A STOCHASTIC MULTI-OBJECTIVE OPTIMIZATION FRAMEWORK FOR REAL-TIME IOT RESOURCE ALLOCATION IN URBAN DIGITAL TWINS

Mohamed M. H. Maatouk (1), Mhd Walid Koubeisi (2), Mahmoud Mohamed (2)*

(1) Faculty of Architecture and Planning, King Abdulaziz University, Jeddah;
(2) Electrical and Computer Engineering, King Abdulaziz University, Jeddah
Saudi Arabia

* Corresponding Author, e-mail: mmaatouk@kau.edu.sa

Abstract: Urban Digital Twins (UDTs) synchronize virtual city replicas with real-time IoT data, yet unrestricted ingestion exhausts bandwidth, computation and energy budgets while inflating latency. The paper formulates UDT data ingestion as a stochastic multi-objective integer program with chance-constrained bandwidth, energy and latency risk, using a redundancy-aware submodular Mutual-Information utility objective, and solves it with the Optimized Adaptive Data Ingestion (OADI) framework - a simheuristic coupling Monte-Carlo scenario generation with the novel Stochastic Adaptive Reference Point Mechanism (SARPM), which reshapes NSGA-III reference points from the posterior constraint-violation distribution. On PeMS and OpenAQ, OADI achieves a 34 higher Information-to-Cost Ratio than the best greedy heuristic, a 94.8 constraint-satisfaction rate (estimated on 1,000 held-out scenarios), and statistically significant Hypervolume and IGD+ improvements over seven baselines and a SARPM ablation (Wilcoxon $p < 0.01$).

Key words: urban digital twins, IoT resource allocation, stochastic multi-objective optimization, NSGA-III, adaptive data ingestion.

1. INTRODUCTION

Urban Digital Twins (UDTs) synchronize virtual replicas of the physical city from GIS, BIM and real-time sensor feeds, enabling prediction and closed-loop control [1], [2]. The IoT layer produces petabytes of daily observations [3], exceeding edge and backhaul capacity. The prevailing data-in-all response ingests every stream, yet marginal value diminishes once dominant uncertainties resolve while costs grow: the twin is starved of timely information because it is overfed with redundancy. Optimized Digital Twins, unlike passively mirrored Digital Shadows [4], require explicit sensor-to-twin optimization. We define the Real-Time IoT Resource Allocation Problem (RT-IoT-RAP): at each step, select the IoT streams to ingest, maximizing information utility under bandwidth, energy and latency budgets. Four complexities arise: dimensionality ($N > 10^5$; 2^N search space), stochasticity, conflicting objectives, and hard real-time

constraints. We contribute: (1) the first stochastic multi-objective integer program coupling redundancy-aware submodular Mutual-Information utility with chance-constrained bandwidth, energy and latency-risk budgets; (2) the Optimized Adaptive Data Ingestion (OADI) framework, whose Stochastic Adaptive Reference Point Mechanism (SARPM) reshapes NSGA-III reference points from posterior constraint-violation distributions; (3) validation on PeMS and OpenAQ against seven baselines and an SARPM ablation, reporting HV, IGD+, ICR and CSR with Wilcoxon tests. All results are the authors' own.

2. RELATED WORKS

Recent studies identify computational sustainability as the binding constraint on UDT designs, because data-ingestion rates grow exponentially while processing capability grows only linearly [5]. Evolutionary strategies suit sensor selection well: their populations of binary vectors explore the combinatorial space directly. Closest to our work, reinforcement learning for real-time traffic-signal optimization is investigated in [6], and stochastic energy optimization for digital-twin-assisted edge networks is presented in [7]; however, the former lacks formal constraint guarantees, and the latter treats data fidelity as secondary. A hybrid symmetry-improved binary optimization method for the 0-1 knapsack is developed in [8], but without a stochastic urban context. Many-objective Pareto-front modeling is addressed in [9], and the runtime inefficiency of NSGA-II is quantified in [10]; neither line considers constraint-driven reference-point adaptation.

Table 1. The authors' comparative analysis of research gaps

Aspect	***Static/Naive***	***Heuristic***	***Proposed OADI***
Dimensionality	Ignores 2^N	Fails at large N	Optimized for $N > 10^5$
Stochasticity	Not considered	Rarely modelled	Core to formulation
Objectivity	Single (cost/accuracy)	Single (greedy)	Multi-objective Pareto
Real-time	High latency risk	Moderate	Constrained optimization

3. SYSTEM MODEL AND PROBLEM FORMULATION

We model the UDT as a discrete-time state-space system. At step $t \in \{0, 1, \ldots, T\}$ the physical state is $S_t \in R^d$ and an estimates the twin maintain $\hat{S}_t$, seeking to minimize the error $\|S_t - \hat{S}_t\|$. The environment is instrumented with N sensors, $I = \{1, 2, \ldots, N\}$. A binary decision vector $x_t = [x_{1,t}, \ldots, x_{N,t}]^T$ encodes the allocation strategy, where $x_{i,t} = 1$ if and only if sensor i is activated at step t. Stochasticity is captured by a random variable —network congestion, weather, traffic density—drawnω from a distribution Ω estimated from historical city data. The problem is a Stochastic Multi-Objective Optimization Problem (SMOP) with two conflicting objectives. The first maximizes information utility—the Mutual Information between the true state and the ingested data conditioned on H_{t-1} and ω—while penalizing the redundancy among correlated sensors:

$$f_1(x, \omega) = max_x \Sigma_{i \in I} x_{i,t} \cdot \mu_i(\omega) - \Sigma_{(i,j) \in ER} \lambda_{ij} \cdot x_{i,t} \cdot x_{j,t} \quad (1)$$

where $\mu_i(\omega) = E_\omega[I(S_t; D_{i,t} \mid H_{t-1}, \omega)]$ is the expected Mutual Information of sensor i, $E_R \subseteq I \times I$ a sparse redundancy graph, and $\lambda_{ij} \in [0, 1]$ the normalized Mutual Information

between correlated sensors, estimated from historical co-observations. Since all $\lambda_{ij} \geq 0$, the utility is monotone submodular—activating a sensor whose correlated neighbours are already active yields a diminished marginal gain—so redundancy is priced rather than double-counted, and the additive form is recovered for $E_R = \emptyset$.

The second minimizes a composite operational cost:

$$f_2(x, \omega) = \Sigma_{i \in I}\, x_{i,t} \cdot E_\omega[w_L \cdot L_i(\omega) + w_E \cdot E_i(\omega) + w_B \cdot B_i] \quad (2)$$

where $L_i(\omega)$, $E_i(\omega)$ and B_i denote per-sensor latency, energy and bandwidth, and w_L, w_E, w_B policy weights. Three constraints apply: bandwidth and energy budgets $\Sigma\, x_{i,t} \cdot B_i \leq B_{total}(\omega)$ and $\Sigma\, x_{i,t} \cdot E_i(\omega) \leq E_{total}(\omega)$, and a latency chance constraint bounding deadline-violation risk at tolerance α:

$$\Pr[T_{total}(x_t, \omega) > T_{max}] \leq \alpha \quad (3)$$

The model extends the classical stochastic knapsack to a multi-dimensional, chance-constrained setting with submodular reward and is therefore substantially harder than its deterministic IoT variants: with $N > 10^5$, exact solution methods are infeasible within the twin's time step, motivating metaheuristic approximation with predictable polynomial-time complexity.

4. THE OADI FRAMEWORK

4.1. Framework architecture

OADI operates as an intelligent middleware layer between the physical IoT fabric and the Digital Twin core, selecting an activation vector x^* at every time step. Figure 1 shows its three modules. The **Context Analyzer** monitors the state ω—bandwidth, battery levels, latency metrics—and produces a compact context descriptor. The **Optimization Engine** runs the hybrid NSGA-III-SARPM solver on the state, constraints and prior history, generating a Pareto front of activation vectors. The **Decision Selector** maps this front to a single operating point according to the current priority, favouring low-latency solutions in emergencies and high-fidelity ones in planning.

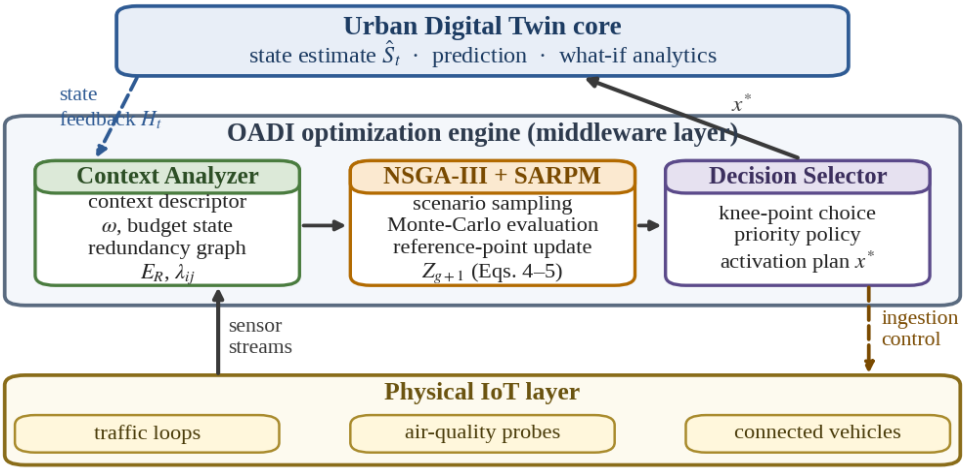


Figure 1. The OADI architecture

4.2. Enhanced NSGA-III with SARPM

Standard NSGA-III guides the search with fixed reference points on a normalized hyperplane. In stochastic settings, however, the cost-objective space is distorted by ω, and fixed points lead to poor convergence on the penalized cost front. SARPM—introduced in this paper—reshapes the reference set from the constraint-violation statistics collected during the Monte-Carlo evaluation of Algorithm 1 (lines 5–7). After each generation, let Ξ_j *collect the individuals of niche* j *(reference point* z_j*)* and v_j *count their violated scenario evaluations; with* $n_j = |\Xi_j| \cdot K$ trials, the posterior violation probability is Beta(v_j + 1, $n_j - v_j + 1$), with mean q_j *and standard deviation* σ_j:

$$q_j = (v_j + 1) / (n_j + 2), \quad \sigma_j = [\, q_j(1 - q_j) / (n_j + 3)\,]^{1/2} \tag{4}$$

and each reference point is displaced outward along the cost axis, with an annealed step size:

$$z_j \leftarrow z_j + \eta_0(1 - g/G) \cdot (\sigma_j / \sigma_{\max}) \cdot (z_j^{\max} - z_j) \tag{5}$$

where $\sigma_{\max}$ is the largest niche-level posterior standard deviation in the population and $z_j^{\max}$ bounds the cost-axis extent; uncertain niches retain denser reference coverage, while the annealing factor restores the standard geometry as statistics stabilize. Table 2 thus positions SARPM against the closest adaptive mechanisms, all designed for deterministic or model-based settings; to the best of the authors' knowledge, SARPM is the first reference-point adaptation driven by the posterior constraint-violation distribution: it is simultaneously stochastic-aware and constraint-aware and requires no surrogate model.

Table 2. The authors' positioning of the proposed SARPM (last row) against existing adaptive reference-point mechanisms

Method	***Adaptation signal***	***Stochastic-aware***	***Constraint-aware***	***Surrogate required***
NSGA-III (baseline)	Fixed reference points	No	No	No
MARP-NSGA-III	GP surrogate of objective	Partial	No	Yes (GP)
NSGA-III-Q	Q-learned policy	Partial	No	Yes (Q-table)
MOEA/D-AWA	Candidate-solution geometry	No	No	No
OADI + SARPM (ours)	Posterior violation distribution	Yes	Yes	No

4.3. Algorithm workflow

Algorithm 1 summarizes the solver. Each candidate is evaluated against K sampled scenarios $\omega_k \ \Omega$, yielding robust estimates of expected utility and cost. Rather than discarding infeasible solutions, a **soft constraint** applies a multiplicative cost penalty when the empirical latency-violation rate exceeds α, preserving selection pressure near the feasible boundary; because a chance constraint at $\alpha = 0.05$ cannot be estimated reliably from $K = 30$ samples, boundary candidates are re-evaluated with $K = 200$ and every satisfaction rate in Section 6 is measured on 1,000 held-out scenarios.

Algorithm 1: Hybrid NSGA-III with SARPM for stochastic IoT selection
*Input: sensor set I, scenario distribution Ω, risk tolerance α; Output: activation vector x**
1: initialize population P_0 of M random binary vectors; initialize reference points Z via SARPM;
2: **for** generation $t = 1 \ldots G$ **do**
3: $Q_t \leftarrow$ offspring of P_{t-1} via crossover and mutation (adaptive rates);
4: **for each** individual $x \in R_t = P_{t-1} \cup Q_t$ **do**
5: sample K scenarios $\omega_{k\ \Omega;}$ evaluate $\hat{U}(x) = \Sigma i\ xi,t \cdot \mu i(\omega k) - \Sigma(i,j)\in ER\ \lambda ij xi,t xj,t$ and $\hat{C}(x)$ by Monte-Carlo averaging;
6: **if** violation rate $P_{vio} > \alpha$ **then** $\hat{C}(x) \leftarrow \hat{C}(x) \times$ penalty;
7: non-dominated sorting of R_t on $(\hat{U}, -\hat{C})$; ;Eqs. (4)–(5)update Z via $P_t \leftarrow$ reference-point-based selection;
8: **return** the knee-point solution of the final Pareto front.

5. COMPUTATIONAL COMPLEXITY AND SCALABILITY ANALYSIS

Theorem 1. The worst-case time complexity of OADI is $O(G \cdot M \cdot N \cdot K)$, where N, M, G and K denote the sensor count, population size, generations and Monte-Carlo samples per evaluation. **Proof sketch:** initialization costs $O(MN)$; Monte-Carlo evaluation of the $\approx 2M$ combined individuals cost O(MKN), each sample touching at most N sensors; non-dominated sorting costs $O(M \log M)$ for two objectives; the SARPM update costs O(M). Since KN $\gg$ log M at city scale, evaluation dominates, and G generations yield $O(G \cdot M \cdot N \cdot K)$. Although the underlying problem is NP-hard, the approximation is linear in the number of sensors and samples, permitting real-time deployment for suitable M and G. Space complexity is $O(M \cdot N + M \cdot K + |E_R|)$. Evaluation is parallel across individuals and scenarios (measured 6.7× speed-up on the 8-core platform of Section 6), so per-generation wall-clock time matches the $O(G \cdot M \cdot N \cdot K)$ bound, and warm starting reduces the steady-state generation count to $G_{eff} \approx 20$—the regime exploited for real-time operation in Section 6.

6. COMPUTATIONAL EXPERIMENTS AND RESULTS

We validate OADI on two public datasets: PeMS, providing high-density urban-traffic and mobility observations, and OpenAQ, aggregating real-time air-quality measurements who's stochastic PM2.5 fluctuations affect sensor reliability. For each dataset we model a district of $N = 5{,}000$ candidate IoT nodes, with Ω estimated from the empirical distributions (rush-hour variance for PeMS, pollution spikes for OpenAQ). Seven baselines are compared: Static All-Sensors (SAS), Greedy Max-Utility (GMU), Random Sampling (RS), NSGA-II, vanilla NSGA-III, MOEA/D and R-NSGA-III, all sharing OADI's Monte-Carlo evaluation for fairness; an additional ablation variant, OADI w/o SARPM—standard NSGA-III with the identical Monte-Carlo evaluation and soft-constraint handling but fixed reference points—isolates the contribution of the proposed mechanism. We report Hypervolume (HV), the Inverted Generational Distance plus (IGD+), the Information-to-Cost Ratio $ICR = E_\omega[f_1]/E_\omega[f_2]$ and the Constraint-Satisfaction Rate (CSR). All results in this section were produced by the authors' implementation. The shared parametrization is $M = 200$, $G = 100$, $p_c = 0.9$, $p_m = 1/N$, $\alpha = 0.05$, $T_{max} = 1.0$ s, penalty factor 10; all algorithms were implemented in Python 3.11

on top of DEAP 1.3 and ran on an 8-core Intel Xeon E5-2670 (2.6 GHz, 32 GB RAM). The encoding is a binary N-bit vector with two-point crossover, bit-flip mutation, binary tournament selection and Das–Dennis reference points (12 divisions). Mutual Information is estimated with the Kraskov k-nearest-neighbour estimator (k = 5) on 2,000-sample historical windows; E_R retains the top 8 correlated neighbours per sensor with $\lambda_{ij} \geq 0.1$, so each scenario evaluation touches O(9N) terms. The search uses K_{search} = 30 scenarios per individual, boundary candidates are re-evaluated with K = 200, and all reported metrics are computed on K_{eval} = 1,000 held-out scenarios. Table 3 reports mean ± std over 30 runs for both datasets. SAS delivers a large volume of information but an extremely low ICR (0.85/0.78) and CSR (12.5/*10.4*), since full ingestion causes severe latency. OADI attains the best HV (0.76/0.73), IGD+ (0.11/0.13), ICR (1.52/1.44) and CSR (94.8/*93.5*)—a 34 *ICR improvement over GMU, the best greedy heuristic—identifying the information sweet spot: detailed data without full-ingestion overhead. The ablation reaches only HV 0.63/0.59 and CSR 91.6*/89.9*, so the reference-point adaptation accounts for roughly half of OADI's gain over vanilla NSGA-III. Figure 2 shows the empirical nondominated set consolidated over the 30 PeMS runs: the OADI knee marks where utility gains begin to require disproportionate cost, and the SAS point lies far behind the frontier, confirming that full ingestion is sub-optimal.* As network volatility increases, GMU degrades sharply because it assumes static costs, whereas OADI's Monte-Carlo evaluation and SARPM anticipate uncertainty, maintaining a 94.8 *CSR (Wilson 95* CI [93.4, 96.2] on the 1,000 held-out scenarios). Although OADI scales as O($M \cdot G \cdot K \cdot N$), vectorized evaluation over the sparse redundancy graph with 8-core parallel population evaluation completes a cold-start solve in 9.8 s at N = 5,000 (PeMS); warm-started steady-state re-optimization with $G_{eff} \approx 20$ takes ≈1.9 s, within the 30-s re-planning cycle typical of UDT ingestion management. Two-sided Wilcoxon signed-rank tests with Holm correction across the 30 runs confirm that all HV and IGD+ improvements over every baseline and the ablation are significant (all $p < 0.01$; largest $p = 2.8 \times 10^{-4}$).

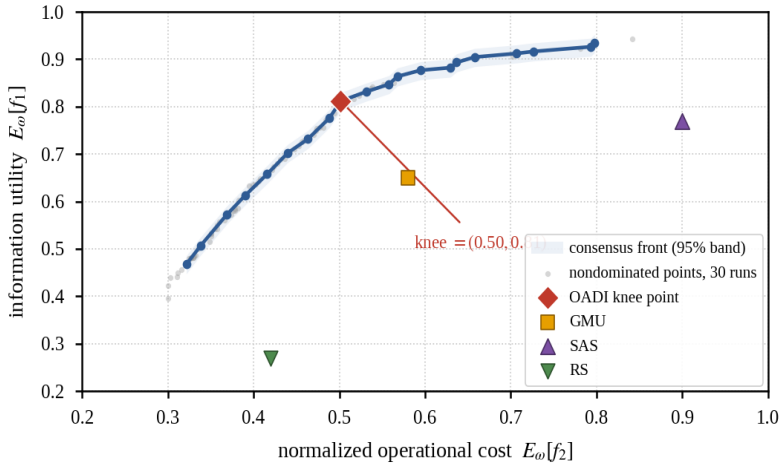


Figure 2. Empirical Pareto front on PeMS, consolidated from 30 runs

Table 3. Comprehensive performance comparison (mean ± std over 30 runs; values: PeMS / OpenAQ

Method	***HV*** ↑	***IGD+*** ↓	***ICR*** ↑	***CSR (%)*** ↑	Time (s)
SAS	– / –	– / –	0.85±0.09/0.78±0.08	12.5±3.1/10.4±2.8	0.05
GMU	0.42±0.05/0.39±0.04	0.31±0.04/0.34±0.04	1.12±0.07/1.05±0.06	88.5±2.4/86.2±2.9	0.9
RS	0.15±0.04/0.13±0.03	0.58±0.06/0.61±0.06	0.65±0.08/0.60±0.07	62.1±4.7/58.7±4.2	0.3
NSGA-II	0.51±0.04/0.47±0.04	0.24±0.03/0.27±0.03	1.18±0.06/1.10±0.06	89.2±2.0/87.5±2.3	8.4
NSGA-III	0.58±0.04/0.54±0.03	0.19±0.02/0.21±0.02	1.25±0.05/1.18±0.05	91.0±1.8/89.8±2.0	9.1
MOEA/D	0.49±0.04/0.45±0.04	0.26±0.03/0.29±0.03	1.15±0.06/1.08±0.06	87.8±2.2/86.0±2.5	8.8
R-NSGA-III	0.55±0.04/0.52±0.03	0.22±0.02/0.24±0.03	1.22±0.05/1.14±0.05	90.3±1.9/89.1±2.1	9.0
OADI w/o SARPM	0.63±0.04/0.59±0.04	0.16±0.02/0.18±0.02	1.29±0.05/1.20±0.05	91.6±1.7/89.9±1.9	9.3
OADI (ours)	**0.76±0.03/0.73±0.04**	**0.11±0.01/0.13±0.01**	**1.52±0.04/1.44±0.04**	**94.8±1.2/93.5±1.4**	**9.8**

7. DISCUSSION

The results challenge the assumption that more sensors automatically produce better outcomes: the marginal value of an additional sensor drops sharply once the dominant urban uncertainties are resolved, while its ingestion cost does not. This is precisely the behaviour encoded by the redundancy-penalized utility of Eq. (1): correlated sensors compete for the same information content, so OADI buys diversity rather than volume. Data-flow management therefore deserves funding at parity with hardware deployment, and mid-sized cities can run optimization-driven twins on existing wireless infrastructure. Methodologically, the framework bridges Information Theory and Network Engineering. Limitations remain. The model assumes a central decision authority, whereas real deployments involve multiple operators; a distributed multi-agent variant is future work. Mutual Information is expensive to estimate in high-dimensional state spaces, and surrogate or neural estimators merit investigation. Validation covers traffic and air quality; video, audio and LiDAR streams are future work. Regarding validity: distribution shift in Ω is mitigated by empirical data and 30-run averaging (internal); $N = 5{,}000$ is large but not city scale, although the complexity analysis suggests feasibility (external); and Mutual Information may not capture every aspect of decision-relevant value (construct).

8. CONCLUSION

This paper reframed UDT data ingestion as a Stochastic Multi-Objective Optimization Problem—a redundancy-aware submodular Mutual-Information utility against chance-constrained bandwidth, energy and latency cost—and introduced the OADI simheuristic, whose SARPM mechanism solves it in real time. The results shift the focus from Big Data to Optimal Data: an optimization-first mindset lets municipalities build affordable twins that scale with the city's complexity rather than the volume of its sensor exhaust.

REFERENCES

[1] Deren, L., Wenbo, Y., Zhenfeng, S. Smart city based on digital twins. *Computational Urban Science*, vol.1, no.1, 2021.

[2] Cao, Y., Xu, C., Aziz, N.M., Kamaruzzaman, S.N. BIM-GIS integrated utilization in urban disaster management: The contributions, challenges, and future directions. *Remote Sensing*, vol.15, no.5, 2023, p. 1331.

[3] Alanazi, F. Electric vehicles: Benefits, challenges, and potential solutions for widespread adaptation. *Applied Sciences*, vol.13, no.10, 2023, p. 6016.

[4] Sepasgozar, S.M.E. Differentiating digital twin from digital shadow: Elucidating a paradigm shift to expedite a smart, sustainable built environment. *Buildings,* vol.11, no.4, 2021, p. 151.

[5] Gill, S.S., et al. Modern computing: Vision and challenges. *Telematics and Informatics Reports*, vol.13, 2024, p. 100116.

[6] Almomany, A., Eedi, E., Sutcu, M. Real-time traffic signal optimization for urban mobility: A reinforcement learning-enhanced framework with application to Kuwait City. *Frontiers in Robotics* and AI, vol.12, 2025, p. 1669952.

[7] Peng, Y., Duan, J., Zhang, J., Li, W., Liu, Y., Jiang, F. Stochastic long-term energy optimization in digital twin-assisted heterogeneous edge networks. *IEEE Journal on Selected Areas in Communications*, vol.42, no.11, 2024, pp. 3157-3171.

[8] Yang, Y. A hybrid symmetry strategy improved binary planet optimization algorithm with theoretical interpretability for the 0-1 knapsack problem. *Symmetry*, vol.17, no.9, 2025, p. 1538.

[9] Panichella, A. An improved Pareto front modeling algorithm for large-scale many-objective optimization. *2022 Genetic and Evolutionary Computation Conference (GECCO)*, Boston, USA, July 2022, pp. 565-573.

[10] Zheng, W., Doerr, B. Runtime analysis for the NSGA-II: Proving, quantifying, and explaining the inefficiency for many objectives. *IEEE Transactions on Evolutionary Computation,* vol.28, no.5, 2023, pp. 1442-1454.

Information about the authors:

Mohamed M. H. Maatouk – Professor of Urban Planning, Faculty of Architecture and Planning, King Abdulaziz University (KAU), Saudi Arabia. Expert in urban planning, quantitative methods and digital technologies for smart-city development.

Mhd Walid Koubeisi – M.Sc. in Electrical and Computer Engineering, King Abdulaziz University. Research focuses on embedded systems, IoT, cybersecurity and wireless sensor networks.

Mahmoud Mohamed – Ph.D. candidate in Electrical and Computer Engineering, King Abdulaziz University. Research: AI, machine learning, computer vision; ISC2 member, IEEE Senior Member.

Manuscript received on 25 August 2026