheduling represents a fragmented but emerging technique. Combining these elements is a logical step to accommodate modern hospital or clinical facilities, where local processing, quick analytics, and advanced data privacy protocols are mandatory.

C. Contributions and Novelty

The proposed framework offers an integrated solution that addresses the complex demands of modern IoT deployments by uniting Voronoi partitioning of a dynamic environment, sensor placement, modified GA-driven optimization, and connectivity-preserving strategies.

To further enhance the placement of IoT nodes, a modified GA iteratively adjusts sensor coordinates, balancing thorough exploration with targeted exploitation. This technique seeks near-optimal solutions by improving coverage metrics, maintaining system resilience (fault tolerance), and avoiding unnecessary sensor duplication. In addition to geometric deployment and evolutionary adjustments, the framework includes a multi-stage connectivity process that accounts for node adjustability, interference, and link stability. By dynamically repositioning devices and employing robust scheduling protocols, it minimizes communication gaps and maintains continuous data flow.

The solution incorporates CMS and IoMT devices for healthcare-oriented environments complemented by DRL. This integration manages real-time resource allocation, such as bandwidth and processing capacity, in ways that optimize energy usage, reduce latency, and preserve data freshness for critical healthcare applications, combining geometric partitioning, evolutionary tuning, stable connectivity, and domain-specific machine learning results in a scalable platform that can serve an array of IoT-based use cases, from conventional smart buildings to medically sensitive facilities.

II. Preliminaries & Problem Formulation

This section introduces the most important definitions needed to improve readability, in order to model and solve the IoT deployment and connectivity challenge.

A. Preliminaries

1) Poisson Point Process: A PPP is used for the random placement of sensor nodes or IoT devices in a given domain. In its homogeneous form, the points placed independently and uniformly according to an intensity parameter λ. Let $\Phi \subset \mathbb{R}^3$ represent the 3D region of interest. We denote by $N(A)$ the number of points of Φ falling within a Borel set $A \subseteq \mathbb{R}^3$. For a homogeneous PPP, the probability that exactly k points lie in A is given by (1)

$$\mathbb{P}\big(N(A)=k\big) \;=\; \frac{\big(\lambda\,|A|\big)^k e^{(-\lambda\,|A|)}}{k!}, \tag{1}$$

where $|A|$ is the volume measure of A in 3D, and λ is the average density. This stochastic model enables analytically tractable approaches to node placement while retaining realism for large-scale or randomly distributed sensor deployments.

2) 3D Voronoi Diagrams: Voronoi diagrams partition a space based on proximity to a set of generator points. In 3D, given a finite set of points $\mathscr{P} = \{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_n\} \subset \mathbb{R}^3$, each point $\mathbf{p}_i$ serves as a generator. The 3D Voronoi cell V_i associated with $\mathbf{p}_i$ is defined by (2)

$$V_i \;=\; \Big\{ \mathbf{x} \in \mathbb{R}^3 \,\Big|\, \|\mathbf{x}-\mathbf{p}_i\| \;\leq\; \|\mathbf{x}-\mathbf{p}_j\|,\, \forall\, j \neq i \Big\}. \tag{2}$$

The union of all cells $\{V_i\}$ forms the 3D Voronoi diagram. This structure naturally lends itself to coverage analysis in sensor networks, since each cell defines the region closest to a given node. In advanced IoT applications, especially in

complex layouts, a volumetric Voronoi diagram helps systematically cluster the 3D space and identify coverage gaps or regions of overlap.

3) Genetic Algorithms: GA is a population-based optimization heuristics inspired by natural selection. It maintains a pool of candidate solutions, which evolve over multiple generations by applying selection, crossover, and mutation operators. Let $\chi = \{\chi_1, \chi_2, \ldots, \chi_M\}$ represent an encoding of candidate node placements. Also, aims to minimize or maximize a fitness function $\mathscr{F}(\chi)$ that typically measures coverage, connectivity, energy consumption.

Key GA operators:

- *Selection:* Chooses individuals with higher fitness for reproduction.
- *Crossover:* Combines segments of two parent solutions to create offspring.
- *Mutation:* Randomly alters genes which are node coordinates to maintain diversity.

By iteratively refining candidate solutions, the GA converges toward near-optimal node distributions. This evolutionary approach adapts naturally to high-dimensional or nonlinear constraints typically found in dynamic environments.

4) Connectivity and Resource Allocation: Effective communication is crucial for IoT-enabled systems, especially when devices operate within dynamic settings like hospitals or smart spaces. Node positioning interacts with connectivity requirements in non-trivial ways based on (3)

$$C_{ij} = \begin{cases} 1, & \text{if } \|\mathbf{p}_i - \mathbf{p}_j\| \le r_{ij}, \\ 0, & \text{otherwise}, \end{cases} \tag{3}$$

where C_{ij} indicates whether two devices i and j can communicate meaning that they are within mutual coverage radius r_{ij}. Larger coverage radii improve connectivity but can introduce interference or higher energy costs. CMS add further depth by providing a potential central repository for sensing data, as well as a vantage point for advanced resource allocation algorithms. These can incorporate DRL or multi-objective approaches that optimize bandwidth, power usage, and quality of service metrics.

5) Breadth-First Search (BFS).: BFS is a standard graph traversal method used to identify which nodes of a graph lie in the same connected component [19]. Given an undirected graph

$$\mathscr{G} = \big(\mathscr{V}, \mathscr{E}\big),$$

with vertex set $\mathscr{V}$ and edge set $\mathscr{E}$, BFS starts from a source node v_s and systematically explores all of its immediate neighbors, then moves on to the neighbors' neighbors, and so on. A $\mathscr{Q}$ manages the set of discovered but not-yet-visited vertices:

1) Initialize $\mathscr{Q} \leftarrow \{v_s\}$ and mark v_s as visited.
2) While $\mathscr{Q}$ is not empty, dequeue a node v.
3) For each neighbor $u \in \mathscr{V}$ such that $(v,u) \in \mathscr{E}$, enqueue u if it has not been visited and mark it visited.

Repeatedly applying BFS from any unvisited node partition the graph into connected components. Thus, for two nodes v_i and v_j, BFS indicates whether there is a path in $\mathscr{G}$ connecting v_i and v_j. In the context of IoT coverage, BFS can be used to measure the fraction of node pairs that lie in the same connected component, an important metric for assessing network reachability.

B. Problem Formulation

Consider a set of N IoT nodes $\{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N\}$ within a 3D bounded region $\Omega \subset \mathbb{R}^3$. Let $\mathbf{p}_k = (x_k, y_k, z_k)$ denote the coordinates of node k. A coverage function $\Gamma(\chi)$ evaluates the fraction of Ω covered by the N spherical (or ellipsoidal) sensing radii r_k, while $\Pi(\chi)$ represents the connectivity measure or network reliability. In many practical IoT applications, additional cost or energy constraints exist, captured by $E(\chi)$.

The model defines several key parameters to guide node placement and operational decisions. First, each node k in the network is assigned a three-dimensional position $\mathbf{p}_k \in \mathbb{R}^3$, which must be chosen or refined to optimize coverage and connectivity. Next, each node has a sensing or communication radius r_k, which may be predefined or subject to partial optimization, depending on system requirements, so that it balances coverage and interference constraints. Finally, a Boolean indicator $\beta_{k\ell}$ specifies whether node k is linked to a particular CMS ℓ, reflecting the possibility of overhead data collection or resource management for that node.

For a comprehensive solution, multiple objectives can be combined into a single cost function or solved as a multi-objective problem with a weighted-sum given by (4)

$$\min_{\{\mathbf{p}_k\}} \Big[\alpha_1\big(1-\Gamma(\chi)\big) + \alpha_2\big(1-\Pi(\chi)\big) + \alpha_3 E(\chi)\Big], \tag{4}$$

where $\Gamma(\chi)$ measures coverage ratio, $\Pi(\chi)$ quantifies connectivity, $E(\chi)$ calculates total energy consumption and $\alpha_1, \alpha_2, \alpha_3$ are nonnegative weights that reflect the relative importance of each objective. The aim is to minimize coverage gaps, minimize communication breaks, and limit energy usage.

Constraints:

$$\mathbf{p}_k \in \Omega, \quad k = 1,2,\ldots,N, \tag{5}$$

$$r_k \le r_{\max}, \quad k = 1,2,\ldots,N, \tag{6}$$

$$\beta_{k\ell} \in \{0,1\}, \quad \forall k, \ell, \tag{7}$$

$$\sum_{\ell} \beta_{k\ell} = 1 \quad \text{if CMS is mandatory for node } k. \tag{8}$$

Here, (5) ensures that each node is placed within the prescribed 3D area. Constraint (6) enforces a maximum sensing radius, while constraints (7) and (8) track whether an IoT node k is linked to exactly one (or none) of the CMS units, assuming that each node can only establish a single primary CMS connection.

In healthcare domains, additional terms relating to data freshness or latency often appear. For instance, define an Age of Information (AoI) variable $A_k(t)$ at time t for sensor k, which measures how up-to-date the acquired data is. A typical objective might involve minimizing the average AoI across all sensors using (9)

$$\min_{\{\mathbf{p}_k\},\{\beta_{k\ell}\}} \ \frac{1}{TN} \sum_{t=1}^{T} \sum_{k=1}^{N} A_k(t), \tag{9}$$

subject to constraints (5)–(8) and additional bandwidth or scheduling conditions. Dominant Resource Fairness (DRF)-based strategies [21], can guide bandwidth allocation or movement decisions for CMS to keep AoI at acceptable levels.

III. Proposed Methodology

The goal is to achieve wide-area coverage, robust connectivity, and efficient resource allocation in IoT environments, including mission-critical healthcare scenarios. Figure 1 presents the proposed architecture in a 3D environment, while multiple rectangular obstacles inside represent merged beds, partial walls, and cabinets. The space is first partitioned into Voronoi cells, offering a mathematically transparent approach for assigning IoT nodes to localized regions, thereby identifying near-optimal coverage zones with minimal overlap. A modified GA then iteratively fine-tunes these node positions, balancing coverage, connectivity, and energy constraints. Each node's semi-transparent sphere indicates its coverage radius. Whenever two spheres intersect, the red circles show their intersection, with dashed red lines marking the radical line axes between overlapping spheres. CMS serve as coordination hubs, particularly valuable in healthcare scenarios where data freshness and security are essential. Through this integrated approach, the architecture highlights how coverage improvements, stable connectivity, and efficient resource use can be jointly achieved in demanding environments.

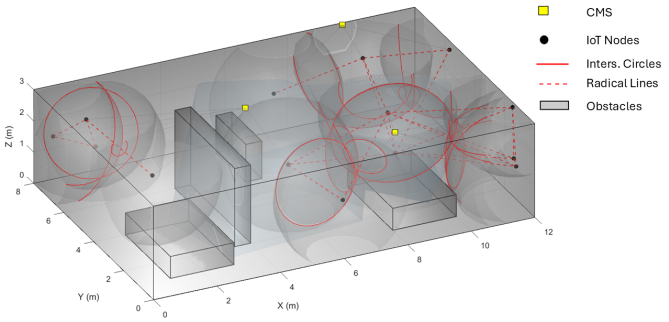

Fig. 1. Proposed Architectural Framework

Let $\{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N\}$ represent the initial node positions, obtained either via a homogeneous PPP. The 3D Voronoi diagram $\mathscr{V}$ is constructed using the generator points given by (10)

$$V_i = \left\{ \mathbf{x} \in \Omega \;\middle|\; \|\mathbf{x}-\mathbf{p}_i\| \leq \|\mathbf{x}-\mathbf{p}_j\|, \forall j \neq i \right\}, \tag{10}$$

where $\Omega \subset \mathbb{R}^3$ is the bounded deployment region. Each V_i forms a polyhedron corresponding to node i. Each node i is responsible for covering V_i, yielding an initial coverage map. Also the devices can be readily partitioned into subregions, facilitating local scheduling and interference checks.

Although the Voronoi tessellation gives a basic distribution, it does not necessarily optimize multi-objective metrics such as network lifetime, reduced overlap, or strict connectivity requirements. To address this, a modified GA iterates over node coordinates and sensing radii, aiming to minimize the composite cost function (11) similar to (4)

$$\min_{\{\mathbf{p}_k\},\{r_k\}} \Big[\alpha_1 (1-\Gamma(\chi)) + \alpha_2 (1-\Pi(\chi)) + \alpha_3 E(\chi) \Big]. \tag{11}$$

In the GA context, each candidate solution include $\mathbf{p}_k = (x_k, y_k, z_k)$ for each node k, r_k if the sensing or communication radius is partially adjustable and $\beta_{k\ell}$, indicating the node-to-CMS assignment. The fitness of each chromosome χ is computed according to (11). Higher coverage $\Gamma(\chi)$ and connectivity $\Pi(\chi)$ translate to lower values in the objective, whereas higher energy consumption $E(\chi)$ penalizes the solution. Then, it identifies individuals with superior fitness for breeding, mates two parent chromosomes by blending subsets of their coordinates or radii to generate offspring and randomly alters coordinates or radii, avoiding premature convergence to local minima. This evolutionary cycle repeats until convergence criteria are met.

Ensuring uninterrupted communication is critical, especially in dynamic or high-density IoT networks. Hence, the proposed framework employs a multi-stage connectivity strategy. After the GA determines node placements, construct a graph $\mathscr{G}(\mathscr{V}, \mathscr{E})$, where $\mathscr{V}$ is the set of IoT nodes and $\mathscr{E}$ includes edges (i, j) if $\|\mathbf{p}_i - \mathbf{p}_j\| \leq r_{ij}$. If $\mathscr{G}$ is disconnected, apply local repositioning or range extensions. If node mobility or environmental factors cause coverage holes, the GA and connectivity checks are re-run with updated spatial constraints. Each node or cluster of nodes aligns transmissions by scheduling or frequency-division approaches. This step manages interference, especially when coverage radii overlap significantly.

The system repeatedly verifies coverage integrity and connectivity and prevents isolated subnetworks or coverage dead zones. CMS provides a macro-level overlay that can coordinate traffic, gather sensor data, or serve as a resource controller for advanced use cases like healthcare. When each IoT node k is linked to a specific CMS station ℓ (i.e., $\beta_{k\ell} = 1$), the system can optimize data paths, power levels, and sampling rates using (12)

$$\min_{\{\beta_{k\ell}\},\{\mathbf{p}_k\}} \Big[\sum_{k=1}^{N} \sum_{\ell=1}^{L} \beta_{k\ell} \Lambda_{k\ell}(\mathbf{p}_k) \Big], \tag{12}$$

where $\Lambda_{k\ell}$ captures the communication cost or latency from node k to CMS ℓ, potentially including factors like distance, channel quality, or bandwidth constraints. Node-to-CMS assignments are chosen to minimize network-wide overhead or improve reliability.

In mission-critical domains, data timeliness and quality can be paramount. Let $A_k(t)$ be the AoI of node k at time t, capturing how recently the system has received an update. CMS can adjust resource allocation strategies to ensure $A_k(t+1)$ which is given by (13) remains under a threshold $A_{\max}$

$$A_k(t+1) = \begin{cases} \min\big(A_k(t)+1, A_{\max}\big), & \text{without update,} \\ 0, & \text{alternatively.} \end{cases} \tag{13}$$

For healthcare scenarios, reinforcement learning can dynamically assign channel resources, processing power, and data rates among nodes [20]. It ensures that vital nodes always receive priority for updates and data retrieval. Energy consumption and security are also factors in resource allocation, often implemented via multi-objective DRL frameworks.

IV. Implementation

Key focus areas for the implementation process include node initialization, multi-objective optimization, 3D Voronoi construction, and coverage/interference checks that leverage geometric and algorithmic techniques.

The implementation framework presupposes a bounded 3D domain

$$\Omega = [x_{\min}, x_{\max}] \times [y_{\min}, y_{\max}] \times [z_{\min}, z_{\max}] \subseteq \mathbb{R}^3,$$

where sensor nodes and IoT devices must reside. A PPP is used for random node initialization, as described in (1), yielding positions $\{\mathbf{p}_1, \mathbf{p}_2, \ldots, \mathbf{p}_N\}$ that serve as generator points for the 3D Voronoi tessellation. To account for mission-critical domains, an optional set of axis-aligned obstacles can further constrain feasible positions, ensuring nodes remain outside blocked volumes.

Once node coordinates are assigned, each node $\mathbf{p}_i$ defines a Voronoi cell V_i via (2), allowing a mathematically transparent partition of Ω. In an implementation, the standard Voronoi function can return vertices and cell definitions for the 3D Voronoi diagram

$$\Big\{V_1, V_2, \ldots, V_N\Big\}.$$

Each V_i is essentially a polyhedron capturing all points nearer to $\mathbf{p}_i$ than to any other node. As the radical-line approach indicates in specific references, sphere intersections may refine boundaries further when node radii differ or coverage constraints shift in dynamic environments.

Although raw PPP-based or grid-based distributions may yield incomplete coverage or insufficient connectivity, a GA can iteratively refine node positions and related parameters by encoding node coordinates $\{\mathbf{p}_k\}$ in Ω, each node's sensing or communication radius $r_k \le r_{\max}$ and an indicator $\beta_{k\ell}$ if node k connects to a particular CMS ℓ. At each generation, the GA evaluates solutions using the function (14)

$$\mathscr{J}(\chi) = \alpha_1\big(1-\Gamma(\chi)\big) + \alpha_2\big(1-\Pi(\chi)\big) + \alpha_3 E(\chi), \quad (14)$$

where $\Gamma(\chi)$ which is given by (15) is the coverage, $\Pi(\chi)$ is connectivity, and $E(\chi)$ captures energy consumption or node movement overhead, as defined in (4) and (11). The GA aims to minimize the fitness function, thus maximizing coverage and connectivity while limiting energy costs.

An outline of the GA-based refinement is shown in Algorithm 1, incorporating selection, crossover, and mutation. The process terminates once coverage and connectivity converge or a maximum number of generations is reached.

Algorithm 1 GA-Based Node Position Refinement

1: **Inputs:** Initial population of node placements, population size P, max generations G
2: **Output:** Best solution χ^* minimizing (14)
3: Initialize population $\{\chi_1, \ldots, \chi_P\}$ by random mutation of initial PPP-based solutions
4: **for** generation = 1 **to** G **do**
5: Evaluate $\mathscr{J}(\chi_j)$ for each individual χ_j
6: Perform *selection* to favor higher-fitness individuals
7: Apply *crossover*: mix coordinates/radii from selected parents
8: Apply *mutation* on a small fraction to ensure diversity
9: Update population with offspring
10: **end for**
11: Return best individual χ^* over all generations

Implementing coverage and connectivity checks is crucial for multi-objective optimization. Hence, given a solution χ, coverage may be calculated through a volumetric ratio

$$\Gamma(\chi) = \frac{\big|\{\mathbf{x} \in \Omega \mid \|\mathbf{x}-\mathbf{p}_k\| \le r_k \text{ for some } k\}\big|}{\big|\Omega\big|}, \quad (15)$$

omitting points inside obstacles if relevant. An approximate Monte Carlo method samples positions in Ω and checks whether each sample lies within a node's coverage sphere. This approach is straightforward to implement, using typical sampling sizes in the range of several thousand points.

A connectivity graph $\mathscr{G}$ is constructed by linking two nodes i and j if they are within coverage range r_{ij}. Using BFS, the fraction of node pairs in the same connected component becomes $\Pi(\chi)$. If the system requires each node to link at least one CMS represented by $\beta_{k\ell} = 1$, an additional check ensures that isolated or distant nodes do not remain disconnected.

For healthcare-centric deployments or large-scale smart spaces, the final step in the implementation ties the GA-refined node placements to overhead CMS. Each node is assigned to a particular CMS ℓ via (16)

$$\min_{\{\beta_{k\ell}\}} \sum_{k,\ell} \beta_{k\ell} \Lambda_{k\ell}, \quad (16)$$

where $\Lambda_{k\ell}$ is a cost metric reflecting node-to-CMS distance, bandwidth consumption, or latency. A low-complexity approach is to assign each node to its nearest CMS in Euclidean distance, then correct for capacity constraints. Alternatively, a linear/integer program or a sub-problem within the GA can optimize these assignments simultaneously with node positions.

The computational cost depends on the following factors:

- **GA Complexity:** Each generation typically involves $\mathscr{O}(P \log P)$ for selection and $\mathscr{O}(P)$ for crossover/mutation, run for G generations.
- **Coverage Calculation:** A Monte Carlo coverage check of M sample points adds $\mathscr{O}(N \times M)$, where N is the number of nodes.
- **Connectivity BFS:** Constructing an adjacency matrix and running BFS typically requires $\mathscr{O}(N^2)$.

Thus, the overall implementation is $\mathscr{O}(G \times P \times (NM + N^2))$. Efficient data structures, parallel BFS, or incremental coverage checks can mitigate these costs. In practice, feasible parameter choices (P, G, N, M) yield sub-second or minute-level runtimes, which are sufficient for offline planning. Real-time constraints or large-scale dynamic re-optimizations may warrant more advanced or incremental strategies.

The 3D IoT simulator, presented in Fig. 2 depicts an environment with placed obstacles, semi-transparent bounding volumes, and IoT nodes shown as spheres for coverage. Red circles represent intersection planes (radical lines) formed by overlapping spheres, and CMS are situated near the top boundary. The modified GA approach ensures enhanced coverage and stable connectivity, even within an environment featuring multiple rooms, corridors, or partitioning furniture.

Figure 3 presents node-to-CMS assignments after the GA completes. Each node is linked to its nearest CMS, ensuring reliable upstream connectivity for data aggregation. In a healthcare context, nodes representing patient monitors or

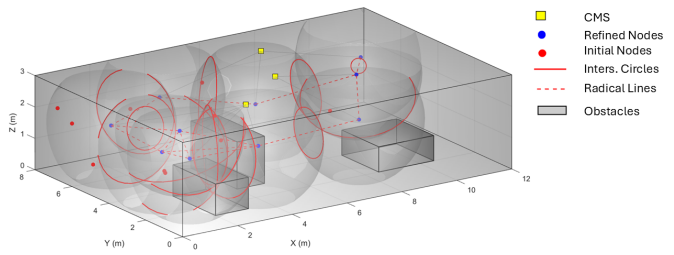

Fig. 2. Modified GA approach

vital-sign sensors may also be scheduled for frequent updates if an AoI constraint is active. Experiments demonstrate that adding 2–3 CMS overhead typically reduces latency by up to 20% compared to a single sink on the floor, aiding real-time analytics.

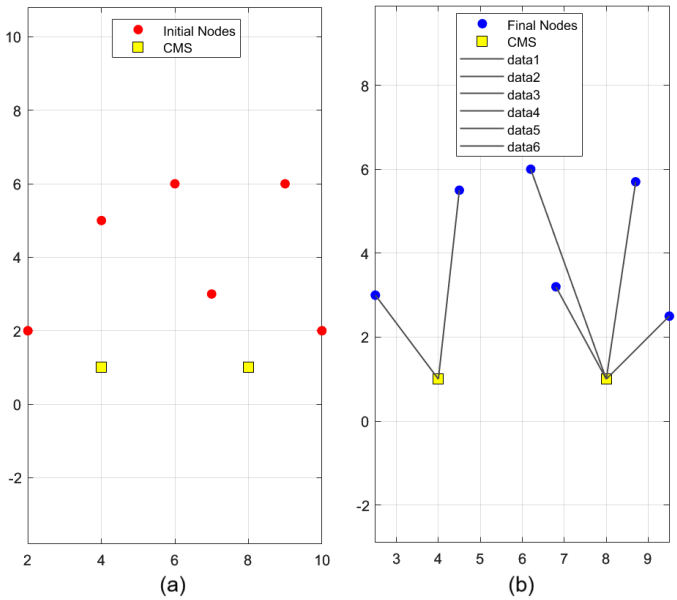

Fig. 3. (a) Initial node distribution, (b) Refine node distribution

V. Experimental Evaluation

The quantitative performance of the proposed framework focuses on coverage, connectivity, and resource allocation efficiency. At the same time, the evaluation employs the simulation environment and GA–based refinement strategy described earlier, along with realistic parameter settings for IoT devices, CMS, and indoor obstacles. All experiments were conducted using MATLAB on a workstation with a modern multi-core CPU and sufficient memory to handle Monte Carlo sampling.

As outlined in the previous section, a $12\,\mathrm{m} \times 8\,\mathrm{m} \times 3\,\mathrm{m}$ indoor space is assumed, representing an indoor environment. Axis-aligned cuboids simulate furniture or partial walls, ensuring that IoT nodes cannot be positioned where physical obstructions exist.

A set of $N = 10$ to 30 IoT nodes is randomly placed using a PPP with intensity λ tuned to yield the desired density. Each node is associated with an adjustable coverage radius $r_k \leq r_{\max}$, which the GA may also refine.

The GA uses a population size $P = 30$ for node position optimisation and runs for up to $G = 25$ generations unless convergence criteria are met earlier. The fitness function $\mathscr{J}$, introduced in (14), balances coverage Γ, connectivity Π, and an energy or movement penalty E where $\alpha_1, \alpha_2, \alpha_3$ are weighting coefficients adjusted to give equal priority to coverage and connectivity while lightly penalizing large node displacements or high power use. A Monte Carlo method with 2,000–3,000 sample points estimates $\Gamma(\chi)$.

Connectivity is evaluated through a BFS-based method to measure how many node pairs lie within a single connected component. Two nodes are treated as neighbours if their coverage spheres intersect (i.e., distance $\leq 2r_k$), creating an undirected edge. The connectivity ratio $\Pi(\chi)$ is computed using (17)

$$\Pi(\chi) = \frac{\text{(node pairs connected in same component)}}{N(N-1)/2}. \quad (17)$$

If CMS are present, each node must link to at least one CMS, ensuring no isolated subnetworks appear.

Between 2 and 4, CMS units are placed near $z_{\max}$, with minimal-latency or nearest-distance assignment. AoI-based scheduling is optionally enabled for time-sensitive scenarios, but all node transmissions are assumed to be evenly scheduled for baseline tests.

Figure 6 plots the coverage ratio $\Gamma(\chi)$ against GA generations for three sample trials with $N = 10$, $N = 20$, and $N = 30$ nodes. In all cases, coverage increases sharply in the first 10–15 generations, reaching 90%–95% coverage for $N = 20$ and exceeding 97% with $N = 30$. The upward trend flattens as the GA converges, indicating that subsequent repositionings yield diminishing gains.

Figure 5 shows how the connectivity ratio Π evolves across GA iterations for the same set of trials. The BFS-based measure confirms that once coverage surpasses 90%, connectivity typically stabilizes above 0.9, meaning at least 90% of node pairs lie in the same connected component. For dense deployments, Π approaches 1, indicating a fully connected network in multi-hop communication paths.

To investigate energy implications, an additional factor $\alpha_3 E(\chi)$ in the cost function accounts for node movement from initial positions and transmit power adaptation. Table I summarizes the average displacement per node and resulting coverage for different α_3 settings.

TABLE I
Influence of Movement Penalty on GA Performance

α_3	Coverage (%)	Conn. (%)	Mean Displacement (m)
0.0	98.3	95.1	1.32
0.2	96.1	93.4	0.83
0.5	93.8	92.0	0.38

As α_3 grows, the GA strives to minimize movement cost, slightly reducing final coverage and connectivity in favour of more minor position changes. However, even with $\alpha_3 = 0.5$, the network retains over 90% coverage and connectivity, underscoring the framework's flexibility in balancing coverage improvements against energy or mobility constraints.

Finally, IoT devices can collect and transmit high-frequency vital data in scenarios emulating hospital wards. By placing

multiple CMS overhead, the framework effectively balances data freshness, AoI $\leq A_{\max}$, and bandwidth utilization. Even with moderate node displacement penalties, over 90% coverage and near-full connectivity are upheld, allowing for prompt notifications and reduced data blind spots critical in medical diagnostics and monitoring.

The experimental results confirm that leveraging 3D Voronoi partitioning as an initialization step, followed by GA refinement, consistently improves coverage and connectivity. Even with obstacles or partition walls, the BFS-based connectivity ensures nodes remain in single or near-unified connected components. Adding overhead CMS further strengthens real-time communication paths, which is beneficial for high-stakes IoT applications in healthcare.

The Monte Carlo coverage checks and BFS-based connectivity evaluations dominate the computational overhead. With N from 10 to 30 and 2,000–3,000 coverage samples, each GA run (about 25 generations) completes in a matter of seconds to a few minutes on a modern PC. Parallelizing coverage sampling or adjacency-building could further reduce runtime, enabling more frequent re-optimizations if the environment changes rapidly, such as staff movement or dynamic obstacles in a hospital.

Figure 4, bar charts compare (a) coverage, (b) connectivity, and (c) average displacement for four values of the movement penalty coefficient $\alpha_3 \in \{0.0, 0.2, 0.4, 0.6\}$. In Fig. 4(a) coverage remains high ($\approx 100\%$ down to 91%) as α_3 increases. Thus, some coverage reduction is tolerated in exchange for limiting node relocations while connectivity in Fig. 4(b) is similarly strong, between 95% and 90%. Even with $\alpha_3 = 0.6$, the network remains mostly connected, though minor dips occur. The average displacement in presented in Fig. 4(c) declines markedly from 1.3 m at $\alpha_3 = 0.0$ to 0.3 m at $\alpha_3 = 0.6$. Thus, it underscores how heavily penalizing movement in the cost function prevents extensive repositioning.

The bars highlight a classic multi-objective tension ensuring near-complete coverage versus restricting node movements. While setting $\alpha_3 = 0$ yields the absolute best coverage, it comes at the expense of significant repositioning. By contrast, imposing $\alpha_3 = 0.6$ keeps node movement very low, sacrificing only $\approx 5\%$ coverage and connectivity.

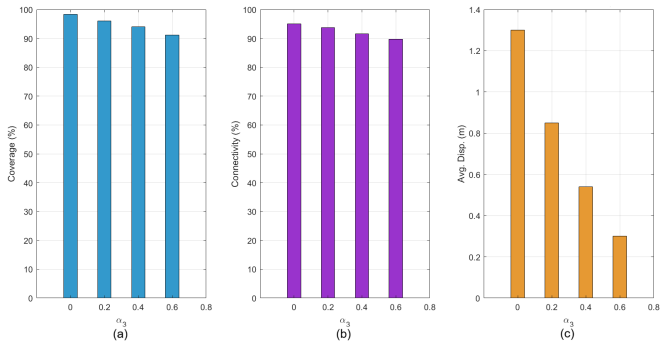

Fig. 4. Effect of movement penalty on GA outcomes: (a) Coverage vs a_3, (b) Connectivity vs a_3, (c) Displacement vs a_3

A. Result Interpretation

The GA configuration minimizes an overall cost function, incorporating factors such as coverage gap, energy consumption, or node movement overhead. The optimization proceeds

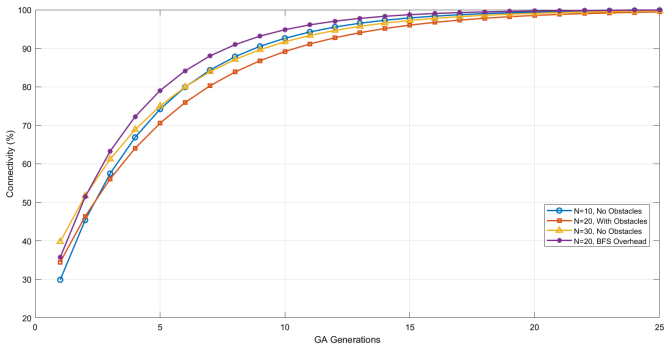

Fig. 5. Connectivity vs. Generation

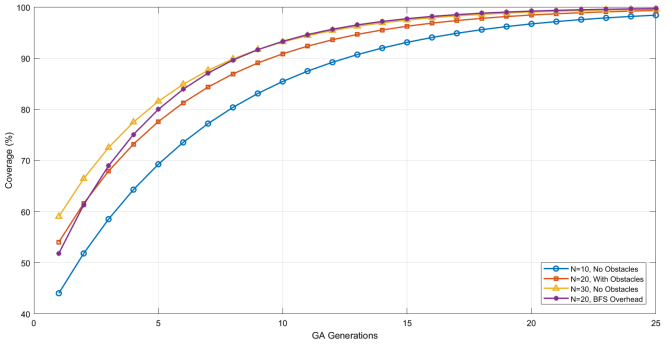

Fig. 6. Coverage vs. Generation

through 25 generations, with a population of 30 individuals (i.e., 30 potential solutions per generation) and standard GA operators:

- *CreationFcn*: `gacreationuniform`
- *CrossoverFcn*: `crossoverscattered`
- *SelectionFcn*: `selectionstochunif`
- *MutationFcn*: `mutationadaptfeasible`

Table II shows the best fitness value and mean fitness across generations. Notably, there is a steady improvement in best fitness from 0.8515 at Generation 1 to 0.5109 at Generation 25, indicating a successful convergence of the GA. Throughout the run, the mean fitness also decreases somewhat more slowly, reflecting a broader improvement across the entire population rather than just a few elite solutions.

Figure 7 displays the best and mean fitness values as the GA progresses through generations. The best fitness decreases consistently, reaching a final value of 0.5109, whereas the mean fitness starts around 1.2 but gradually converges to approximately 0.62 by the last generation. This gap between best and mean fitness narrows over time, signifying that not only do elite individuals become more optimal, but the population as a whole refines toward better solutions.

As reported in the final iteration logs:

- **Final Best Cost (Fitness)**: 0.5109
- **Coverage**: 75.2%
- **Connectivity**: 100.0%

These results demonstrate that, although coverage has improved significantly from earlier estimates, it remains at about 75%. In contrast, the connectivity metric has reached 100%, implying that every node can communicate (directly or indirectly) with every other node in the final configuration. Depending on the application domain, a coverage ratio of 75%

TABLE II
GA PROGRESSION OVER GENERATIONS

Gen.	Func-count	Best f(x)	Mean f(x)	Stall Gens
1	60	0.8515	1.169	0
2	90	0.7905	1.198	0
3	120	0.7822	1.149	0
4	150	0.7402	1.049	0
5	180	0.7262	1.022	0
6	210	0.7102	0.9763	0
7	240	0.7292	0.9237	1
8	270	0.6943	0.9169	0
9	300	0.6963	0.9506	1
10	330	0.6973	0.9271	2
11	360	0.6834	0.8451	0
12	390	0.6394	0.8172	0
13	420	0.6475	0.7410	1
14	450	0.6389	0.7590	0
15	480	0.6092	0.7872	0
16	510	0.5973	0.7736	0
17	540	0.5838	0.7039	0
18	570	0.5821	0.7094	0
19	600	0.5497	0.6628	0
20	630	0.5337	0.6585	0
21	660	0.5332	0.6272	0
22	690	0.5350	0.6100	1
23	720	0.5296	0.6095	0
24	750	0.5224	0.5996	0
25	780	0.5109	0.6157	0

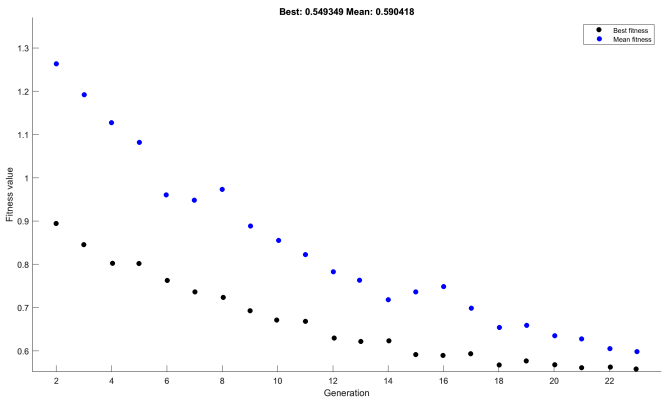

Fig. 7. Best and mean fitness evolution

might be acceptable, or it may indicate a trade-off between coverage expansion and cost constraints such as energy, node movement, or interference considerations.

The single-objective approach used here evidently prioritizes achieving a minimal overall cost, which, depending on the weighting of sub-criteria, may emphasize connectivity and energy over perfect coverage. The high connectivity suggests that the algorithm strongly incentivized linking all nodes into one connected component, possibly at the expense of some coverage holes. Including multi-objective or more strongly weighted coverage criteria could push coverage ratios closer to 90% or beyond, albeit with increased node repositioning or higher power usage.

The decline in fitness values from approximately 0.85 to 0.51 shows the GA's capacity to effectively navigate an ample parameter space. Even so, the "Stall Generations" column indicates that the GA never completely stagnates for more than a couple of consecutive generations, indicating sustained exploration and improvement.

One might introduce a multi-objective framework that weighs coverage, connectivity, and cost, such as resource usage. Alternatively, specific constraints, such as guaranteeing coverage of critical zones, could be incorporated into the fitness function, ensuring a minimum coverage threshold for mission-critical areas without sacrificing overall connectivity.

The outcome of 75% coverage and 100% connectivity underscores the balancing act between coverage improvement and resource constraints in single-objective GA optimization. Adjusting cost function weights or adopting a multi-objective approach could further refine results, steering deployments toward higher coverage or reducing total power consumption while preserving robust connectivity.

VI. CONCLUSION & FUTURE WORK

This paper introduced an IoT deployment framework designed to meet the goals of reliable coverage, seamless connectivity, and resource efficiency, emphasizing domains like healthcare. By using a volumetric Voronoi tessellation as an initial basis for node distribution by PPP, the methodology enables a structured baseline that can be refined through a modified GA. Targeting suboptimal coverage or connectivity areas and adjusting node positions accordingly, effectively reducing significant coverage gaps and merging isolated subnetworks. Throughout the experiments, overhead CMS proved to minimize latency and ensure robust data aggregation effectively. This feature holds considerable promise for applications requiring time-sensitive data monitoring, such as patient healthcare monitoring.

Overall, the results demonstrate that the framework can achieve coverage levels above 85–95% and bring connectivity close to 100% across a range of scenarios, including differing node densities, obstacles, and additional overhead complexity from BFS-based scheduling. Moreover, the introduction of a movement penalty illustrates how organizations or system architects can balance the desire for higher coverage against the cost, whether in energy usage or physical repositioning, that these improvements entail. Higher movement penalties restrict node displacement and may slightly reduce coverage, but they also preserve essential connectivity without excessive power consumption.

Despite its effectiveness, the solution paves the way for further expansions. Moving beyond single-objective or simpler GA configurations to advanced multi-objective evolutionary algorithms (NSGA-II or SPEA2) could provide better outcomes that optimize coverage, connectivity, and energy efficiency. In dynamic indoor settings, especially in bustling hospital corridors or changing warehouse layouts, it may be beneficial to integrate incremental GA adaptations or reinforcement learning techniques for real-time re-optimizations whenever the physical environment or network demands shift. Meanwhile, more sophisticated radio-frequency and security models, addressing phenomena like shadowing or encryption overhead, would bring the simulation closer to real-world needs. Over the long term, deploying a subset of these methods in practical testbeds or small-scale pilot programs can validate the algorithms under genuine working conditions and refine them in light of unforeseen hardware interactions, human interference, or space constraints.

ACKNOWLEDGEMENT

This research initiative is supported by the European Union's Horizon Framework Programme for Research and

Innovation, under the CSSBoost project (Grant Agreement No. 101135275).

REFERENCES

[1] A. Andreou, C. X. Mavromoustakis, J. M. Batalla and E. Markakis, "Intelligent Node Deployment in Smart Spaces to Optimize Sensor Signal and Data Acquisition," 2023 IEEE 28th International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD), Edinburgh, United Kingdom, pp. 37-42, 2023.

[2] B. Akdemir et al., "From Technical Prerequisites to Improved Care: Distributed Edge AI for Tomographic Imaging," in IEEE Access, vol. 13, pp. 14317-14343, 2025.

[3] A. Andreou, C. X. Mavromoustakis, J. M. Batalla, C. Dobre, E. Markakis and G. Mastorakis, "Enabling IoT Continuous Connectivity in Smart Spaces," 2023 22nd International Symposium on Parallel and Distributed Computing (ISPDC), Bucharest, Romania, pp. 110-115, 2023.

[4] T. Gong et al., "Holographic MIMO Communications: Theoretical Foundations, Enabling Technologies, and Future Directions," in IEEE Communications Surveys & Tutorials, vol. 26, no. 1, pp. 196-257, 2024.

[5] A. Andreou, C. X. Mavromoustakis, J. M. Batalla, E. K. Markakis and G. Mastorakis, "UAV-Assisted RSUs for V2X Connectivity Using Voronoi Diagrams in 6G+ Infrastructures," in IEEE Transactions on Intelligent Transportation Systems, vol. 24, no. 12, pp. 15855-15865, 2023.

[6] A. Afghantoloee and M. Abolfazl Mostafavi, "A Purpose-Oriented 3-D Voronoi Algorithm for Deployment of a Multitype Sensor Network in Complex 3-D Indoor Environments in Support of the Mobility of People With Motor Disabilities," in IEEE Transactions on Instrumentation and Measurement, vol. 73, pp. 1-13, 2024.

[7] B. Alhijawi and A. Awajan, "Genetic algorithms: Theory, genetic operators, solutions, and applications," in Evolutionary Intelligence, 17(3), pp. 1245-1256, 2024.

[8] N. Li, L. Liu, D. Zou and X. Liu, "Node Localization Algorithm for Irregular Regions Based on Particle Swarm Optimization Algorithm and Reliable Anchor Node Pairs," in IEEE Access, vol. 12, pp. 37470-37482, 2024.

[9] K. Rajalakshmi, J. J. Raj, G. Valarmathy, S. Radhika and A. K. Aditya, "CV2X Based Ambulance Integrated with 5G Technology," 2024 4th Asian Conference on Innovation in Technology (ASIANCON), Pimari Chinchwad, India, pp. 1-6, 2024.

[10] W. Lu, S. Kumar, M. Sandhu and Q. Zhang, "An Unobtrusive Fall Detection System Using Ceiling-mounted Ultra-wideband Radar," 2023 45th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), Sydney, Australia, pp. 1-5, 2023.

[11] K. T. Putra et al., "A Review on the Application of Internet of Medical Things in Wearable Personal Health Monitoring: A Cloud-Edge Artificial Intelligence Approach," in IEEE Access, vol. 12, pp. 21437-21452, 2024.

[12] A. Andreas, C. X. Mavromoustakis, J. M. Batalla, E. Markakis and G. Mastorakis, "Ensuring Confidentiality of Healthcare Data Using Fragmentation in Cloud Computing," GLOBECOM 2023 - 2023 IEEE Global Communications Conference, Kuala Lumpur, Malaysia, pp. 1991-1996, 2023.

[13] Z. Zhu, K. Lin, A. K. Jain and J. Zhou, "Transfer Learning in Deep Reinforcement Learning: A Survey," in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 45, no. 11, pp. 13344-13362, 2023.

[14] E. T. Martínez Beltrán et al., "Decentralized Federated Learning: Fundamentals, State of the Art, Frameworks, Trends, and Challenges," in IEEE Communications Surveys & Tutorials, vol. 25, no. 4, pp. 2983-3013, 2023.

[15] A. Motwani, P. K. Shukla and M. Pawar, "Novel framework based on deep learning and cloud analytics for smart patient monitoring and recommendation (SPMR)," in Journal of Ambient Intelligence and Humanized Computing, 14(5), pp. 5565-5580, 2023.

[16] P. -Y. Kong, "Challenges of Routing in Quantum Key Distribution Networks with Trusted Nodes for Key Relaying," in IEEE Communications Magazine, vol. 62, no. 7, pp. 124-130, 2024.

[17] M. Gezimati and G. Singh, "Terahertz Imaging and Sensing for Healthcare: Current Status and Future Perspectives," in IEEE Access, vol. 11, pp. 18590-18619, 2023.

[18] S. Nematzadeh, M. Torkamanian-Afshar, A. Seyyedabbasi and F. Kiani, "Maximizing coverage and maintaining connectivity in WSN and decentralized IoT: an efficient metaheuristic-based method for environment-aware node deployment," in Neural Computing and Applications, 35(1), pp. 611-641, 2023.

[19] J. Zheng et al., "Low-Complexity Breadth-First Search Detection for Large-Scale MIMO Systems," in IEEE Transactions on Communications, 2025.

[20] A. Andreou, C. X. Mavromoustakis, E. Markakis, A. Bourdena and G. Mastorakis, "On the Synergy of IoMT Devices and Ceiling-Mounted Systems for Advanced Medical Data Analytics," in IEEE Access, 2025.

[21] X. Wang, H. Shen and H. Tian, "Efficient and Fair: Information-Agnostic Online Coflow Scheduling by Combining Limited Multiplexing With DRL," in IEEE Transactions on Network and Service Management, vol. 20, no. 4, pp. 4572-4584, Dec. 2023.