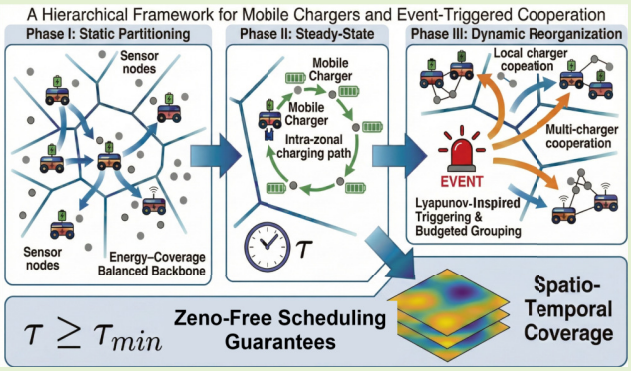
SCOPE: A Hierarchical Event-Triggered Framework for Spatio-Temporal Coverage Maximization in WRSNs With Zeno-Free Scheduling Guarantees

Pei Xu, *Member, IEEE*, Shilong Han, Zijun Wu, Chih-Yung Chang, *Member, IEEE*, and Diptendu Sinha Roy, *Senior Member, IEEE*

***Abstract*— Spatio-temporal coverage quality is a critical performance metric in wireless rechargeable sensor networks (WRSNs). Existing studies predominantly focus on survival-oriented scheduling or centralized optimization, often relying on Boolean sensing abstractions or simplified coverage metrics. Such approaches fail to capture the stochastic nature of physical sensing under the probabilistic sensing model (PSM), overlook the diminishing returns of cooperative coverage, and exhibit limited scalability in large-scale dynamic networks with mobile chargers (MCs). This article proposes spatio-temporal coverage via partitioning and events (SCOPE), a hierarchical event-triggered framework for quality-aware coverage scheduling in WRSNs. SCOPE uses an *N*-aware hierarchy that transforms large-scale sensor-level charging demands into three coordinated stages: energy–coverage density construction and static partitioning, zone-level low-latency scheduling, and Lyapunov-inspired event-triggered cooperative reorganization. The theoretical analysis is restricted to a fixed-cost single-knapsack budgeted PSM grouping kernel with a constant-factor approximation guarantee, while the full multi-MC route-cost implementation is evaluated empirically. Simulation results under the evaluated settings show that SCOPE maintains an average coverage-quality ratio of approximately 0.93 under the default setting. In the dense-network case with $N = 600$, SCOPE improves travel-normalized charging efficiency by approximately 18% over the multicharger cooperative charging mechanism (MC^3), the strongest baseline in this metric, and maintains an online computation time of approximately 25–45 ms/decision for $N = 200$–600. Theoretical analysis further establishes bounded solution quality for the scoped kernel and a scheduling-level Zeno-freeness guarantee for the event-triggered reorganization process.**



***Index Terms*— Event-triggered control, mobile charging scheduling, probabilistic sensing model (PSM), submodular maximization, wireless rechargeable sensor networks (WRSNs).**

Received 18 July 2026; accepted 12 August 2026. Date of publication 24 August 2026; date of current version 2 October 2026. This work was supported in part by the National Natural Science Foundation of China through the General Program under Grant 72271085 and in part by the Industry-Sponsored Research Projects under Grant 902/22050124126 and Grant 902/22050124063. The associate editor coordinating the review of this article and approving it for publication was Dr. Yang Li. *(Corresponding author: Chih-Yung Chang.)*

Pei Xu and Shilong Han are with the School of Artificial Intelligence and Big Data, Hefei University, Hefei 230601, China (e-mail: xupei@hfuu.edu.cn; 18856698650@163.com).

Zijun Wu is with the School of Automotive and Transportation Engineering, Hefei University of Technology, Hefei 230009, China (e-mail: zijun.wu@hfut.edu.cn).

Chih-Yung Chang is with the Department of Computer Science and Information Engineering, Tamkang University, New Taipei City 25137, Taiwan (e-mail: cychang@mail.tku.edu.tw).

Diptendu Sinha Roy is with the Department of Computer Science and Engineering, National Institute of Technology Meghalaya, Shillong 793003, India (e-mail: diptendu.sr@nitm.ac.in).

Digital Object Identifier 10.1109/JSEN.2026.3724657

I. Introduction

WIRELESS rechargeable sensor networks (WRSNs) have emerged as a critical infrastructure for long-term industrial monitoring and surveillance applications, ranging from coverage and connectivity-aware networks [1], [7] to systems employing range-adjustable sensors [16]. Unlike traditional battery-powered networks constrained by finite lifetimes [27], WRSNs leverage mobile chargers (MCs) to replenish energy wirelessly, aiming to extend network operation over long durations [2]. Compared with ambient energy harvesting techniques such as solar charging, which are inherently intermittent and weather-dependent [29], MC-based wireless charging provides a more reliable and controllable energy supply, making it particularly suitable for long-term and quality-critical monitoring applications. In such dynamic environments, the primary performance metric has shifted from merely delaying the first node death to maximizing

1558-1748 © 2026 IEEE. All rights reserved, including rights for text and data mining, and training of artificial intelligence and similar technologies. Personal use is permitted, but republication/redistribution requires IEEE permission.
See https://www.ieee.org/publications/rights/index.html for more information.
Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

the quality of service (QoS), specifically the spatio-temporal coverage quality of targets or areas [22].

Fig. 1 illustrates the considered WRSN scenario, where static sensors are deployed in a monitoring region Ω and a set of MCs provides on-demand energy replenishment. Once the residual energy of a sensor drops below a prescribed threshold, the sensor issues a charging request carrying its location and a service deadline, and an MC is dispatched to perform wireless energy transfer. Without timely replenishment, energy-depleted sensors may create coverage holes. As discussed in [10], prioritizing the recharging of nodes critical to data quality is crucial to repair these voids and maintain surveillance integrity. This tight coupling between sensing quality, charging timeliness, and charger mobility motivates a coordinated scheduling framework that is both performance-aware and designed to avoid scheduling chattering under event-triggered reorganization.

Early studies of MC scheduling focused mainly on the "survival" of the network, employing algorithms to minimize node failure rates or service cost [4], [11]. While effective for keeping nodes alive, these approaches often rely on binary availability assumptions (alive or dead), effectively treating coverage as a Boolean state [4]. However, this binary abstraction fails to capture the stochastic nature of physical sensing, where detection probability decays with distance and environmental noise [22]. Under the more realistic probabilistic sensing model (PSM), coverage exhibits a diminishing returns property (submodularity) [28]. Recent studies have explored dynamic probabilistic sensing capabilities [13], yet traditional time-window-based schedules [9] often ignore the nonlinear coupling between energy and coverage, leading to "coverage saturation" in redundant regions while leaving critical areas poorly monitored.

To address these limitations, recent studies have explored cooperative charging schemes [5], [22] and advanced scheduling techniques. For instance, Wen et al. [14] introduced hierarchical path reduction strategies to optimize charger lifetimes, while Hu et al. [21] proposed collaborative partitioning strategies to coordinate multiple chargers. Nevertheless, applying these concepts to large-scale, event-driven scenarios introduces a *computational dilemma*. Maximizing submodular coverage functions is generally NP-hard. Global optimization approaches, such as multiobjective evolutionary algorithms [23], suffer from high computational latency, making them unsuitable for real-time responses to sudden energy bursts or topological changes. Centralized schedulers also face communication bottlenecks, motivating a hybrid framework that can maintain a persistent "static backbone" for routine operations while possessing the agility to dynamically reorganize resources when local demand surges.

A critical challenge in designing such event-driven scheduling systems is preventing scheduling-level Zeno behavior, i.e., the occurrence of infinitely many scheduling or reorganization triggers within a finite-time interval. In WRSNs, unconstrained event triggering caused by fluctuating energy demands can lead to *computational chattering*, resulting in excessive communication, computation, and charger route reconfiguration. To prevent scheduling chattering, it is necessary to enforce a strictly positive minimum interevent interval, $\tau_{\min}$. By incorporating a Lyapunov-inspired event-trigger mechanism with a dwell-time gate, the proposed framework remains responsive to real-time demand changes without being repeatedly reorganized by transient noise. Accordingly, the guarantee established in this article concerns only scheduling-level Zeno-freeness of the event-triggered reorganization process; it does not constitute the type of closed-loop Lyapunov-stability or domain-of-attraction guarantee studied in classical nonlinear-control analysis [19].

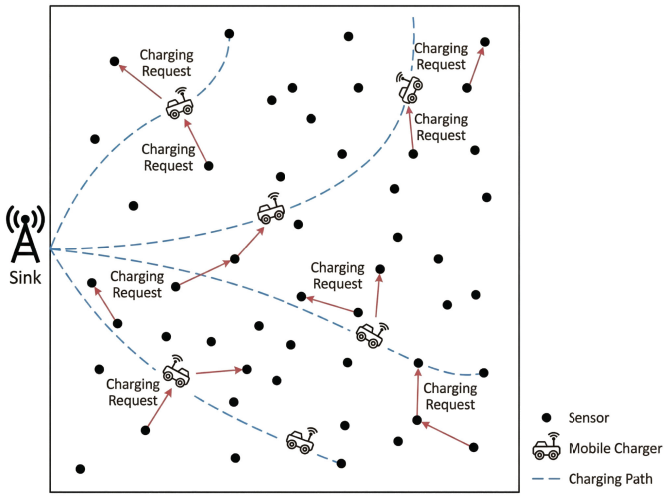


Fig. 1. Illustration of the considered WRSN with MCs. Sensors generate charging requests when their residual energy falls below a threshold, and MCs travel from depots to replenish energy to prevent coverage holes.

These observations motivate us to study the coverage maximization problem in WRSNs under the PSM, with the goal of jointly optimizing surveillance quality and response latency. To this end, we propose a hierarchical, event-triggered scheduling framework. Our approach, termed spatio-temporal coverage via partitioning and events (SCOPE), departs from the traditional "one-to-one" charging paradigm. Instead, it employs a gradient-based static partitioning to balance the inherent coverage potential, followed by an event-triggered mechanism that dynamically groups charging requests. By exploiting the submodularity of the PSM coverage objective, SCOPE derives a scoped approximation guarantee for the fixed-cost single-knapsack budgeted PSM grouping kernel used inside event-triggered cooperation, while the full multi-MC route-dependent scheduling implementation is handled as a practical extension and evaluated empirically.

Relative to the multicharger cooperative charging mechanism (MC^3) [22], the closest PSM-aware cooperative baseline, SCOPE, does not claim novelty from probabilistic sensing or multi-MC cooperation alone. MC^3 dynamically adjusts charging groups according to charger positions and sensor choices to balance workloads and reduce waiting and travel costs. SCOPE differs in three specific respects. First, it introduces an N-aware hierarchy that aggregates the states of all N sensors into an energy–coverage demand field, partitions the network into K service zones, and confines most online decisions to local request sets; cross-zone cooperation is activated only for the triggered coalition. Second, cooperative reorganization is governed by a Lyapunov-inspired sentinel and an explicit dwell-time gate, which yields a strictly pos-

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

itive lower bound on scheduling interevent times. Third, the approximation claim is deliberately restricted to the fixed-cost single-knapsack PSM coverage-selection kernel after feasibility pruning, and no approximation guarantee is claimed for the complete route-dependent multi-MC implementation. Accordingly, PSM-based utility, coverage-aware prioritization, request grouping, route-aware costs, and multi-MC cooperation are integrated and extended components, whereas the N-aware hierarchy, the scheduling-level Zeno-freeness mechanism, and the scoped kernel guarantee constitute the principal new contributions.

In this process, the number of sensors N is not treated merely as a simulation parameter. It determines the active sensor population in the PSM-based coverage objective, the number of residual-energy states and charging requests, and the scale of the pending request set. The novelty of SCOPE in handling N lies in its hierarchical decomposition: sensor-level energy consumption and marginal PSM coverage contribution are first aggregated into an energy–coverage-aware demand density using kernel density estimation (KDE) [20], the resulting N-sensor demand field is partitioned into K charger service zones, and online decisions are then made mainly on local zonal request sets. Cross-zone cooperation is invoked only when the sentinel detects persistent regional stress, thereby avoiding repeated global optimization over all N sensors while preserving coverage quality and real-time responsiveness.

To make the scope of this study explicit, this work addresses three research questions. *Q1:* How can the energy states, request-generation load, and marginal probabilistic coverage contributions of an N-sensor WRSN be converted into a scalable energy–coverage demand representation for K MCs? *Q2:* How can each MC make low-latency local scheduling decisions that jointly account for marginal PSM coverage gain, request urgency, travel cost, deadlines, and energy-safety constraints (ESCs)? *Q3:* When persistent regional stress occurs, how can cooperative reorganization be triggered and accepted so that the budgeted submodular grouping subproblem has a clearly scoped approximation guarantee and the event-triggered scheduler avoids scheduling-level Zeno behavior? These questions define the scope of the proposed framework and motivate the contributions summarized as follows.

Accordingly, the main contributions of this article are summarized as follows. To distinguish the principal novelty from supporting integration, the first three items describe the core architectural and theoretical contributions. The fourth item describes a supporting implementation contribution that combines established coverage-gain, urgency, and travel-cost factors for online execution; it is not claimed as a new general approximation or stability result.

1) *Hierarchical Event-Triggered Framework With Zeno-Free Scheduling Guarantees:* We propose a multiphase scheduling architecture that decouples macroscopic topology initialization from microscopic task scheduling. By integrating gradient-flow-based static partitioning [25] with a Lyapunov-inspired event trigger [17], the framework suppresses unnecessary charger reorganization during routine operation. Moreover, by enforcing a strictly positive interevent lower bound, the proposed trigger is provably free of scheduling-level Zeno behavior at the event-triggered scheduling layer.
2) *N-Aware Hierarchical Decomposition for Scalable Scheduling:* We explicitly account for the sensor population size N by converting sensor-level energy states, charging demands, and marginal PSM coverage contributions into an energy–coverage demand density. Instead of repeatedly optimizing over all N sensors, SCOPE decomposes the large-scale charging problem into K local service-zone scheduling subproblems and rare event-triggered cooperative reorganization. This process-level design is distinct from centralized optimization, fixed region assignment, and proximity-based greedy scheduling and enables scalable coverage-aware operation as the sensor population grows.
3) *Submodularity-Based Cooperative Grouping With a Scoped Approximation Guarantee:* Under the PSM, we formulate the core coverage-selection step of cooperative grouping as a fixed-cost single-knapsack budgeted PSM coverage-selection kernel after feasibility pruning. The grouped coverage objective is evaluated over the union of selected requests, so PSM monotonicity and submodularity are not affected by MC labels. For this scoped kernel, the modified greedy strategy with a max-singleton check provides a constant-factor approximation guarantee. The complete multi-MC assignment with route-dependent costs, deadlines, and gated acceptance is treated as a practical scheduling extension and evaluated empirically.
4) *Low-Latency Local Scheduling as a Supporting Implementation Component:* Motivated by multiobjective optimization principles, we design a lightweight local adaptive scoring mechanism for steady-state scheduling. The mechanism enables each MC to switch online between coverage-aggressive and energy-conservative modes according to real-time battery states and request urgency, providing robust performance under heterogeneous energy consumption while avoiding the high latency of global evolutionary optimization. This component supports the architecture but is not claimed as a separate approximation or stability theorem.

II. Related Works

The scheduling of MCs in WRSNs has evolved from ensuring basic network survival to optimizing high-level service quality. We categorize existing studies into three streams: survival-oriented scheduling, coverage-aware scheduling, and cooperative multi-MC coordination.

A. Survival-Oriented Scheduling

Early research primarily focused on the "survival" problem, aiming to minimize the number of dead nodes or extend network lifetime [4], [11]. Building on these foundations, Xu et al. [11] formulated the problem of maximizing sensor lifetime with minimal service cost. To address the constraints of charger capacity, Wei et al. [9] introduced time-window constraints into the scheduling model.

More recently, research has shifted toward optimizing charging efficiency and overhead within these survival constraints. For instance, Wen et al. [14] proposed the joint perpetual lifetime and path reduction (PLPR) algorithm, which introduces a

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

hierarchical charging architecture using super MCs to recharge regular MCs, thereby maintaining the perpetual lifetime of the charging system while reducing movement paths. Similarly, Kan et al. [1] incorporated network connectivity into the energy replenishment decision, ensuring that the critical relay nodes essential for data backhaul are prioritized to prevent network partitioning.

Limitations and Transition: While these survival-oriented strategies effectively delay node failures, they typically rely on binary availability assumptions (alive or dead), effectively treating coverage as a Boolean state [4]. They ignore the heterogeneous "value" of data generated by different sensors. In surveillance applications, maintaining a node in a redundant area is far less valuable than saving a node in a coverage hole. This agnosticism to sensing quality necessitates a shift toward coverage-aware scheduling.

B. Coverage-Aware and Quality-Oriented Scheduling

Recognizing the shortcomings of survival-first approaches, recent studies have integrated coverage quality into the scheduling objective. The core idea is to prioritize nodes that contribute most to the global sensing utility. Yu et al. [12] proposed the coverage-aware energy replenishment mechanism (CAERM), which quantifies the coverage contribution of each sensor. Similarly, Dande et al. [15] introduced a workload-aware recharge scheduling mechanism to further optimize surveillance quality. To further enhance data quality, Hu et al. [10] and Dande et al. [26] explored mechanisms to prioritize recharging for nodes critical to data quality, thereby dynamically repairing sensing voids and improving spatio-temporal data accuracy.

In the context of the more realistic PSM, research has begun to address the nonlinear relationship between energy and coverage. Zhang et al. [13] proposed a dynamic probabilistic sensing scheme for solar-powered sensors, optimizing target coverage under unstable energy harvesting conditions. Similarly, Xu et al. [16] developed the joint monitoring quality and network lifetime for range-adjustable sensors (JMLA) algorithm, which adjusts the sensing radii of sensors to balance monitoring quality against network lifetime. Yu et al. [22] proposed MC^3, which uses a cooperative mechanism to maximize surveillance quality under PSM.

More recently, Xue et al. [30] proposed MaxCov and the request-grouping-based MaxCov-RG for maximizing target coverage with multiple chargers while balancing coverage performance and computational complexity. Their formulation focuses on selecting and grouping charging requests for target-coverage maximization, whereas SCOPE addresses PSM-based spatio-temporal coverage through an N-aware service-zone hierarchy and event-triggered regional cooperation.

Among these methods, MC^3 [22] is the closest to SCOPE because both adopt PSM-based surveillance quality and cooperative multicharger operation. MC^3 updates charging groups according to the current positions of the MCs and permits sensors to select a suitable charger, with the objectives of workload balancing and reducing waiting and travel costs. CAERM [12] prioritizes coverage contribution; the charging scheduling for maximizing coverage of targets (CSCT) method [5] addresses target-coverage-oriented multicharger scheduling and request grouping, and the collaborative scheme in [21] uses partitioning and batch recharging to improve data quality and lifetime. These studies establish the importance of coverage-aware and cooperative charging, but they do not report the same combination of: 1) an N-aware density–zone–local-set hierarchy; 2) event-triggered regional cooperation with a positive scheduling interevent lower bound; and 3) a formally scoped approximation guarantee for a fixed-cost PSM selection kernel. Conversely, SCOPE does not claim that PSM modeling, multicharger cooperation, dynamic service-region adjustment, request grouping, or coverage-aware prioritization are individually new.

Limitations and Transition: Coverage-aware, probabilistic-sensing, and target-coverage-oriented methods improve charging priorities by accounting for sensing utility or target importance [12], [13], [16], [22], [30]. Request-grouping approaches further improve coordination among multiple chargers [5], [30]. However, these studies mainly optimize coverage or grouping decisions at the sensor/request level and do not jointly provide an N-aware zonal decomposition, low-latency intrazonal scheduling, and persistent-stress-triggered regional cooperation. This motivates a hierarchical framework that preserves coverage awareness while restricting most online decisions to local request sets. First, the coverage function under PSM exhibits submodularity, making the optimization problem NP-hard [12]. Most existing solutions rely on centralized heuristic searches or complex evolutionary algorithms [23], which incur high computational latency and are unsuitable for real-time response. Second, simple greedy strategies based on coverage contribution often fail to guarantee global efficiency due to the "diminishing returns" effect. This creates a need for scheduling algorithms that are both theoretically grounded and computationally efficient for real-time execution.

C. Cooperative Coordination and Optimization Techniques

To handle the complexity of large-scale networks, multi-MC cooperation and advanced optimization techniques have been explored. Lin et al. [6] proposed game-theoretic collaborative charging scheduling (GTCCS) to model the competition and cooperation among multiple MCs. To enhance collaboration efficiency, Hu et al. [21] proposed a collaborative recharging scheme that optimizes charger coordination to enhance data quality and sustainable lifetime.

Recent work has also extended WRSN coordination toward probabilistic and integrated sensing and communication (ISAC)-assisted charging, edge-assisted service, and mobility-constrained scheduling. Qaisar et al. [31] proposed Poised, which models multiple charging-priority attributes probabilistically to balance the loads of multiple mobile charging vehicles and uses ISAC-assisted coordination to reduce travel and charging conflicts. Its probabilistic formulation concerns charging attributes and ISAC coordination, whereas SCOPE evaluates sensing utility directly under the PSM and invokes cross-zone cooperation only under persistent regional stress. Qiu et al. [32] jointly optimized wireless charging and fresh-data retrieval using a speed-adjustable mobile edge

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

device. Their objective is charging–freshness co-optimization, while SCOPE focuses on sustaining PSM-based spatio-temporal coverage with multiple chargers. Zhang et al. [33] jointly designed charger trajectories and charging schedules in obstacle-constrained WRSNs, emphasizing mobility feasibility and obstacle-aware trajectory planning. In comparison, SCOPE uses an N-aware zonal hierarchy for routine scheduling and event-triggered local cooperation, with obstacle-aware travel costs included in the extended evaluation.

Recent data-driven approaches have also been applied to complex scheduling problems. Liang et al. [18] used digital-twin-assisted deep reinforcement learning (DRL) to optimize resource allocation and offloading decisions in autonomous systems. On the theoretical front, Xue et al. [5] emphasized the importance of submodular maximization in coverage tasks, proposing request grouping strategies to approximate the optimal solution.

Limitations and Motivation: These recent studies provide complementary advances in probabilistic and ISAC-assisted charging, edge-enabled charging and data retrieval, obstacle-aware trajectory design, and multicharger cooperation [31], [32], [33]. Nevertheless, they address different primary objectives and do not jointly combine PSM-based spatio-temporal coverage, an N-aware zonal scheduling hierarchy, and persistent-stress-triggered local reorganization. SCOPE is designed to connect these three elements within one interpretable scheduling process. DRL-based methods [18], while promising, operate as closed-box models with limited interpretability of their internal scheduling decisions and high training overhead, which is risky for safety-critical industrial applications. On the other hand, existing cooperative schemes [8], [21] often lack a rigorous mechanism to handle the tradeoff between static-topology persistence and the agility of dynamic response. The remaining gap is therefore not the absence of PSM-based quality metrics or multicharger cooperation, both of which have been established in prior work. Rather, it is the absence of a process-level architecture that simultaneously localizes routine online decisions as N grows, activates cross-zone cooperation only under persistent stress, and states theoretical guarantees at precisely delimited layers. SCOPE addresses this gap through the N-aware hierarchy, the dwell-time-gated event process, and the fixed-cost single-knapsack PSM selection kernel. The complete multi-MC route-cost assignment remains an empirical implementation layer and is not included in the approximation claim.

III. System Model and Problem Formulation

A. Network Architecture and Entities

As shown in Fig. 2, we consider a WRSN deployed over a bounded 2-D monitoring domain $\Omega \subset \mathbb{R}^2$. The system comprises N stationary rechargeable sensors, K MCs initially located at depots, and a sink that collects charging requests and coordinates scheduling decisions. The domain boundaries are known and, where needed, the region is discretized into a grid of points for computation.

Each sensor $s_i \in \mathcal{S}$ is located at static coordinates $q_i \in \Omega$ and powered by a rechargeable battery with capacity $E_s^{\max}$. Sensors consume energy continuously for sensing and data transmission.

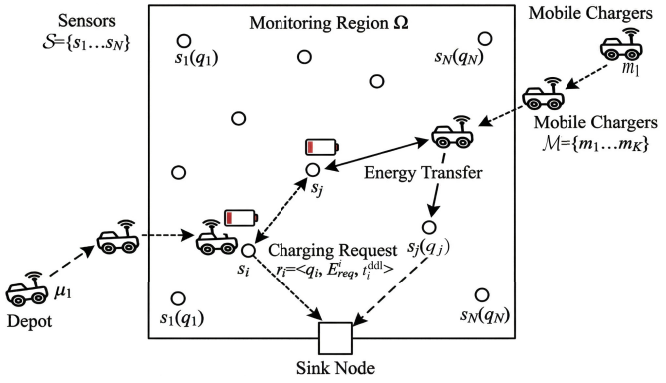


Fig. 2. Network architecture and entities. Static sensors $\mathcal{S} = \{s_1, \ldots, s_N\}$ are deployed in Ω and report charging requests to the sink; MCs $\mathcal{M} = \{m_1, \ldots, m_K\}$ depart from depots $\{\mu_k\}$ to replenish sensors via wireless energy transfer.

The residual energy of sensor s_i is denoted by $E_i(t)$ and measured in joules. When $E_i(t)$ falls below the request threshold E_{th}, the sensor generates a charging request $r_i = \langle q_i, E^i_{\text{req}}, t_i^{\text{ddl}} \rangle$, where E^i_{req} is the requested energy and t_i^{ddl} is the service-completion deadline. Let t_i^{gen} be the request-generation time, and let $\widehat{P}_i(t_i^{\text{gen}}) > 0$ be the current or window-averaged discharge rate, measured in joules per second. This time-local estimate is distinguished from the long-term average $\bar{P}_i$ used in Phase I. The request-generation slack is

$$\ell_i^{\text{gen}} = \frac{E_i(t_i^{\text{gen}}) - E_{\min}}{\widehat{P}_i(t_i^{\text{gen}})}.$$

The numerator is measured in joules and the denominator in joules per second; therefore, ℓ_i^{gen} is measured in seconds. The absolute service-completion deadline is

$$t_i^{\text{ddl}} = t_i^{\text{gen}} + \ell_i^{\text{gen}}.$$

Thus, t_i^{ddl} is the latest time by which the charging service should be completed before the sensor is predicted to reach $E_{\min}$. At a later time t, the remaining slack is

$$\ell_i(t) = \max\{t_i^{\text{ddl}} - t, 0\}.$$

The MCs serve as energy replenishment agents. Each MC $m_k \in \mathcal{M}$ is initially stationed at a specific home depot μ_k (determined by the static partitioning in Phase I, inspired by [25]) and is equipped with a high-capacity battery $B^{\max}$. The MCs travel with a velocity $v \in [0, v_{\max}]$ and can wirelessly transfer energy to sensors upon arrival. In the simulations, we assume reliable state exchange among sensors, MCs, and the sink so that the scheduling logic can be evaluated under controlled conditions. The impact of packet loss, communication delay, and deployment-level coordination constraints is discussed in Section V-D.

Role of the Sensor Population Size N: The number of sensors $N = |\mathcal{S}|$ is explicitly considered throughout the formulation. First, N determines the maximum cardinality of the active sensor set $\mathcal{S}^{\pi}_{active}(t)$ in the spatio-temporal coverage objective, and therefore affects both the achievable PSM coverage utility and the degree of sensing redundancy. Second, each sensor s_i contributes an individual residual-energy state, discharge process, charging request, and marginal PSM coverage contribution, so increasing N changes both the spatial

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

distribution of coverage-critical nodes and the online request-generation load. Third, let

$$\mathcal{R}(t) = \{r_i(t) \mid s_i \in \mathcal{S},\ r_i(t) \ \text{ is pending at time } \ t\}$$

be the pending request set. Since each sensor generates at most one active request at a decision epoch, $|\mathcal{R}(t)| \le N$, implying that N directly controls the worst-case online scheduling scale. SCOPE handles this scaling effect by partitioning the N-sensor network into K charger service zones and making most online decisions locally within each zone.

B. Probabilistic Sensing and Submodularity

Let

$$d_i(x) = \|x - q_i\|$$

be the Euclidean distance, in meters, between sensor s_i at q_i and target point $x \in \Omega$. Unlike the deterministic Boolean disk model, the sensing capability of realistic sensors exhibits stochastic decay. Following [22], the per-sensor detection probability is modeled by the truncated exponential-power PSM

$$c_i(x) = \begin{cases} \exp\left[-\eta_1 d_i(x)^{\eta_2}\right], & 0 \le d_i(x) < R_s \\ 0, & d_i(x) \ge R_s \end{cases} \tag{1}$$

where R_s (m) is the physical sensing cutoff radius, $\eta_2 > 0$ is a dimensionless shape parameter, and $\eta_1 > 0$ has units of $\mathrm{m}^{-\eta_2}$, so that $\eta_1 d_i(x)^{\eta_2}$ is dimensionless. A larger η_1 produces faster attenuation, whereas η_2 controls the curvature of the distance-decay profile. The reported experiments use homogeneous values of R_s, η_1, and η_2 shared by all sensors.

In the present simulator, η_1 and η_2 are effective parameters that aggregate hardware sensitivity and environment-dependent sensing attenuation. No additional additive measurement-noise process is sampled on top of (1). A failed sensor or a sensor sleeping at or below $E_{\min}$ is assigned zero sensing availability and therefore contributes zero detection probability. Conditional on the target location and sensor states, detection events from different contributing sensors are assumed independent.

Coverage quality is defined by the collaborative detection probability. For a subset of active sensors $\mathcal{A} \subseteq \mathcal{S}$, the joint detection probability at point x is given by the complement of the failure probability, expressed as $Q(x, \mathcal{A}) = 1 - \prod_{s_i \in \mathcal{A}}(1 - c_i(x))$. To quantify the macroscopic quality of surveillance, we define the *global coverage utility* $\Phi(\mathcal{A})$ as the weighted integral of detection probabilities over the entire domain

$$\Phi(\mathcal{A}) = \int_{\Omega} w(x) \left(1 - \prod_{s_i \in \mathcal{A}} (1 - c_i(x))\right) dx \tag{2}$$

where $w(x)$ represents the spatial weight or importance density of location x.

A critical theoretical property of this model is *diminishing returns*. Let $\Delta\Phi(s|\mathcal{A}) = \Phi(\mathcal{A} \cup \{s\}) - \Phi(\mathcal{A})$ be the marginal gain of adding sensor s to set $\mathcal{A}$. The coverage utility function $\Phi(\cdot)$ is monotone nondecreasing and *submodular*. That is, for any subsets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{S}$ and a sensor $s \notin \mathcal{B}$, the following inequality holds:

$$\Phi(\mathcal{A} \cup \{s\}) - \Phi(\mathcal{A}) \ge \Phi(\mathcal{B} \cup \{s\}) - \Phi(\mathcal{B}). \tag{3}$$

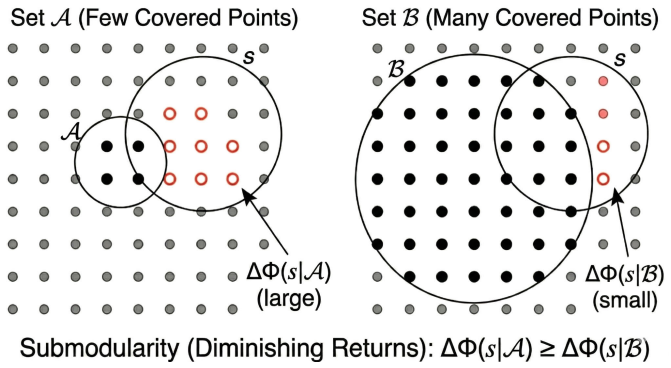


Fig. 3. Submodularity (diminishing returns) intuition of probabilistic coverage. Adding the same sensor s yields a larger marginal gain on a smaller covered set $\mathcal{A}$ than on a larger covered set $\mathcal{B}$, i.e., $\Delta\Phi(s|\mathcal{A}) \ge \Delta\Phi(s|\mathcal{B})$.

Fig. 3 illustrates the diminishing-returns behavior that underlies the submodularity of the coverage-quality objective. Intuitively, when only a few points are already covered by the sensor set $\mathcal{A}$, adding sensor s contributes substantial new coverage, leading to a large marginal gain $\Delta\Phi(s \mid \mathcal{A})$. In contrast, when the active sensor set becomes larger, i.e., $\mathcal{B} \supseteq \mathcal{A}$, more locations within the sensing footprint of s are already covered, and the marginal gain $\Delta\Phi(s \mid \mathcal{B})$ becomes smaller. This property motivates greedy-style selection with provable approximation guarantees.

This submodularity property provides the theoretical basis for the scoped approximation guarantee of the single-budget grouping kernel analyzed in Section IV-C [3].

Model-Scope Clarification: The implementation uses the truncated exponential-power profile in (1), but the SCOPE scheduling logic is not tied to the numerical values of η_1, η_2, or R_s. For any fixed per-sensor detection probabilities $c_i(x) \in [0, 1]$ combined through the same conditional-independence rule, the union-based coverage utility remains monotone and submodular. However, the numerical marginal gains, sensing redundancy, and therefore the selected charging priorities can vary with the decay scale, shape, and cutoff radius. This numerical sensitivity is evaluated in Section V-B.

C. Energy Dynamics and Mobility Model

Fig. 4(a) depicts a typical residual-energy trajectory of a sensor. A charging request is generated once the residual energy falls below E_{th}, and request r_i is associated with a service deadline t_i^{ddl} to prevent energy depletion. Fig. 4(b) illustrates the mobility of an MC traveling from a depot to a requesting sensor with bounded speed. These coupled energy and mobility dynamics impose feasibility constraints on charging timeliness and motivate the scheduling formulation in Section III-D.

The operation of MCs is constrained by a strict energy budget. The total energy consumption of an MC m_k comprises two components: movement energy expended in mobility and the replenishment energy transferred to sensors. Let $\mathcal{T}_k(t)$ be the trajectory of MC m_k up to time t. The movement power is modeled as a convex function of velocity, $P_{\text{move}}(v) = \xi v^2 + \beta$, where ξ and β are mechanical coefficients [9]. In the simulations, ξ and β are calibrated with respect to the default fixed-speed motion at $v_0 = 2$ m/s. Let c_{move}^0 be the default

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

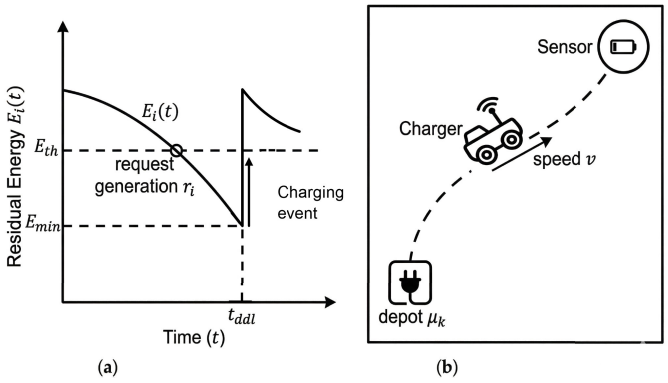


Fig. 4. Energy dynamics and charger mobility. (a) Sensor generates a request when $E_i(t)$ falls below E_{th}. The deadline t_i^{ddl} is the latest service-completion time before the sensor is predicted to reach E_{min}. (b) MC departs from depot μ_k and travels with speed v to serve a requesting sensor.

constant movement-energy coefficient used in M0, so that the reference movement power is $P_0 = c^0_{\text{move}} v_0$. To preserve consistency between M0 and M1 at the nominal speed, the coefficients are selected to satisfy

$$P_{\text{move}}(v_0) = P_0.$$

A second calibration point is imposed at $v_{\max}$ using a high-speed penalty factor γ_v

$$P_{\text{move}}(v_{\max}) = \gamma_v P_0.$$

Thus,

$$\xi = \frac{(\gamma_v - 1)P_0}{v_{\max}^2 - v_0^2}, \qquad \beta = P_0 - \xi v_0^2.$$

In the reported simulations, $c^0_{\text{move}} = 1.0$ J/m, $v_{\max} = 3.0$ m/s, and $\gamma_v = 1.8$, and the same calibrated ξ and β are used for all algorithms in M1–M3.

For compact notation, we define the operational travel distance between two locations q_a and q_b as

$$d(q_a, q_b) = \begin{cases} \|q_a - q_b\|, & \text{in obstacle-free scenarios} \\ d_{\mathcal{F}}(q_a, q_b), & \text{in obstacle-constrained scenarios} \end{cases}$$

where $d_{\mathcal{F}}(q_a, q_b)$ denotes the shortest feasible path distance in the free space $\mathcal{F}$. All travel-time, movement-energy, and operational travel-distance calculations in SCOPE use $d(\cdot,\cdot)$. The movement energy from q_a to q_b is computed as

$$E_{\text{move}}(q_a, q_b) = \int_0^{d(q_a,q_b)/v} P_{\text{move}}(v)\, d\tau$$

which reduces to $c^0_{\text{move}} d(q_a, q_b)$ under the default fixed-speed M0 setting.

For an MC m_k to accept a sequence of tasks, it must satisfy the *ESC* [11]. This constraint ensures that at any decision epoch, the MC retains sufficient residual energy to complete the assigned tasks and return to its designated depot μ_k or a charging station. Formally, let $\mathcal{R}_{plan}$ be the set of planned requests. The residual battery $B_k^{\text{res}}(t)$ must satisfy

$$B_k^{\text{res}}(t) - \underbrace{\sum_{r_j \in \mathcal{R}_{plan}} \left(\int_{path} P_{\text{move}}(v)\, d\tau + E_{\text{req}}^j \right)}_{\text{Service Cost}}$$

$$-E_{return}(q_{last}, \mu_k) \geq B_{safe} \tag{4}$$

where E_{return} is the estimated movement energy required to return to the depot from the last task location q_{last}, and B_{safe} is a reserved buffer. The subset of state space satisfying (4) forms the feasible operating manifold for the scheduling algorithm.

D. Problem Statement

The overarching objective is to maintain the highest possible level of spatio-temporal coverage quality. However, due to the finite number of MCs and their energy constraints, not all charging requests can be served immediately, leading to potential sensor outages and coverage holes. This necessitates a scheduling policy that trades off between maximizing coverage gain and minimizing service latency.

We formulate the dynamic scheduling problem as a *constrained spatio-temporal coverage-maximization problem*. Let $\pi = \{\pi_1, \ldots, \pi_K\}$ be a joint scheduling policy for the K MCs, where $\pi_k(t)$ specifies the action (move, charge, or wait) of MC m_k at time t. The set of active sensors at time t under policy π, denoted as $\mathcal{S}^{\pi}_{active}(t)$, varies as sensors deplete energy or get replenished. Let $\mathcal{R}_{\text{served}}$ denote the set of charging requests completed under policy π. The objective is to maximize the time-averaged global coverage utility over an infinite horizon

$$\max_{\pi \in \Pi} \lim_{T \to \infty} \frac{1}{T} \int_0^T \Phi\left(\mathcal{S}^{\pi}_{active}(t)\right) dt \tag{5}$$

subject to the following operational constraints:

$$\text{C1 (Deadline): } t_j^{\text{arr}} + \tau_j^{\text{svc}} \leq t_j^{\text{ddl}} \quad \forall r_j \in \mathcal{R}_{\text{served}} \tag{6a}$$

$$\text{C2 (Energy): (4) holds for all } m_k \in \mathcal{M} \text{ at all } t \tag{6b}$$

$$\text{C3 (Flow): } \frac{dp_k(t)}{dt} = v_k(t), \quad \|v_k(t)\| \leq v_{\max}. \tag{6c}$$

Constraint C1 is completion-based: the predicted arrival time plus the service duration must not exceed the absolute deadline. All three quantities are measured in seconds. Constraint C2 imposes the energy safety guarantee for the MCs. Constraint C3 governs the kinematic limits of the MCs. This problem formulation combines characteristics of the vehicle-routing problem with time windows (VRPTW) and the submodular function-maximization problem, both of which are NP-hard [12]. The computational difficulty increases with N, because a larger sensor population increases the number of possible active-sensor states, potential charging requests, and feasible service sequences under deadline and MC energy constraints. Consequently, solving (5) exactly in real time is computationally intractable. To address this, we propose a hierarchical, event-triggered framework in Section IV, which decomposes the problem into static partitioning, steady-state independent scheduling, and dynamic regional reorganization.

IV. Proposed Algorithm

Fig. 5 summarizes SCOPE's three-phase architecture. Phase I constructs a static energy–coverage-balanced partition and assigns charger depots. Phase II performs low-latency intra-zonal scheduling on local request sets. Phase III is activated only when persistent regional stress is detected, enabling local cooperative reorganization without repeated full-network optimization.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

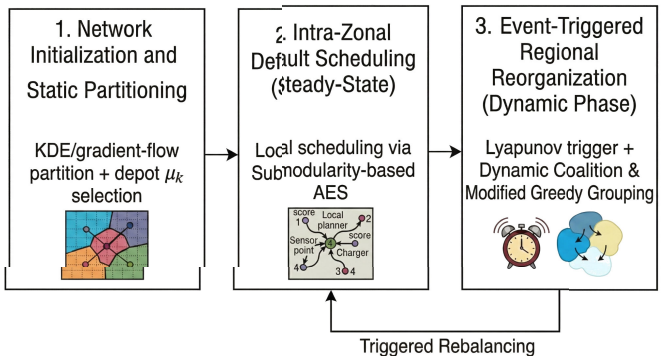


Fig. 5. Overall framework of the proposed SCOPE algorithm. Phase I performs network initialization and static partitioning with depot selection. Phase II executes intrazonal default scheduling in steady state. Phase III performs event-triggered regional reorganization for load rebalancing.

Sensor-Number-Aware Novelty of SCOPE: The novelty of SCOPE is not that it simply varies N in simulations, but that it handles the resulting N-sensor scheduling complexity through a closed-loop hierarchical process. Phase I compresses sensor-level energy–coverage information from all N sensors into a continuous nominal demand density and partitions it into K service zones. Let

$$\mathcal{S}_k = \{s_i \in \mathcal{S} : q_i \in \mathcal{V}_k\}, \quad N_k = |\mathcal{S}_k|, \quad \sum_{k=1}^{K} N_k = N.$$

Phase II operates on the local request set

$$\mathcal{R}_k(t) = \{r_i(t) \in \mathcal{R}(t) \mid s_i \in \mathcal{S}_k\}$$

with

$$\sum_{k=1}^{K} |\mathcal{R}_k(t)| \le N$$

rather than repeatedly searching over the entire network. Under a balanced partition, the expected online scheduling load of each MC scales mainly with the local sensor population N_k, approximately N/K, instead of the full N. Phase III is invoked only when the sentinel detects persistent regional stress, and the cooperative grouping is restricted to the triggered coalition rather than all sensors. This N-aware decomposition distinguishes SCOPE from centralized global optimization, fixed assignment, and distance-greedy scheduling and enables scalable coverage-aware responsiveness.

A. Phase I: Network Initialization and Static Partitioning

The primary objective of Phase I is to establish a robust "static backbone" for the WRSN by partitioning the continuous monitoring domain $\Omega \subset \mathbb{R}^2$ into K disjoint subregions $\mathcal{V} = \{\mathcal{V}_1, \ldots, \mathcal{V}_K\}$. To make the static backbone insensitive to instantaneous network fluctuations, we adopt a *time-scale separation principle*. Rather than responding to transient energy deficits, this phase focuses on the macroscopic energy demand distribution to initialize a topology that equilibrates the "potential coverage gain" across the MCs.

1) Energy–Coverage Density Construction: We consider a set of sensors $\mathcal{S} = \{s_1, \ldots, s_N\}$ stochastically deployed in Ω. To capture the realistic sensing quality described in [22], we adopt the truncated exponential-power PSM $c_i(x)$ specified in (1). The collective coverage quality at point x provided by the active set $\mathcal{S}_{active}$ is $Q(x, \mathcal{S}_{active}) = 1 - \prod_{s_i \in \mathcal{S}_{active}} (1 - c_i(x))$.

To quantify the topological importance of each sensor, we define the *coverage criticality weight* ω_i^{cov} as the marginal integral gain over the domain Ω

$$\omega_i^{cov} = \Phi(\mathcal{S}) - \Phi(\mathcal{S} \setminus \{s_i\}) = \int_{\Omega} w(x) \left[Q(x, \mathcal{S}) - Q(x, \mathcal{S} \setminus \{s_i\}) \right] dx. \tag{7}$$

This metric ensures that nodes covering sparse or unique areas are prioritized.

Definition 1 (Energy–Coverage Density): Crucially, to construct a time-invariant reference field, let $\bar{P}_i$ be the long-term average power-consumption rate of sensor s_i. We first define the dimensionless energy–coverage weight

$$a_i = \alpha \frac{\bar{P}_i}{\max_j \bar{P}_j} + (1 - \alpha) \frac{\omega_i^{cov}}{\max_j \omega_j^{cov}}$$

$$\widetilde{\rho}_{\mathrm{EC}}(x) = \sum_{i=1}^{N} a_i K_{\mathbf{H}}(x - q_i)$$

$$\rho_{\mathrm{EC}}(x) = \frac{\widetilde{\rho}_{\mathrm{EC}}(x)}{\int_{\Omega} \widetilde{\rho}_{\mathrm{EC}}(z)\, dz}, \quad \int_{\Omega} \rho_{\mathrm{EC}}(x)\, dx = 1. \tag{8}$$

Here, $\widetilde{\rho}_{\mathrm{EC}}(x)$ is the unnormalized kernel field and $\rho_{\mathrm{EC}}(x)$ is its normalized form, provided that $\int_{\Omega} \widetilde{\rho}_{\mathrm{EC}}(z)\, dz > 0$. The density represents long-term energy demand weighted by each sensor's marginal PSM coverage contribution, while the sum over all N sensors explicitly embeds the sensor population in the Phase-I backbone.

For a fixed partition $\mathcal{V} = \{\mathcal{V}_1, \ldots, \mathcal{V}_K\}$, define the energy–coverage mass of zone k by

$$\Lambda_k^{\mathrm{EC}}(\mathcal{V}) = \int_{\mathcal{V}_k} \rho_{\mathrm{EC}}(x)\, dx, \quad \sum_{k=1}^{K} \Lambda_k^{\mathrm{EC}}(\mathcal{V}) = 1.$$

These Phase-I aggregates can therefore be balanced across the K service zones. They compress the N-sensor energy demand and marginal coverage information into K zone loads without discarding the underlying individual service requests. This nominal zone mass is distinct from the online request workload $\Lambda_k(t)$ used in Phase III.

At initialization, the nominal energy–coverage density used in the subsequent Phase-I validity monitoring and sentinel construction is defined as

$$\rho_0(x) \equiv \rho_{\mathrm{EC}}(x).$$

After a validated persistent distribution shift, $\rho_0(x)$ is updated according to (17).

The continuous density is used solely for spatial aggregation in Phase I; it is not itself a submodular set function. Submodularity applies to the discrete PSM utility over active sensors, or equivalently over feasible service requests. Once the zones are fixed, they only group requests, whereas the utility continues to be evaluated over the underlying union of individual requests.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

In all Phase-I KDE computations, $K_{\mathbf{H}}(\cdot)$ is implemented as an isotropic Gaussian kernel

$$K_{\mathbf{H}}(z) = \frac{1}{2\pi h_0^2} \exp\left(-\frac{\|z\|^2}{2h_0^2}\right) \tag{9}$$

where $\mathbf{H} = h_0^2 I_2$. In the simulations, we set $h_0 = 40$ m, consistent with the KDE bandwidth used for the online request-density field. The same kernel definition is also used when constructing the online empirical density $\widehat{\rho}_t(x)$ for Phase-I validity monitoring. Thus, $\rho_{\mathrm{EC}}(x)$ provides a normalized, time-smoothed energy–coverage input for the partitioning algorithm rather than a stationary probability model of future coverage gains.

2) Variational Partitioning via Gradient Flow: The partitioning problem is formulated as a *variational coverage–energy balancing problem*. We seek a partition that minimizes a *generalized distortion-energy functional* $\mathcal{J}(\mathcal{V}, \boldsymbol{\mu})$, ensuring both geometric compactness and load balance relative to the energy–coverage density $\rho_{\mathrm{EC}}(x)$. Let $\mathcal{C}_k$ be the nominal service capacity of MC m_k. The optimization is defined as

$$\min_{\mathcal{V},\boldsymbol{\mu}} \mathcal{J}(\mathcal{V}, \boldsymbol{\mu}) = \sum_{k=1}^{K} \int_{\mathcal{V}_k} \Bigg(\underbrace{\|x - \mu_k\|^2}_{\text{Compactness}} + \zeta \underbrace{\left(\int_{\mathcal{V}_k} \rho_{\mathrm{EC}}(z) dz - \mathcal{C}_k \right)^2}_{\text{Load Balance}} \Bigg) \times \rho_{\mathrm{EC}}(x) dx \tag{10}$$

where $\boldsymbol{\mu} = \{\mu_1, \ldots, \mu_K\}$ denotes the set of generating centroids. The nonnegative coefficient ζ controls the penalty imposed on energy–coverage load imbalance.

To solve (10), we employ a *continuous-time gradient-descent flow* to iteratively evolve the generators. Since $\rho_{\mathrm{EC}}(x)$ is treated as time-invariant during Phase-I optimization, the gradient-flow iteration is run until a stationary partition is reached in implementation. The control law governing the motion of the anchors is derived as

$$\dot{\mu}_k(\tau) = -\kappa_{prop} \left(\mu_k(\tau) - \frac{\int_{\mathcal{V}_k} x \rho_{\mathrm{EC}}(x) dx}{\Lambda_k^{\mathrm{EC}}(\mathcal{V})} \right) - \kappa_{int} \left(\Lambda_k^{\mathrm{EC}}(\mathcal{V}) - \mathcal{C}_k \right) \nabla_{\mu_k} \Lambda_k^{\mathrm{EC}} \tag{11}$$

where τ represents the fictitious optimization time, and $\Lambda_k^{\mathrm{EC}}(\mathcal{V}) = \int_{\mathcal{V}_k} \rho_{\mathrm{EC}}(x) dx$ is the aggregate energy–coverage load. When the stopping criterion is satisfied, the final centroid μ_k^* is designated as the static *home depot* for MC m_k.

Remark 1 (Time-Scale Separation): The validity of this static partitioning relies on the assumption that the statistical distribution of energy demand [captured by $\rho_{\mathrm{EC}}(x)$] changes slowly compared with the daily discharge cycles of sensors. This separation allows the derived backbone $\mathcal{V}^*$ to remain optimal for long periods, while instantaneous deviations are handled by the dynamic scheduling in subsequent phases.

3) Validity Monitor and Rare Phase-I Relearning: The time-scale separation assumption in Phase I is valid only when the long-term statistical demand density remains close to the nominal energy–coverage density field $\rho_0(x)$. In practice, this assumption may become invalid under persistent consumption changes, sensor redeployment, large-scale sensor failures, or long-term environmental changes. We introduce a lightweight Phase-I validity monitor to detect persistent deviations from the nominal energy–coverage density field.

Let $\widehat{\rho}_t(x)$ be the online empirical demand density estimated over a sliding monitoring window of length T_{mon}. It is constructed using the same KDE structure as in Phase I but with the recently observed average sensor consumption and request statistics

$$\widehat{\rho}_t(x) = \sum_{i=1}^{N} \left[\alpha \frac{\widehat{P}_i(t)}{\max_j \widehat{P}_j(t)} + (1-\alpha) \frac{\omega_i^{\mathrm{cov}}(t)}{\max_j \omega_j^{\mathrm{cov}}(t)} \right] \times K_{\mathbf{H}}(x - q_i) \tag{12}$$

where $\widehat{P}_i(t)$ is the window-averaged consumption rate of sensor s_i, and $\omega_i^{\mathrm{cov}}(t)$ is the current PSM marginal coverage contribution. Both $\rho_0(x)$ and $\widehat{\rho}_t(x)$ are normalized so that their integrals over Ω equal one.

We quantify the distribution drift by the total-variation distance

$$D_\rho(t) = \frac{1}{2} \int_\Omega \left| \widehat{\rho}_t(x) - \rho_0(x) \right| dx. \tag{13}$$

In addition, we monitor the induced regional load imbalance under the current static partition

$$D_\Lambda(t) = \max_k \frac{\left| \int_{\mathcal{V}_k} \widehat{\rho}_t(x)\, dx - \mathcal{C}_k \right|}{\mathcal{C}_k + \varepsilon_C}. \tag{14}$$

Finally, to capture topological changes caused by redeployment or failures, we define

$$D_F(t) = \frac{|\mathcal{S}_{\mathrm{fail}}(t) \cup \mathcal{S}_{\mathrm{redeploy}}(t)|}{N} \tag{15}$$

where $\mathcal{S}_{\mathrm{fail}}(t)$ and $\mathcal{S}_{\mathrm{redeploy}}(t)$ denote sensors that have failed or have been redeployed since the last Phase-I update.

The time-scale separation assumption is considered to be broken when

$$D_\rho(t) > D_\rho^{\mathrm{th}} \quad \text{or} \quad D_\Lambda(t) > D_\Lambda^{\mathrm{th}} \quad \text{or} \quad D_F(t) > D_F^{\mathrm{th}} \tag{16}$$

for M_{hold} consecutive validity checks. Each check is evaluated every Δ_{chk} seconds using the most recent sliding monitoring window of length T_{mon}. Thus, T_{mon} controls the temporal support used to estimate $\widehat{\rho}_t(x)$, while Δ_{chk} controls the frequency at which the break condition is tested. This sliding-window implementation prevents short-lived bursts from triggering costly global repartitioning, while still allowing persistent distribution shifts to be detected within a few check intervals. Such transient events are still handled by Phase III.

When the break condition in (16) is satisfied, SCOPE performs a rare Phase-I relearning step. The nominal density is updated by

$$\rho_0^+(x) = (1 - \beta_\rho) \rho_0(x) + \beta_\rho \widehat{\rho}_t(x) \tag{17}$$

where $\beta_\rho \in (0, 1]$ controls the update rate. The variational partitioning in (10) and (11) is then rerun using the current centroids as warm-starts

$$\mu_k^+(0) = \mu_k^{\mathrm{old}} \tag{18}$$

which substantially reduces the computation and avoids unnecessary large-scale region reshaping.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

The new partition is accepted only if it provides a sufficient reduction of the distortion functional

$$\mathcal{J}\left(\mathcal{V}^{\text{old}}, \boldsymbol{\mu}^{\text{old}}; \rho_0^+\right) - \mathcal{J}\left(\mathcal{V}^{\text{new}}, \boldsymbol{\mu}^{\text{new}}; \rho_0^+\right) \geq \eta_J \, \mathcal{J}\left(\mathcal{V}^{\text{old}}, \boldsymbol{\mu}^{\text{old}}; \rho_0^+\right) \tag{19}$$

and the time since the last Phase-I relearning satisfies

$$t - t_{\text{last}}^P \geq T_{\min}^P. \tag{20}$$

The dwell-time $T_{\min}^P$ ensures that Phase-I relearning remains rare and cannot degenerate into frequent global reclustering. During relearning, ongoing charging tasks are not interrupted; only newly generated requests are assigned according to the updated partition. MC depots are updated virtually first and physically relocated only when the corresponding MC becomes idle. This design preserves scheduling continuity while allowing the static backbone to adapt to persistent distribution shifts. The default values of T_{mon}, Δ_{chk}, M_{hold}, and the drift thresholds are summarized in Section V-A.

B. Phase II: Intrazonal Default Scheduling (Steady State)

Following the macroscopic partitioning in Phase I, the system transitions to the steady-state execution phase. In this phase, each MC m_k operates as an autonomous agent within its designated zone $\mathcal{V}_k^*$. To maintain consistency with the global coverage maximization objective in Phase III, the local scheduling policy is formulated not as a simple weighted sum, but as a *maximization of spatio-temporal efficiency density*. This decentralized mechanism involves feasibility pruning, adaptive efficiency scoring, and lookahead trajectory optimization.

From the perspective of sensor-number scalability, Phase II avoids a full N-sensor search and instead schedules only the locally pending requests $\mathcal{R}_k(t)$, which makes the online decision load depend on the zonal request density rather than the entire network size.

1) Feasibility Masking and Manifold Pruning: Let $\mathcal{R}_k(t)$ be the set of pending requests in zone k. To ensure operational safety, we first apply a *feasibility mask* $\Psi_k : \mathcal{R}_k \to \{0, 1\}$ to prune the candidate space. An action a_j (servicing request r_j) is admissible if and only if it satisfies the strict ESC and the real-time deadline constraint (RDC) defined in [9]. Considering the nonlinear energy consumption model $P_{\text{move}}(v)$ and the operational distance $d(\cdot,\cdot)$ defined in Section III-C, the feasible action manifold is rigorously defined as

$$\tilde{\mathcal{A}}_k(t) = \Big\{ a_j \in \mathcal{R}_k(t) \;\Big|\; \mathbb{I}\left(B_k^{\text{res}}(t) - E_{\text{move}}\left(p_k(t), q_j\right) - E_{\text{req}}^j \geq B_{safe}\right) \cdot \mathbb{I}\left(t_{\text{now}} + \frac{d\left(p_k(t), q_j\right)}{v} + \tau_{\text{svc}}^j \leq t_{\text{ddl}}^j\right) = 1 \Big\}. \tag{21}$$

This hard filtering restricts the subsequent optimization to a physically realizable subspace, preventing any safety violations.

2) Adaptive Marginal Efficiency Scoring: To align with the submodular coverage maximization logic of Phase III, we define the "value" of a task not by its absolute priority, but by its **marginal efficiency density**—the ratio of coverage gain to resource consumption. Let $\Delta\Phi_{\text{PSM}}(a_j)$ be the marginal coverage gain of task j calculated under the PSM [22], and let $\Delta D_{\text{travel}}(a_j) = d(p_k(t), q_j)$ be the operational travel-distance feature, measured in meters, that serves as the travel-cost component in the local scheduling score. Inspired by the adaptive operator logic in [23], we introduce state-dependent exponents to dynamically balance the tradeoff between coverage greediness and energy conservation. The scoring function $R(s_t, a_j)$ is formulated as

$$R\left(s_t, a_j\right) = \frac{\left(\Delta\Phi_{\text{PSM}}\left(a_j\right)\right)^{\lambda\left(\bar{B}_k(t)\right)}}{\left(\Delta D_{\text{travel}}\left(a_j\right)\right)^{\gamma\left(\bar{B}_k(t)\right)}} \times \exp\left(-\beta_{time} \frac{t_{\text{ddl}}^j - t_{\text{arr}}^j}{T_{span}}\right). \tag{22}$$

Here, $E_{\text{move}}(q_a, q_b)$ is the physical movement energy, measured in joules, whereas $\Delta D_{\text{travel}}(a_j)$ is the operational travel-distance feature, measured in meters, used in the local scheduling score. These quantities serve different roles and are not used interchangeably. The normalized MC residual-battery state is

$$\bar{B}_k(t) = \frac{B_k^{\text{res}}(t)}{B^{\max}}.$$

The exponents $\lambda(\bar{B}_k(t))$ and $\gamma(\bar{B}_k(t))$ are designed as adaptive control functions using exponential decay and saturation to smooth the mode transition. Specifically, they are defined as

$$\lambda\left(\bar{B}_k(t)\right) = 1 + \alpha_\lambda\left(1 - e^{-\beta_\lambda \bar{B}_k(t)}\right) \tag{23}$$

$$\gamma\left(\bar{B}_k(t)\right) = 1 + \alpha_\gamma e^{-\beta_\gamma \bar{B}_k(t)} \tag{24}$$

where $\alpha_\lambda, \alpha_\gamma > 0$ determine the magnitude of the weighting adjustment, and $\beta_\lambda, \beta_\gamma > 0$ control the sensitivity of the adaptation. As energy depletes ($\bar{B}_k(t) \to 0$), $\gamma(\bar{B}_k(t))$ approaches $1+\alpha_\gamma$ to significantly penalize long-distance travel; conversely, when the buffer is sufficiently charged, $\gamma(\bar{B}_k(t))$ decays toward unity while $\lambda(\bar{B}_k(t))$ increases, allowing the coverage gain to dominate the efficiency score. This adaptive formulation enables the MC to autonomously shift its behavioral mode from "coverage-aggressive" to "energy-conservative" without central intervention.

3) Lookahead Trajectory Optimization via Truncated-Horizon Dynamic Programming: To avoid myopic decisions common in pure greedy strategies, we employ a *truncated-horizon dynamic programming* approach (equivalent to an n-path search). The decision problem is to find the optimal next-hop a^* that maximizes the cumulative efficiency density over a finite lookahead horizon n. The value function $V_n(s_t, a_j)$ is defined recursively by the Bellman equation [8]

$$V_n\left(s_t, a_j\right) = R\left(s_t, a_j\right) + \beta_{disc} \cdot \max_{a' \in \tilde{\mathcal{A}}_k(t+1)} V_{n-1}\left(s_{t+1}, a'\right) \tag{25}$$

with the terminal condition $V_0(\cdot) = 0$. The optimal action is selected as $a^* = \arg\max_{a \in \tilde{\mathcal{A}}_k(t)} V_n(s_t, a)$. If $\tilde{\mathcal{A}}_k(t) = \emptyset$, the MC enters *parking mode* and retreats to the home depot μ_k^* derived in Phase I. Since μ_k^* is the centroid of the coverage potential field, hovering there minimizes the expected response latency for future stochastic arrivals.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

C. Phase III: Event-Triggered Regional Reorganization (Dynamic Phase)

While the autonomous scheduling in Phase II ensures scalability under steady-state conditions, the system must remain responsive to substantial spatio-temporal disturbances, such as sudden demand bursts or topological disruptions. Phase III introduces an *event-triggered cross-region adjustment (CRA)* mechanism. To address demand imbalance without incurring the high overhead of global reclustering, we reformulate the dynamic reorganization as a **cooperative request grouping (CRG)** problem governed by the **SCOPE** framework and a Lyapunov-inspired sentinel.

The novelty of Phase III is that cross-zone cooperation is not performed periodically or globally over all N sensors; it is triggered only by persistent sentinel-detected stress and is limited to the affected regional coalition.

1) Lyapunov-Inspired Sentinel: We construct a Lyapunov-inspired sentinel $W_k(t)$ to quantify the operational stress of zone k. Let $\Lambda_k(t)$ be the online pending-request workload defined below, which is distinct from the Phase-I mass $\Lambda_k^{\mathrm{EC}}(\mathcal{V})$. Let $\mathcal{C}_k$ be the nominal service capacity.

Here, $E_{\mathrm{res}}^k(t)$ denotes the arithmetic mean residual energy of the sensors assigned to zone $\mathcal{V}_k$, i.e.,

$$E_{\mathrm{res}}^k(t) = \frac{1}{|\mathcal{S}_k|} \sum_{s_i \in \mathcal{S}_k} E_i(t)$$

where $\mathcal{S}_k$ is the set of sensors assigned to zone $\mathcal{V}_k$. This zone-level sensor statistic is distinct from the residual battery $B_k^{\mathrm{res}}(t)$ of MC m_k.

The sentinel is defined as the Hessian-weighted norm of load saturation combined with energy criticality

$$W_k(t) = \frac{1}{2} \left\| \frac{\Lambda_k(t)}{\mathcal{C}_k} - 1 \right\|_{H_k}^2 + \frac{\kappa}{2} \exp\left(\frac{E_{\mathrm{crit}} - E_{\mathrm{res}}^k(t)}{\sigma_e} \right). \quad (26)$$

The quantities E_{crit}, $E_{\mathrm{res}}^k(t)$, and σ_e are all measured in joules; hence, the exponential argument in (26) is dimensionless.

To make the sentinel fully reproducible, we explicitly define the Hessian-weighted norm in (26). Let

$$e_k(t) = \frac{\Lambda_k(t)}{\mathcal{C}_k} - 1 \quad (27)$$

be the normalized workload saturation of zone k. Since $e_k(t)$ is a scalar load-imbalance variable, the Hessian-weighted norm in (26) is implemented as a curvature-weighted scalar norm

$$\|e_k(t)\|_{H_k}^2 \triangleq H_k e_k^2(t) \quad (28)$$

where $H_k \geq 1$ is a zone-dependent Hessian weight computed from the nominal demand density field $\rho_0(x)$ constructed in Phase I.

Specifically, we first normalize the KDE-based nominal demand field as

$$\sum_{x_g \in \mathcal{G}} \rho_0(x_g) \Delta A = 1 \quad (29)$$

where $\mathcal{G}$ is the density-evaluation grid and ΔA is the grid-cell area. For a grid point $x_g = (u, v)$, the Hessian matrix of $\rho_0(x)$ is approximated by second-order finite differences

$$\rho_{xx}(u,v) = \frac{\rho_0(u+\Delta_g, v) - 2\rho_0(u,v) + \rho_0(u-\Delta_g, v)}{\Delta_g^2} \quad (30)$$

$$\rho_{yy}(u,v) = \frac{\rho_0(u, v+\Delta_g) - 2\rho_0(u,v) + \rho_0(u, v-\Delta_g)}{\Delta_g^2} \quad (31)$$

$$\begin{aligned} \rho_{xy}(u,v) = \frac{1}{4\Delta_g^2} \Big[& \rho_0(u+\Delta_g, v+\Delta_g) \\ & - \rho_0(u+\Delta_g, v-\Delta_g) \\ & - \rho_0(u-\Delta_g, v+\Delta_g) \\ & + \rho_0(u-\Delta_g, v-\Delta_g) \Big]. \end{aligned} \quad (32)$$

Boundary grid points are handled using one-sided finite differences. The local curvature magnitude is then measured by the Frobenius norm

$$h(x_g) = \left\| \nabla^2 \rho_0(x_g) \right\|_F = \sqrt{\rho_{xx}^2(x_g) + 2\rho_{xy}^2(x_g) + \rho_{yy}^2(x_g)}. \quad (33)$$

For each zone $\mathcal{V}_k$, the average curvature is

$$\bar{h}_k = \frac{1}{|\mathcal{G}_k|} \sum_{x_g \in \mathcal{G}_k} h(x_g), \quad \mathcal{G}_k = \mathcal{G} \cap \mathcal{V}_k. \quad (34)$$

The final Hessian weight is normalized as

$$H_k = \mathrm{clip}_{[1, H_{\max}]} \left(1 + \chi \frac{\bar{h}_k}{\frac{1}{K} \sum_{\ell=1}^{K} \bar{h}_\ell + \varepsilon_h} \right) \quad (35)$$

where χ controls the influence of geometric curvature, ε_h avoids division by zero, and $H_{\max}$ prevents the isolated grid-level noise from producing excessive trigger sensitivity.

The online workload intensity $\Lambda_k(t)$ is computed from the current pending request set rather than from instantaneous sensor energy alone. Let $\mathcal{R}(t)$ be the set of pending charging requests at time t. We construct an online request-density field

$$\rho_t(x) = \sum_{r_j \in \mathcal{R}(t)} \varpi_j(t) K_{h_\rho}(x - q_j) \quad (36)$$

where $K_{h_\rho}(\cdot)$ is a Gaussian kernel with bandwidth h_ρ, and $\varpi_j(t)$ is the urgency-weighted coverage importance of request r_j

$$\begin{aligned} \varpi_j(t) &= \Delta\Phi_j \left(1 + \frac{\eta_u}{\ell_j(t) + \varepsilon_\ell} \right) \\ \ell_j(t) &= \max\left\{ t_j^{\mathrm{ddl}} - t, 0 \right\}. \end{aligned} \quad (37)$$

Here, $\Delta\Phi_j$ is the PSM marginal coverage contribution of sensor s_j. The remaining slack $\ell_j(t)$, urgency time-scale coefficient η_u, and regularizer ε_ℓ are measured in seconds, so the urgency ratio is dimensionless. The aggregate workload of zone k is then

$$\Lambda_k(t) = \int_{\mathcal{V}_k} \rho_t(x)\, dx \approx \sum_{x_g \in \mathcal{G}_k} \rho_t(x_g) \Delta A. \quad (38)$$

The nominal service capacity is inherited from the Phase-I nominal partitioning

$$\mathcal{C}_k = \int_{\mathcal{V}_k} \rho_0(x)\, dx \approx \sum_{x_g \in \mathcal{G}_k} \rho_0(x_g) \Delta A. \quad (39)$$

For homogeneous MCs and a balanced Phase-I partition, $\mathcal{C}_k$ is close to $1/K$.

In implementation, the drift of $W_k(t)$ is estimated by a smoothed finite difference. Let

$$\bar{W}_k(t) = \frac{1}{L_W} \sum_{\ell=0}^{L_W - 1} W_k(t - \ell \Delta t) \quad (40)$$

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

be the moving average of $W_k(t)$ over a window of length L_W. The numerical drift estimate is

$$\widehat{\dot{W}}_k(t) = \frac{\bar{W}_k(t) - \bar{W}_k(t - \Delta_W)}{\Delta_W}. \tag{41}$$

A reorganization event is generated only when both the stress-growth condition and the dwell-time condition are satisfied

$$\mathcal{E}_{\text{trig}} : \widehat{\dot{W}}_k(t) > \delta e^{-\lambda_{\text{decay}}(t - t_{\text{last}})} \quad \text{and} \quad t - t_{\text{last}} \geq \tau_{\min}. \tag{42}$$

The first condition detects persistent growth in regional workload stress, while the second condition explicitly prevents scheduling chattering. The minimum dwell time is chosen as

$$\tau_{\min} = \max\left\{\tau_{\text{ctrl}}, \frac{\epsilon_{\text{th}}}{L_{\text{rate}}}\right\} \tag{43}$$

where τ_{ctrl} is the minimum control-cycle duration required for information exchange and task redistribution. Theorem 2 assumes that L_{rate} is a conservative upper bound on the actual sentinel drift $|\dot{W}_k(t)|$ during continuous operation. In the discrete-time implementation, its numerical value is calibrated from the maximum observed magnitude of the smoothed finite-difference drift $|\widehat{\dot{W}}_k(t)|$, multiplied by a safety factor.

a) Response to bursty or localized stress: A sudden regional increase in sensor discharge produces additional charging requests and reduces their remaining slack. Through (36)–(38), these requests increase the urgency-weighted request density $\rho_t(x)$ and the regional workload $\Lambda_k(t)$. Persistent growth of the resulting sentinel $W_k(t)$ satisfies the trigger condition in (42), after which the sink forms a local coalition containing the stressed zone and its participating neighbors. The coalition collects only the corresponding pending requests, performs feasibility pruning and cooperative grouping through Algorithm 1, evaluates the gated acceptance conditions in (47), and then either commits the new MC task sets or retains the current routes if the tentative reallocation is rejected. Thus, localized stress is handled by local event-triggered route reallocation rather than periodic full-network replanning.

Let t_{det} be the decision epoch at which the trigger condition in (42) is detected. Let t_{com} be the time at which state collection, coalition formation, cooperative grouping, gate evaluation, assignment dissemination, and the commit/acknowledgment exchange have been completed, so that each participating MC either installs the accepted task set or explicitly retains its current route. We define the scheduling reaction time as

$$T_{\text{react}} = t_{\text{com}} - t_{\text{det}}.$$

This quantity measures trigger-to-route-commitment latency. The reaction clock starts at trigger detection and therefore excludes the preceding stress accumulation and moving-average detection delay. It also excludes subsequent physical travel to a reassigned sensor and the charging-service duration.

2) CRG via Scope Strategy: Upon triggering, the initiator m_k establishes a dynamic coalition $\mathcal{N}_k$ with its neighbors. Unlike the pairwise bargaining in prior works, we adopt a holistic *request-grouping* approach to redistribute the union of pending tasks $\mathcal{T}_{union} = \mathcal{R}_k \cup \bigcup_{j \in \mathcal{N}_k} \mathcal{R}_j$.

Algorithm 1 SCOPE CRG Strategy (Modified Greedy)

Require: Union of pending tasks $\mathcal{T}_{union}$; Candidate MCs $\mathcal{N}_k \cup \{m_k\}$; Initial available budget B_i for each MC m_i;
Ensure: Task partition $\mathbf{G} = \{\mathcal{G}_i\}_{i=1}^{|\mathcal{N}_k|+1}$;
Initialize: $\mathcal{G}_i \leftarrow \emptyset$, $B_i^{\text{rem}} \leftarrow B_i$ for all MCs; $\mathcal{T}_{rem} \leftarrow \mathcal{T}_{union}$;
// Phase 1: Density-Based Greedy Construction
while $\mathcal{T}_{rem} \neq \emptyset$ **do**
  Step 1: Calculate Marginal Density
  for each task $r \in \mathcal{T}_{rem}$ and MC m_i **do**
    Let $\mathcal{G}_{\text{all}} \leftarrow \bigcup_h \mathcal{G}_h$. Compute global gain $\Delta\Phi = \Phi_{\text{PSM}}(\mathcal{G}_{\text{all}} \cup \{r\}) - \Phi_{\text{PSM}}(\mathcal{G}_{\text{all}})$;
    Compute cost $\Delta\mathcal{C}_{\text{cost}}(\mathcal{G}_i, r) = \mathcal{C}_{\text{cost}}(\mathcal{G}_i \cup \{r\}) - \mathcal{C}_{\text{cost}}(\mathcal{G}_i)$;
    Set density $\rho_{i,r} \leftarrow \Delta\Phi / (\Delta\mathcal{C}_{\text{cost}}(\mathcal{G}_i, r) + \varepsilon_c)$;
  end for
  Step 2: Greedy Selection
  Find $(r^*, m_{best}) = \arg\max_{r \in \mathcal{T}_{rem}, m_i} \rho_{i,r}$;
  Let $\Delta C^* \leftarrow \Delta\mathcal{C}_{\text{cost}}(\mathcal{G}_{m_{best}}, r^*)$;
  Step 3: Feasibility Check
  if $\Delta C^* \leq B_{m_{best}}^{\text{rem}}$ **then**
    $\mathcal{G}_{m_{best}} \leftarrow \mathcal{G}_{m_{best}} \cup \{r^*\}$; {assign task}
    $B_{m_{best}}^{\text{rem}} \leftarrow B_{m_{best}}^{\text{rem}} - \Delta C^*$; {update budget}
    $\mathcal{T}_{rem} \leftarrow \mathcal{T}_{rem} \setminus \{r^*\}$;
  else
    $\mathcal{T}_{rem} \leftarrow \mathcal{T}_{rem} \setminus \{r^*\}$; {Discard infeasible task}
  end if
end while
// Phase 2: Max-singleton check
Step 4: Comparison with Single Best Task
Find the best feasible singleton $r_{\max}$:
$r_{\max} = \arg \max_{\substack{r \in \mathcal{T}_{union} \\ \exists m_i : \mathcal{C}_{\text{cost}}(\{r\}) \leq B_i}} \Phi_{\text{PSM}}(\{r\})$;
Let $\mathbf{G}_{single}$ be the partition containing only $r_{\max}$;
Step 5: Final Selection
if $\Phi_{\text{PSM}}(\{r_{\max}\}) > \Phi_{\text{PSM}}(\bigcup_h \mathcal{G}_h)$ **then**
  return $\mathbf{G} \leftarrow \mathbf{G}_{single}$;
else
  return $\mathbf{G}$; {Return Greedy Solution}
end if

a) Scope of monotonicity and submodularity: The monotonicity and submodularity used in the theoretical analysis refer to the PSM coverage utility itself, not to a cost-penalized routing objective. Travel distance, route-dependent service cost, deadline feasibility, and MC energy safety are handled through feasibility pruning, nonnegative budget constraints, and the gated acceptance rule. If a path-cost penalty were subtracted from the coverage objective, monotonicity would not be guaranteed in general. Therefore, the approximation guarantee is stated only for the budgeted PSM coverage-selection kernel with fixed feasible candidates and fixed nonnegative costs.

b) Scope of the approximation guarantee: The implemented CRG in Phase III is a multi-MC assignment problem with route-dependent incremental service costs, per-MC energy budgets, deadline and energy-safety pruning, and the gated acceptance rule in (47). These practical components are necessary for executable MC scheduling, but they are not covered by the constant-factor bound in Theorem 1. The formal approximation guarantee is therefore stated for the

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

fixed-cost single-knapsack budgeted PSM coverage-selection kernel obtained after fixing the candidate set and using fixed nonnegative service costs. The full assignment layer is written as follows:

$$\begin{aligned} \max_{\{\mathcal{G}_i\}} \quad & \mathcal{J}_{\text{cov}}\left(\{\mathcal{G}_i\}\right) = \Phi_{\text{PSM}}\left(\bigcup_{i\in\{k\}\cup\mathcal{N}_k} \mathcal{G}_i\right) \\ \text{s.t.} \quad & \bigcup_i \mathcal{G}_i \subseteq \mathcal{T}_{union}, \quad \mathcal{G}_i \cap \mathcal{G}_j = \emptyset \ \ (\forall i \neq j) \\ & \mathcal{C}_{\text{cost}}\left(\mathcal{G}_i\right) \leq B_i \quad \forall i \in \{k\} \cup \mathcal{N}_k \end{aligned} \tag{44}$$

where $\Phi_{\text{PSM}}(\cdot)$ is the PSM coverage utility evaluated over the union of served requests, $\mathcal{C}_{\text{cost}}(\mathcal{G}_i)$ represents the energy cost of serving the group assigned to MC m_i, and B_i denotes the available energy budget of MC m_i (e.g., remaining battery minus safety threshold). Using the union avoids double-counting coverage contributions across MC groups.

Lemma 1 (Submodularity Under Fixed-Zone Aggregation): Fix a decision epoch and a partition $\mathcal{V} = \{\mathcal{V}_1, \ldots, \mathcal{V}_K\}$. For each zone k, let $\mathcal{R}_k^{\text{feas}}$ be the fixed feasible request set obtained after deadline and energy-safety pruning, and assign every request to exactly one zone. For any zone-index set $\mathcal{A} \subseteq \{1, \ldots, K\}$, define

$$\mathcal{R}(\mathcal{A}) = \bigcup_{k\in\mathcal{A}} \mathcal{R}_k^{\text{feas}}, \qquad F_{\text{zone}}(\mathcal{A}) = \Phi_{\text{PSM}}(\mathcal{R}(\mathcal{A})).$$

If $w(x) \geq 0$ and $0 \leq c_r(x) \leq 1$ for every feasible request r, then F_{zone} is monotone nondecreasing and submodular on the fixed set of zones.

Proof: Let $\mathcal{A} \subseteq \mathcal{B} \subseteq \{1, \ldots, K\}$ and let $z \notin \mathcal{B}$. The marginal contribution of the fixed request bundle in zone z is

$$F_{\text{zone}}\left(\mathcal{A} \cup \{z\}\right) - F_{\text{zone}}\left(\mathcal{A}\right) = \int_\Omega w\left(x\right) \prod_{u\in\mathcal{R}(\mathcal{A})} \left(1 - c_u\left(x\right)\right) \times \left[1 - \prod_{r\in\mathcal{R}_z^{\text{feas}}} \left(1 - c_r\left(x\right)\right)\right] dx.$$

The bracketed term is nonnegative, so every such marginal is nonnegative and F_{zone} is monotone. Moreover, $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{R}(\mathcal{B})$ and $0 \leq c_u(x) \leq 1$ imply

$$\prod_{u\in\mathcal{R}(\mathcal{B})} \left(1 - c_u(x)\right) \leq \prod_{u\in\mathcal{R}(\mathcal{A})} \left(1 - c_u(x)\right).$$

Multiplying both sides by the nonnegative zone-bundle factor and $w(x)$, and then integrating over Ω, gives

$$F_{\text{zone}}(\mathcal{A} \cup \{z\}) - F_{\text{zone}}(\mathcal{A}) \geq F_{\text{zone}}(\mathcal{B} \cup \{z\}) - F_{\text{zone}}(\mathcal{B})$$

which is the diminishing-returns condition. The result concerns the PSM utility of the underlying request union; it does not treat the continuous Phase-I density as a submodular function. □

Proposition 1 (Approximation-Factor Preservation Under a Fixed-Zone Coalition): Fix a zone coalition $\mathcal{C} \subseteq \{1, \ldots, K\}$ and the coalition-restricted candidate ground set

$$\mathcal{R}_{\text{cand}}(\mathcal{C}) = \bigcup_{k\in\mathcal{C}} \mathcal{R}_k^{\text{feas}}.$$

Apply the fixed-cost single-knapsack PSM kernel defined below in (45), with $\mathcal{V}_{\text{cand}} = \mathcal{R}_{\text{cand}}(\mathcal{C})$, and the same modified density-greedy procedure with a max-singleton check as in Theorem 1. If $G_{\mathcal{C}}^{\text{alg}}$ is the returned solution and $G_{\mathcal{C}}^{\star}$ is the optimum over this same coalition-restricted candidate set, then

$$\Phi_{\text{PSM}}(G_{\mathcal{C}}^{\text{alg}}) \geq \alpha_{\text{ker}}\, \Phi_{\text{PSM}}(G_{\mathcal{C}}^{\star})$$

where α_{ker} is the approximation factor already stated in Theorem 1. Fixed-zone aggregation introduces no additional multiplicative factor because it only partitions and restricts the candidate ground set; it does not change the PSM objective, fixed request costs, single budget, or selection rule. This guarantee is relative only to the optimum over the same fixed coalition. It excludes coalition formation, route-dependent multi-MC assignment, deadline dynamics, and gated acceptance in the complete implementation.

Proposition 2 (Trigger-to-Route-Commitment Bound Under One-Cycle Coordination): Suppose that a sink-assisted implementation guarantees that state collection, coalition formation, cooperative grouping, gate evaluation, assignment dissemination, and commit/acknowledgment exchange are completed within one fixed coordination interval τ_{ctrl}. Then, for every detected trigger

$$T_{\text{react}} = t_{\text{com}} - t_{\text{det}} \leq \tau_{\text{ctrl}}.$$

In the reported configuration, $\tau_{\text{ctrl}} = 50$ s is the nominal coordination interval used in the simulations. Hence, under this one-cycle coordination abstraction, an accepted reallocation is committed, or a rejected tentative reallocation is resolved without changing the current routes, within the nominal bound of 50 s after trigger detection.

Proof: Under the stated one-cycle assumption, all control stages defining t_{com} finish no later than $t_{\text{det}} + \tau_{\text{ctrl}}$. Therefore, $t_{\text{com}} \leq t_{\text{det}} + \tau_{\text{ctrl}}$, and subtracting t_{det} gives the claimed bound. □

c) Scope of physical execution: After route commitment, the physical time required to reach and serve a reassigned request remains task-dependent. For an assignment of request r_i to MC m_k, feasibility continues to be governed by the deadline condition in (21), which at the commitment time requires

$$t_{\text{com}} + \frac{d(p_k(t_{\text{com}}), q_i)}{v} + \tau_{\text{svc}}^i \leq t_i^{\text{ddl}}.$$

Accordingly, Proposition 2 bounds only the scheduling decision under the stated coordination assumption, while the existing deadline and energy-safety checks govern physical execution. Packet-level exchange and acknowledgment timing are abstracted rather than explicitly simulated, so 50 s is not a measured end-to-end hardware latency. A hardware implementation must replace the nominal interval with a verified worst-case communication-and-computation latency, including any additional coordination cycles.

For the theoretical guarantee, we consider the following single-budget coverage-selection kernel:

$$\begin{aligned} \max_{G\subseteq\mathcal{V}_{\text{cand}}} \quad & \Phi_{\text{PSM}}\left(G\right) \\ \text{s.t.} \quad & \sum_{r\in G} \bar{C}_r \leq B \end{aligned} \tag{45}$$

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

where $\mathcal{V}_{\text{cand}}$ is the feasible candidate request set after energy-safety and deadline pruning, $\bar{C}_r \geq 0$ is the fixed nonnegative service cost of candidate request r, and B is a single available energy budget. Path-dependent assignment, multi-MC budget coupling, and gated acceptance are used in the executable Phase-III implementation but are not included in this scoped kernel.

To guide the cooperative redistribution, we use a density-based modified greedy strategy with a max-singleton check. The formal constant-factor bound is established for the single-knapsack coverage-selection kernel in (45). In the full implementation, the same density principle is embedded in a multi-MC assignment procedure with feasibility pruning, route-dependent incremental costs, and gated acceptance. The algorithm iteratively assigns the unassigned task r^* with the highest *marginal gain density* to the MC m_{best} that maximizes

$$\begin{aligned}(r^*, m_{best}) &= \arg \max_{r \in \mathcal{T}_{rem},\, m_i} \frac{\Delta\Phi_{\text{PSM}}(r \mid \mathcal{G}_{\text{all}})}{\Delta\mathcal{C}_{\text{cost}}(\mathcal{G}_i, r) + \varepsilon_c} \\ \Delta\Phi_{\text{PSM}}(r \mid \mathcal{G}_{\text{all}}) &= \Phi_{\text{PSM}}(\mathcal{G}_{\text{all}} \cup \{r\}) - \Phi_{\text{PSM}}(\mathcal{G}_{\text{all}}) \\ \mathcal{G}_{\text{all}} &= \bigcup_h \mathcal{G}_h \end{aligned} \tag{46}$$

where $\mathcal{G}_{\text{all}}$ is the union of all tasks currently selected across participating MCs, so the numerator measures the global marginal PSM coverage gain without double-counting coverage across MC groups. The denominator remains MC-specific: $\Delta\mathcal{C}_{\text{cost}}(\mathcal{G}_i, r)$ is the incremental group cost of adding task r to the current task set of MC m_i. The sets $\mathcal{G}_i$ store the tasks assigned to each MC, B_i denotes the initial available budget of m_i, and B_i^{rem} tracks the remaining budget during Algorithm 1. In particular, $\mathcal{G}_{m_{best}}$ and $B_{m_{best}}^{\text{rem}}$ represent the current task set and remaining budget of the selected MC m_{best}. If adding r^* to m_{best} satisfies $\Delta\mathcal{C}_{\text{cost}}(\mathcal{G}_{m_{best}}, r^*) \leq B_{m_{best}}^{\text{rem}}$, it is included in $\mathcal{G}_{m_{best}}$, and $B_{m_{best}}^{\text{rem}}$ is decreased by the same incremental cost. The small constant ε_c prevents numerical instability when the incremental cost is close to zero.

This density-based greedy strategy, together with the max-singleton check (Step 4 in Algorithm 1), guides the cooperative redistribution. The formal bound in Theorem 1 applies to the scoped single-knapsack kernel after candidate-set fixing and feasibility pruning, whereas the implemented multi-MC route-cost assignment is treated as a practical extension and evaluated empirically.

3) Gated Acceptance and Theoretical Analysis: The derived grouping strategy $\{\mathcal{G}_i\}^*$ is committed only if it satisfies the following physical and scheduling-level gatekeeping criteria:

$$\Delta\mathcal{U}_{\text{net}} \geq \epsilon_U, \qquad \Delta\mathcal{C}_{\text{cost}} \leq \Theta, \qquad \Delta W_k \geq \epsilon_{\text{th}}. \tag{47}$$

Here, $\Delta\mathcal{U}_{\text{net}}$ denotes the net PSM coverage gain produced by the reassignment, and ϵ_U is the minimum normalized net-utility-gain threshold. The implementation uses $\epsilon_U = 0.5$, preserving the existing numerical utility gate. The quantity $\Delta\mathcal{C}_{\text{cost}}$ is the additional travel/charging cost introduced by the cooperative grouping. The third condition is the sentinel stress-decrease gate, where

$$\Delta W_k = W_k\left(t_\ell^-\right) - W_k\left(t_\ell^+\right). \tag{48}$$

The tentative reassignment is evaluated before commitment by recomputing the Lyapunov-inspired sentinel W_k under the updated workload distribution. Thus, an event is accepted only when it provides a nonnegligible coverage benefit, respects the cost-surge budget, and guarantees a measurable decrease in the sentinel stress metric. The stress-decrease gate in (48) is the decrease condition used in Theorem 2 to establish scheduling-level Zeno-freeness.

In all obstacle-constrained scenarios, the incremental cost $\Delta\mathcal{C}_{\text{cost}}(m_i, r)$ and the group cost $\mathcal{C}_{\text{cost}}(\mathcal{G}_i)$ are computed using the operational distance $d(\cdot,\cdot)$ defined in Section III-C.

The event-trigger threshold δ acts on the sentinel drift, whereas ϵ_{th} acts on the accepted sentinel decrease. To ensure dimensional consistency, the minimum accepted stress decrease is parameterized as

$$\epsilon_{\text{th}} = \eta_\epsilon \delta \tag{49}$$

where $\eta_\epsilon > 0$ is a threshold-conversion time scale. In the reported simulations, $\delta = 1.0$ normalized stress/s and $\eta_\epsilon = 0.5$ s, yielding $\epsilon_{\text{th}} = 0.5$ normalized stress.

The cost-surge threshold is defined as an aggregate budget-normalized cap over the participating MCs

$$\Theta = \sum_{m_i \in \{m_k\} \cup \mathcal{N}_k} \eta_\Theta B_i \tag{50}$$

where B_i is the available energy budget of MC m_i, as defined for the grouping problem above, and $\eta_\Theta \in (0, 1)$ controls the maximum additional cost allowed by one accepted cooperative reorganization. In all simulations, we set $\eta_\Theta = 0.2$, meaning that an accepted cooperative reorganization cannot introduce additional travel/charging overhead exceeding 20% of the aggregate available budget of the participating MCs. This gated design ensures that every accepted event produces a nonnegligible coverage benefit and sentinel stress decrease while keeping the extra physical cost bounded, which is required by the proof of scheduling-level Zeno-freeness in Theorem 2.

Although the union-based grouped coverage objective remains monotone and submodular, the full implementation also contains multiple MC-specific budgets, route-dependent incremental costs, deadline pruning, and gated acceptance. Hence, the constant-factor approximation guarantee is restricted to the fixed-cost single-knapsack PSM coverage-selection kernel stated below, while the complete multi-MC route-cost reorganization implementation is treated as a practical extension evaluated empirically.

By Proposition 1, applying the scoped kernel within a fixed triggered-zone coalition preserves the approximation factor stated in Theorem 1 relative to the optimum over the same coalition-restricted candidate set. No additional approximation factor is introduced by the fixed-zone labels.

Theorem 1 (Fixed-Cost PSM Kernel Guarantee): Consider the single-knapsack budgeted PSM coverage-selection kernel in (45), where $\Phi_{\text{PSM}}(\cdot)$ is monotone nondecreasing and submodular, each candidate request has a fixed nonnegative cost $\bar{C}_r$, and B is a single energy budget. The modified density-greedy procedure with a max-singleton check returns a solution G_{ker} satisfying

$$\Phi_{\text{PSM}}(G_{\text{ker}}) \geq \left(1 - 1/\sqrt{e}\right) \Phi_{\text{PSM}}\left(G_{\text{ker}}^\star\right) \tag{51}$$

where $G_{\text{ker}}^\star$ is the optimal solution of the scoped kernel. This guarantee applies only to the single-knapsack selection kernel. It is not claimed for the full implementation, which includes

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

multi-MC route-cost assignment, feasibility pruning, deadline constraints, and gated acceptance.

Proof: The PSM coverage utility remains monotone because adding a feasible request cannot decrease the collaborative detection probability. It is submodular because the marginal detection gain decreases as the already selected active coverage set grows. Formally, for any subsets $\mathcal{A} \subseteq \mathcal{B} \subseteq \mathcal{V}_{\text{cand}}$ and a candidate request $s \notin \mathcal{B}$

$$\Phi_{\text{PSM}}(\mathcal{A} \cup \{s\}) - \Phi_{\text{PSM}}(\mathcal{A}) \geq \Phi_{\text{PSM}}(\mathcal{B} \cup \{s\}) - \Phi_{\text{PSM}}(\mathcal{B}). \quad (52)$$

Since the service cost $\bar{C}_r$ is imposed only through the single budget constraint and is not subtracted from the objective, the objective of (45) remains monotone and submodular. The scoped kernel therefore matches the budgeted coverage-selection setting addressed by the modified density-greedy algorithm with a max-singleton check. Following the standard budgeted maximum-coverage result of Khuller et al. [24], the returned solution satisfies the stated $(1 - 1/(e)^{1/2})$ bound with respect to the optimum of (45). The full multi-MC implementation is a practical route-cost assignment extension and is evaluated empirically rather than covered by this theorem. □

Remark 2 (Physical Zeno Behavior Versus Scheduling-Level Zeno Behavior): It is crucial to distinguish between physical Zeno behavior and *scheduling-level Zeno behavior*. In WRSNs, physical Zeno behavior, i.e., infinitely many physical movements in finite time, is naturally precluded by the bounded velocity $v_{\max}$ of MCs. However, the primary threat in our event-triggered framework is *scheduling-level Zeno behavior* or chattering, where the controller repeatedly generates reallocation events within a finite interval due to small state fluctuations. Such chattering can cause excessive communication, computation, and route reconfiguration overhead. Therefore, the dwell-time gate in (42) explicitly enforces a positive interevent lower bound and guarantees feasibility of communication and computation cycles.

Scope of the Scheduling-Level Zeno-Freeness Guarantee: Accordingly, the guarantee established in this article concerns only scheduling-level Zeno-freeness of the event-triggered reorganization process; it does not constitute a Lyapunov-stability guarantee for the complete WRSN closed-loop system, including all sensor energy states, charger battery states, request queues, and physical trajectories. The result instead ensures that, under the bounded-drift assumption and the enforced dwell-time gate, the scheduler cannot generate infinitely many reorganization events within a finite-time interval.

Theorem 2 (Absence of Scheduling-Level Zeno Behavior): Consider the event-triggered scheduling process induced by the stress-growth condition and the dwell-time gate in (42). Suppose that the drift of the sentinel is bounded during continuous operation, i.e.,

$$\left|\dot{W}_k(t)\right| \leq L_{\text{rate}} < \infty$$

and that every accepted reorganization satisfies the sentinel stress-decrease gate $\Delta W_k = W_k(t_\ell^-) - W_k(t_\ell^+) \geq \epsilon_{\text{th}} > 0$ in (47). Then the event-triggered scheduling process is free of scheduling-level Zeno behavior. In particular, the interevent interval is lower bounded by

$$t_{\ell+1} - t_\ell \geq \tau_{\min} = \max\left\{\tau_{\text{ctrl}}, \frac{\epsilon_{\text{th}}}{L_{\text{rate}}}\right\} > 0. \quad (53)$$

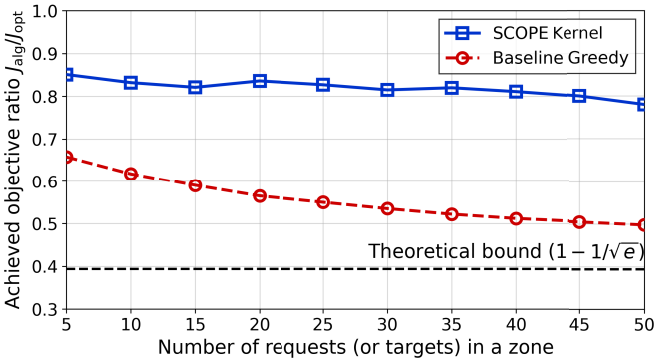


Fig. 6. Empirical approximation ratio for the fixed-cost single-knapsack budgeted PSM grouping kernel. The dashed line indicates the theoretical lower bound $1 - 1/(e)^{1/2}$ for the scoped kernel. This experiment checks the theoretical kernel used in the cooperative grouping step, rather than claiming a constant-factor bound for the complete multi-MC route-dependent scheduling implementation.

Proof: By the gated acceptance rule in (47), each committed event satisfies the sentinel stress-decrease condition $\Delta W_k = W_k(t_\ell^-) - W_k(t_\ell^+) \geq \epsilon_{\text{th}}$. This condition is used below to derive the positive interevent lower bound.

Let $\{t_\ell\}_{\ell=1}^{\infty}$ be the event-triggering sequence. We prove that the sequence cannot accumulate within a finite-time interval. By the dwell-time gate in (42), a new event cannot be generated unless

$$t - t_{\text{last}} \geq \tau_{\min}.$$

Therefore, for any two consecutive events t_ℓ and $t_{\ell+1}$

$$t_{\ell+1} - t_\ell \geq \tau_{\min}.$$

Since $\tau_{\min} > 0$, infinitely many events cannot occur within a finite-time interval.

We further show that this enforced dwell time is consistent with the sentinel decrease produced by the gated acceptance test. At an accepted event time t_ℓ, the reorganization yields a nonnegligible reduction in the sentinel stress metric

$$W_k\left(t_\ell^+\right) \leq W_k\left(t_\ell^-\right) - \epsilon, \quad \epsilon \geq \epsilon_{\text{th}}. \quad (54)$$

During the subsequent continuous phase, the growth rate of the sentinel is bounded by L_{rate}. Hence, the minimum time required for the sentinel to recover the removed stress amount is lower bounded by

$$\Delta t \geq \frac{\epsilon}{L_{\text{rate}}} \geq \frac{\epsilon_{\text{th}}}{L_{\text{rate}}}. \quad (55)$$

Together with the minimum control-cycle duration τ_{ctrl}, this yields the lower bound in (53). Thus, the event-triggered mechanism strictly precludes scheduling-level Zeno behavior and prevents computational chattering. □

a) Implementation-oriented summary: In implementation, Phase III is executed as a local event-triggered procedure rather than as a full-network optimization. When the sentinel detects persistent regional stress, the sink-assisted scheduler first forms the triggered local coalition and collects the pending requests in the corresponding union set. It then removes requests that violate the deadline or ESCs, evaluates the marginal PSM coverage gain over the union of already selected requests, applies the density-based greedy assignment rule, and commits the tentative reorganization only when the

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

TABLE I
SIMULATION PARAMETERS

Parameter	Value
Network Environment	
Monitoring area (Ω)	500×500 m^2
Number of sensors N	200–600
Number of MCs K	3–8
Deployment distribution	Uniform random
Communication range	20 m
Sensing & Energy Model	
Per-sensor PSM $c_i(x)$	Truncated PSM; (1)
Sensing cut-off radius R_s	40 m (default)
Attenuation scale η_1	2.5×10^{-3} m^{-2}
Decay-shape exponent η_2	2.0 (dimensionless, default)
Multi-sensor fusion	Conditional independence
Additional sensing-noise sampling	None
Sensor battery $E_s^{\max}$	1000 J
Sensor residual energy $E_i(t)$	$0 \le E_i(t) \le E_s^{\max}$ [J]
Initial sensor energy $E_i(0)$	$U(0.25, 0.65)E_s^{\max}$
Coverage grid step	10 m
Spatial weight $w(x)$	Uniform, $w(x) = 1$
Discharge rate $\widehat{P}_i(t)$	0.1 J/s (default)
Charging rate (MC-to-Sensor)	2 J/s
MC velocity v	2 m/s
MC battery $B^{\max}$	3.0×10^4 J
MC safety reserve B_{safe}	$0.1B^{\max}$
Charging target E_{tar}	$0.9E_s^{\max}$
Service energy E_{req}^i	$E_{\text{tar}} - E_i(t)$
Service time τ_{svc}^i	$E_{\text{req}}^i/P_{\text{ch}}$
Request-generation slack ℓ_i^{gen}	$[E_i(t_i^{\text{gen}}) - E_{\min}]/\widehat{P}_i(t_i^{\text{gen}})$ [s]
Remaining slack $\ell_i(t)$	$\max\{t_i^{\text{ddl}} - t, 0\}$ [s]
Service-completion deadline t_i^{ddl}	$t_i^{\text{gen}} + \ell_i^{\text{gen}}$ [s]
Default movement-energy coefficient c_{move}^0	1.0 J/m
High-speed penalty factor γ_v	1.8
Request threshold E_{th}	$0.3E_s^{\max}$
Sleep threshold $E_{\min}$	$0.1E_s^{\max}$
SCOPE Algorithm Settings	
Time slot Δt	1 s
Phase I: KDE Weight Factor α	0.5
Phase I: Load Penalty Weight ζ	10^{-3}
Phase I: Gradient Gains $\kappa_{prop}, \kappa_{int}$	0.8, 0.2
Phase II: Time Urgency Factor β_{time}	2.5
Phase II: Reward Gain Factor α_λ	2.0
Phase II: Reward Decay Rate β_λ	3.0
Phase II: Cost Penalty Factor α_γ	4.0
Phase II: Cost Decay Rate β_γ	5.0
Phase II: Lookahead horizon n	5 requests
Phase II: Discount factor β_{disc}	0.9
Phase II: Planning span T_{span}	300 s
Phase III: Energy-criticality weight κ	1.0
Phase III: Energy Sensitivity σ_e	50 J
Phase III: Event-trigger threshold δ	1.0 normalized stress/s
Phase III: Sentinel stress-decrease threshold ϵ_{th}	0.5 normalized stress
Simulation Config	
Simulation time	2000 s / 6000 s*
Monte Carlo runs	50

PSM details: The per-sensor model is the truncated exponential-power profile in (1). The default values are $R_s = 40$ m, $\eta_1 = 2.5 \times 10^{-3}$ m^{-2}, and $\eta_2 = 2.0$; Fig. 12 varies R_s from 20 to 40 m. Multi-sensor detections are conditionally independent, with $Q = 1 - \prod_i (1 - c_i)$; no additional sensing noise is sampled because it is represented by the effective PSM parameters.

Note: * The 6000-s horizon is used only for the finite-horizon coverage-retention analysis in Fig. 10; all other default and robustness experiments use 2000 s.

coverage, cost-surge, and sentinel stress-decrease gates are satisfied. Thus, route-dependent costs and deadlines enter the executable pruning and acceptance stages, whereas the formal approximation statement is restricted to the fixed-cost single-knapsack PSM coverage-selection kernel.

TABLE II
EXTENDED ROBUSTNESS SCENARIOS

Scenario	Setting	Purpose
S0	Uniform deployment without obstacles	Default reference case
S1	Gaussian-mixture deployment: 70% clustered sensors from four truncated Gaussian clusters and 30% uniform background sensors	Non-uniform/clustered topology evaluation
S2	S1 plus the fixed rectangular obstacle map in Table III; sensors and MCs are excluded from obstacle regions	Obstacle-constrained mobility evaluation
S3	S2 plus restricted buffer zones obtained by a 10-m dilation around the obstacles in Table III	Joint obstacle/restricted-area evaluation
Burst	Moving Gaussian demand hotspot from $b_0 = (125, 370)$ to $b_1 = (390, 140)$ during 500–1100 s	Spatio-temporally correlated bursty-demand evaluation
M0	Constant speed, homogeneous load, fixed charging efficiency	Default mobility/energy reference
M1	M0 plus variable trip speed and nonlinear movement power $P_{\text{move}}(v) = \xi v^2 + \beta$	Nonlinear mobility evaluation
M2	M1 plus heterogeneous and higher-average sensor loads with bounded temporal variation	Heterogeneous/high-load robustness evaluation
M3	M2 plus service-level time-varying wireless charging efficiency	Simulated mobility/high-load/efficiency evaluation

TABLE III
OBSTACLE MAP USED IN S2 AND S3

Obstacle	x-range (m)	y-range (m)
O_1	[90, 130]	[350, 425]
O_2	[170, 190]	[220, 275]
O_3	[205, 245]	[130, 215]
O_4	[285, 325]	[370, 465]
O_5	[325, 400]	[360, 390]
O_6	[390, 455]	[255, 285]
O_7	[430, 465]	[285, 330]
O_8	[335, 380]	[135, 170]
O_9	[420, 445]	[30, 85]

D. Empirical Checks of Scoped Theoretical Results

To empirically check the scoped approximation guarantee in Theorem 1, we evaluate the fixed-cost budgeted PSM coverage-selection kernel on small-scale instances where the optimal solution can be computed exactly. Fig. 6 reports the approximation ratio $J_{\text{alg}}/J_{\text{opt}}$ as a function of the number of candidate requests in the scoped kernel. The horizontal dashed line indicates the theoretical bound $1 - 1/(e)^{1/2}$ derived for the fixed-cost single-knapsack budgeted PSM grouping kernel.

We further examine Theorem 2 and its proof by plotting the Lyapunov-inspired sentinel and the interevent time induced by the event-triggered reorganization mechanism. As shown in Fig. 7(a), the sentinel $W(t)$ decreases at accepted event instants, indicating that each accepted reorganization enforces a measurable reduction in the sentinel stress metric. The vertical dashed lines in Fig. 7(a) correspond to accepted event instants that satisfy the stress-decrease gate $\Delta W_k \ge \epsilon_{\text{th}}$. Fig. 7(b)

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

TABLE IV
REPRODUCIBILITY SETTINGS FOR THE EVENT-TRIGGERED SENTINEL

Symbol	Meaning	Value / Selection Rule
h_0	KDE bandwidth for Phase-I nominal density $K_{\mathbf{H}}$	40 m
h_ρ	KDE bandwidth for online request density K_{h_ρ}	40 m
Δ_g	Grid step for Hessian calculation	10 m
χ	Hessian-curvature influence factor	0.5
ε_h	Hessian normalization regularizer	10^{-6}
$H_{\max}$	Upper bound of Hessian weight	3.0
L_W	Moving-average window for $W_k(t)$	10 s
Δ_W	Finite-difference interval for drift	10 s
η_u	Urgency time-scale coefficient	1 s
ε_ℓ	Remaining-slack regularizer	1 s
ε_c	Marginal-density cost regularizer	10^{-6}
δ	Sentinel-drift event-trigger threshold	1.0 normalized stress/s
λ_{decay}	Event-trigger-threshold decay rate	$\ln 2/300\ \text{s}^{-1}$
τ_{ctrl}	Fixed sink-assisted coordination interval; also used as the minimum control-cycle term in the dwell-time gate	50 s (nominal)
T_{react}	Conditional trigger-to-route-commitment scheduling latency	$T_{\text{react}} \le \tau_{\text{ctrl}}$ under one-cycle coordination; $\tau_{\text{ctrl}} = 50$ s (nominal)
L_{rate}	Assumed conservative upper bound on the actual sentinel drift $\lvert\dot{W}_k(t)\rvert$; numerically calibrated from the smoothed finite-difference drift	0.015 normalized stress/s
$\tau_{\min}$	Minimum inter-event dwell time	$\max\{\tau_{\text{ctrl}}, \epsilon_{\text{th}}/L_{\text{rate}}\} = 50$ s
κ	Energy-criticality weight	1.0
σ_e	Energy sensitivity coefficient	50 J
E_{crit}	Critical zone-average sensor-energy reference	$E_{\text{crit}} = E_{\text{th}} = 300$ J
η_ϵ	Sentinel threshold-conversion time scale	0.5 s
ϵ_U	Minimum normalized net-utility-gain threshold	0.5
ϵ_{th}	Minimum sentinel stress-decrease threshold	$\eta_\epsilon \delta = 0.5$ normalized stress
η_Θ	Cost-surge threshold ratio	0.2
Θ	Maximum accepted cost surge	$\sum_{m_i \in \{m_k\} \cup \mathcal{N}_k} \eta_\Theta B_i$

TABLE V
REPRODUCIBILITY SETTINGS FOR PHASE-I VALIDITY MONITORING AND RELEARNING

Symbol	Meaning	Value / Selection Rule
T_{mon}	Sliding monitoring window for empirical density estimation	300 s
Δ_{chk}	Validity-check interval for drift-score evaluation	100 s
M_{hold}	Consecutive validity checks required before re-learning	3
D_ρ^{th}	Demand-distribution drift threshold	0.15
D_Λ^{th}	Regional load-imbalance threshold	0.25
D_F^{th}	Failure/re-deployment ratio threshold	0.10
β_ρ	Density re-learning update rate	0.2
η_J	Minimum distortion-reduction ratio for accepting a new partition	0.03
$T_{\min}^P$	Minimum dwell time between two Phase-I re-learning operations	1000 s

compares the empirically observed minimum interevent time $\min_k T_{\text{IE}}^{(k)}$ with the analytical dwell-time lower bound

$$\tau_{\min} = \max\left\{\tau_{\text{ctrl}}, \frac{\epsilon_{\text{th}}}{L_{\text{rate}}}\right\}.$$

In the plotted simulation range, the measured $\min_k T_{\text{IE}}^{(k)}$ remains no smaller than this lower bound for all tested ϵ_{th}. These results

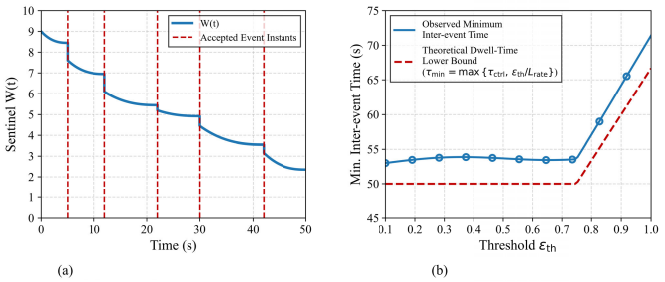


Fig. 7. Validation of scheduling-level Zeno-freeness. (a) Evolution of the Lyapunov-inspired sentinel with accepted event instants marked by vertical dashed lines. (b) Minimum interevent time versus the minimum sentinel stress-decrease threshold ϵ_{th}; the dashed line indicates the theoretical dwell-time lower bound $\tau_{\min} = \max\{\tau_{\text{ctrl}}, \epsilon_{\text{th}}/L_{\text{rate}}\}$.

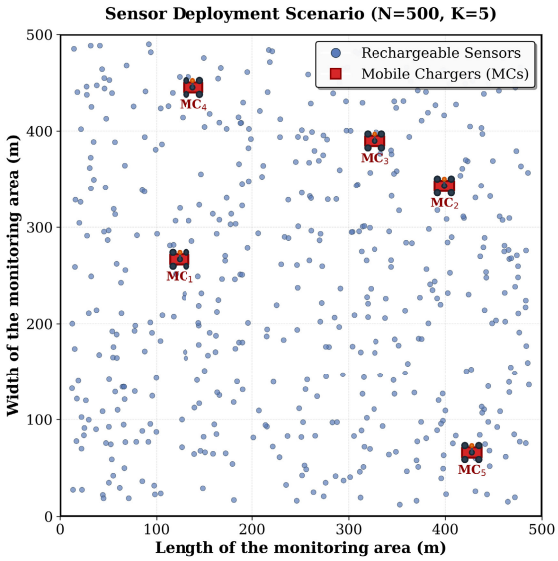


Fig. 8. Simulation scenario and illustrative charger-region partition (sensor deployment, MC depots, and region boundaries).

support the scheduling-level Zeno-freeness guarantee of the proposed event-triggered reorganization mechanism under the stated dwell-time and bounded-drift assumptions.

V. PERFORMANCE EVALUATION

In this section, we evaluate the proposed **SCOPE** framework and compare it with three representative baselines, namely, **MC^3**, **CAERM**, and **Dist-Greedy**. MC^3 is a recent multi-MC cooperative charging scheme with dynamic service-region adjustment, CAERM is a coverage-aware method with relatively static charger assignment, and Dist-Greedy represents a myopic heuristic that makes charging decisions based on instantaneous information. The evaluation focuses on coverage effectiveness, charging efficiency, network survivability, and charging latency under identical simulated system settings. Table I summarizes the simulation parameters, Table II summarizes the extended robustness scenarios, Table III specifies the obstacle map used in S2 and S3, and Table IV lists the reproducibility settings for the event-triggered sentinel.

The event-triggered sentinel parameters are selected using a normalization-first procedure. The Phase-I KDE uses the Gaussian kernel in (9) with $h_0 = 40$ m, followed by normalization of $\rho_0(x)$ over the monitoring region. First, the

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

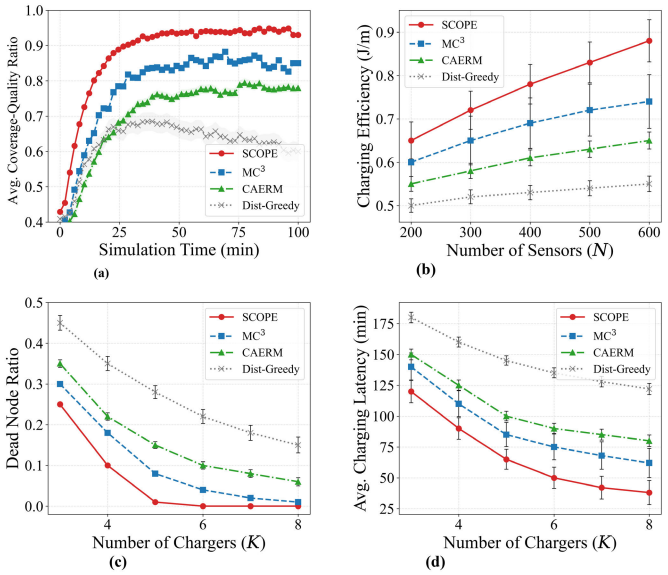


Fig. 9. Performance comparison among SCOPE, MC^3, CAERM, and Dist-Greedy. (a) Average coverage-quality ratio versus time. (b) Travel-normalized charging efficiency versus the number of sensors. (c) Dead-node ratio versus the number of chargers. (d) Average charging latency versus the number of chargers. Shaded bands in (a) and error bars in (b)–(d) indicate 95% CIs over 50 paired Monte Carlo runs.

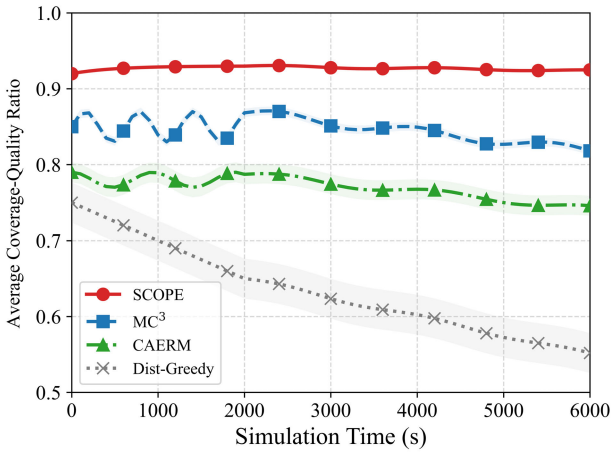


Fig. 10. Finite-horizon coverage retention: average coverage-quality ratio versus time under an extended 6000-s simulation horizon. Shaded bands indicate 95% CIs over 50 paired Monte Carlo runs.

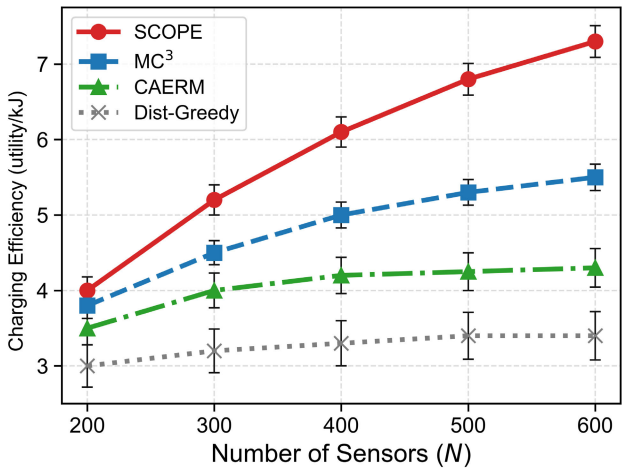


Fig. 11. Utility-normalized charging efficiency versus the number of sensors. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

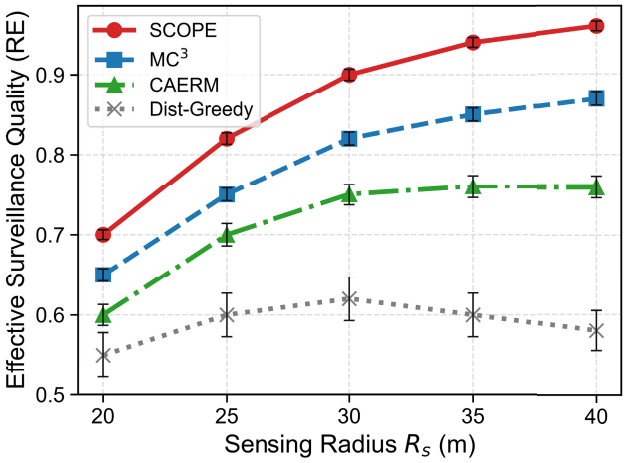


Fig. 12. Effective surveillance quality (RE) versus sensing radius R_s. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

normalized nominal demand density makes the nominal zone capacity C_k comparable across different deployments. Second, the Hessian weights H_k are normalized by the average regional curvature and clipped to $[1, H_{\max}]$, preventing isolated KDE artifacts from dominating the event trigger. Third, the event-trigger threshold δ is chosen from the sensitivity analysis in Section V-B8. A smaller δ increases the event frequency and control overhead, whereas an excessively large δ reduces responsiveness. We therefore use $\delta = 1.0$ normalized stress/s as the default tradeoff. Finally, the minimum sentinel stress-decrease threshold is set to $\epsilon_{\text{th}} = \eta_\epsilon \delta$, with $\eta_\epsilon = 0.5$ s, ensuring that every accepted event yields a measurable reduction in the sentinel stress metric.

In the simulations and in the validation of Fig. 7, τ_{ctrl} is set to 50 s, corresponding to one complete information-exchange, coalition-formation, and task-redistribution cycle in the 1-s slotted simulator. Under the one-cycle coordination assumption in Proposition 2, the same nominal control period gives $T_{\text{react}} \leq 50$ s from trigger detection to route commitment or explicit rejection of the tentative reallocation. Packet-level dissemination and acknowledgment delays are abstracted rather than measured in the simulator. In the pilot calibration, each candidate parameter setting was evaluated over 50 no-drift reference runs using an independent seed pool disjoint from the evaluation runs. For implementation, L_{rate} was numerically calibrated from the maximum observed magnitude of the smoothed finite-difference drift $|\widehat{W}_k(t)|$, multiplied by a safety factor of 1.2, and was conservatively fixed at 0.015 normalized stress/s. With $\epsilon_{\text{th}} = 0.5$ normalized stress, the analytical recovery term is $\epsilon_{\text{th}}/L_{\text{rate}} \approx 33.3$ s, and thus $\tau_{\min} = \max\{\tau_{\text{ctrl}}, \epsilon_{\text{th}}/L_{\text{rate}}\} = 50$ s, which is used as a common dwell-time gate for all tested scenarios.

The Phase-I validity thresholds in Table V are selected using the same independent pilot-calibration protocol, with a seed pool disjoint from the evaluation runs. Specifically, D_ρ^{th}, D_Λ^{th}, and D_F^{th} were selected above the 95th percentile of normal short-term fluctuations observed in the no-drift reference runs,

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

TABLE VI
COMMUNICATION OVERHEAD OF THE EVENT-TRIGGERED COOPERATIVE REORGANIZATION UNDER THE DEFAULT SETTING

Metric	Value
Event-trigger threshold δ	1.0 normalized stress/s
Minimum dwell time $\tau_{\min}$	50 s
Dwell-time upper-bound frequency	72 events/h
Accepted reorganization events	$\approx$ 8.3 per 2000-s run
Reorganization frequency	$\approx$ 14.9 events/h
Average coalition size	$\approx$ 3 MCs/event
Average local union request set	$\approx$ 8–10 requests/event
Logical messages per accepted CRA	$\approx$ 11 messages/event
Logical messages per run	$\approx$ 91 messages/2000-s run
Control payload per accepted CRA	$\approx$ 1.0 KB/event
CRA control payload per run	$\approx$ 8–9 KB per 2000-s run

Note: The values are compact scheduling-control payload accounting values derived from the released configuration. They include scheduling-control messages only and exclude sensing-data traffic, MAC/PHY headers, retransmissions, encryption/authentication fields, and hardware-specific protocol overhead.

and then verified to be exceeded by the persistent distribution shifts in B1–B3 for consecutive validity checks. The choice $T_{\text{mon}} = 300$ s, $\Delta_{\text{chk}} = 100$ s, and $M_{\text{hold}} = 3$ requires persistent violation over multiple consecutive checks, while $\eta_J = 0.03$ and $T^P_{\min} = 1000$ s reject numerically insignificant repartitioning and ensure that Phase-I relearning remains rare.

A. Simulation Model

We consider a square monitoring domain of 500×500 m^2, in which N rechargeable sensors are deployed uniformly at random. A fleet of MCs is initialized at designated depots and travels at a constant velocity. A sensor triggers a recharge request when its residual energy falls below a predefined threshold E_{th}; should the energy further drop below a critical minimum $E_{\min}$, the sensor is deemed dead (i.e., unable to perform sensing tasks). The PSM detailed in Section III is adopted for all comparative methods. Fig. 8 depicts a snapshot of the simulation environment, illustrating sensor distribution, MC depot locations, and the partition of charger service regions. Table I summarizes the simulation parameters. Unless otherwise stated, all reported data points are averages over 50 independent Monte Carlo trials, and the terms deployment, operating condition, and scenario refer to controlled simulations rather than hardware or field-test validation.

To emulate an ongoing long-term WRSN rather than a cold-start network with all sensors fully charged, the initial residual energy of each sensor is sampled as

$$E_i(0) \sim U(0.25, 0.65)E_s^{\max}.$$

This warm-start energy distribution ensures that a subset of sensors may generate charging requests during the 2000-s evaluation horizon, while still leaving sufficient heterogeneity for the scheduling policies to exhibit a different coverage and survivability behavior.

All stochastic deployments, initial energy states, demand perturbations, mobility/load/efficiency realizations, and random seeds are generated once per Monte Carlo trial and shared by all compared algorithms. This common-randomness protocol ensures that the reported differences are caused by the scheduling policies rather than by different environmental realizations.

Reproducibility Artifact Statement: To support reproducibility, the source code, random-seed lists, simulation scripts, and configuration files used in this study are publicly available at https://github.com/shilonghan244-maker/Scope. The repository contains the seed pool used for the 50 paired Monte Carlo trials, scripts for generating random sensor deployments, initial energy states, demand perturbations, mobility/load/charging-efficiency realizations, and configuration files corresponding to Tables I–V and the S0–S3, Burst, and M0–M3 scenarios. Because most experiments in this work are based on random deployment and stochastic operating realizations, the reported results should be interpreted as statistical Monte Carlo results rather than deterministic single-run outcomes. A single run or a small number of runs may not reproduce the reported mean curves exactly; reproducibility is instead achieved by using the released seed pool and configuration files to regenerate the same paired trial set and recompute the reported means, standard deviations, confidence intervals (CIs), and statistical tests. The released package is intended for key statistical reproducibility: it regenerates the 50 paired Monte Carlo trial records, Tables VI–VIII, and figure-concordance checks for the main trends and method rankings. It is not intended as a pixel-level or point-by-point reproduction of every panel in Figs. 9–20.

Statistical Analysis: Since all compared algorithms are evaluated using the same Monte Carlo seed pool and the same environmental realizations, pairwise comparisons are performed using paired samples. For each reported key metric in Table VIII, SCOPE is compared with the strongest or representative non-SCOPE baseline under the corresponding setting, as specified in the table. The paired difference is defined so that a positive value favors SCOPE, except where the table explicitly reports a latency tradeoff or a statistically comparable near-zero result. The Wilcoxon signed-rank test is used as the primary test, and Holm correction is applied to control multiple comparisons across the reported statistical tests. We report the mean paired difference, 95% CI, Holm-adjusted p-value, and paired effect size d_z. The effect size d_z is reported as a paired standardized difference and is used only to describe the magnitude of the paired effect; it is interpreted together with the corrected p-value and CI, rather than as an independent significance criterion. When the adjusted p-value is not below 0.05, or the CI includes zero, the result is interpreted as comparable performance rather than as a statistically supported separation.

In all figures reporting Monte Carlo means, error bars or shaded bands denote 95% CIs over the same 50 paired trials; illustrative deployment/path panels, representative trajectory panels, and theoretical single-trace validations are excluded.

Charging-Latency Statistic for Table VII: Let $\mathcal{R}_{\text{lat}} \subseteq \mathcal{R}_{\text{served}}$ be the postinitialization requests that reach charging service start and are included in this statistic; requests pending at initialization and requests that do not reach service start within the evaluation horizon are excluded. For each $r_i \in \mathcal{R}_{\text{lat}}$, let t_i^{asg} be the time at which the request is first accepted and assigned to an MC by the scheduler. Let τ_{off} be the fixed control-cycle offset used by the simulator latency metric. In the reported experiments, $\tau_{\text{off}} = 25$ s. Initial backlog requests remain included in the coverage and survivability evaluation.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

TABLE VII
STATISTICAL SUMMARY OF THE MAIN SIMULATION RESULTS OVER 50 PAIRED MONTE CARLO RUNS

Metric	Evaluation setting	SCOPE	MC^3	CAERM	Dist-Greedy
Average coverage-quality ratio C_{avg}	Default, $N = 500$, $K = 5$, $T = 2000$ s	0.933 ± 0.010 [0.931, 0.936]	0.882 ± 0.020 [0.877, 0.888]	0.796 ± 0.046 [0.783, 0.810]	0.712 ± 0.092 [0.686, 0.738]
Travel-normalized charging efficiency (J/m)	Dense network, $N = 600$	0.876 ± 0.170 [0.828, 0.925]	0.725 ± 0.218 [0.664, 0.787]	0.558 ± 0.068 [0.538, 0.577]	0.546 ± 0.063 [0.528, 0.563]
Dead-node ratio (%)	Default charger setting, $K = 5$	4.43 ± 0.74 [4.22, 4.64]	8.56 ± 1.70 [8.08, 9.05]	16.92 ± 3.33 [15.97, 17.86]	23.07 ± 6.42 [21.24, 24.89]
Average charging latency (s)	Default charger setting, $K = 5$	68.8 ± 32.5 [59.6, 78.1]	76.3 ± 39.1 [65.1, 87.4]	81.1 ± 15.5 [76.7, 85.5]	81.8 ± 15.0 [77.5, 86.1]
Utility-normalized charging efficiency (utility/kJ)	Dense network, $N = 600$	7.96 ± 0.75 [7.75, 8.17]	6.24 ± 0.62 [6.06, 6.41]	4.92 ± 0.90 [4.66, 5.17]	2.80 ± 1.13 [2.48, 3.12]
Computation time per decision (ms)	Dense network, $N = 600$	44.7 ± 2.9 [43.9, 45.5]	724.2 ± 37.5 [713.5, 734.8]	971.8 ± 56.3 [955.8, 987.8]	8.4 ± 1.1 [8.0, 8.7]

Note: Each entry is reported as mean $\pm$ standard deviation [95% CI] over 50 paired Monte Carlo runs. The CI is computed using $\bar{x} \pm t_{0.975,49} s/\sqrt{50}$.

TABLE VIII
SUMMARY OF PAIRED STATISTICAL TESTS FOR REPRESENTATIVE KEY COMPARISONS

Comparison	Metric / baseline	Diff. (95% CI)	p_{adj}/d_z	Interpretation
Fig. 9(b), $N = 600$	Travel-normalized efficiency vs MC^3	$+0.151$ J/m $[+0.077, +0.225]$	$0.002/0.58$	Statistically supported efficiency increase over the strongest baseline.
Fig. 17(b), S3	PSM coverage retention vs MC^3	$+0.769$ pp $[+0.683, +0.855]$	$< 0.001/2.53$	Statistically supported mean increase under S3.
Fig. 17(c), S3	Dead-node ratio vs MC^3	$+0.727$ pp reduction $[+0.607, +0.847]$	$< 0.001/1.72$	Statistically supported mean reduction in node mortality.
Fig. 17(d), S3	Average charging latency vs Dist-Greedy	-97.956 s $[-114.198, -81.714]$	$< 0.001/-1.71$	Raw latency favors the distance-oriented baseline; no latency superiority is claimed for SCOPE.
Fig. 18(c), burst window	Peak coverage loss vs CAERM	$+0.580$ pp reduction $[+0.267, +0.892]$	$0.004/0.53$	Statistically supported peak-loss reduction; the absolute effect size remains modest-to-moderate.
Fig. 18(d), burst window	Burst-critical outage ratio vs MC^3	$+0.362$ pp reduction $[+0.255, +0.469]$	$< 0.001/0.96$	Statistically supported reduction in burst-critical low-energy severity.
Fig. 19(a), M3	PSM coverage retention vs MC^3	$+0.048$ pp $[-0.013, +0.109]$	$0.297/0.22$	Comparable robustness; no statistical separation is claimed.
Fig. 19(b), M3	Coverage-critical outage ratio vs MC^3	$+0.078$ pp reduction $[+0.032, +0.125]$	$0.004/0.48$	Statistically supported coverage-critical outage reduction; the absolute effect size remains small-to-moderate.
Fig. 19(c), M0–M3	Coverage-critical outage degradation vs MC^3	$+0.153$ pp reduction $[+0.089, +0.216]$	$< 0.001/0.69$	Statistically supported reduction in coverage-critical outage degradation; the absolute effect size remains modest-to-moderate.

Note: Table VIII summarizes paired statistical tests for representative key comparisons. Reported paired differences, CIs, Holm-adjusted p-values (p_{adj}), and paired effect sizes are computed from the released paired Monte Carlo trial records for the specified comparison in each row. The paired effect size d_z is interpreted jointly with p_{adj} and the CI. When the CI includes zero or $p_{\text{adj}} \geq 0.05$, the result is interpreted as comparable performance rather than as a statistically supported separation. "pp" denotes percentage points.

The default comparison experiments in Fig. 9 and the extended robustness tests are conducted over a 2000-s simulation horizon. The finite-horizon coverage-retention test in Fig. 10 extends the horizon to 6000 s.

We further evaluate SCOPE under three groups of extended robustness experiments involving nonuniform obstacle-constrained simulated deployments, spatio-temporally correlated bursty demands, and simulated mobility/energy dynamics. The extended experiments use the same basic network scale as the default setting, i.e., a 500×500 m^2 monitoring area, $N = 500$ rechargeable sensors, $K = 5$ MCs, a simulation horizon of 2000 s, and 50 independent Monte Carlo trials.

For the extended robustness experiments in Figs. 17–20, the reported PSM coverage-related percentage is interpreted as normalized PSM coverage retention with respect to the fully active reference network under the corresponding monitoring domain and scenario configuration. Therefore, these values are used to compare robustness degradation within each extended scenario, rather than to directly replace the absolute coverage-time trajectories in Figs. 9 and 10. For obstacle-constrained scenarios, the PSM utility is evaluated over feasible monitoring grid cells in $\mathcal{F}$ and normalized by the fully active PSM utility on the same feasible domain.

For the simulated mobility/energy-dynamics scenarios M1–M3, $v_0 = 2$ m/s, $v_{\min} = 1.2$ m/s, $v_{\max} = 3.0$ m/s, and $\sigma_v = 0.35$ m/s. The load-variation process uses $\Delta_{\text{load}} = 60$ s,

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

$\rho_P = 0.85$, $\sigma_P = 0.03$, and $\epsilon_P^{\max} = 0.15$. The charging-efficiency factor is sampled at the beginning of each charging service from $U(0.72, 1.0)$. The common-randomness protocol described above is also applied to the M1–M3 mobility, load, and efficiency realizations.

In obstacle-constrained scenarios, gray blocks denote physical obstacles such as buildings or blocked structures, while red dashed regions denote restricted buffer zones. Sensors are not deployed inside these regions, and MCs are not allowed to traverse them. Therefore, the Euclidean travel distance is replaced by the shortest feasible path distance $d_{\mathcal{F}}(\cdot,\cdot)$ in the free space $\mathcal{F}$, computed using a grid-based shortest path search. The same obstacle-aware distance model is applied to SCOPE and all baselines.

The nonuniform deployment in S1 is generated by a truncated Gaussian-mixture model. Specifically, 70% of the sensors are sampled from four spatial clusters, and the remaining 30% are sampled uniformly over the monitoring region. The cluster centers are sampled uniformly in Ω with a minimum intercenter distance of 80 m, and clustered sensors are sampled from $\mathcal{N}(c_\ell, \sigma_c^2 I_2)$ with $\sigma_c = 35$ m. Rejection sampling is used to ensure that all sensors lie inside Ω and outside infeasible regions.

For S2 and S3, the physical obstacle set is represented as $\mathcal{O} = \cup_{\ell=1}^{9} O_\ell$, where each O_ℓ is an axis-aligned rectangular building block. The coordinates of all obstacle blocks are listed in Table III, match the scripts used in the obstacle-constrained robustness experiments, and are kept identical for all Monte Carlo trials and all compared algorithms. The restricted buffer set in S3 is constructed by a 10-m dilation around $\mathcal{O}$, i.e.,

$$\mathcal{B} = \{x \in \Omega : \mathrm{dist}(x, \mathcal{O}) \le 10 \quad \mathrm{m}\}.$$

Thus, the feasible space is $\mathcal{F} = \Omega \setminus \mathcal{O}$ for S2 and $\mathcal{F} = \Omega \setminus (\mathcal{O} \cup \mathcal{B})$ for S3. All candidate sensor positions falling in infeasible regions are rejected, and MC operational travel distances are computed on a 10-m grid using a shortest path search over feasible cells.

1) Default Evaluation Metrics: Let $\mathcal{G}$ be the coverage-evaluation grid and ΔA the grid-cell area. The average coverage-quality ratio used in Figs. 9(a) and 10 is

$$C_{\mathrm{avg}} = \frac{1}{T}\int_0^T \frac{\sum_{x_g \in \mathcal{G}} w\left(x_g\right) Q\left(x_g, \mathcal{A}\left(t\right)\right) \Delta A}{\sum_{x_g \in \mathcal{G}} w\left(x_g\right) \Delta A} dt. \tag{56}$$

The coverage entry in Table VII is reported as an average coverage-quality ratio computed using the same Fig. 9(a) coverage-quality reporting pipeline and the released coverage-evaluation configuration.

The travel-normalized charging efficiency in Fig. 9(b) is defined as

$$\eta_{\mathrm{travel}} = \frac{\sum_{r_i \in \mathcal{R}_{\mathrm{served}}} E_{\mathrm{del}}^i}{\sum_{m_k \in \mathcal{M}} d_k^{\mathrm{move}} + \varepsilon_d} \tag{57}$$

where E_{del}^i is the energy delivered to sensor s_i, d_k^{move} is the travel distance of MC m_k, and ε_d avoids division by zero. The utility-normalized charging efficiency in Fig. 11 is

$$\eta_{\mathrm{util}} = 1000 \frac{\sum_{r_i \in \mathcal{R}_{\mathrm{served}}} \Delta\Phi_i}{\sum_{r_i \in \mathcal{R}_{\mathrm{served}}} E_{\mathrm{del}}^i + \sum_{m_k \in \mathcal{M}} E_k^{\mathrm{move}}} \tag{58}$$

reported in utility/kJ. The average dead-node ratio and the Table VII charging-latency statistic are computed as

$$D_{\mathrm{avg}} = 100 \frac{1}{TN} \int_0^T \sum_{i=1}^{N} \mathbb{I}\{E_i(t) \le E_{\min}\} dt \tag{59}$$

$$L_{\mathrm{avg}} = \frac{1}{|\mathcal{R}_{\mathrm{lat}}|} \sum_{r_i \in \mathcal{R}_{\mathrm{lat}}} \left(t_i^{\mathrm{start}} - t_i^{\mathrm{asg}} + \tau_{\mathrm{off}}\right). \tag{60}$$

Finally, the effective surveillance quality RE in Fig. 12 is evaluated as the normalized PSM utility under the corresponding sensing radius R_s

$$RE\left(R_s\right) = \frac{1}{T}\int_0^T \frac{\Phi_{R_s}\left(\mathcal{A}\left(t\right)\right)}{\Phi_{R_s}\left(\mathcal{S}_{\mathrm{full}}\right)} dt. \tag{61}$$

2) Metric Definitions for Extended Robustness Tests: For the extended robustness tests, the PSM coverage-related percentage is reported as a normalized coverage retention metric. Let $\mathcal{A}(t)$ be the active sensor set at time t, and let $\mathcal{S}_{\mathrm{full}}$ be the fully active reference network under the same deployment and feasible monitoring domain. For a feasible domain $\mathcal{F} \subseteq \Omega$, we define

$$R_{\mathrm{cov}} = 100 \frac{\frac{1}{T}\int_0^T \Phi_{\mathcal{F}}\left(\mathcal{A}\left(t\right)\right) dt}{\Phi_{\mathcal{F}}\left(\mathcal{S}_{\mathrm{full}}\right)} \tag{62}$$

where $\Phi_{\mathcal{F}}(\cdot)$ is the PSM coverage utility evaluated over $\mathcal{F}$. For unconstrained cases, $\mathcal{F} = \Omega$; for obstacle-constrained cases, $\mathcal{F}$ excludes obstacles and restricted regions. This metric is used in Figs. 17, 19, and 20.

For the burst-demand experiment, the risk-weighted PSM coverage loss at time t is computed over the burst-affinity field. Let

$$\psi_b(x,t) = \exp\left(-\frac{\|x - b(t)\|^2}{2\sigma_b^2}\right)$$

be the spatial burst-affinity weight centered at the moving hotspot $b(t)$. The corresponding weighted PSM utility is

$$\Phi_b\left(\mathcal{A}\left(t\right), t\right) = \int_\Omega \psi_b\left(x,t\right) Q\left(x, \mathcal{A}\left(t\right)\right) dx. \tag{63}$$

The risk-weighted coverage loss is then defined as

$$L_b\left(t\right) = 100\left[1 - \frac{\Phi_b\left(\mathcal{A}\left(t\right), t\right)}{\Phi_b\left(\mathcal{S}_{\mathrm{full}}, t\right)}\right] \tag{64}$$

and is reported in percentage points.

The peak coverage loss is computed over the early burst-response window $\mathcal{T}_{\mathrm{br}} = [500, 1300]$ s

$$L_{\mathrm{peak}}^{\mathrm{br}} = \max_{t \in \mathcal{T}_{\mathrm{br}}} L_b\left(t\right). \tag{65}$$

Let $\mathbb{I}_i^{\mathrm{out}}(t)$ be an outage indicator that equals one when sensor s_i is below the sleep threshold and zero otherwise. The burst-critical outage ratio is defined as

$$R_{\mathrm{out}}^{\mathrm{burst}} = 100 \frac{\sum_{i=1}^{N} \int_{\mathcal{T}_{\mathrm{burst}}} \mathbb{I}_i^{\mathrm{out}}\left(t\right) \omega_i^{\mathrm{cov}} \psi_b\left(q_i, t\right) dt}{\sum_{i=1}^{N} \int_{\mathcal{T}_{\mathrm{burst}}} \omega_i^{\mathrm{cov}} \psi_b\left(q_i, t\right) dt} \tag{66}$$

where $\mathcal{T}_{\mathrm{burst}} = [500, 1100]$ s is the burst-active window. This metric weights outage duration by both PSM marginal coverage contribution and burst-region affinity.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

For the simulated mobility/energy-dynamics experiment, the coverage-critical outage ratio is defined over the full simulation horizon as

$$R_{\text{out}} = 100\frac{\sum_{i=1}^{N}\int_0^T \mathbb{I}_i^{\text{out}}(t)\,\omega_i^{\text{cov}}dt}{T\sum_{i=1}^{N}\omega_i^{\text{cov}}}. \tag{67}$$

The critical-outage degradation from M0 to M3 is then

$$\Delta R_{\text{out}} = R_{\text{out}}(\text{M3}) - R_{\text{out}}(\text{M0}). \tag{68}$$

B. Simulation Results and Discussion

1) Average Coverage-Quality Ratio and Short-Term Convergence: Fig. 9(a) illustrates the temporal evolution of the average coverage-quality ratio. SCOPE demonstrates rapid convergence and sustains the highest steady-state coverage among all compared methods. Dist-Greedy exhibits a sharp initial rise but subsequently degrades due to its myopic decision-making, which fails to prevent coverage holes under energy scarcity. CAERM outperforms Dist-Greedy by incorporating coverage awareness, yet its limited adaptability leads to a lower performance plateau. While MC^3 improves upon CAERM by leveraging cooperative multi-MC coordination and dynamic region adjustment, it consistently underperforms SCOPE, suggesting that purely heuristic region reshaping and proximity-based strategies may not fully align with the global quality-aware objective under PSM.

2) Charging Efficiency Versus Network Size (Travel-Normalized): Fig. 9(b) presents the charging efficiency as a function of the number of sensors, measured in a travel-normalized metric (J/m) to quantify energy delivered per unit of operational travel distance. SCOPE yields the highest plotted mean performance as N increases, indicating that its scheduling decisions generate higher effective charging outcomes per unit of operational travel distance in the evaluated mean trends. Although MC^3 improves efficiency over CAERM via dynamic service-region adjustment, frequent region reshaping incurs additional movement overhead, thereby reducing travel-normalized efficiency. Dist-Greedy exhibits the poorest performance due to a lack of global coordination, frequently resulting in inefficient charging trajectories. At $N = 600$, SCOPE improves the travel-normalized charging efficiency by approximately 18% over MC^3, the strongest baseline in this metric.

3) Network Survivability and Charging Latency: Fig. 9(c) compares the dead-node ratio across varying charger quantities. SCOPE achieves the lowest dead-node ratio across all tested values of K, thereby improving survivability under the default moderate-resource setting. Fig. 9(d) evaluates the average charging latency versus K. SCOPE attains the lowest average charging latency in all settings by jointly prioritizing request urgency (energy safety) and coverage-quality contribution.

4) Charging Efficiency (Utility-Normalized): Complementing the travel-normalized efficiency in Fig. 9(b), Fig. 11 reports the utility-normalized charging efficiency (utility/kJ) versus the number of sensors. SCOPE also has the highest plotted mean utility return per unit of energy expenditure under this metric, followed by MC^3, while CAERM and Dist-Greedy show lower mean values.

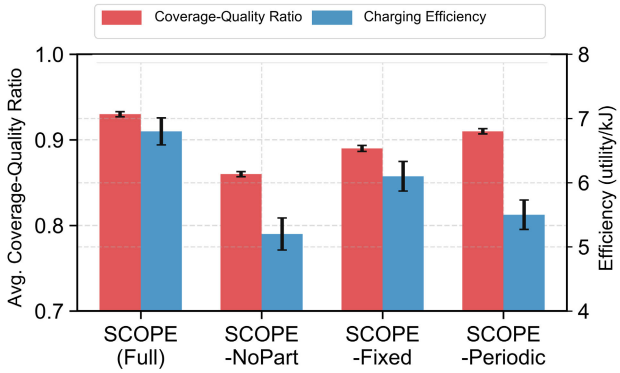


Fig. 13. Ablation study of SCOPE: average coverage-quality ratio and charging efficiency for SCOPE-Full and three variants. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

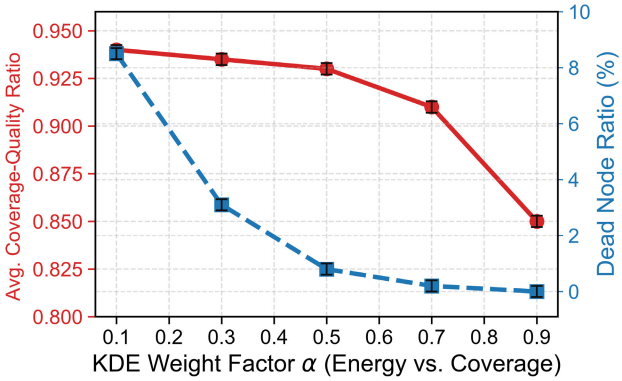


Fig. 14. Sensitivity to KDE weight factor α: average coverage-quality ratio and dead-node ratio. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

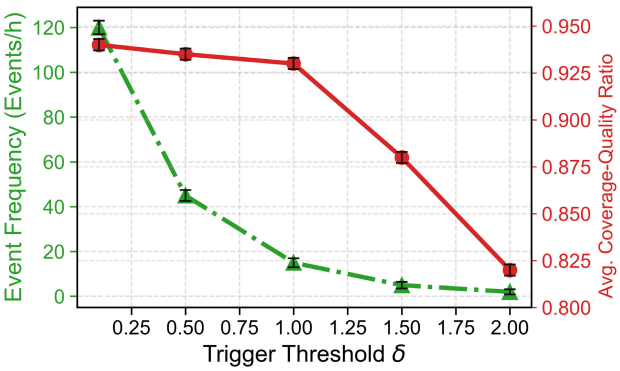


Fig. 15. Sensitivity to event-trigger threshold δ (normalized stress/s): event frequency and average coverage-quality ratio. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

5) Finite-Horizon Coverage Retention: Fig. 10 examines finite-horizon coverage retention over an extended 6000-s simulation period. This experiment is intended as a coverage-retention stress test rather than a proof of indefinite energy-neutral operation. Under the default setting, the aggregate sensing consumption of all sensors can exceed the aggregate instantaneous charging rate of the MC fleet; therefore, the objective is not to keep every sensor alive indefinitely, but to preserve the PSM coverage quality by prioritizing coverage-critical sensors under redundancy. SCOPE

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

TABLE IX
SENSITIVITY OF THE AVERAGE COVERAGE-QUALITY RATIO TO THE CALIBRATED PSM DECAY SHAPE

PSM profile	η_2	η_1 [$\mathrm{m}^{-\eta_2}$]	SCOPE C_{avg} mean [95% CI]	MC^3 C_{avg} mean [95% CI]
Gradual decay	1.5	1.1180×10^{-2}	0.9429 [0.9404, 0.9454]	0.8923 [0.8873, 0.8972]
Default	2.0	2.5000×10^{-3}	0.9334 [0.9306, 0.9362]	0.8824 [0.8768, 0.8880]
Steeper decay	2.5	5.5902×10^{-4}	0.9271 [0.9247, 0.9295]	0.8757 [0.8704, 0.8811]

Note: Each entry is the mean [95% CI] over 50 paired Monte Carlo runs, with the interval computed as $\bar{x} \pm t_{0.975,49} s/\sqrt{50}$. The cut-off radius is fixed at $R_s = 40$ m, and η_1 is calibrated so that $c(20\ \mathrm{m}) = e^{-1}$. SCOPE and MC^3 use the same PSM profile and paired environmental realization in every trial.

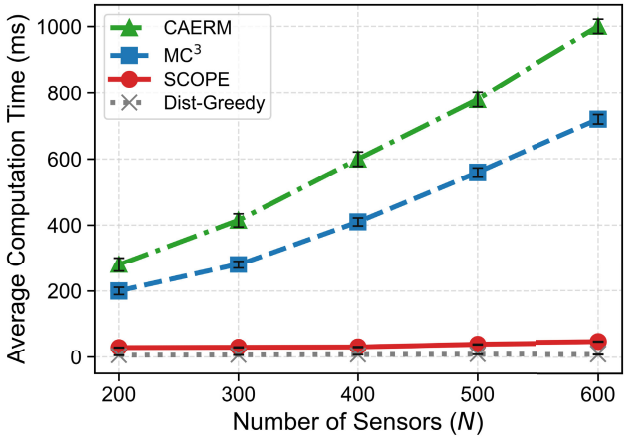


Fig. 16. Comparison of average computation time per decision epoch versus sensor population size N. SCOPE maintains low computation time because most online decisions are made on local zonal request sets rather than by repeated global optimization over all sensors. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

maintains a relatively high average coverage-quality ratio over the evaluated finite-horizon setting, whereas the baselines suffer larger oscillations or accumulated coverage degradation.

Sensor-Number Scalability and Process Novelty: To directly answer how the proposed framework considers the number of sensors, we vary N from 200 to 600 and evaluate SCOPE from both performance and computational perspectives. First, Fig. 9(b) reports the travel-normalized charging efficiency versus the number of sensors. SCOPE consistently yields the highest delivered-energy return per unit of operational travel distance as N increases, indicating that the energy–coverage-aware scheduling decisions remain effective in denser networks. Second, Fig. 11 reports the utility-normalized charging efficiency versus N, confirming that SCOPE also achieves the highest PSM utility return per unit charging energy. Third, the computation-overhead analysis in Section V-B9 reports the average computation time per decision epoch as (N) increases from 200 to 600. SCOPE maintains a controlled online computation time of approximately 25–45 ms because its online scheduler operates mainly on local zonal request sets, whereas baselines involving global reclustering or iterative optimization scale more rapidly with N. These observations support the proposed N-aware hierarchical scheduling process, which transforms large-scale sensor-level requests into density-based partitioning, local scheduling, and event-triggered cooperation rather than relying on empirical parameter tuning alone.

6) Effective Surveillance Quality Under Different Sensing Radii: Fig. 12 investigates the effective surveillance quality under varying sensing radii R_s. While all algorithms benefit from extended sensing ranges, SCOPE yields the highest plotted mean RE and the largest observed mean gain as R_s increases. This indicates that SCOPE more effectively exploits enhanced sensing capabilities while mitigating redundancy and energy waste inherent in the PSM. Fig. 12 specifically evaluates sensitivity to the PSM cutoff radius R_s: it shows how the normalized PSM utility changes as the sensing footprint expands, while η_1 and η_2 are held at their default values.

To separately evaluate sensitivity to the distance-decay shape, we use the default setting $N = 500$, $K = 5$, $T = 2000$ s, fix $R_s = 40$ m, and test $\eta_2 \in \{1.5, 2.0, 2.5\}$ with the same 50 paired Monte Carlo seeds and common-randomness protocol. For each profile, η_1 is calibrated as

$$\eta_1(\eta_2) = \frac{-\ln p_{\mathrm{ref}}}{d_{\mathrm{ref}}^{\eta_2}}, \quad d_{\mathrm{ref}} = 20 \ \ \mathrm{m}, \quad p_{\mathrm{ref}} = e^{-1}$$

so that every tested profile satisfies $c(d_{\mathrm{ref}}) = e^{-1}$ while the decay shape changes. All non-PSM paired realizations, including the deployment, initial energy, and base-load field, are held fixed. Table IX reports the resulting average coverage-quality ratios for SCOPE and MC^3, the strongest PSM-aware baseline. Taken together, Fig. 12 and Table IX cover two distinct PSM sensitivities: the cutoff radius R_s and the calibrated distance-decay shape. The framework and its union-based coverage formulation remain applicable under all tested profiles, whereas the absolute coverage level and the magnitude of the performance gap remain conditional on the selected PSM parameters.

7) Ablation Study: Fig. 13 presents an ablation study quantifying the contributions of key SCOPE components. We compare the full framework (SCOPE-Full) against three variants: SCOPE-NoPart (excluding adaptive partitioning), SCOPE-Fixed (using fixed charger assignments), and SCOPE-Periodic (employing periodic rescheduling rather than event-triggered control). The full version simultaneously achieves the highest observed average coverage-quality ratio and charging efficiency among the tested variants, supporting the roles of both adaptive partitioning and event-triggered control.

8) Sensitivity Analysis of Key Parameters: Fig. 14 analyzes the KDE weight factor α, reporting both average coverage-quality ratio and dead-node ratio. An intermediate range is shown to provide a desirable tradeoff between coverage quality and survivability. Fig. 15 evaluates the event-trigger threshold δ in terms of event frequency and average coverage-quality ratio. A smaller δ triggers frequent events (increasing control overhead) yielding marginally higher coverage, whereas an excessively large δ compromises responsiveness and degrades coverage. These results offer simulation-based guidance for parameter selection under the evaluated settings.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

9) Computation Overhead Analysis: We measure the average computation time per decision epoch for all algorithms as N increases from 200 to 600. Simulations were conducted on a PC equipped with an Intel Core i7-10700 CPU at 2.90 GHz and 16-GB RAM.

Fig. 16 shows that Dist-Greedy achieves the lowest computation time (<10 ms) because of its simple scheduling logic. SCOPE maintains a slightly higher but controlled computation overhead (approximately 25–45 ms/decision in the tested range of N = 200–600), which is negligible relative to the physical travel and charging time scales. This efficiency stems from its hierarchical design: the computationally intensive partitioning (Phase I) is performed offline, whereas the online scheduling (Phase II) relies on efficient local calculations. MC^3 and CAERM exhibit higher computational costs that increase rapidly with N, primarily because of complex iterative optimization and global reclustering processes.

The practical consequence of this architectural distinction is quantified by the paired results. At N = 600, SCOPE requires 44.7 ± 2.9 ms/decision, whereas MC^3 requires 724.2 ± 37.5 ms (see Table VII). Under the same dense-network setting, the paired travel-normalized-efficiency difference is +0.151 J/m (95% CI: +0.077 to +0.225 and Holm-adjusted p = 0.002) in favor of SCOPE (see Table VIII). These values quantify the empirical consequence of the hierarchy over the tested range; they are not presented as an asymptotic complexity proof or a universal dominance claim.

10) Communication Overhead Analysis: In addition to computation time, we account for the scheduling-control communication overhead introduced by SCOPE. The reported overhead includes request-report, trigger, coalition-formation, state-summary, assignment, and commit/acknowledgment messages used for scheduling and event-triggered cooperative reorganization. Regular sensing-data traffic, medium access control/physical layer (MAC/PHY) headers, retransmissions, encryption/authentication fields, and hardware-specific protocol overhead are not included. Thus, the numbers below should be interpreted as compact scheduling-control payload accounting under the tested simulation settings, rather than as a complete wireless-link-layer measurement.

For an accepted cooperative reorganization event e, let $\mathcal{C}_e$ be the triggered charger coalition, $q_e = |\mathcal{C}_e|$ be the number of participating MCs, and $R_e = |\mathcal{T}_{\text{union},e}|$ be the number of pending requests considered in the local union set. With batched task assignment, the number of logical control messages per accepted event is

$$M_e = 1 + 2\,(q_e - 1) + q_e + q_e = 4q_e - 1 \tag{69}$$

where the four terms correspond to the trigger report, coalition invitation/state-summary exchange, assignment packets, and commit/acknowledgment messages, respectively. Using a compact binary payload format, the trigger report, coalition invitation, MC state summary, assignment header, task descriptor, and commit/acknowledgment message require 48, 40, 64, 32, 64, and 24 B, respectively. Therefore, the payload size of one accepted event is estimated as

$$\begin{aligned} B_e &= 48 + (q_e - 1)\,(40 + 64) \\ &\quad + \sum_{m_i \in \mathcal{C}_e} (32 + 64|\mathcal{G}_i|) + 24q_e \\ &= 160q_e - 56 + 64R_e \text{ bytes.} \end{aligned} \tag{70}$$

Table VI summarizes the resulting event frequency, logical-message count, and compact control-payload estimates under the default setting.

Under the default setting with δ = 1.0 normalized stress/s and $\tau_{\min}$ = 50 s, SCOPE triggers about 8.3 accepted cooperative reorganizations over a 2000-s run, corresponding to approximately 14.9 events/h. This frequency is well below the dwell-time upper bound frequency of 3600/50 = 72 events/h. With an average local coalition size of about three MCs and about 8–10 pending requests in the triggered union set, each accepted reorganization requires about 11 logical control messages and approximately 1.0 kB of control payload. Therefore, the additional scheduling-control communication cost remains controlled under the tested simulation settings, while avoiding periodic global reorganization.

Table VII summarizes the trial-level means, standard deviations, and 95% CIs over the 50 paired Monte Carlo runs. Table VIII reports representative paired comparisons under the statistical protocol defined in Section V-A.

Table IX examines the sensitivity to the PSM decay shape while keeping the detection probability at d_{ref} = 20 m fixed. The absolute coverage-quality values change with the decay profile because the marginal PSM gains and sensing redundancy change. Nevertheless, SCOPE retains the same relative mean ranking against MC^3 across the tested profiles: the paired SCOPE–MC^3 difference is positive for every profile, with 95% CIs $[0.0470, 0.0542]$, $[0.0464, 0.0556]$, and $[0.0473, 0.0556]$ for η_2 = 1.5, 2.0, and 2.5, respectively. This indicates that the observed scheduling advantage is not caused solely by the default choice η_2 = 2, although the numerical performance remains parameter-dependent.

C. Robustness Under Simulated Deployment and Dynamic Operating Conditions

1) Nonuniform and Obstacle-Constrained Simulated Deployments: To evaluate the proposed framework beyond the standard uniform simulated deployment, we construct four simulated deployment scenarios with increasing spatial complexity. S0 is the default uniform simulated deployment without obstacles. S1 is generated by a Gaussian-mixture topology with four mildly separated clusters and a uniform background. S2 further includes the fixed physical-obstacle map listed in Table III and illustrated in Fig. 17(a), and S3 combines the same obstacles with restricted buffer zones. In S2 and S3, the operational travel distance of each MC is computed using the shortest feasible path distance $d_{\mathcal{F}}(\cdot,\cdot)$ in the free space, rather than the Euclidean distance. All compared algorithms use the same obstacle map, the same feasible-path distance model, and the same initial network realizations.

Fig. 17(a) illustrates an example simulated S3 deployment using the fixed obstacle map in Table III. The coverage-retention results in Fig. 17(b) are computed using (62). As shown in Fig. 17(b), SCOPE has the highest mean PSM coverage retention in all four simulated scenarios. In the most constrained S3 scenario, SCOPE maintains an average PSM coverage retention of 98.75%, compared with 97.91% for the best baseline. Fig. 17(c) further shows that SCOPE achieves

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

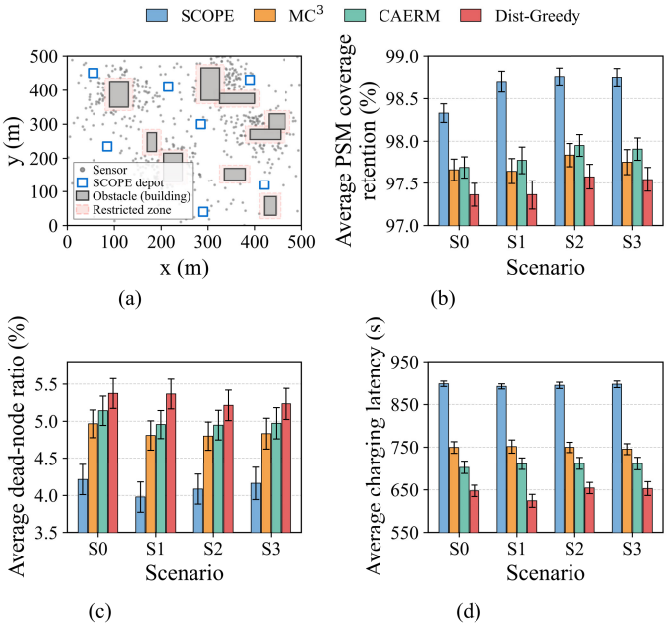


Fig. 17. Robustness under nonuniform and obstacle-constrained simulated deployments. (a) Example S3 simulated deployment with physical obstacles, restricted buffer zones, and SCOPE depots. (b) Average PSM coverage retention under S0–S3. (c) Average dead-node ratio under S0–S3. (d) Average charging latency. Error bars in (b)–(d) indicate 95% CIs over 50 paired Monte Carlo runs.

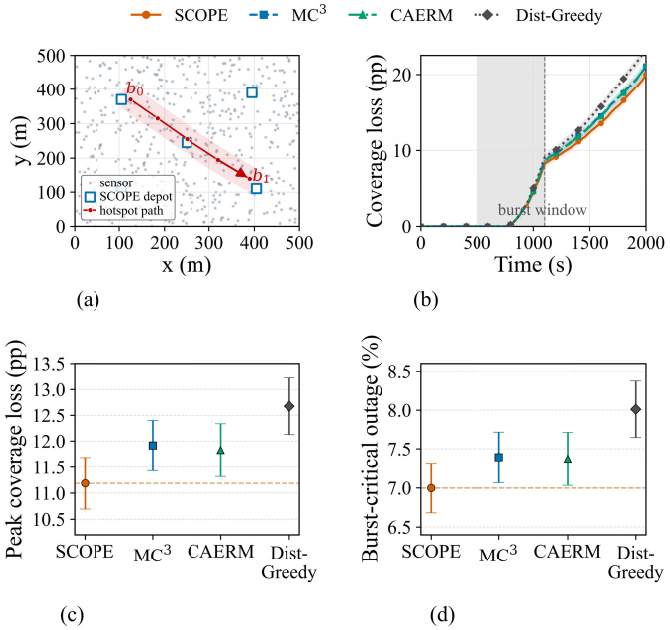


Fig. 18. Response to simulated spatio-temporally correlated bursty demands. (a) Moving burst-demand corridor with $b_0 = (125, 370)$ and $b_1 = (390, 140)$. (b) Risk-weighted coverage loss during and after the burst window. (c) Peak coverage loss within the early burst-response window $t \in [500, 1300]$ s. (d) Burst-critical outage ratio. Shaded bands in (b) and error bars in (c) and (d) indicate 95% CIs over 50 paired Monte Carlo runs.

the lowest average dead-node ratio. In S3, the dead-node ratio is reduced to 4.16%, whereas the best baseline remains at 4.83%.

Fig. 17(d) reports the raw average charging latency. Since distance-oriented baselines tend to serve nearby requests first, they may obtain lower raw latency in some obstacle-constrained cases. However, this comes at the cost of higher coverage degradation and higher node mortality. In contrast, SCOPE prioritizes coverage-critical requests under the PSM and therefore shows improved mean coverage robustness and survivability under the constrained simulated topologies. The paired tests in Table VIII support the S3 mean improvements in PSM coverage retention and dead-node ratio, with statistically supported paired differences of +0.769 percentage points and +0.727 percentage points, respectively. Table VIII also confirms that SCOPE does not optimize raw charging latency in this obstacle-constrained setting. This is consistent with the intended coverage-critical scheduling objective.

2) Spatio-Temporally Correlated Bursty Demands: We next evaluate the dynamic response of SCOPE under spatially and temporally correlated bursty demands. A moving Gaussian demand hotspot is introduced from $t = 500$ s to $t = 1100$ s. The hotspot center moves linearly from $b_0 = (125, 370)$ to $b_1 = (390, 140)$. The burst-induced discharge increment of sensor s_i is modeled as

$$\Delta P_i(t) = \lambda_b P_i^0 \exp\left(-\frac{\|q_i - b(t)\|^2}{2\sigma_b^2}\right) \tag{71}$$

where P_i^0 is the nominal discharge rate and $b(t)$ is the time-varying hotspot center. In the experiments, we set $\lambda_b = 8.0$ and $\sigma_b = 115$ m, producing a localized but persistent demand surge along the moving corridor rather than a network-wide load increase. This setting is more challenging than independent random request arrivals because it creates both spatial correlation and temporal persistence in the generated charging requests.

The moving-hotspot experiment evaluates the burst-handling mechanism described in Section IV-C. The formal bound in Proposition 2 concerns the scheduling interval from trigger detection to route commitment or explicit rejection; the delayed physical coverage response in Fig. 18 additionally includes request generation, sensor-energy depletion, charger travel, and charging-service time.

To emphasize the affected region, we evaluate the risk-weighted PSM coverage loss along the burst corridor. The risk-weighted coverage loss, early window peak loss, and burst-critical outage ratio are computed using (64)–(66). Fig. 18(b) shows that all algorithms experience a delayed coverage degradation after the burst starts, because sensors first consume energy faster, then generate requests, and eventually may fall below the sleep threshold if not served in time. SCOPE consistently yields the smallest mean coverage loss throughout the burst-response process.

The peak coverage loss in Fig. 18(c) is computed within the early burst-response window $t \in [500, 1300]$ s, rather than over the entire 0–2000-s horizon. SCOPE achieves the lowest mean peak loss of 11.19 percentage points, compared with 11.83 percentage points for CAERM, 11.91 percentage points for MC^3, and 12.68 percentage points for Dist-Greedy. Fig. 18(d) further reports the burst-critical outage ratio, where sensor outage durations are weighted by both PSM marginal coverage contribution and burst-region affinity. SCOPE achieves the lowest mean burst-critical outage ratio of 7.00%, compared

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

with 7.37% for the best baseline. These mean results suggest improved response under spatio-temporally correlated bursty demands, but the statistical interpretation depends on the CIs, corrected p-values, and effect sizes in Table VIII. Table VIII reports statistically supported but modest-to-moderate reductions in the burst window: 0.580 percentage points for peak coverage loss relative to CAERM and 0.362 percentage points for the burst-critical outage ratio relative to MC^3. We therefore describe the result as an improved mean response to correlated bursts, not as broad dominance.

3) Nonlinear Mobility, Heterogeneous Loads, and Variable Charging Efficiency: Finally, we evaluate SCOPE under simulated mobility and energy dynamics. Four simulated operating conditions are considered. M0 is the default reference case with fixed-speed motion at $v_0 = 2$ m/s, homogeneous sensor load, and fixed charging efficiency. In M0, movement energy is computed using the default constant per-distance coefficient c^0_{move}.

M1 activates the calibrated nonlinear movement-power model $P_{\text{move}}(v) = \xi v^2 + \beta$ under variable trip speeds. For each trip, the speed is sampled as $v_\ell = \text{clip}(v_0 + \epsilon^v_\ell, v_{\min}, v_{\max})$, where $\epsilon^v_\ell \sim \mathcal{N}(0, \sigma^2_v)$, $v_{\min} = 1.2$ m/s, $v_{\max} = 3.0$ m/s, and $\sigma_v = 0.35$ m/s. The corresponding movement energy is computed as $E_{\text{move},\ell} = P_{\text{move}}(v_\ell) d_\ell / v_\ell$, where d_ℓ is the feasible path length.

M2 further introduces heterogeneous and higher average sensor loads with bounded temporal variation. The baseline discharge rate is sampled as $P^0_i \sim U(0.085, 0.185)$ J/s, and the time-varying load is $P_i(t) = P^0_i(1 + \epsilon^P_i(t))$. The perturbation $\epsilon^P_i(t)$ is updated every $\Delta_{\text{load}} = 60$ s using a bounded autoregressive process with $\rho_P = 0.85$, $\sigma_P = 0.03$, and $\epsilon^{\max}_P = 0.15$.

M3 additionally introduces service-level wireless charging efficiency. At the beginning of each charging service, the effective charging efficiency is sampled as $\eta_i(t) \sim U(0.72, 1.0)$, and the received charging power is $P^{\text{eff}}_{\text{ch},i}(t) = \eta_i(t) P_{\text{ch}}$. The common-randomness protocol described above is applied to all M1–M3 mobility, load, and efficiency realizations.

The coverage-critical outage ratio and M0–M3 degradation are computed using (67) and (68).

Fig. 19 shows the results. In M0 and M1, SCOPE remains comparable to the strongest baseline because the simulated operating conditions are still relatively mild. In terms of average PSM coverage retention, the leading methods remain very close under the most demanding simulated M3 setting. Specifically, SCOPE obtains 98.29%, compared with 98.26% for MC^3, 98.21% for CAERM, and 97.57% for Dist-Greedy.

The mean robustness trend is more clearly reflected by the coverage-critical outage ratio and the M0–M3 coverage-critical outage degradation. Fig. 19(b) reports the coverage-critical outage ratio. In M3, SCOPE reduces this ratio to 4.02%, while MC^3, CAERM, and Dist-Greedy obtain 4.11%, 4.26%, and 5.93%, respectively. Moreover, Fig. 19(c) compares the increase in coverage-critical outage from M0 to M3. SCOPE has the smallest degradation of 3.12 percentage points, compared with 3.24 for MC^3, 3.25 for CAERM, and 4.41 for Dist-Greedy. Under M3, the coverage-retention values of SCOPE and MC^3 are statistically comparable. Table VIII nevertheless reports small-to-moderate reductions for SCOPE in the coverage-critical outage ratio and its M0–M3 degradation. Accordingly, the evidence supports comparable coverage retention with modest outage-robustness improvements.

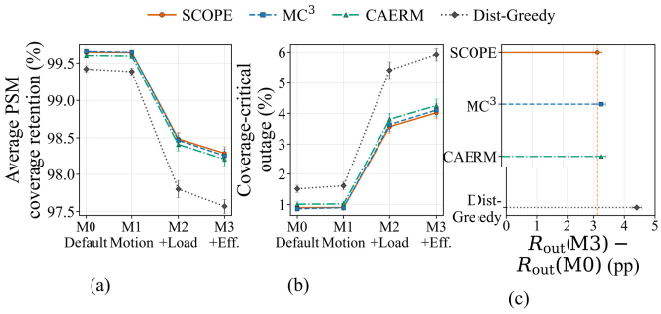


Fig. 19. Robustness under simulated nonlinear mobility, heterogeneous loads, and variable charging efficiency. (a) Average PSM coverage retention. (b) Coverage-critical outage ratio. (c) Increase in coverage-critical outage from M0 to M3. Error bars indicate 95% CIs over 50 paired Monte Carlo runs.

4) Failure of Time-Scale Separation and Rare Phase-I Relearning: The static partitioning in Phase I relies on a time-scale separation assumption: the long-term statistical demand density should evolve more slowly than the short-term request and energy fluctuations handled by Phases II and III. This assumption is appropriate under routine operation, where transient energy deficits, local request bursts, and short-lived workload fluctuations can be absorbed by the local scheduling and event-triggered reorganization mechanisms. However, it may become invalid when the long-term demand distribution changes persistently, such as after a sudden regional consumption shift, sensor redeployment, large-scale sensor failure, or sustained environmental changes.

Using the Phase-I validity monitor in Section IV-A, we examine when the time-scale separation assumption breaks and how rare Phase-I relearning restores the static backbone. The monitor jointly evaluates the online demand-distribution drift, the regional load mismatch, and the failure/redeployment ratio, and triggers Phase-I relearning only when the persistent break condition in (16) remains true for consecutive validity checks. This section focuses on the resulting stress-test behavior under persistent distribution shifts.

Three persistent-shift scenarios are considered: B1, a regional consumption shift; B2, sensor redeployment into a new spatial cluster; and B3, permanent failures of coverage-critical sensors. These shifts are injected at $t = 800$ s and are fundamentally different from short-lived bursty requests, because they alter the long-term statistical demand distribution used by the static Phase-I backbone. We compare three variants: SCOPE-Static, which keeps the original Phase-I partition unchanged; SCOPE-Periodic, which reruns Phase I periodically; and SCOPE-Triggered, which uses the Phase-I validity monitor in Section IV-A, including the persistent break condition in (16), to perform rare trigger-based relearning only when the persistent drift condition is satisfied.

The three persistent break scenarios are instantiated as follows. In B1, at $t = 800$ s, sensors within a regional shift center $c_B = (360, 360)$ experience a persistent consumption increase

$$P_i(t) \leftarrow P_i(t)\left[1 + \lambda_B \exp\left(-\frac{\|q_i - c_B\|^2}{2\sigma^2_B}\right)\right]$$

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

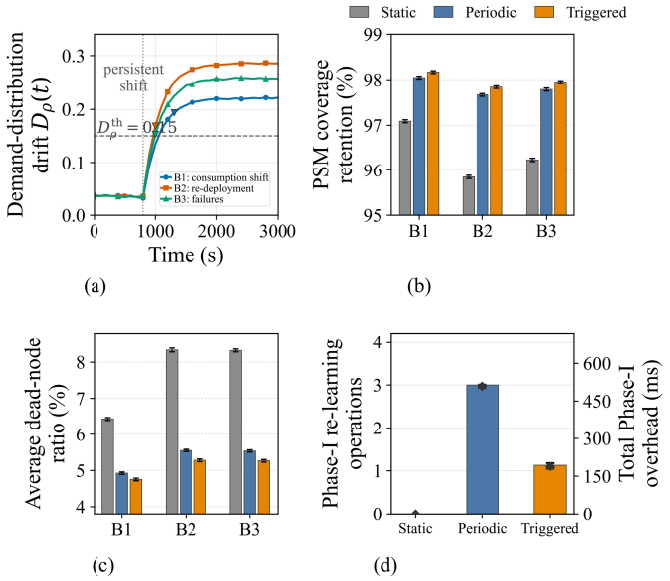


Fig. 20. Time-scale separation break and rare Phase-I relearning. (a) Representative drift-score trajectories under the three persistent-shift scenarios. Downward markers indicate the detected break time after $M_{\text{hold}} = 3$ consecutive validity checks evaluated every $\Delta_{\text{chk}} = 100$ s using sliding $T_{\text{mon}} = 300$ s monitoring windows. (b) PSM coverage retention, normalized by the no-drift reference coverage under the same network realization. (c) Average dead-node ratio. (d) Number of Phase-I relearning operations and total Phase-I overhead. Error bars in (b)–(d) indicate 95% CIs over 50 paired Monte Carlo runs. In (d), the bars show the average number of Phase-I relearning operations, and the diamond markers show the average total Phase-I computational overhead.

with $\lambda_B = 0.8$ and $\sigma_B = 100$ m. In B2, 15% of sensors are redeployed at $t = 800$ s into a new truncated Gaussian cluster centered at $c_R = (380, 140)$ with standard deviation 35 m, while avoiding obstacles and restricted zones. In B3, 15% of sensors are selected from the highest coverage-criticality quartile and permanently marked as failed at $t = 800$ s. These perturbations are held fixed after injection and follow the common-randomness protocol described in the simulation model.

To avoid mixing absolute coverage levels with recovery-oriented stress-test statistics, Fig. 20(b) reports PSM coverage retention, defined as the achieved PSM coverage after the persistent shift normalized by the no-drift reference coverage under the same network realization. Thus, the values in Fig. 20(b) should be interpreted as the retained fraction of nominal PSM coverage after a persistent distribution shift, rather than as the raw average coverage level reported in the main baseline comparison.

Fig. 20(a) shows that the drift score remains below the validity threshold before the persistent shift and exceeds it afterward. The detection markers follow the sliding-window validity-check protocol summarized in Table V. Fig. 20(b) and (c) shows that SCOPE-Static suffers from reduced PSM coverage retention and increased dead-node ratio after persistent shifts, whereas SCOPE-Triggered restores the coverage-retention level to around 98% while keeping the dead-node ratio close to the periodic relearning baseline. Fig. 20(d) further shows that SCOPE-Triggered achieves this adaptation with fewer Phase-I relearning operations and substantially lower Phase-I overhead than periodic repartitioning.

a) Summary: The extended tests show that SCOPE preserves coverage-aware robustness across obstacle-constrained, burst-demand, mobility/energy, and persistent-shift scenarios. Statistical claims are limited to the paired comparisons in Table VIII; small absolute differences are described as modest improvements, and comparisons without statistical separation are reported as comparable performance.

D. Practical Deployment Implications

The results above should be interpreted as simulation-based evidence under controlled and reproducible settings rather than as hardware or field-test validation. In practical WRSN deployments, several assumptions used in the simulator may affect the achieved performance. First, the simulation uses reliable state exchange as a baseline abstraction rather than modeling perfect communication at the packet level. Packet losses, communication delays, and stale state reports may affect request deadlines, residual-energy estimates, and workload imbalance measurements. A practical implementation should therefore use timestamped requests, acknowledgments, retry policies, and conservative slack margins in the deadline and energy-safety checks. The current framework also uses a sink-assisted coordination layer for request aggregation, event monitoring, and cooperative reorganization. This architecture is suitable for deployments with an edge gateway or base station, but fully decentralized deployments would require local leader election, distributed consensus, or redundant gateway support to avoid a single coordination bottleneck.

Second, the dwell-time gate and the sentinel assume that trigger checks and deadline evaluations are performed on a sufficiently consistent time base. Clock-synchronization errors should therefore be bounded and incorporated into the control-cycle duration, dwell-time margin, and service-deadline feasibility test. Global Positioning System (GPS) and indoor-positioning errors, as well as charger localization errors, can also affect travel-distance estimation, obstacle-aware path planning, and charger-to-sensor alignment. In deployment, feasible-path distances, travel-time estimates, and charging-alignment tolerances should be inflated by a localization-error margin, especially in indoor or GPS-denied environments.

Third, battery degradation changes sensor capacity, discharge behavior, and available charger budgets over long horizons. Periodic recalibration of sensor energy models, MC battery states, and safety reserves is therefore required. Practical wireless power transfer (WPT) may also be affected by distance variation, coil or antenna misalignment, environmental interference, and time-varying charging efficiency. These factors can be handled only conservatively in a simulation-oriented study, for example through online efficiency estimation, longer charging-time allocation, and larger energy-safety buffers. These deployment factors delimit the conditions under which the simulation-based performance and the scheduling-level guarantees should be interpreted and transferred to hardware-oriented implementations.

VI. Conclusion

This article studied spatio-temporal coverage maximization in WRSNs under probabilistic sensing and mobile-charging

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

constraints. To address the coupling among coverage quality, sensor energy depletion, charger mobility, and online request generation, we proposed SCOPE, a hierarchical event-triggered scheduling framework. SCOPE decomposes the scheduling process into energy–coverage demand partitioning, local adaptive scheduling, and event-triggered cooperative grouping, so that most online decisions are made on local zonal request sets rather than by repeated global optimization over all sensors.

The novelty claim is therefore limited to the N-aware hierarchical process, the scheduling-level Zeno-freeness mechanism, and the scoped PSM kernel guarantee; probabilistic sensing, coverage-aware prioritization, and multi-MC cooperation are treated as established ingredients integrated into the proposed framework.

The theoretical results are scoped to two components of the proposed framework. For the Phase-III fixed-cost single-knapsack budgeted PSM coverage-selection kernel after feasibility pruning, the modified greedy selection with a max-singleton check provides a constant-factor approximation guarantee. For the event-triggered reorganization process, the bounded-drift assumption together with the dwell-time gate ensures scheduling-level Zeno-freeness by enforcing a strictly positive lower bound on the interevent interval. These results do not claim Lyapunov stability of the complete WRSN closed-loop system, including all sensor energy states, charger battery states, request queues, and physical trajectories.

Simulation results under the evaluated finite-horizon settings indicate that SCOPE improves coverage-aware charging performance over representative baselines. Under the default setting, SCOPE maintains an average coverage-quality ratio of approximately 0.93. In the dense-network case with N = 600, SCOPE improves travel-normalized charging efficiency by approximately 18% over MC^3, the strongest baseline in this metric, and keeps the online computation time within approximately 25–45 ms/decision for N = 200–600. The extended simulated scenarios further suggest that SCOPE maintains robust coverage-aware behavior under nonuniform deployments, obstacle-constrained travel, bursty demands, and heterogeneous mobility/energy conditions. When the empirical separation is small or CIs overlap, these results should be interpreted as robustness consistency rather than large statistical dominance.

Several limitations remain. The current validation is simulation-based and does not include hardware field tests. In practical deployments, communication delays and packet losses can render request states stale; localization errors can distort feasible-path and charging-alignment estimates; battery aging can alter effective capacity and discharge behavior; and WPT uncertainty can change service times and delivered energy. Future work will incorporate these effects into hardware-oriented validation and scheduling design.

REFERENCES

[1] Y. Kan, C.-Y. Chang, C.-H. Kuo, and D. S. Roy, "Coverage and connectivity aware energy charging mechanism using mobile charger for WRSNs," *IEEE Syst. J.*, vol. 16, no. 3, pp. 3993–4004, Sep. 2022.

[2] A. Kaswan, P. K. Jana, and S. K. Das, "A survey on mobile charging techniques in wireless rechargeable sensor networks," *IEEE Commun. Surveys Tuts.*, vol. 24, no. 3, pp. 1750–1783, 3rd Quart., 2022.

[3] P. Zhou, C. Wang, and Y. Yang, "Static and mobile target k-coverage in wireless rechargeable sensor networks," *IEEE Trans. Mobile Comput.*, vol. 18, no. 10, pp. 2430–2445, Oct. 2019.

[4] C. Lin, J. Zhou, C. Guo, H. Song, G. Wu, and M. S. Obaidat, "TSCA: A temporal-spatial real-time charging scheduling algorithm for on-demand architecture in wireless rechargeable sensor networks," *IEEE Trans. Mobile Comput.*, vol. 17, no. 1, pp. 211–224, Jan. 2018.

[5] H. Xue, H. Chen, Q. Dai, K. Lin, J. Li, and Z. Li, "CSCT: Charging scheduling for maximizing coverage of targets in WRSNs," *IEEE Trans. Computat. Social Syst.*, vol. 11, no. 3, pp. 3049–3061, Jun. 2024.

[6] C. Lin et al., "GTCCS: A game theoretical collaborative charging scheduling for on-demand charging architecture," *IEEE Trans. Veh. Technol.*, vol. 67, no. 12, pp. 12124–12136, Dec. 2018.

[7] M. Majid et al., "Applications of wireless sensor networks and Internet of Things frameworks in the industry revolution 4.0: A systematic literature review," *Sensors*, vol. 22, no. 6, p. 2087, Mar. 2022.

[8] J. Xu, S. Hu, S. Wu, K. Zhou, H. Dai, and L. Xu, "Cooperative charging as service: Scheduling for mobile wireless rechargeable sensor networks," in *Proc. IEEE 41st Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2021, pp. 685–695.

[9] Z. Wei, M. Li, Q. Zhao, Z. Lyu, S. Zhu, and Z. Wei, "Multi-MC charging schedule algorithm with time windows in wireless rechargeable sensor networks," *IEEE Access*, vol. 7, pp. 156217–156227, 2019.

[10] Z. Hu, C.-Y. Chang, R. Shen, C.-H. Kuo, and D. S. Roy, "A recharging mechanism for improving data quality and maintaining sustainable lifetime in WRSNs," *IEEE Sensors J.*, vol. 23, no. 22, pp. 28244–28260, Nov. 2023.

[11] W. Xu, W. Liang, X. Jia, Z. Xu, Z. Li, and Y. Liu, "Maximizing sensor lifetime with the minimal service cost of a mobile charger in wireless sensor networks," *IEEE Trans. Mobile Comput.*, vol. 17, no. 11, pp. 2564–2577, Nov. 2018.

[12] H. Yu, C.-Y. Chang, Y. Wang, D. S. Roy, and X. Bai, "CAERM: Coverage aware energy replenishment mechanism using mobile charger in wireless sensor networks," *IEEE Sensors J.*, vol. 21, no. 20, pp. 23682–23697, Oct. 2021.

[13] C. Zhang, Y. Li, P. Xu, C.-Y. Chang, and D. S. Roy, "Dynamic probabilistic sensing for enhanced target coverage in WRSNs," *IEEE Sensors J.*, vol. 24, no. 16, pp. 26699–26715, Aug. 2024.

[14] W. Wen, C.-Y. Chang, Q. Zhang, and D. Sinha Roy, "PLPR: Joint perpetual lifetime and path reduction for energy recharging using multiple mobile chargers in WRSNs," *IEEE Sensors J.*, vol. 24, no. 19, pp. 31157–31172, Oct. 2024.

[15] B. Dande, C.-Y. Chang, S.-J. Wu, and D. S. Roy, "WLARS: Workload-aware recharge scheduling mechanism for improving surveillance quality in wireless rechargeable sensor networks," *IEEE Sensors J.*, vol. 23, no. 11, pp. 12237–12250, Jun. 2023.

[16] P. Xu, X. Wang, C.-Y. Chang, Y. Li, and D. S. Roy, "JMLA: Joint monitoring quality and network lifetime for range-adjustable sensors in WRSNs," *IEEE Sensors J.*, vol. 24, no. 20, pp. 33522–33543, Oct. 2024.

[17] P. Tabuada, "Event-triggered real-time scheduling of stabilizing control tasks," *IEEE Trans. Autom. Control*, vol. 52, no. 9, pp. 1680–1685, Sep. 2007.

[18] K. Liang, Y. Wang, Z. Li, G. Zheng, K.-K. Wong, and C.-B. Chae, "Digital twin assisted deep reinforcement learning for computation offloading in AAV systems," *IEEE Trans. Veh. Technol.*, vol. 74, no. 5, pp. 8466–8471, May 2025.

[19] G. Chesi, *Domain of Attraction: Analysis and Control via SOS Programming*. London, U.K.: Springer, 2011.

[20] B. W. Silverman, *Density Estimation for Statistics and Data Analysis*. London, U.K.: Chapman & Hall, 1986.

[21] Z. Hu, X. Gui, C.-Y. Chang, X. Zhang, and D. Sinha Roy, "A multiple mobile chargers collaboratative recharging scheme for enhanced data quailty and sustainable lifetime in WRSNs," *IEEE Sensors J.*, vol. 25, no. 17, pp. 34000–34016, Sep. 2025.

[22] H. Yu, C.-Y. Chang, W. Guan, Y. Wang, and D. Sinha Roy, "MC^3: A multicharger cooperative charging mechanism for maximizing surveillance quality in WRSNs," *IEEE Sensors J.*, vol. 24, no. 10, pp. 17080–17091, May 2024.

[23] Q. Guo, Y. Qiu, Y. Shi, Q. Zheng, and H. Guo, "An enhanced dual-population NSGA-II for solving multi-objective emergency supplies dispatch problem based on truck-drone-dispatch vehicle collaboration," *Comput. Oper. Res.*, vol. 187, Mar. 2026, Art. no. 107324.

[24] S. Khuller, A. Moss, and J. Naor, "The budgeted maximum coverage problem," *Inf. Process. Lett.*, vol. 70, no. 1, pp. 39–45, Apr. 1999.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.

[25] J. Cortes, S. Martinez, T. Karatas, and F. Bullo, "Coverage control for mobile sensing networks," *IEEE Trans. Robot. Autom.*, vol. 20, no. 2, pp. 243–255, Apr. 2004.
[26] B. Dande, C.-Y. Chang, C.-H. Kuo, and D. S. Roy, "RS-STQ: Recharge scheduling algorithm for maximizing spatial and temporal data accuracy," *IEEE Sensors J.*, vol. 22, no. 23, pp. 23565–23580, Dec. 2022.
[27] P. Xu, J. Wu, C. Shang, and C.-Y. Chang, "GSMS: A barrier coverage algorithm for joint surveillance quality and network lifetime in WSNs," *IEEE Access*, vol. 7, pp. 159608–159621, 2019.
[28] P. Xu, J. Wu, C.-Y. Chang, C. Shang, and D. S. Roy, "MCDP: Maximizing cooperative detection probability for barrier coverage in rechargeable wireless sensor networks," *IEEE Sensors J.*, vol. 21, no. 5, pp. 7080–7092, Mar. 2021.
[29] P. Xu, Y. Li, C.-Y. Chang, C.-H. Kuo, and D. S. Roy, "Dynamic weather-adaptive enhanced barrier coverage with adjustable-range sensors for WRSNs," *J. Internet Technol.*, vol. 26, no. 4, pp. 523–545, Jul. 2025.
[30] H. Xue, H. Chen, Z. Ni, X. Liu, and F. Xia, "Towards maximizing coverage of targets for WRSNs by multiple chargers scheduling," *IEEE Trans. Mobile Comput.*, vol. 23, no. 10, pp. 9959–9970, Oct. 2024, doi: 10.1109/TMC.2024.3369054.
[31] M. U. F. Qaisar et al., "Poised: Probabilistic on-demand charging scheduling for ISAC-assisted WRSNs with multiple mobile charging vehicles," *IEEE Trans. Mobile Comput.*, vol. 23, no. 12, pp. 10818–10834, Dec. 2024, doi: 10.1109/TMC.2024.3382668.
[32] X. Qiu et al., "Joint device charging and fresh data retrieval with mobile edge device in wireless-powered IoT systems," *IEEE Trans. Consum. Electron.*, vol. 70, no. 4, pp. 7385–7397, Nov. 2024, doi: 10.1109/TCE.2024.3419128.
[33] X. Zhang, R. Jia, Q. Yin, Z. Zheng, and M. Li, "Intelligent trajectory design and charging scheduling in wireless rechargeable sensor networks with obstacles," *IEEE Trans. Mobile Comput.*, vol. 23, no. 9, pp. 8664–8679, Sep. 2024, doi: 10.1109/TMC.2024.3350075.



Pei Xu (Member, IEEE) received the Ph.D. degree in computer science and technology from Anhui University, Hefei, China, in 2020.

He is currently a Lecturer with the School of Artificial Intelligence and Big Data, Hefei University, Hefei. His current research interests include smart sensing technology, the Internet of Things, and deep learning.



Shilong Han received the B.S. degree in software engineering from Bengbu University, Bengbu, China, in 2022. He is currently pursuing the M.S. degree with the School of Artificial Intelligence and Big Data, Hefei University, Hefei, China.

His research interests include wireless sensor networks, the Internet of Things (IoT), and optimization theory.



Zijun Wu received the Ph.D. degree from Clausthal University of Technology, Clausthal-Zellerfeld, Germany, in 2015.

He is currently a Full Professor with Hefei University of Technology, Hefei, China. His research focuses mainly on the field of operations research, including algorithm design and analysis, algorithmic game theory, transportation, and machine learning.



Chih-Yung Chang (Member, IEEE) received the Ph.D. degree in computer science and information engineering from National Central University, Taoyuan City, Taiwan, in 1995.

He is currently a Distinguished Professor with the Department of Computer Science and Information Engineering and the Department of Artificial Intelligence, Tamkang University, New Taipei City, Taiwan. His current research interests include artificial intelligence, deep learning and machine learning, the Internet of Things, and wireless sensor networks.

Dr. Chang is the Co-Chair of ACM SIGMOBILE, Taiwan. He served/has been serving as an Associate Guest Editor for several SCI-indexed journals, including *Journal of Internet Technology* (JIT) from 2004 to 2008, *Journal of Information Science and Engineering* (JISE) in 2008, *Telecommunication Systems* (TS) in 2010, *IET Communications* in 2011, *International Journal of Distributed Sensor Networks* (IJDSN) from 2012 to 2014, *International Journal of Ad Hoc and Ubiquitous Computing* (IJAHUC) since 2011, and *Journal of Applied Science and Engineering* (JASE) since 2018.



Diptendu Sinha Roy (Senior Member, IEEE) received the Ph.D. degree in engineering from the Birla Institute of Technology, Mesra, Ranchi, India, in 2010.

In 2016, he joined the Department of Computer Science and Engineering, National Institute of Technology (NIT) Meghalaya, Shillong, India, as an Associate Professor, where he has been the Chair of the Department of Computer Science and Engineering since January 2017. Prior to his stint at NIT Meghalaya, he was with the Department of Computer Science and Engineering, National Institute of Science and Technology, Brahmapur, India. His current research interests include software reliability, distributed and cloud computing, and the Internet of Things (IoT), specifically applications of artificial intelligence/machine learning for smart integrated systems.

Dr. Roy is a member of the IEEE Computer Society.

Authorized licensed use limited to: Tamkang Univ.. Downloaded on October 06,2026 at 06:26:38 UTC from IEEE Xplore. Restrictions apply.