

Article

WPT-JCCO: Co-Optimisation of Communication and Computation Cost Through Advanced Wireless-Power Transfer Strategies for Swarm Robotics

Amir Ijaz ^{1,*}, Hashem Haghbayan ¹, Ethiopia Nigussie ²  and Juha Plosila ¹¹ Department of Computing, University of Turku, FI-20014 Turku, Finland² Faculty of Technology, University of Turku, FI-20014 Turku, Finland

* Correspondence: amir.ijaz@utu.fi

Abstract

Wireless-power mobile edge computing, SWIPT-MEC, priority-aware WPT scheduling and swarm resource allocation already solve important parts of the energy-management problem. The novelty of WPT-JCCO is not any one of those elements; it is a single swarm-supervisory feasible set that couples decisions which the three adjacent method classes normally separate. Each epoch-level action jointly selects the robot to charge and one of three physically distinct WPT modalities: far-field radio-frequency, resonant near-field and directional lightwave transfer, together with the SWIPT split, local/edge task placement, CPU frequency, bandwidth and transmit power. Relative to SWIPT-MEC, the formulation adds discrete recipient-modality selection with pose, alignment, blockage and dwell-dependent feasibility. Relative to conventional WPT scheduling, charging is not a separate priority or routing stage but is solved jointly with computation and radio allocation. Relative to swarm resource-allocation methods, energy replenishment is endogenous and an individual minimum-battery constraint protects the weakest robot. A fourth coupling makes the centrally generated resource vector admissible only when the complete sense-compute-actuate age fits the one-second supervisory epoch; otherwise a previously feasible or local-safe action is applied. Nonlinear harvesting, partial offloading, priority scoring and augmented-Lagrangian primal-dual updates are treated as established techniques. This paper derives the continuous block updates, keeps the WPT variables binary through candidate screening, and declares convergence only when stationarity, feasibility, merit-change and binary-hold tests are jointly satisfied. Normalised primal steps are safeguarded by backtracking, dual and penalty updates are bounded, and a local tracking bound plus divergence monitor delimit real-time operation without claiming global mixed-integer optimality or closed-loop motion stability. Numerical evaluation over a 20-robot swarm and 30 Monte Carlo runs shows that WPT-JCCO reduces net energy depletion by 23.8% relative to communication-computation optimisation with static WPT and by 49.7% relative to local-only execution, while increasing task success from 93.5% to 97.3%. A released common-trace comparison shows normalised-cost reductions of 11.1%, 11.3% and 5.8% relative to two-stage WPT+CCO, fixed-SWIPT dynamic offloading and an offline Q-learning scheduler. Convergence and one-factor-at-a-time sensitivity studies further examine swarm size, task load, WPT budget, bandwidth, edge capacity, mobility and channel margin. The headline values remain scoped to the nominal independent-task case; mode-specific RF, near-field and lightwave operating envelopes, robust pose/CSI, WPT-safety and task-DAG extensions are formulated but not presented as hardware-validated results.



Academic Editors: Pablo García Triviño, Cancan Rong, Xiaosheng Wang, Jiayu Zhou and Yibo Wang

Received: 3 June 2026

Revised: 17 June 2026

Accepted: 24 June 2026

Published: 26 June 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: swarm robotics; wireless power transfer; simultaneous wireless information and power transfer; mobile edge computing; computation offloading; robust resource allocation; task graphs; alignment uncertainty; primal-dual optimisation

1. Introduction

Swarm robotics relies on coordinated behaviour among multiple robots that share sensing, computation, communication and energy resources under decentralised or partially decentralised control [1]. A swarm can cover a larger area than a single robot, tolerate individual failures and execute distributed sensing or inspection missions. These advantages diminish when low-energy robots can no longer communicate, process assigned data or maintain motion. Energy management therefore affects not only average consumption but also mission continuity, task completion and the number of robots that remain available to the collective.

Mobile edge computing (MEC) and adaptive wireless networking reduce the burden on embedded processors, but they create a coupled allocation problem. Offloading a perception task can save local CPU energy while consuming radio energy and edge capacity; local execution avoids radio contention but may drain a robot with a computation-intensive workload. Communication–computation co-optimisation (CCO) addresses this coupling by jointly controlling offloading, transmit power, bandwidth and processor frequency [2–5]. In most swarm-oriented MEC studies, however, residual energy remains a state to be conserved rather than a resource that can also be replenished during operation.

The integration of WPT and MEC is itself a mature research direction. WP-MEC studies have jointly optimised energy beamforming, local computing and task offloading under energy and task causality [6]; coordinated wireless charging with partial or binary offloading under latency constraints [7]; maximised residual energy with mixed offloading [8]; and combined communication and computation cooperation in NOMA-based wireless-power networks [9]. A comprehensive survey documents the breadth of this literature [10]. More recently, UAV-assisted SWIPT-MEC has coupled terminal battery sustainability, nonlinear harvesting and computational-resource management [11]. Consequently, neither “joint WPT and offloading” nor “minimum net energy depletion” is claimed here as a new research question.

The remaining difficulty is specific to online swarm operation. A mobile robot is simultaneously a task source, a computing node, a radio terminal, an actuator and a mission-critical member whose loss can alter collective capability. Existing WP-MEC formulations primarily optimise network-level energy, computation rate, latency or energy-transmitter expenditure for generic wireless devices. They normally assume one wireless-power architecture or one harvest-then-compute protocol and do not decide among physically different charging modes. Conversely, swarm-MEC and cloud-robotics studies model robot mobility, cooperation and control traffic but do not actively replenish robot batteries [4,5]. The unresolved control question is therefore: how should a supervisory controller jointly select the robot to charge, the applicable WPT mode, the SWIPT split, the local/offloaded workload and the radio/CPU allocation so that urgent tasks are completed without allowing a weak robot or a stale base-station command to become a mission bottleneck?

Three non-ideal effects materially change this decision. First, a robot can translate or rotate during a charging interval, reducing magnetic coupling or moving a directional lightwave receiver outside the beam footprint. Second, communication and WPT gains are estimated from finite pilot and pose information; an optimistic estimate can produce an infeasible rate, decoding split or harvested-energy prediction. Third, robotic perception and

planning applications are often pipelines rather than divisible data blocks: a successor stage cannot start before its predecessors complete, and when consecutive stages execute at different locations their intermediate data must also be transmitted. The revised formulation therefore separates the nominal model used for the reported Monte Carlo comparison from a deployment-oriented robust counterpart with pose/alignment uncertainty, conservative channel gains and directed acyclic graph (DAG) precedence constraints.

This paper addresses that question through WPT-JCCO, a WPT-aware joint communication–computation–charging framework for energy-constrained swarm robotics. Far-field RF, resonant near-field and directional lightwave WPT are represented as alternative actions with different range, alignment and efficiency profiles. Their scheduling is coupled with SWIPT power splitting, local CPU frequency, offloading or stage-placement decisions, radio bandwidth and transmit power. The objective is not only to reduce aggregate depletion but also to protect the minimum robot battery, satisfy task and packet-delivery constraints, and ensure that the resulting resource action remains fresh enough to be applied within the one-second supervisory epoch.

Direct novelty statement: WPT-JCCO differs from the three closest method classes in the decision that is optimised. SWIPT-MEC normally optimises energy–information splitting, power and offloading inside one RF transfer architecture; WPT-JCCO additionally chooses the recipient robot and the physical WPT modality, and it makes that binary choice conditional on predicted pose, alignment, blockage, safety and charging dwell. Priority-aware WPT scheduling normally selects who or when to charge before or independently of computation and communication allocation; WPT-JCCO places the charging action in the same constrained problem as task placement, CPU allocation, bandwidth, transmit power and SWIPT decoding. Swarm resource-allocation methods normally conserve a finite battery while assigning tasks or communication resources; WPT-JCCO makes replenishment an endogenous control variable and couples it to an individual minimum-battery constraint so that the aggregate optimum cannot repeatedly sacrifice the same robot. The additional control-age constraint rejects a centrally computed action when its state information or command delivery has become stale.

The genuinely new object is therefore the coupled mixed discrete–continuous feasible set, not the separate ingredients. Partial offloading, CPU-frequency control, bandwidth and power allocation are inherited from MEC-CCO; power splitting and nonlinear harvesting are inherited from SWIPT-MEC; priority-weighted charging and projected primal–dual iteration are established. The algorithmic contribution is a problem-specific decomposition that screens integral recipient–mode candidates, updates the continuous MEC/SWIPT blocks, enforces minimum-battery and timing feasibility, and invokes a previous-feasible or local-safe fallback. It is not presented as a new general-purpose optimisation family.

The contributions are as follows:

1. A direct novelty boundary is established against SWIPT-MEC, WPT scheduling and swarm resource allocation. The paper does not claim novelty for offloading, SWIPT splitting, nonlinear harvesting, priority scoring or primal–dual optimisation; it claims the joint swarm-supervisory decision that makes charging modality, computation, communication, weakest-robot battery safety and command freshness mutually constraining.
2. A mixed robot–modality action is introduced. The binary decision selects both the recipient and one of three physically distinct charging modes. A mode-specific operating envelope accounts for translation and rotation, coil or beam misalignment, path clearance, hardware-interlock state, protected-zone exposure and charging dwell. RF may continue at reduced gain when a directional mode is rejected; near-field and lightwave actions are inhibited when coupling or tracking cannot be maintained.

The accepted charging action is then optimised jointly with task placement and radio/CPU variables rather than in a separate charging stage.

3. Aggregate depletion is coupled to two swarm-specific feasibility mechanisms: an individual minimum-battery constraint with a fairness multiplier, and an end-to-end sense–compute–actuate age constraint with previous-feasible and local-safe fallbacks. These mechanisms prevent the optimiser from improving the mean by repeatedly depleting one robot or applying a stale central command.
4. A tailored mixed discrete–continuous decomposition is derived from an inequality-augmented Lagrangian. Binary WPT variables remain integral through candidate screening and merit comparison; continuous SWIPT, offloading, CPU and radio blocks use projected first-order or closed-form updates, and multipliers use projected residual ascent. The revision specifies a four-part convergence declaration, normalised parameter-selection rules, a local time-varying tracking bound and a residual-divergence monitor. Conditional stationarity is stated for the fixed-binary continuous subproblem, while neither global mixed-integer optimality nor safety-critical motion stability is claimed.
5. A 20-robot, 30-seed numerical study, stronger common-trace baselines, component ablation, convergence sweeps and a seven-factor sensitivity study evaluated the proposed coupling. The released data and code cover swarm size, task intensity, WPT power, bandwidth, edge capacity, mobility and channel uncertainty. Robust pose/CSI, WPT-safety and DAG extensions were formulated for deployment, while the headline energy and success values remained scoped to the nominal independent-task model.

The paper is organised as follows. Section 2 reviews MEC-CCO, SWIPT-MEC and WPT-enabled robotics, then provides component-level novelty attribution. Section 3 presents the system model. Section 4 formulates the optimisation problem. Section 5 describes the proposed algorithm. Section 6 gives the simulation setup. Section 7 analyses the results. Section 8 discusses interpretation, limitations and deployment. Section 9 concludes the paper.

2. Related Work

2.1. Communication–Computation Control in MEC and Robotics

Joint communication and computation management is well established in MEC. Stochastic radio and computational-resource allocation has been used to coordinate local execution, offloading and transmission under channel and task-buffer uncertainty [2], while broader surveys classify offloading, radio allocation and green MEC as central design problems [3]. In robot and UAV systems, mobility adds time-varying channels, task priorities and coordination requirements. The broader UAV-communication literature establishes mobility, trajectory and link variation as coupled design variables [12]. UAV-swarm MEC work has subsequently addressed task offloading and data compression [4], while wireless cloud-robotics research has coupled communication and control decisions through multi-agent learning [5]. These methods reduce consumed energy or delay, but they do not schedule energy replenishment as part of the robot-control action.

Mobility can also be learned rather than treated only as an exogenous trace. Cao et al. study resource-constrained small-cell networks in which user routes and content preferences are initially unknown [13]. Their federated routing and popularity learning procedure allows neighbouring small-cell base stations to learn mobility-conditioned service demand without centralising raw user observations, after which a greedy cache-placement step approximates the service-cost minimiser. This work is relevant to the prediction layer of mobile edge systems because it shows how distributed mobility and demand estimates can guide resource placement under limited bandwidth and storage. Its decision variables are

nevertheless routing and caching, not robot charging, SWIPT, computation offloading or battery safety. WPT-JCCO presently consumes reported or predicted robot pose, queue and battery states; it does not claim a new federated-learning method.

For online robotic supervision, solver complexity alone does not establish deployability. State reports may be stale because of medium access, queueing and retransmissions, and the resulting command may arrive after the state used to compute it has changed. Practical networked-control scheduling studies show that communication-stack delay and information freshness affect closed-loop performance [14]. Distributed multi-robot control can similarly lose safety margins under delayed communication [15], and 5G service requirements treat latency, availability and reliability as coupled properties of cyber-physical communication [16]. These findings motivate the explicit sense–compute–actuate constraint used in WPT-JCCO.

Recent safety-critical multi-agent control illustrates a complementary design direction. Luo et al. applied an improved proximal-policy-optimisation controller to adaptive orbit–attitude takeover by cellular satellites [17]. That study addressed safety-critical continuous control of cooperating autonomous agents, whereas the present work addresses a slower supervisory layer that allocates charging, communication and computation resources. The comparison is important because an energy-aware supervisor cannot replace the inner safety controller: WPT-JCCO must provide a timely feasible resource action while collision avoidance, stabilisation and actuator-level safety remain local to the robots.

2.2. Wireless-Power MEC, SWIPT and Residual-Energy Optimisation

Wireless-power communication established the basic opportunities and constraints of supplying energy through the radio interface [18]. WP-MEC then integrated wireless charging with computation offloading, and its scope is considerably broader than passive energy harvesting. Representative formulations jointly optimise energy beamforming, local computation and remote task execution under energy and task causality [6]; wireless charging and binary or partial offloading under a latency deadline [7]; and mixed offloading to maximise residual energy [8]. Of particular relevance, Zeng et al. formulated a three-node NOMA WP-MEC system in which a near user cooperatively computes and forwards part of a far user’s workload, jointly optimising AP transmit power, offloading coefficients and stage durations to mitigate the double-near–far effect [9]. The survey by Wang et al. organised this mature body of work around energy transfer, offloading modes, resource allocation, multiple access and learning-based control [10]. These studies demonstrate that joint charging and offloading is not, by itself, a sufficient novelty claim.

SWIPT adds a direct energy–information trade-off. UAV-assisted SWIPT-MEC has recently combined directional transmission, nonlinear harvesting, dynamic tasks, terminal battery sustainability and computational-resource allocation through deep reinforcement learning [11]. WPT technologies also differ substantially in range, alignment and conversion efficiency: resonant near-field transfer can be efficient over short gaps, far-field RF provides broader coverage at lower received power, and directional optical or microwave links require line of sight and beam control [19,20]. Most WP-MEC models select power, time or beamforming within one transfer architecture; they do not select among heterogeneous charging modalities according to a robot’s location, urgency and battery condition.

Physical uncertainty is consequential in both halves of the problem. Coil displacement and angular error change mutual coupling and therefore the output power of resonant WPT, while directional optical transfer is sensitive to pointing error, beam footprint and blockage [20]. Robust wireless-power resource allocation has consequently been studied under imperfect channel-state information and nonlinear harvesting [21]; imperfect channel estimation has also been incorporated directly into UAV-assisted MEC offloading and band-

width allocation [22]. On the computation side, many applications comprise dependent stages. DAG-based MEC studies show that ignoring precedence and intermediate-data transfer can create infeasible placements or understate application completion time [23]. These results motivate the uncertainty sets and precedence constraints introduced below.

Priority-aware charging is also established outside MEC. Temporal–distantal priority scheduling ranks charging requests using residual operating time and charger distance [24], while deadline-constrained multi-node charging maximises effective charging utility under request deadlines [25]. Related real-time charging schemes similarly combine energy state, urgency or travel cost when selecting the next receiver. These studies make clear that a weighted score formed from battery condition, queue pressure and deadline information is a familiar scheduling device. Accordingly, the priority index used later in WPT-JCCO is not presented as a theoretical novelty or as an optimal solution of the mixed-integer problem. Its purpose is to provide an interpretable, low-overhead recipient-ranking rule inside the wider cross-layer controller.

2.3. Direct Novelty Boundary Against Adjacent Method Classes

Table 1 shows that the three adjacent classes stop at different boundaries. SWIPT-MEC couples information decoding, harvesting and offloading but generally assumes one RF transfer architecture and does not choose a mobile robot–modality pair. WPT scheduling selects a receiver, order, path or charging power but normally treats task placement and shared communication/computation resources as external or subsequent decisions. Swarm resource-allocation methods jointly assign tasks, routes, bandwidth or processing capacity, but generally treat battery energy as a finite state to conserve rather than a replenishment action to optimise. WPT-JCCO crosses these boundaries by defining one supervisory action over charging recipient and modality, SWIPT split, task placement, CPU, bandwidth and transmit power.

Table 1. Representative adjacent studies and the swarm-control requirements not jointly covered by them.

Study and Setting	Established Contribution	Requirement Remaining for WPT-Aware Swarm Supervision
Wang et al. [6], multi-user WP-MEC	Real-time joint energy beamforming, local computing and offloading under task and energy causality	Generic terminals and a single RF charging architecture; no charging-mode choice, robot-level mission role, SWIPT split or end-to-end control-age constraint
Malik and Vu [7], multi-user MEC	Joint wireless charging and partial/binary offloading under latency constraints	No mobile-swarm model, mode-dependent WPT selection, packet-delivery requirement or minimum-robot-battery protection
Wu et al. [8], WP-MEC	Residual-energy maximisation with mixed offloading	Residual energy is optimised for wireless devices without task urgency, heterogeneous WPT modes or freshness of a supervisory robot command
Lin et al. [24] and Yang et al. [25], rechargeable sensor networks	Priority- and deadline-aware charging-request scheduling using residual operating time, distance, utility and deadline information	No joint local/edge computation, SWIPT splitting, heterogeneous WPT-mode decision, robot-level battery-fairness constraint or supervisory control-age model
Zeng et al. [9], NOMA WP-MEC	Three-node harvest-then-offload protocol with NOMA-based communication/computation cooperation; joint AP power, offloading and stage-duration optimisation mitigates the double-near-far effect	Cooperative fixed wireless users and one RF energy-transfer protocol; no mobile-swarm role, hybrid WPT-mode selection, minimum-robot-battery constraint or supervisory control-age model
Luo et al. [17], cellular-satellite control	Improved-PPO-based adaptive safety-critical orbit–attitude takeover for cooperating autonomous spacecraft	Addresses low-level safety-critical dynamics and control, not wireless charging, MEC offloading, SWIPT or supervisory energy/resource allocation; complementary rather than a direct energy baseline

Table 1. Cont.

Study and Setting	Established Contribution	Requirement Remaining for WPT-Aware Swarm Supervision
Chen et al. [11], UAV-assisted SWIPT-MEC	Joint UAV energy, terminal battery sustainability, nonlinear harvesting and computational-resource control	A UAV serves ground IoT terminals through one directional SWIPT architecture; no multi-mode robot charging, minimum-battery fairness dual or base-station control-age/fallback model
Cao et al. [13], mobility-aware small-cell routing/caching	Federated learning of user routing and content popularity under unknown mobility/preferences, followed by service-cost-oriented cache placement	Learns mobility-conditioned demand but does not allocate robot computation, radio/WPT resources, enforce per-robot battery safety or model supervisory control age; complementary prediction layer rather than a direct charging baseline
Hu et al. [4] and Xu et al. [5], swarm/cloud robotics	Mobility-aware offloading, compression, and communication–control coordination	Robot dynamics and communication are represented, but active WPT recipient/mode selection and SWIPT are absent
WPT-JCCO	Recipient-modality selection across RF/NF/lightwave WPT, coupled SWIPT/MEC allocation, minimum-battery dual and age-gated supervisory action	The joint feasible-set coupling is the claimed contribution; robust pose/CSI and DAG extensions still require measured hardware and workload validation

The novelty is carried by three constraints and their coupling. First, the integral variable $u_{i,m}$ selects a robot-mode pair and is bounded by the mode-availability gate $z_{i,m}$; therefore, a high-priority robot cannot receive a near-field or lightwave command that will violate alignment, blockage, safety or dwell requirements. Second, the individual minimum-battery residual and its fairness multiplier feed back into charging, offloading, CPU and radio allocation, so weakest-robot protection is not a post-processing rule. Third, $T_{\text{loop}} \leq \Delta t - T_{\text{guard}}$ makes state collection, optimisation and command delivery part of resource feasibility and activates a defined fallback when the central decision is stale. These are the paper-specific modelling mechanisms. The familiar constituent methods—partial offloading, SWIPT splitting, nonlinear harvesting, priority ranking and augmented-Lagrangian updates remain attributed to the prior literature.

The federated mobility-learning contribution of Cao et al. is adjacent but distinct: it estimates mobility-conditioned routing and demand statistics, whereas WPT-JCCO maps an available state estimate to a constrained charging–communication–computation action. A future implementation may combine these layers, but the present work neither learns the mobility model federatively nor attributes its resource-allocation contribution to federated learning. Tables 1 and 2 state the study-level and component-level boundaries, respectively. The positioning relies on independently authored adjacent studies; prior publications by the present authors are not used to establish the gap or the performance claim.

Table 2. Direct distinction between WPT-JCCO and the three adjacent method classes.

Adjacent Class	Typical Optimisation Boundary	Difference Introduced by WPT-JCCO
SWIPT-MEC	Joint power splitting, harvesting, offloading and communication/computation allocation, usually within one RF or directional SWIPT architecture [6,9,11]	Adds an integral robot–modality choice across RF, resonant near-field and lightwave WPT; the choice must satisfy pose, alignment, blockage, safety and dwell feasibility and is coupled to weakest-robot battery protection and command age.
Priority-aware or mobile WPT scheduling	Selects receiver, charging order, route, dwell or transmit power from residual energy, distance, urgency or deadline information [24,25]	Does not stop after choosing who to charge: charging is solved in the same constrained problem as task placement, CPU frequency, edge allocation, bandwidth, transmit power and SWIPT decoding.

Table 2. Cont.

Adjacent Class	Typical Optimisation Boundary	Difference Introduced by WPT-JCCO
Swarm resource allocation and cloud robotics	Allocates tasks, communication, compression, routes or edge resources while conserving a finite onboard battery [4,5]	Makes energy replenishment an endogenous action and couples aggregate depletion to an individual minimum-battery residual, preventing repeated sacrifice of a mission-critical robot.
Established optimisation primitives	Partial offloading, nonlinear harvesting, priority scores, projection and primal–dual or augmented-Lagrangian updates are established	Reused without standalone novelty claims. The algorithmic contribution is the problem-specific candidate screening, coupled continuous updates, feasibility/timing acceptance and fallback sequence.
Paper-specific contribution	The adjacent classes contain individual ingredients or pairwise couplings	A single mixed discrete–continuous supervisory feasible set links recipient/modality availability, communication–computation allocation, minimum-battery fairness and end-to-end action freshness.

3. System Model

Consider a swarm of N robots indexed by $i \in \mathcal{N} = \{1, \dots, N\}$ operating in discrete epochs $t \in \{1, \dots, T\}$ of duration Δt . Each robot has a battery state $B_i(t)$, position $\mathbf{q}_i(t)$, orientation $\boldsymbol{\theta}_i(t)$ and a queue of sensing or perception applications. The nominal model represents an application by $\mathcal{J}_i(t) = \{D_i(t), L_i(t), c_i(t), \tau_i^{\max}(t)\}$, where $D_i(t)$ is the input data size in bits, $L_i(t)$ is the output or result size, $c_i(t)$ is the required CPU cycles per bit and $\tau_i^{\max}(t)$ is the latency deadline. A dependent application is represented more generally by a DAG $\mathcal{G}_i(t) = (\mathcal{V}_i(t), \mathcal{E}_i(t))$. Each stage $v \in \mathcal{V}_i(t)$ has descriptor $\mathcal{J}_{i,v}(t) = \{D_{i,v}, L_{i,v}, c_{i,v}\}$, and an edge $(u, v) \in \mathcal{E}_i(t)$ means that stage v cannot start before stage u and any required intermediate-data transfer have completed. The scalar task model is the one-node or fully parallelisable special case of this graph representation.

The system contains an edge access point with computational capacity $F^{\text{edge}}(t)$ and a hybrid WPT hub. The WPT hub supports a finite set of modes $\mathcal{M} = \{\text{RF}, \text{NF}, \text{LW}\}$, representing far-field RF WPT, resonant near-field WPT and directional lightwave WPT. The access point can also support SWIPT in the downlink, where a robot splits received power between information decoding and energy harvesting. Figure 1 summarises the layered architecture.




Figure 1. Layered WPT-JCCO architecture. The energy-transfer, communication/SWIPT, computation, and optimisation layers exchange state feedback and validated resource actions among the robots, access point, edge processor, and WPT hub.

3.1. Mobility, Pose Prediction and WPT Availability

Robot motion is not assumed to preserve charger alignment throughout an epoch. A first-order predictor is

$$\mathbf{q}_i(t+1) = \mathbf{q}_i(t) + \mathbf{v}_i(t)\Delta t + \mathbf{w}_i^q(t), \quad (1)$$

$$\boldsymbol{\vartheta}_i(t+1) = \boldsymbol{\vartheta}_i(t) + \boldsymbol{\omega}_i(t)\Delta t + \mathbf{w}_i^\vartheta(t), \quad (2)$$

where $\mathbf{v}_i$ and $\boldsymbol{\omega}_i$ are the estimated translational and angular velocities and $\mathbf{w}_i^q$ and $\mathbf{w}_i^\vartheta$ collect unmodelled acceleration, localisation error and attitude error. During the candidate charging interval $\mathcal{I}_t = [t\Delta t, (t+1)\Delta t]$, the true pose is assumed to lie in

$$\mathcal{U}_i(t) = \left\{ (\mathbf{q}, \boldsymbol{\vartheta}) : \|\mathbf{q} - \hat{\mathbf{q}}_i(\tau)\|_2 \leq \varepsilon_i^q, \|\boldsymbol{\vartheta} - \hat{\boldsymbol{\vartheta}}_i(\tau)\|_2 \leq \varepsilon_i^\vartheta, \tau \in \mathcal{I}_t \right\}. \quad (3)$$

The bounds can be obtained from localisation covariance, maximum acceleration, controller tracking error or measured prediction residuals. Obstruction is represented separately from pose error. Let $\chi_{i,m}^{\text{clr}}(\tau) \in [0, 1]$ denote the fraction of the nominal transfer path that remains usable at time τ : it is one for a clear path, zero for complete blockage and may take an intermediate value for partial optical occlusion, metallic shielding or multipath attenuation. From a perception or occupancy predictor, the controller forms the lower confidence value

$$\underline{\chi}_{i,m}^{\text{clr}}(t) = \max \left\{ \min_{\tau \in \mathcal{I}_t} \hat{\chi}_{i,m}^{\text{clr}}(\tau) - \varepsilon_{i,m}^\chi, 0 \right\}. \quad (4)$$

A mobility-sensitive WPT mode m is declared available only if its blockage-adjusted worst-case transfer gain remains above a calibrated threshold throughout the scheduled dwell interval:

$$z_{i,m}(t) = \mathbf{1} \left\{ \underline{\chi}_{i,m}^{\text{clr}}(t) \min_{(\mathbf{q}, \boldsymbol{\vartheta}) \in \mathcal{U}_i(t)} g_{i,m}^{\text{wpt}}(\mathbf{q}, \boldsymbol{\vartheta}) \geq g_m^{\min} \right\}, \quad u_{i,m}(t) \leq z_{i,m}(t). \quad (5)$$

When a calibrated blockage distribution is available, the same gate can be written as the chance constraint $\Pr\{g_{i,m}^{\text{eff}} < g_m^{\min}\} \leq \epsilon_m^{\text{blk}}$. Thus, a robot predicted to leave the near-field coupling region, lose lightwave line of sight, enter an occluded corridor or exceed the allowed pointing error is not scheduled with that mode. RF WPT can remain available at a lower transfer gain when the directional modes are rejected.

Mode-Specific Operating Envelopes on Moving Robots:

The three charging modes as outlined in Table 3 do not react to robot motion in the same way. To distinguish geometric feasibility from transfer-gain uncertainty, let $v_m^{\max}$ and $\omega_m^{\max}$ denote calibrated translational- and angular-speed limits for mode m , and let $s_{i,m}^{\text{hw}}(t) \in \{0, 1\}$ denote the status of the mode-specific hardware protection chain. The final operational gate is

$$\bar{z}_{i,m}(t) = z_{i,m}(t) \mathbf{1} \left\{ \|\mathbf{v}_i(t)\|_2 \leq v_m^{\max}, \|\boldsymbol{\omega}_i(t)\|_2 \leq \omega_m^{\max}, s_{i,m}^{\text{hw}}(t) = 1 \right\}, \quad u_{i,m}(t) \leq \bar{z}_{i,m}(t). \quad (6)$$

For RF, the speed bounds can be loose when a broad beam or electronically steered array maintains coverage; the dominant penalties are antenna-orientation loss, shadowing and exposure. Resonant near-field charging normally requires the robot to enter a short-range coupling zone and hold a sufficiently stable pose, so the admissible speed, coil gap, lateral offset and angular error are substantially tighter. Directional lightwave charging requires continuous line of sight and closed-loop beam tracking; the gate is removed immediately when the receiver leaves the beam footprint or an obstacle enters the protected path. The hardware state $s_{i,m}^{\text{hw}}$ covers the checks that must remain outside the optimiser, including

RF or magnetic-field protection, foreign-object and thermal detection for near-field coupling, and the independent optical shutter or beam-stop interlock for lightwave transfer.

Table 3. Operational interpretation of the three WPT modes for moving swarm robots. Thresholds are hardware- and site-specific and must be calibrated before deployment.

Mode	Motion and Misalignment	Blockage and Fallback	Safety and Hardware Enforcement
Far-field RF	A broad or steered beam can tolerate moderate translation. Robot attitude and polarisation enter $A_{i,RF}$; the scheduler uses the lowest predicted gain over the dwell interval.	Walls, metallic structures, people or other robots can introduce shadowing rather than a purely binary LOS loss. If the lower gain bound fails, the controller reduces power, changes beam/recipient, defers charging or selects no charge.	Protected-location electric/magnetic-field or incident-power-density limits are hard constraints. Antenna power, duty cycle, leakage and electromagnetic compatibility require measured calibration; the transmitter-power budget alone does not establish compliance.
Resonant near-field	The robot enters a coupling or docking zone and slows or holds position. Coil separation, lateral offset, angular error, vibration and angular speed determine $A_{i,NF}$ and the permitted dwell.	Nearby conductive objects, shielding, another robot or departure from the coupling zone can reduce efficiency or cause unintended heating. A failed coupling, foreign-object or thermal check aborts the action and returns to RF, deferred charging or no charge.	Magnetic-field exposure, induced currents, receiver temperature and foreign-object heating are enforced by independent charger/receiver protection. The optimiser may request a transfer only while these protection states remain healthy.
Directional light-wave	A tracking unit points the beam at the photovoltaic receiver using predicted pose and feedback. Translation, rotation and vibration must keep the receiver within the beam footprint and angular acceptance cone.	Continuous LOS is required. A person, robot, obstacle, dust/smoke event or tracking loss causes immediate beam inhibition; the scheduler then chooses RF, postpones charging or applies the no-charge fallback.	Accessible optical irradiance, exclusion zones and product classification must satisfy the applicable IEC 60825 safeguards. A fail-safe shutter or beam stop operates independently of network and optimisation software.

3.2. Communication Model

Let $p_i(t)$ denote the uplink transmit power, $b_i(t)$ the allocated bandwidth and $h_i(t)$ the channel power gain between robot i and the access point. The state report provides the estimate

$$\hat{h}_i(t) = G_0 \left(\frac{d_0}{d_i(t)} \right)^\gamma 10^{-\hat{\xi}_i(t)/10}, \quad (7)$$

where G_0 is the reference gain at distance d_0 , $d_i(t)$ is the robot–access-point distance, γ is the path-loss exponent and $\hat{\xi}_i(t)$ is the estimated shadowing term. The actual channel is $h_i(t) = \hat{h}_i(t) + e_i^h(t)$ with $|e_i^h(t)| \leq \varepsilon_i^h(t)$. The robust controller uses the lower confidence bound

$$\underline{h}_i(t) = \max\{\hat{h}_i(t) - \varepsilon_i^h(t), 0\} \quad (8)$$

and reserves resources according to the conservative rate

$$\underline{R}_i(t) = b_i(t) \log_2 \left(1 + \frac{p_i(t) \underline{h}_i(t)}{N_0 b_i(t) + I_i(t)} \right), \quad (9)$$

where N_0 is the noise spectral density and $I_i(t)$ is aggregate interference. In a deployment with simultaneous robot transmissions or concurrent energy beams, a conservative interference bound is

$$\bar{I}_i^{\text{tot}}(t) = \sum_{j \neq i} v_{ji}(t) p_j(t) \bar{h}_{ji}(t) + \delta_{\text{coex}}(t) \sum_k \sum_{j \neq i} \sum_m u_{k,j,m}(t) P_{k,m}^{\text{wpt}}(t) \bar{\ell}_{k,j \rightarrow i,m}(t), \quad (10)$$

where ν_{ji} indicates time–frequency overlap, $\bar{h}_{ji}$ bounds co-channel coupling, k indexes charging hubs or simultaneous beams, $\bar{\ell}_{k,j \rightarrow i,m}$ bounds energy-beam leakage into robot i 's communication receiver, and δ_{coex} is one only when WPT and data reception overlap. The robust rate uses $I_i = \bar{I}_i^{\text{tot}}$. The nominal FDMA and time-separated experiment sets $\nu_{ji} = 0$ and $\delta_{\text{coex}} = 0$, hence $I_i = 0$. Setting $\varepsilon_i^h = 0$ recovers the nominal perfect-estimate model. A probabilistic implementation can replace Equation (8) by a channel quantile chosen to satisfy a prescribed outage probability.

If $\alpha_i(t) \in [0, 1]$ is the fraction of the task offloaded to the edge, uplink transmission time and energy are

$$T_i^{\text{tx}}(t) = \frac{\alpha_i(t)D_i(t)}{\underline{R}_i(t)}, \quad E_i^{\text{tx}}(t) = p_i(t)T_i^{\text{tx}}(t). \quad (11)$$

For SWIPT reception, $\rho_i(t) \in [0, 1]$ denotes the fraction of received downlink power assigned to energy harvesting; $1 - \rho_i(t)$ is assigned to information decoding.

3.3. Online Control Exchange and Timing Model

At the start of epoch t , each robot transmits a timestamped state report containing its battery level, position, queue state, task descriptor and channel-quality indicator. Let S_i^{ul} and S_i^{dl} denote the report and command sizes in bits, respectively, and let $R_i^{\text{ul},c}$ and $R_i^{\text{dl},c}$ denote their effective control-plane rates after coding and scheduling. A conservative serial-access bound for the uplink and downlink exchange is

$$T_c^{\text{ul}}(t) = \sum_{i \in \mathcal{N}} \frac{\chi_i^{\text{ul}}(t)S_i^{\text{ul}}}{R_i^{\text{ul},c}(t)} + T_{\text{MAC}}^{\text{ul}}(t) + T_q^{\text{ul}}(t), \quad (12)$$

$$T_c^{\text{dl}}(t) = \sum_{i \in \mathcal{N}} \frac{\chi_i^{\text{dl}}(t)S_i^{\text{dl}}}{R_i^{\text{dl},c}(t)} + T_{\text{MAC}}^{\text{dl}}(t) + T_q^{\text{dl}}(t), \quad (13)$$

where χ_i^{ul} and χ_i^{dl} are retransmission-inflation factors and T_{MAC} and T_q collect medium-access and queueing delay. The sums provide a conservative bound for contention-free time-division access; they can be replaced by the corresponding scheduled completion time when orthogonal transmissions occur in parallel.

The complete age of a control action is represented by the sense–compute–actuate interval

$$T_{\text{loop}}(t) = T_{\text{samp}}(t) + T_c^{\text{ul}}(t) + T_{\text{BS}}(t) + T_{\text{opt}}(t) + T_c^{\text{dl}}(t) + T_{\text{act}}(t), \quad (14)$$

where T_{BS} is base-station parsing and state-assembly time, T_{opt} is the measured wall-clock execution time of WPT-JCCO and T_{act} is command validation and actuation time at the robot. Propagation delay is included in the communication terms; over the considered 50 m area it is negligible compared with medium-access, queueing and retransmission delay. The online supervisory update is accepted only if

$$T_{\text{loop}}(t) \leq \Delta t - T_{\text{guard}}, \quad (15)$$

where $T_{\text{guard}} > 0$ reserves time for clock uncertainty and execution jitter. The age of the state used by the new command is therefore upper-bounded by $T_{\text{loop}}(t)$. If Equation (15) is violated or if a required report is missing then the base station does not transmit a partially updated decision. Each robot retains the previous feasible resource allocation for one epoch or switches to a local safe mode with bounded transmit power, no new offloading, and battery-preserving CPU frequency. Low-level motion stabilisation, collision avoidance and emergency stopping remain local and are not closed through the 1 s wireless loop.

When control and task traffic share the same time-frequency resource, the control-plane airtime fraction is

$$\omega_c(t) = \frac{T_c^{\text{ul}}(t) + T_c^{\text{dl}}(t)}{\Delta t}, \quad 0 \leq \omega_c(t) < 1, \quad (16)$$

and the data plane can use at most $(1 - \omega_c(t))B^{\text{tot}}$. The corresponding communication overhead energy of robot i is

$$E_i^{\text{ctrl}}(t) = p_i^{\text{ul},c}(t) \frac{\chi_i^{\text{ul}}(t) S_i^{\text{ul}}}{R_i^{\text{ul},c}(t)} + p_i^{\text{rx},c}(t) \frac{\chi_i^{\text{dl}}(t) S_i^{\text{dl}}}{R_i^{\text{dl},c}(t)} + E_i^{\text{codec}}(t). \quad (17)$$

This term exposes the implementation-dependent radio and packet-processing cost instead of assuming that control exchange is free.

3.4. Computation Model

The local computation workload is $(1 - \alpha_i(t))D_i(t)c_i(t)$ CPU cycles. If $f_i(t)$ is the local CPU frequency then

$$T_i^{\text{loc}}(t) = \frac{(1 - \alpha_i(t))D_i(t)c_i(t)}{f_i(t)}, \quad E_i^{\text{loc}}(t) = \kappa_i(1 - \alpha_i(t))D_i(t)c_i(t)f_i^2(t), \quad (18)$$

where κ_i is the effective switched-capacitance coefficient. If the offloaded task is assigned $F_i(t)$ CPU cycles/s at the edge then the edge execution time is

$$T_i^{\text{edge}}(t) = \frac{\alpha_i(t)D_i(t)c_i(t)}{F_i(t)}. \quad (19)$$

The end-to-end latency is

$$T_i(t) = \max\{T_i^{\text{loc}}(t), T_i^{\text{tx}}(t) + T_i^{\text{edge}}(t) + T_i^{\text{dl}}(t)\}, \quad (20)$$

where $T_i^{\text{dl}}(t)$ is the downlink result-return time. This formulation allows the local and offloaded components to execute in parallel when the workload is partitionable.

3.5. Dependency-Aware Task Extension

For a dependent application, let $s_{i,v}(t)$ and $C_{i,v}(t)$ denote the start and completion times of stage v , and let $y_{i,v}(t) \in \{0, 1\}$ indicate edge execution ($y_{i,v} = 1$) or local execution ($y_{i,v} = 0$). Its stage execution time is selected from the corresponding local or edge model. For every precedence edge $(u, v) \in \mathcal{E}_i(t)$,

$$s_{i,v}(t) \geq C_{i,u}(t) + T_{i,u,v}^{\text{int}}(t), \quad (21)$$

where $T_{i,u,v}^{\text{int}}$ is zero when both stages execute at the same location and otherwise includes the transmission time of the intermediate output $L_{i,u}$. The application latency is

$$T_i^{\text{DAG}}(t) = \max_{v \in \mathcal{V}_i^{\text{exit}}(t)} C_{i,v}(t) - r_i(t), \quad (22)$$

where $r_i(t)$ is the release time and $\mathcal{V}_i^{\text{exit}}$ is the set of exit stages. In the dependency-aware formulation, the scalar offloading ratio α_i is replaced by stage-placement variables $y_{i,v}$, and Equation (21) prevents the optimiser from executing dependent stages in parallel. The fractional task model used in the numerical comparison remains appropriate only for independently partitionable workloads such as image tiles, frame batches or separable sensor windows.

3.6. Hybrid Wireless Power Transfer Model

Let $u_{i,m}(t) \in \{0, 1\}$ indicate whether robot i is served by WPT mode m in epoch t . At most one primary WPT mode is selected per robot per epoch:

$$\sum_{m \in \mathcal{M}} u_{i,m}(t) \leq 1. \quad (23)$$

The received WPT power under mode m is

$$P_{i,m}^{\text{rx}}(t) = P_m^{\text{wpt}}(t) g_{i,m}^{\text{wpt}}(t), \quad (24)$$

where $P_m^{\text{wpt}}(t)$ is the WPT transmit power. To expose mobility-induced alignment loss, the transfer gain is written as a distance term multiplied by a normalised alignment factor:

$$g_{i,\text{RF}}^{\text{wpt}}(t) = G_{\text{RF}} \left(\frac{d_0}{d_i^c(t)} \right)^{\gamma_{\text{RF}}} A_{i,\text{RF}}(t) \chi_{i,\text{RF}}^{\text{clr}}(t), \quad (25)$$

$$g_{i,\text{NF}}^{\text{wpt}}(t) = \frac{G_{\text{NF}}}{1 + (d_i^c(t)/d_{\text{NF}})^6} A_{i,\text{NF}}(t) \chi_{i,\text{NF}}^{\text{clr}}(t), \quad (26)$$

$$g_{i,\text{LW}}^{\text{wpt}}(t) = G_{\text{LW}} \exp(-\zeta d_i^c(t)) A_{i,\text{LW}}(t) \chi_{i,\text{LW}}^{\text{clr}}(t). \quad (27)$$

Here, $d_i^c(t)$ is the charger distance, $\chi_{i,m}^{\text{clr}}(t)$ represents mode-dependent path clearance or blockage and ζ is the lightwave attenuation coefficient. The factors $A_{i,m} \in [0, 1]$ capture lateral and angular misalignment. Engineering surrogates that can be calibrated from measurements are

$$A_{i,\text{NF}}(t) = \exp \left[-k_r^{\text{NF}} \|\mathbf{r}_{i,\perp}(t)\|_2^2 - k_\theta^{\text{NF}} \phi_i^2(t) \right], \quad (28)$$

$$A_{i,\text{LW}}(t) = \exp \left[-2 \|\mathbf{r}_{i,\perp}(t)\|_2^2 / w_b^2 \right] [\cos \phi_i(t)]_+^{k_\theta^{\text{LW}}}, \quad (29)$$

where $\mathbf{r}_{i,\perp}$ is lateral displacement from the coil or beam axis, ϕ_i is angular error and w_b is the lightwave beam radius. RF transfer may use $A_{i,\text{RF}} \approx 1$ or a measured antenna-orientation factor.

The controller does not assume that the estimated transfer gain is exact. With calibration error $|e_{i,m}^g| \leq \varepsilon_{i,m}^g$ and the pose set in Equation (3), the conservative gain and received power are

$$\underline{g}_{i,m}^{\text{wpt}}(t) = \max \left\{ \min_{(\mathbf{q}, \boldsymbol{\theta}) \in \mathcal{U}_i(t)} \hat{g}_{i,m}^{\text{wpt}}(\mathbf{q}, \boldsymbol{\theta}) - \varepsilon_{i,m}^g, 0 \right\}, \quad \underline{P}_{i,m}^{\text{rx}}(t) = P_m^{\text{wpt}}(t) \underline{g}_{i,m}^{\text{wpt}}(t). \quad (30)$$

Nonlinear harvested power for robust feasibility is represented as

$$\underline{P}_i^{\text{harv}}(t) = \sum_{m \in \mathcal{M}} u_{i,m}(t) \left(\frac{P_i^{\text{sat}}}{1 + \exp[-a_i(\underline{P}_{i,m}^{\text{rx}}(t) - b_i^{\text{eh}})]} - P_i^{\text{off}} \right)^+. \quad (31)$$

The nominal formulation is recovered by setting the pose, blockage and gain uncertainty radii to zero. Only the commanded beam is credited in Equation (31). Uncommanded RF, magnetic or optical leakage is treated as interference or exposure rather than useful harvested energy; coherent multi-transmitter combining would require a complex-valued phase and coupling model that is outside the present single-hub formulation.

3.7. Safety-Constrained WPT and Multi-Robot Coexistence

A charger-power cap alone does not establish safe operation because exposure depends on frequency, waveform, beam geometry, distance, averaging time and the presence of people or non-target robots. Let $\mathcal{L}(t)$ be a set of protected locations obtained from the site

map and perception system. For exposure metric r , for example electric-field strength, magnetic-field strength, incident power density or optical irradiance, define

$$\mathcal{X}_\ell^{(r)}(t) = \sup_{\tau \in \mathcal{I}_t} \sum_k \sum_i \sum_m u_{k,i,m}(t) P_{k,m}^{\text{wpt}}(t) \psi_{\ell,k,i,m}^{(r)}(\tau), \quad \ell \in \mathcal{L}(t), \quad (32)$$

where $\psi_{\ell,k,i,m}^{(r)}$ is a calibrated transfer coefficient from beam k to protected location ℓ . The hard safety condition is

$$\mathcal{X}_\ell^{(r)}(t) \leq \mathcal{X}_{\ell,\max}^{(r)}, \quad \forall \ell, r, t. \quad (33)$$

The limits $\mathcal{X}_{\ell,\max}^{(r)}$ are not objective weights: they must be selected from the applicable jurisdiction, operating frequency and exposure category. RF and magnetic-field operation should be assessed against recognised human-exposure guidance such as IEEE C95.1 and the ICNIRP radiofrequency guidelines [26,27]. Lightwave hardware requires product classification, accessible-emission control and engineering safeguards under the IEC 60825 series; the moving-platform specification is directly relevant when the emitting aperture is mounted on a robot or mobile charger [28,29]. A directional optical action is therefore admissible only when its safety interlock, exclusion-zone and beam-stop checks are healthy; optimisation software is not a substitute for a hardware interlock.

For multiple simultaneous charging actions, the deployment model also imposes receiver and infrastructure coexistence limits,

$$\bar{I}_i^{\text{tot}}(t) \leq I_i^{\max}, \quad \sum_{k,i,m} u_{k,i,m}(t) \leq K_{\text{beam}}, \quad (34)$$

where $I_i^{\max}$ is a calibrated communication-receiver or electronics-immunity threshold and K_{beam} is the certified number of concurrent beams. These constraints capture cross-beam leakage, charging-receiver saturation and shared-hub conflicts at the resource-allocation level. Waveform-specific rectifier intermodulation, coherent beam interaction and electromagnetic-compatibility testing still require hardware measurements.

Battery dynamics are

$$B_i(t+1) = \min \{ B_i^{\max}, B_i(t) - E_i^{\text{loc}}(t) - E_i^{\text{tx}}(t) - E_i^{\text{ctrl}}(t) - E_i^{\text{mov}}(t) - E_i^{\text{idle}}(t) + \Delta t P_i^{\text{harv}}(t) \}. \quad (35)$$

4. Problem Formulation

The objective is to minimise net energy depletion while preserving task feasibility and battery fairness. Define the decision vector

$$\mathbf{x}(t) = \{ \alpha_i(t), f_i(t), p_i(t), b_i(t), F_i(t), \rho_i(t), u_{i,m}(t), P_m^{\text{wpt}}(t) \}_{i \in \mathcal{N}, m \in \mathcal{M}}. \quad (36)$$

The per-epoch cost is

$$\begin{aligned} \mathcal{C}(t) = \sum_{i \in \mathcal{N}} & \left[E_i^{\text{loc}}(t) + E_i^{\text{tx}}(t) + E_i^{\text{ctrl}}(t) + E_i^{\text{mov}}(t) + E_i^{\text{idle}}(t) - \eta_B \Delta t P_i^{\text{harv}}(t) \right] \\ & + \lambda_T \sum_{i \in \mathcal{N}} (T_i(t) - \tau_i^{\max}(t))^+ + \lambda_F \sum_{i \in \mathcal{N}} (\bar{B}(t) - B_i(t))^+, \end{aligned} \quad (37)$$

where η_B weights harvested energy, λ_T penalises deadline violation, λ_F penalises battery imbalance and $\bar{B}(t)$ is the swarm-average battery state.

The optimisation problem is

$$\begin{aligned}
\mathbf{P0} : \quad & \min_{\{\mathbf{x}(t)\}_{t=1}^T} \sum_{t=1}^T \mathcal{C}(t) \\
\text{s.t.} \quad & 0 \leq \alpha_i(t) \leq 1, \\
& 0 \leq p_i(t) \leq p_i^{\max}, \quad 0 \leq b_i(t), \quad \sum_{i \in \mathcal{N}} b_i(t) \leq (1 - \omega_c(t)) B^{\text{tot}}, \\
& f_i^{\min} \leq f_i(t) \leq f_i^{\max}, \quad \sum_{i \in \mathcal{N}} F_i(t) \leq F^{\text{edge}}(t), \\
& 0 \leq \rho_i(t) \leq 1, \quad u_{i,m}(t) \in \{0, 1\}, \quad \sum_{m \in \mathcal{M}} u_{i,m}(t) \leq 1, \quad u_{i,m}(t) \leq z_{i,m}(t), \\
& \sum_{i \in \mathcal{N}} \sum_{m \in \mathcal{M}} u_{i,m}(t) \leq 1, \quad 0 \leq P_m^{\text{wpt}}(t) \leq P^{\text{wpt,max}} \sum_{i \in \mathcal{N}} u_{i,m}(t), \\
& \sum_{m \in \mathcal{M}} P_m^{\text{wpt}}(t) \leq P^{\text{wpt,max}}, \\
& T_i(t) \leq \tau_i^{\max}(t) + \epsilon_T, \quad B_i^{\min} \leq B_i(t) \leq B_i^{\max}, \\
& T_{\text{loop}}(t) \leq \Delta t - T_{\text{guard}}, \quad 0 \leq \omega_c(t) < 1, \\
& \mathcal{X}_\ell^{(r)}(t) \leq \mathcal{X}_{\ell,\max}^{(r)}, \quad \bar{I}_i^{\text{tot}}(t) \leq I_i^{\max}, \\
& \text{PDR}_i(t) \geq \text{PDR}_i^{\min}, \quad i \in \mathcal{N}, \ell \in \mathcal{L}(t), r \in \mathcal{R}, t = 1, \dots, T.
\end{aligned} \tag{38}$$

For the single-primary-beam hybrid hub used in the numerical study, the global binary constraint in Problem (38) permits either one robot–mode pair or the no-charge action in each epoch; a multi-beam charger would replace this constraint by its measured beam-count and interference limits. For deployment under bounded uncertainty, Problem (38) is interpreted as the robust counterpart **P0-R**. Communication latency and decoding constraints use $\underline{l}_i$ and $\underline{R}_i$ from Equations (8) and (9); battery and charging constraints use $\underline{g}_{i,m}^{\text{wpt}}$ and $\underline{p}_i^{\text{harv}}$ from Equations (30) and (31); and $u_{i,m} \leq z_{i,m}$ enforces the predicted alignment, blockage and dwell gate in Equation (5). Equations (33) and (34) additionally remove unsafe or electromagnetically incompatible actions from the feasible set. For a dependent application, $\mathbf{x}(t)$ is augmented by $\{y_{i,v}, s_{i,v}, C_{i,v}\}$ and the latency constraint is replaced by Equations (21) and (22). The nominal formulation is the special case $\epsilon_i^q = \epsilon_i^\theta = \epsilon_i^h = \epsilon_{i,m}^g = \epsilon_{i,m}^x = 0$, one clear path, one beam, non-overlapping WPT/data phases and a single-node task graph.

The paper-specific structure of **P0** lies in its feasible set rather than in the use of a weighted energy objective. Conventional MEC-CCO variables $\{\alpha_i, f_i, p_i, b_i, F_i\}$ and the established SWIPT variable ρ_i are coupled to three additional mechanisms: the robot–modality gate $u_{i,m} \leq z_{i,m}$; the individual battery requirement $B_i \geq B_i^{\min}$ together with the fairness term; and the action-age constraint $T_{\text{loop}} \leq \Delta t - T_{\text{guard}}$. Removing the first recovers a fixed-architecture or exogenous-charging MEC problem; removing the second reduces mission preservation to aggregate energy optimisation; removing the third yields a conventional epoch optimiser whose solution may be stale when applied. These three constraints and their simultaneous interaction with the MEC/SWIPT variables are the optimisation-model contribution claimed in this paper.

Problem (38) is non-convex because of binary WPT decisions, nonlinear harvesting, logarithmic rate expressions and products between task fractions and continuous resources. Instead of claiming global optimality, this paper seeks a feasible supervisory update with interpretable resource decisions. Feasibility now has two distinct meanings: the resource vector must satisfy the energy, latency and capacity constraints, and the complete state-report/optimisation/command cycle must satisfy Equation (15). A low iteration count does not, by itself, establish the second condition because T_{opt} depends on processor, software

and memory implementation, while the remaining terms depend on the wireless stack and current network load.

5. Proposed WPT-JCCO Algorithm

The solver is an alternating mixed discrete–continuous procedure. The binary block selects a robustly admissible robot–mode pair, while the continuous block allocates offloading, local and edge computation, radio resources, SWIPT splitting and WPT transmit power. Alternating minimisation, augmented-Lagrangian penalties and projected primal–dual updates are standard tools; the contribution here is the particular decomposition and the acceptance logic required by the swarm battery, physical WPT and control-age constraints. This section derives the implemented updates and states the guarantees that can and cannot be made.

5.1. Continuous Augmented-Lagrangian Block

Within one epoch, the time index is suppressed. Let

$$\mathbf{x}_c = \{\alpha_i, f_i, p_i, b_i, F_i, \rho_i, P_m^{\text{wpt}}\}_{i \in \mathcal{N}, m \in \mathcal{M}} \quad (39)$$

collect the continuous variables and let $\mathcal{X}(\mathbf{u})$ denote their box, simplex and binary-linking constraints for a fixed WPT selector $\mathbf{u}$. The nonlinear inequalities are written as $g_j(\mathbf{x}_c, \mathbf{u}) \leq 0, j \in \mathcal{J}$. In particular,

$$g_i^{\text{loc}} = \frac{(1 - \alpha_i)D_i c_i}{f_i} - \tau_i^{\text{max}}, \quad (40)$$

$$g_i^{\text{off}} = \frac{\alpha_i D_i}{R_i} + \frac{\alpha_i D_i c_i}{F_i} + T_i^{\text{dl}} - \tau_i^{\text{max}}, \quad (41)$$

$$g^{\text{bw}} = \sum_i b_i - (1 - \omega_c)B^{\text{tot}}, \quad g^{\text{edge}} = \sum_i F_i - F^{\text{edge}}, \quad (42)$$

$$g^{\text{wpt}} = \sum_m P_m^{\text{wpt}} - P^{\text{wpt,max}}, \quad g_i^{\text{bat}} = B_i^{\text{min}} - B_i^+, \quad (43)$$

$$g_i^{\text{dec}} = \Gamma_i^{\text{min}}(N_0 b_i + I_i) - (1 - \rho_i)P_i^{\text{dl}} l_i, \quad (44)$$

$$g_{\ell,r}^{\text{safe}} = \mathcal{X}_{\ell}^{(r)} - \mathcal{X}_{\ell,\text{max}}^{(r)}, \quad g_i^{\text{int}} = \bar{I}_i^{\text{tot}} - I_i^{\text{max}}. \quad (45)$$

Here, B_i^+ is the next-epoch battery value produced by Equation (35). PDR, DAG-precedence, blockage-outage, safety-exposure and any active robust-gain requirements are appended to the same residual vector. Let $\mathcal{C}_0$ denote the differentiable part of the per-epoch cost, including a smooth approximation of the hinge penalties when they are used. The Powell–Hestenes–Rockafellar form of the inequality augmented Lagrangian is

$$\mathcal{L}_{\beta}(\mathbf{x}_c, \mathbf{u}, \boldsymbol{\lambda}) = \mathcal{C}_0(\mathbf{x}_c, \mathbf{u}) + \sum_{j \in \mathcal{J}} \frac{[\lambda_j + \beta g_j(\mathbf{x}_c, \mathbf{u})]_+^2 - \lambda_j^2}{2\beta}, \quad (46)$$

where $\beta > 0, \lambda_j \geq 0$ and $[a]_+ = \max\{a, 0\}$. Defining the effective multiplier

$$\hat{\lambda}_j^{(k)} = [\lambda_j^{(k)} + \beta_k g_j(\mathbf{x}_c^{(k)}, \mathbf{u}^{(k)})]_+ \quad (47)$$

gives $\nabla \mathcal{L}_{\beta} = \nabla \mathcal{C}_0 + \sum_j \hat{\lambda}_j \nabla g_j$. For block q , a first-order strongly convex local model yields

$$\begin{aligned} \mathbf{x}_q^{(k+1)} &= \arg \min_{\mathbf{z} \in \mathcal{X}_q} \left\langle \nabla_q \mathcal{L}_{\beta}(\mathbf{x}_q^{(k)}, \mathbf{u}^{(k)}, \boldsymbol{\lambda}^{(k)}), \mathbf{z} - \mathbf{x}_q^{(k)} \right\rangle + \frac{1}{2s_q^{(k)}} \|\mathbf{z} - \mathbf{x}_q^{(k)}\|_{H_q}^2 \\ &= \Pi_{\mathcal{X}_q}^{H_q} \left(\mathbf{x}_q^{(k)} - s_q^{(k)} H_q^{-1} \nabla_q \mathcal{L}_{\beta} \right), \end{aligned} \quad (48)$$

where $H_q \succ 0$ is diagonal scaling and $s_q^{(k)}$ is selected by backtracking so that the augmented-Lagrangian merit does not increase. Euclidean projection is used when $H_q = I$. The bandwidth and edge-CPU blocks are projected onto their respective non-negative simplexes, so their shared-capacity constraints are satisfied at every primal update.

The principal scalar updates follow directly from Equation (48). Let $\ell_i = D_i c_i$, let $\hat{v}_i^{\text{loc}}$ and $\hat{v}_i^{\text{off}}$ be the effective multipliers of Equations (40) and (41), and let $a_i = 1 + \hat{\xi}_i$, where $\hat{\xi}_i$ is the effective multiplier of the minimum-battery constraint. Holding the other blocks fixed gives

$$\frac{\partial \mathcal{L}_\beta}{\partial \alpha_i} = -a_i \kappa_i \ell_i f_i^2 + a_i \frac{p_i D_i}{R_i} - \hat{v}_i^{\text{loc}} \frac{\ell_i}{f_i} + \hat{v}_i^{\text{off}} \left(\frac{D_i}{R_i} + \frac{\ell_i}{F_i} \right) + d_i^{\text{soft}}, \quad (49)$$

where d_i^{soft} collects differentiable fairness or task-penalty terms. Therefore,

$$\alpha_i^{(k+1)} = \Pi_{[0,1]} \left(\alpha_i^{(k)} - s_{\alpha,i}^{(k)} \frac{\partial \mathcal{L}_\beta}{\partial \alpha_i} \right). \quad (50)$$

For a non-zero local workload, minimising the local energy and local-latency terms with respect to f_i gives the cube-root update

$$f_i^{(k+1)} = \Pi_{[f_i^{\min}, f_i^{\max}]} \left[\left(\frac{\hat{v}_i^{\text{loc}}}{2a_i \kappa_i + \epsilon_f} \right)^{1/3} \right], \quad (51)$$

while $f_i = f_i^{\min}$ is used when $(1 - \alpha_i) \ell_i = 0$. Similarly, the unconstrained edge-share stationary point is

$$\tilde{F}_i = \left(\frac{\hat{v}_i^{\text{off}} \alpha_i \ell_i}{\hat{\lambda}_{\text{edge}} + \epsilon_F} \right)^{1/2}, \quad \mathbf{F}^{(k+1)} = \Pi_{\{\mathbf{F} \geq 0: \sum_i F_i \leq F_{\text{edge}}\}}(\tilde{\mathbf{F}}). \quad (52)$$

For the radio block, set $q_i = N_0 b_i + I_i$ and $\gamma_i = p_i \underline{h}_i / q_i$. The rate derivatives used by the projected step are

$$\frac{\partial R_i}{\partial p_i} = \frac{b_i \underline{h}_i}{\ln(2)(q_i + p_i \underline{h}_i)}, \quad (53)$$

$$\frac{\partial R_i}{\partial b_i} = \log_2(1 + \gamma_i) - \frac{b_i p_i \underline{h}_i N_0}{\ln(2) q_i (q_i + p_i \underline{h}_i)}. \quad (54)$$

If $A_i = \alpha_i D_i (a_i p_i + \hat{v}_i^{\text{off}})$, the transmission-energy and offloading-latency contribution satisfies

$$\frac{\partial \mathcal{L}_\beta}{\partial p_i} = \alpha_i D_i \left[\frac{a_i}{R_i} - \frac{a_i p_i + \hat{v}_i^{\text{off}}}{R_i^2} \frac{\partial R_i}{\partial p_i} \right] + d_{p,i}, \quad (55)$$

$$\frac{\partial \mathcal{L}_\beta}{\partial b_i} = -\frac{A_i}{R_i^2} \frac{\partial R_i}{\partial b_i} + \hat{\lambda}_{\text{bw}} + d_{b,i}, \quad (56)$$

where $d_{p,i}$ and $d_{b,i}$ contain the active decoding, PDR and proximal-penalty derivatives. Equation (48) then projects p_i onto $[0, p_i^{\max}]$ and $\mathbf{b}$ onto the available-bandwidth simplex.

The decoding constraint gives the explicit upper bound

$$\bar{p}_i = \Pi_{[0,1]} \left(1 - \frac{\Gamma_i^{\min}(N_0 b_i + I_i)}{P_i^{\text{dl}} \underline{h}_i + \epsilon_p} \right). \quad (57)$$

Accordingly, ρ_i is updated by Equation (48) on $[0, \bar{\rho}_i]$. When harvested energy is the only term that varies with ρ_i and is monotone increasing, the scalar solution is simply $\rho_i^{(k+1)} = \bar{\rho}_i$.

For a fixed selected pair (i, m) , let $H_{i,m}(P)$ denote the nonlinear harvested-power function in Equation (31). The WPT-power derivative is

$$\frac{\partial \mathcal{L}_\beta}{\partial P_m^{\text{wpt}}} = -\Delta t (\eta_B + \hat{\xi}_i) H'_{i,m}(P_m^{\text{wpt}}) + \hat{\lambda}_{\text{wpt}} + d_{P,m}, \quad (58)$$

followed by projection onto the binary-linked interval in Problem (38). Finally, the multipliers are updated by projected ascent,

$$\lambda_j^{(k+1)} = \left[\lambda_j^{(k)} + s_\lambda^{(k)} g_j(\mathbf{x}_c^{(k+1)}, \mathbf{u}^{(k+1)}) \right]_+. \quad (59)$$

The penalty is increased only when the maximum positive constraint residual fails to contract by the prescribed factor; otherwise it is retained. This safeguard avoids increasing β at every iteration.

5.2. Binary Robot–Mode Decision

The binary WPT variables are not relaxed and rounded. The robust candidate set is

$$\mathcal{A}_t = \left\{ (i, m) : z_{i,m}(t) = 1, B_i(t) < B_i^{\max}, \underline{g}_{i,m}^{\text{wpt}}(t) \geq g_m^{\min}, \max_{\ell,r} g_{\ell,r}^{\text{safe}}(\mathbf{u}^{(i,m)}, P_m^{\text{probe}}) \leq 0, g_i^{\text{int}} \leq 0 \right\} \cup \{\emptyset\}, \quad (60)$$

where $\emptyset$ denotes no charging. Candidate screening therefore occurs before the continuous solve and excludes predicted blockage, unsafe exposure and excessive receiver-interference cases at the probe power. For each $(i, m) \in \mathcal{A}_t$, the associated vector $\mathbf{u}^{(i,m)}$ has one unit entry; all other entries are zero. The conventional score

$$\Psi_i(t) = w_B \frac{B_i^{\max} - B_i(t)}{B_i^{\max}} + w_Q \frac{Q_i(t)}{Q_i^{\max}} + w_T \frac{1}{\tau_i^{\max}(t) + \epsilon} + w_C \frac{c_i(t) D_i(t)}{C^{\max}} \quad (61)$$

is used only to order the candidates. A mode-aware ordering score is

$$s_{i,m} = \frac{\Psi_i(t) H_{i,m}(P_m^{\text{probe}})}{P_m^{\text{probe}} + \epsilon_P}. \quad (62)$$

The reported low-overhead implementation evaluates the highest-ranked feasible pair, the incumbent pair and the no-charge action. For every evaluated candidate c , the continuous block is warm-started and updated to produce $\tilde{\mathbf{x}}_c^c$. The binary decision is then

$$c^{(k+1)} = \arg \min_{c \in \mathcal{S}_t} \mathcal{L}_\beta(\tilde{\mathbf{x}}_c^c, \mathbf{u}^c, \boldsymbol{\lambda}^{(k)}), \quad (63)$$

where $\mathcal{S}_t$ contains the evaluated candidates. A new pair is accepted only if its merit is lower than that of the incumbent by at least ϵ_{bin} and it passes the feasibility test below. Thus, $\mathbf{u}$ remains binary throughout and no fractional-mode artefact is introduced. If all elements of $\mathcal{A}_t$ are evaluated, Equation (63) gives the best one-primary-action binary block conditional on the accuracy of the continuous solves. With the implemented shortlist, it is a heuristic neighbourhood search and does not guarantee the globally best robot–mode pair.

5.3. Stopping Tests, Feasibility and Convergence Scope

The implementation distinguishes numerical stationarity from feasibility. For the current binary vector, define

$$r_{\text{stat}}^{(k)} = \left\| \mathbf{x}_c^{(k)} - \Pi_{\mathcal{X}(\mathbf{u}^{(k)})} \left(\mathbf{x}_c^{(k)} - \nabla_{\mathbf{x}_c} \mathcal{L}_\beta \right) \right\|_\infty, \quad (64)$$

$$r_{\text{feas}}^{(k)} = \left\| [\mathbf{g}(\mathbf{x}_c^{(k)}, \mathbf{u}^{(k)})]_+ \right\|_\infty. \quad (65)$$

Two additional tests prevent a numerically flat but infeasible or switching solution from being labelled converged:

$$r_{\text{merit}}^{(k)} = \frac{|\mathcal{L}_\beta^{(k)} - \mathcal{L}_\beta^{(k-1)}|}{\max\{1, |\mathcal{L}_\beta^{(k-1)}|\}}, \quad (66)$$

$$r_{\text{bin}}^{(k)} = \mathbf{1}\{\mathbf{u}^{(k)} \neq \mathbf{u}^{(k-1)}\}. \quad (67)$$

Numerical convergence is declared only when $r_{\text{stat}}^{(k)} \leq \epsilon_{\text{stat}}$, $r_{\text{feas}}^{(k)} \leq \epsilon_{\text{feas}}$, $r_{\text{merit}}^{(k)} \leq \epsilon_{\text{merit}}$, and the binary action has remained unchanged for H_{stop} consecutive sweeps. The reference settings are $\epsilon_{\text{stat}} = \epsilon_{\text{feas}} = 10^{-4}$, $\epsilon_{\text{merit}} = 10^{-5}$ and $H_{\text{stop}} = 3$. Reaching the iteration or wall-clock limit without satisfying all four tests is reported as an early stop, not as convergence.

For fixed $\mathbf{u}$ and fixed (β, λ) , the primal block is a block-successive upper-bound method. If (i) $\mathcal{X}(\mathbf{u})$ is non-empty, compact and convex; (ii) the smooth functions have Lipschitz-continuous gradients on that set; (iii) each local model in Equation (48) is strongly convex, gradient-consistent and accepted by a sufficient-decrease line search; and (iv) every block is updated then every accumulation point is stationary for the fixed augmented-Lagrangian subproblem, following standard block-successive minimisation results [30]. If the inner residuals are driven to zero, the penalty/multiplier sequence is safeguarded, a feasible accumulation point exists and a standard constraint qualification holds there, augmented-Lagrangian theory gives KKT stationarity for the fixed-binary continuous problem [31]. These are conditional guarantees; the finite 80-iteration online truncation may stop before the asymptotic conditions are reached.

For the complete mixed-integer procedure, global optimality is not claimed. Because the shortlist does not enumerate all binary configurations and the continuous problem remains non-convex, the returned point is at most stationary with respect to the continuous block and locally accepted with respect to the evaluated binary neighbourhood. If each binary change must decrease the merit by at least $\epsilon_{\text{bin}} > 0$ and the merit is bounded below, only finitely many binary changes can be accepted. This prevents binary cycling but does not establish the global optimum of Problem (38).

Simple box and simplex constraints are satisfied exactly by projection. Nonlinear resource, latency, decoding and battery constraints are enforced by the residual acceptance test. A command is issued only when $r_{\text{feas}} \leq \epsilon_{\text{feas}}$ and Equation (15) is satisfied. If the new candidate fails then the previous action is re-evaluated under the current robust state; if it is no longer admissible then the controller invokes the predefined local safe mode and defers or drops tasks as needed. The safe mode preserves hard radio, CPU and battery-protection bounds but does not guarantee that every task deadline is met. Robust feasibility is also conditional on the actual pose and channel errors lying inside the uncertainty sets used to construct $\underline{h}_i$ and $\underline{g}_{i,m}^{\text{wpt}}$.

5.4. Parameter Selection and Update Safeguards

The numerical parameters are selected after scaling, not in their heterogeneous physical units. For a bounded scalar block $x_q \in [\ell_q, u_q]$, the solver uses $\tilde{x}_q = (x_q - \ell_q)/(u_q - \ell_q)$; a constraint residual is divided by a characteristic physical limit $G_j > 0$, giving $\tilde{g}_j = g_j/G_j$. The stationarity and feasibility tolerances therefore refer to dimensionless residuals. The di-

agonal matrix H_q in Equation (48) contains the same block scaling and, where available, a local curvature upper bound. The listed primal step is an initial trial value.

The initial primal steps were chosen conservatively after normalisation: 0.05 for bounded fractions and shared-resource shares and 0.02 for frequency, power and SWIPT blocks, whose gradients are more strongly affected by inverse-latency, logarithmic-rate and nonlinear-harvesting terms. These values are common to all seeds and baselines and were not retuned to improve a particular result. The initial penalty $\beta_0 = 1$ makes the normalised objective and constraint terms comparable. It is doubled only if the feasibility residual does not contract by 20% within five sweeps and is clipped at 10^4 ; this residual-balancing rule avoids making the primal model unnecessarily stiff. Multipliers are likewise projected onto $[0, 10^4]$.

The reported 80-sweep implementation uses $s_\lambda^{(k)} = 0.05/\sqrt{k+1}$ as a finite-horizon damping schedule. This schedule is part of the disclosed numerical implementation, but it is not used to justify an infinite-iteration theorem. The asymptotic KKT statement requires a safeguarded sequence satisfying $s_\lambda^{(k)} > 0$, $\sum_k s_\lambda^{(k)} = \infty$ and $\sum_k (s_\lambda^{(k)})^2 < \infty$; for example, $0.05/(k+1)^{0.6}$ satisfies those conditions. This distinction separates the bounded online implementation from the assumptions of the asymptotic convergence result.

5.5. Local Real-Time Tracking and Stability Scope

Real-time operation involves a time-varying optimisation problem because queues, batteries, channels and feasible WPT modes change between epochs. Let $\mathbf{z}_t = (\mathbf{x}_{c,t}, \lambda_t)$, let $\mathbf{z}_t^*$ be a regular fixed-binary KKT point for epoch t , and let Φ_t denote one accepted projected primal–dual sweep with fixed penalty. If, in a neighbourhood of $\mathbf{z}_t^*$, Φ_t is contractive with factor $q < 1$, and if $d_t = \|\mathbf{z}_t^* - \mathbf{z}_{t-1}^*\|$ is the inter-epoch optimiser drift then a warm start followed by K_t sweeps gives the local tracking bound

$$e_t \equiv \|\mathbf{z}_t^{(K_t)} - \mathbf{z}_t^*\| \leq q^{K_t} (e_{t-1} + d_t). \quad (68)$$

If $K_t \geq K_{\min}$ and $d_t \leq \bar{d}$ then with $a = q^{K_{\min}} < 1$,

$$\limsup_{t \rightarrow \infty} e_t \leq \frac{a}{1-a} \bar{d}. \quad (69)$$

Equations (68) and (69) are sufficient local tracking conditions, not a global property of the non-convex mixed-integer problem. Large state jumps, a changed binary active set or loss of regularity can invalidate the contraction assumption.

The implementation therefore monitors contraction rather than assuming it. With $\epsilon_0 = 10^{-12}$, define

$$\hat{q}^{(k)} = \max \left\{ \frac{r_{\text{stat}}^{(k+1)}}{\max\{r_{\text{stat}}^{(k)}, \epsilon_0\}}, \frac{r_{\text{feas}}^{(k+1)}}{\max\{r_{\text{feas}}^{(k)}, \epsilon_0\}} \right\}. \quad (70)$$

If $\hat{q}^{(k)} \geq 1$ for $H_{\text{div}} = 3$ consecutive sweeps, the line search reaches its minimum step, a multiplier remains clipped, or the remaining timing budget cannot accommodate another sweep then the current candidate is not transmitted. The controller revalidates the last feasible allocation and otherwise invokes the local safe mode. Projection, backtracking and clipping keep the numerical iterates bounded; the residual and timing gates provide recursive resource-action feasibility when a valid fallback exists. They do not prove Lyapunov stability of formation, navigation or collision-avoidance dynamics, which remain under the robots' local safety-critical controllers.

Table 4 summarises the distinction.

Table 4. Meaning and scope of the algorithm stability statements.

Property	Condition or Safeguard	Claim Made
Fixed-epoch numerical convergence	Four residual/hold tests and the regularity assumptions in Section 5.3	Conditional stationarity/KKT convergence for the fixed-binary continuous problem; no global mixed-integer guarantee
Time-varying optimiser tracking	Local contraction, bounded optimiser drift and enough sweeps before the deadline	Local bound in Equations (68) and (69); not guaranteed after large mode/state changes
Online resource feasibility	Projection, nonlinear residual test, timestamp/age check and revalidated fallback	No unconverged, infeasible or stale central allocation is intentionally issued
Robot motion stability	Local stabilisation, obstacle avoidance and emergency control remain on board	Not proved by the primal–dual resource allocator and not claimed in this paper

Table 5 summarises the revised update.

Table 5. Derived alternating primal–dual WPT–JCCO supervisory update.

Input: Timestamped robot reports; task or DAG descriptors; pose/channel uncertainty bounds; resource limits; tolerances ϵ_{stat}, ϵ_{feas}, ϵ_{merit}, ϵ_{bin}; hold lengths H_{stop}, H_{div}; epoch deadline $\Delta t - T_{\text{guard}}$. 1. Validate report age; construct pose, blockage and gain uncertainty sets; evaluate protected-zone exposure and cross-beam leakage; compute $z_{i,m}$, $\underline{h}_i$, $\underline{g}_{i,m}^{\text{wpt}}$ and form $\mathcal{A}_t$. 2. Warm-start continuous variables and multipliers from the last accepted epoch; revalidate the incumbent action. 3. Rank feasible robot–mode pairs with Equation (62); form the evaluated set $\mathcal{S}_t$ including the incumbent and no-charge candidates. 4. For each $c \in \mathcal{S}_t$, fix $\mathbf{u}^c$ and alternate Equations (50)–(58), using simplex projection for $\mathbf{b}$ and $\mathbf{F}$. 5. Update effective constraint multipliers and apply Equation (59); increase β only if the feasibility residual fails to contract. 6. Compute r_{stat}, r_{feas}, r_{merit} and $\hat{q}$; declare convergence only after all residual tests and the binary–hold condition pass for H_{stop} sweeps. 7. Apply Equation (63); accept a binary change only when the merit decrease exceeds ϵ_{bin}, then restart the hold counter. 8. Monitor the remaining wall-clock budget and the divergence conditions after every sweep; stop before the timing guard or after H_{div} non-contracting sweeps. 9. Issue the timestamped command only if the latest converged/accepted iterate passes blockage, alignment, exposure, interference, resource and timing checks; otherwise inhibit the beam, revalidate the previous action or invoke the local safe mode. Output: A binary robot–mode decision, continuous resource allocation and validity timestamp, or the safe fallback action.
--

With a one-candidate shortlist, the binary screening cost is $\mathcal{O}(N|\mathcal{M}|)$. One continuous sweep is $\mathcal{O}(N \log(1/\epsilon_r))$ when scalar rate/power searches use tolerance ϵ_r and the simplex projections are implemented by sorting. Evaluating L robot–mode candidates multiplies the continuous-block cost by L . These are operation counts, not wall-clock guarantees. If $T_{\text{opt}}^{\text{max}}(t)$ is the solver budget left by Equation (15) then the real-time sweep cap is

$$K_t^{\text{rt}} = \min \left\{ 80, \left\lceil \frac{T_{\text{opt}}^{\text{max}}(t) - T_{\text{init}}}{\bar{T}_{\text{iter}}} \right\rceil \right\}. \quad (71)$$

A new allocation can be transmitted only if the four convergence/acceptance tests pass by K_t^{rt} ; otherwise the controller uses the revalidated fallback. The iteration cap therefore converts a measured per-sweep time into an admissible online horizon, but it cannot be evaluated from asymptotic complexity alone.

6. Simulation Setup

Our numerical study evaluated whether WPT-JCCO improves energy balance and task success. The nominal experiment contained $N = 20$ robots in a $50 \text{ m} \times 50 \text{ m}$ square, with the access point and hybrid WPT hub co-located at $(25, 25) \text{ m}$. Each run contained 100 one-second epochs. The initial positions were sampled uniformly over the square, the initial battery states were sampled uniformly from 65% to 100% of capacity, and all stochastic quantities were generated from a seed-specific PCG64 stream. The tasks arrived according to a Poisson process with mean 0.10 jobs per robot per epoch; an arrived job had input size $D_i \sim \mathcal{U}(0.5, 3.0) \text{ MB}$, workload density $c_i \sim \mathcal{U}(500, 1500) \text{ cycles/bit}$, and result size $L_i = 0.05D_i$. Low-, medium-, and high-priority jobs occurred with probabilities 0.35, 0.45, and 0.20 and used deadlines 1.20, 0.80, and 0.45 s, respectively. Unfinished jobs remained in an eight-job FIFO queue; a job was successful only if its completion time did not exceed its assigned deadline. Tables 6–8 provide the complete nominal scenario and solver settings.

Table 6. Nominal run, task-generation, mobility, communication and computation settings.

Parameter	Value	Implementation Detail
Robots, area and epochs Randomisation	$N = 20$; $50 \times 50 \text{ m}^2$; 100 epochs PCG64; seeds 0–29	AP/WPT hub at $(25, 25) \text{ m}$; $\Delta t = 1 \text{ s}$ Identical seed-specific traces for every compared scheme
Initial robot position	$x_i, y_i \sim \mathcal{U}(0, 50) \text{ m}$	Independently sampled at the start of each run
Initial battery	$B_i(0)/B_i^{\max} \sim \mathcal{U}(0.65, 1.00)$	$B_i^{\max} = 1800 \text{ J}$; safety threshold $0.10B_i^{\max}$
Task arrivals	Poisson, $\lambda_0 = 0.10 \text{ jobs/robot/epoch}$	FIFO queue; maximum eight waiting jobs
Task input and output	$D_i \sim \mathcal{U}(0.5, 3.0) \text{ MB}$; $L_i = 0.05D_i$	Independently partitionable task in the nominal experiment
Workload density	$c_i \sim \mathcal{U}(500, 1500) \text{ cycles/bit}$	Independently sampled for each arrived job
Priority/deadline classes	$(0.35, 1.20 \text{ s})$, $(0.45, 0.80 \text{ s})$, $(0.20, 0.45 \text{ s})$	Entries give class probability and deadline for low/medium/high priority
Mobility model	Correlated random direction	$v_i \sim \mathcal{U}(0.2, 1.2) \text{ m/s}$; heading innovation $\mathcal{N}(0, 15^\circ)$; specular boundary reflection
Movement power	$P_i^{\text{mov}} = 12 + 3v_i^2 \text{ W}$	Exogenous and common to all schemes; not optimised
Reference channel gain	$G_0 = -40 \text{ dB}$ at $d_0 = 1 \text{ m}$	Path-loss exponent sampled once per run from $\mathcal{U}(2.2, 3.0)$
Shadowing	4-dB log-normal, correlation 0.90	First-order Gauss–Markov update once per epoch
Noise model	-174 dBm/Hz ; 7-dB receiver noise figure	Thermal noise and receiver front-end contribution
Interference/coexistence assumption	Orthogonal FDMA; one primary beam; $I_i = 0$	No inter-robot co-channel overlap, no simultaneous energy beams, and non-overlapping WPT/data subslots; hence $\nu_{ji} = \delta_{\text{coex}} = 0$
Bandwidth and power	$B^{\text{tot}} = 8 \text{ MHz}$; $p_i \in [0.01, 0.50] \text{ W}$	Data-plane bandwidth excludes the control-plane fraction
CPU parameters	$f_i \in [0.2, 1.4] \text{ GHz}$; $F^{\text{edge}} = 24 \text{ GHz}$	$\kappa_i = 2.0 \times 10^{-28} \text{ J/(cycle Hz}^2)$
Downlink/SWIPT	$P_i^{\text{dl}} = 1 \text{ W}$; $\Gamma_i^{\text{min}} = 6 \text{ dB}$	Orthogonal downlink; ρ_i projected onto the decoding-feasible interval
PDR requirement	0.95	Task accepted only when latency, decoding and PDR requirements are satisfied
Control messages	256-B report; 128-B command	Timing audit uses 0.25–2 Mbps and a 20 ms guard

Table 7. Nominal WPT, nonlinear harvesting and objective/priority settings.

Parameter	Value	Implementation Detail
WPT modes	RF, resonant near-field, lightwave	At most one primary robot–mode pair per epoch
Total WPT power	$p_{\text{wpt}, \text{max}} = 15 \text{ W}$	Candidate probe power $P_m^{\text{probe}} = 5 \text{ W}$; one-second maximum dwell
Nominal mode availability	RF: $d \leq 25 \text{ m}$; NF: $d \leq 2.5 \text{ m}$; LW: $d \leq 25 \text{ m}$ and LOS	Epoch-start distance with $A_{i,m} = 1$, $\chi_{i,m}^{\text{clr}} = 1$, no dynamic blocker, and zero gain-estimation error
RF transfer parameters	$G_{\text{RF}} = 4.0 \times 10^{-3}$; $\gamma_{\text{RF}} = 2.2$	Reference gain at 1 m in the transfer model
Near-field parameters	$G_{\text{NF}} = 0.10$; $d_{\text{NF}} = 1.8 \text{ m}$	Sixth-order distance roll-off

Table 7. Cont.

Parameter	Value	Implementation Detail
Lightwave parameters	$G_{\text{LW}} = 0.08; \zeta = 0.06 \text{ m}^{-1}$	Nominal LOS; beam radius $w_b = 0.35 \text{ m}$ in the robust extension
Alignment coefficients	$k_r^{\text{NF}} = 2.0 \text{ m}^{-2}; k_\theta^{\text{NF}} = 4.0 \text{ rad}^{-2}; k_\theta^{\text{LW}} = 8$	Used by the deployment-oriented alignment model; nominal numerical comparison sets alignment loss to zero
Harvester saturation	$P_i^{\text{sat}} = 1.2 \text{ W}$	Common nonlinear receiver model for all robots
Logistic harvester slope/turn-on	$a_i = 150 \text{ W}^{-1}; b_i^{\text{eh}} = 0.014 \text{ W}$	$P_i^{\text{off}} = P_i^{\text{sat}} / [1 + \exp(a_i b_i^{\text{eh}})]$ gives zero output at zero RF input
Harvested-energy weight	$\eta_B = 1.0$	One joule harvested offsets one joule of depletion in the nominal cost
Deadline penalty	$\lambda_T = 50 \text{ J/s}$	Applied to positive deadline violation
Battery-imbalance penalty	$\lambda_F = 0.02$	Multiplies the positive battery deficit in joules
Priority weights	$(w_B, w_Q, w_T, w_C) = (0.40, 0.25, 0.25, 0.10)$	Fixed across all runs; weights sum to one and are not learnt
Priority stabilisers	$\epsilon = 10^{-3} \text{ s}; C^{\text{max}} = 3.0 \text{ MB} \times 1500 \text{ cycles/bit}$	Prevents division by zero and normalises the computation term

Table 8. Alternating primal–dual and binary-selection settings used in the reference implementation.

Parameter	Value	Implementation Detail
Warm start	First epoch: $\alpha = 0.35, f = 0.8 \text{ GHz}, p = 0.10 \text{ W}, \rho = 0.25$	Equal bandwidth/edge shares and $P_m^{\text{wpt}} = 5 \text{ W}$ initially; later epochs use the last accepted allocation and multipliers
Initial multipliers/penalty	$\lambda_j^{(0)} = 0; \beta_0 = 1$	Multiplier clipping interval $[0, 10^4]$; $\beta_{\text{max}} = 10^4$
Dual step	$s_\lambda^{(k)} = 0.05 / \sqrt{k+1}$	Finite 80-sweep reference schedule; projected residual ascent with multiplier clipping. For an untruncated asymptotic run, use a square-summable schedule such as $0.05 / (k+1)^{0.6}$
Penalty update	$\beta \leftarrow \min(2\beta, \beta_{\text{max}})$	Applied if r_{feas} has not contracted by factor 0.80 over five sweeps
Normalised primal steps	$s_\alpha = 0.05;$ $s_f = s_p = s_\rho = s_P = 0.02;$ $s_b = s_F = 0.05$	Initial trials after variable/residual normalisation; Armijo backtracking, rather than the listed value alone, determines the accepted step
Backtracking	factor 0.5; Armijo coefficient 10^{-4} ; minimum step 10^{-6}	Retains the first non-increasing augmented-Lagrangian merit step
Binary shortlist	top-ranked pair, incumbent, no-charge	Maximum of three candidates per epoch; no relaxation or rounding
Binary acceptance	$\epsilon_{\text{bin}} = 10^{-5}$	Strict augmented-merit decrease plus feasibility/timing acceptance
Stationarity tolerance	$\epsilon_{\text{stat}} = 10^{-4}$	Infinity norm in Equation (64); dimensionless after scaling
Feasibility tolerance	$\epsilon_{\text{feas}} = 10^{-4}$	Maximum positive normalised constraint residual
Merit tolerance/hold	$\epsilon_{\text{merit}} = 10^{-5}; H_{\text{stop}} = 3$	Merit, stationarity and feasibility tests must pass while the binary action remains unchanged for three sweeps
Divergence monitor	$\hat{q} \geq 1$ for $H_{\text{div}} = 3$ sweeps	Reject current candidate after repeated non-contraction, minimum-step saturation or persistent multiplier clipping
Scalar-search tolerance	$\epsilon_r = 10^{-6}$	Rate/power bisection or scalar line-search tolerance
Outer-iteration cap	80	Online truncation limit for every candidate
Convergence stress test	$N = 10, 20, 40, 60;$ $\lambda = 0.25, 0.50, 0.75, 1.00$	One background task plus additional Poisson arrivals; 30 seeds

The reported numerical comparison used the nominal special case of the generalised model. Communication and WPT gains were evaluated from the epoch-start state without injected estimation error; near-field and lightwave availability was checked from distance and line of sight at that state rather than from measured intra-epoch pose traces; and each workload was independently partitionable rather than a multi-stage DAG. These choices preserved comparability with the original baseline data but made the resulting energy and success values optimistic with respect to dynamic blockage, alignment loss, channel-estimation outage, exposure-driven power reduction, multi-beam interference and precedence-induced waiting. The uncertainty-aware equations in Sections 3.1 and 3.5

specify how a deployment study should be performed once pose-error, pointing-error, channel-estimation and task-graph data are available; the present paper does not claim that the nominal values quantify those effects.

The same priority-index definition in Equation (61) was used throughout the numerical study. Its weights were treated as fixed design parameters rather than optimisation variables, and the study does not claim that their additive combination was superior to TADP, earliest-deadline-first or other priority schedulers. The numerical comparison evaluated the complete WPT-JCCO controller against the stated baselines; it should not be read as a standalone validation of the priority formula.

The convergence study used a separate controlled stress test of the continuous allocation core inside the WPT-JCCO outer loop. The plotted relative objective change was retained for comparability with the earlier experiment; it was not substituted for the stationarity and feasibility residuals in Equations (64) and (65), for which no retrospective trace was generated. This design isolated the numerical effect of resource coupling from discrete WPT-mode switching and from hardware communication delay. Each trial started from the same feasible uniform bandwidth, edge-compute and charging-share allocation. Every robot had one background task, and additional arrivals followed a Poisson distribution with mean λ tasks per robot per epoch. Two sweeps were performed over the same 30 deterministic seeds: (i) $N \in \{10, 20, 40, 60\}$ in the fixed 50 m $\times$ 50 m area at $\lambda = 0.50$, corresponding to densities from 0.004 to 0.024 robots/m²; and (ii) $\lambda \in \{0.25, 0.50, 0.75, 1.00\}$ at $N = 20$. The maximum number of outer iterations was 80.

For objective value $J^{(k)}$ at outer iteration k , the convergence residual was

$$\delta^{(k)} = \frac{|J^{(k)} - J^{(k-1)}|}{\max\{J^{(k-1)}, 10^{-12}\}}. \quad (72)$$

A run was classified as stabilised at the first iteration for which $\delta^{(k)} < 10^{-3}$ held for three consecutive updates. The plotted curves report the seed-wise mean residual. This stress test evaluated the iteration scaling of the tolerance-controlled continuous update map. It did not measure solver wall-clock time, discrete mode-switching behaviour or communication delay, which were addressed separately by the timing audit.

6.1. Simulation Implementation Environment

The numerical simulator and the result-processing workflow were implemented in Python 3.13.5 and executed in a 64-bit Linux environment on the x86_64 architecture. Vectorised numerical operations, nonlinear model evaluation and pseudo-random sampling were performed with NumPy 2.3.5. Per-epoch and per-run records were organised with pandas 2.2.3, and the publication figures were generated with Matplotlib 3.10.8. The implementation was CPU-based and did not require MATLAB R2025b, GPU acceleration or a commercial optimisation package. Double-precision floating-point arithmetic was used throughout.

WPT-JCCO was implemented as a discrete-time Monte Carlo simulator. At the beginning of each 1 s epoch, the simulator updated the robot positions, task arrivals, channel states and residual battery levels. It then applied the selected baseline or the proposed resource-allocation policy and recorded net energy depletion, task latency, completion status, charging utilisation and the updated battery state. A separate NumPy pseudo-random generator based on PCG64 was initialised for each run using the deterministic seed set $\{0, 1, \dots, 29\}$. The same seed was used across all the compared schemes in a given run so that each method experienced the same mobility, workload and channel realisa-

tions. Package versions and the seed policy were recorded with the simulation files to support repeatability.

The revision package released a machine-readable configuration file, deterministic generators for task, mobility and channel traces, the convergence stress-test implementation, aggregate-result and figure scripts, and SHA-256 checksums.

Running the supplied scenario generator regenerated the nominal exogenous traces for seeds 0–29 and wrote a compact NumPy archive plus a CSV audit summary, as highlighted in Table 9. The figure script treated the supplied aggregate-results CSV as reported study output; it did not manufacture optimiser trajectories. The released code therefore made the scenario construction, convergence experiment and post-processing auditable. Target-specific radio drivers and hardware timing measurements remained outside the present simulation release.

The reported software environment describes the numerical implementation and post-processing pipeline; it is not used as evidence of real-time execution on a deployed base station. As discussed in Section 8, compliance with the 1 s online-control budget still requires wall-clock measurements on the intended processor, operating system, radio driver and protocol stack.

The interference and WPT coexistence assumptions in Tables 6 and 7 are deliberately explicit. The nominal experiment used orthogonal radio allocation and time-separated WPT/data phases, so it did not quantify co-channel swarm interference, adjacent-channel leakage, nonlinear rectifier coupling under simultaneous multi-beam charging, or inter-hub interference. Equations (10)–(34) provide the deployment extension, Table 3 states the operational response of each charging modality, and Table 10 identifies the coefficients that remain uncalibrated. The nominal results therefore do not hide these effects in an unspecified I_i term or infer compliance from the transmitter-power cap.

Table 9. Treatment of the three non-ideal effects and scope of the reported numerical results.

Effect	Generalised Model or Control Response	Status in the 30-Seed Results
Robot mobility and WPT alignment	Pose/orientation uncertainty set, mode-dependent alignment factor, lower-bound transfer gain and dwell-interval availability gate; reject near-field/lightwave actions that cannot maintain the minimum gain.	Bounded position updates were simulated, but gain was evaluated from the epoch-start pose without calibrated lateral/angular error traces.
Channel estimation error	Bounded CSI and WPT-gain errors; conservative rate, decoding and harvested-power constraints; stale or missing reports trigger the existing fallback action.	Perfect epoch-start gain estimates were used. No outage or calibration-error sensitivity was inferred from the headline values.
Task dependencies	DAG stages, local/edge placement variables, precedence constraints and intermediate-output transfer delay; application latency is the completion time of the exit stage.	Tasks were independently partitionable. The reported latency is not a makespan result for dependent perception or planning pipelines.

Table 10. Deployment-oriented treatment of WPT blockage, safety and multi-robot coexistence.

Effect	Model or Operational Constraint	Status of the Reported Numerical Results
Dynamic blockage/LOS	Lower confidence clearance factor $\chi_{i,m}^{\text{clr}}$ over the complete dwell; deterministic gate or calibrated blockage-outage constraint.	Clear-path epoch-start gate; no dynamic people, obstacles, dust, smoke or optical occlusion traces.
Alignment/orientation	Pose uncertainty set and mode-specific lateral/angular factors; reject a mode when worst-case gain drops below g_m^{min} .	Nominal $A_{i,m} = 1$; no measured coil-offset, vibration, receiver-attitude or beam-tracker error.
RF/NF exposure	Protected-zone field or power-density constraint using site-calibrated coefficients and applicable exposure limits.	The 15 W transmitter cap was a resource bound, not evidence of exposure compliance.

Table 10. Cont.

Effect	Model or Operational Constraint	Status of the Reported Numerical Results
Lightwave safety	Irradiance/accessibility constraint, product classification, exclusion-zone monitoring and fail-safe beam interlock.	No eye-safety classification or interlock response was simulated; lightwave results assume an admissible clear path.
Multi-robot/multi-beam coupling	Upper-bound communication leakage $\bar{I}_i^{\text{tot}}$, beam-count limit, receiver-saturation and protected-location exposure constraints.	One primary beam and time-separated WPT/data; no cross-beam, inter-hub or coherent-combining effects.
Calibration required	Pose-to-efficiency maps, blockage traces, leakage matrices, exposure maps and interlock latency.	Released schema marks these hardware/site-specific quantities as unavailable rather than assigning synthetic compliance values.

The headline comparison evaluated local-only execution without WPT, edge offloading without WPT, CCO with static round-robin WPT, greedy WPT-aware charging and the proposed WPT-JCCO method. The main metrics were net energy depletion per epoch, mean task latency, task success ratio, minimum residual battery after 100 epochs, charger utilisation and convergence of the relative objective-change residual. The reported energy values used mean $\pm$ one standard deviation over 30 seeds. To answer the concern that these comparators were relatively simple, Section 6.2 adds a separate released common-trace stress test with two-stage WPT+CCO, fixed-SWIPT dynamic offloading and offline Q-learning scheduling. The original task-level Monte Carlo evaluation did not emulate a specific MAC implementation or base-station processor. Consequently, a separate analytical timing audit is reported below. It used the packet sizes in Table 6, a conservative serial exchange for 20 robots, a 25% airtime inflation for headers and retransmission allowance, 15 ms for sampling, parsing, synchronisation and actuation, and the 20 ms guard. The remaining interval was the maximum permissible measured solver time; it is not presented as a hardware benchmark. Because the receiver-circuit power, packet-processing energy and retransmission energy depended on the selected hardware and protocol stack, the headline energy values set $E_i^{\text{ctrl}} = 0$ and therefore did not claim a measured control-plane energy saving. Equation (17) is retained so that this term can be calibrated in the planned hardware-in-the-loop evaluation.

6.2. Stronger Baseline Implementations

Additional baselines were evaluated in a separate reduced-order policy stress test because the historical aggregate results do not contain per-epoch optimiser trajectories that can be replayed retrospectively. The test nevertheless used the same 30 released task, mobility, channel and initial-battery traces for every policy, and WPT-JCCO was re-evaluated inside the same evaluator. The resulting comparison is therefore internally matched, but its normalised cost should not be mixed numerically with the original headline energy table.

The action grid contained offloading fractions $\{0, 0.25, 0.50, 0.75, 1\}$ and uplink powers $\{0.15, 0.30, 0.45\}$ W. A service slice equal to 6.5% of each generated task was processed per supervisory update. The reduced-order battery update included a movement-energy scale of $0.028P_i^{\text{mov}}$, 0.06 J of control energy per robot and 0.15 J of sensing energy per active robot. These stress-test-only coefficients, together with the policy definitions below, are recorded in `simulation_config_v12.json` and were not used to alter the headline results.

The three stronger alternatives were implemented as follows.

- *Two-stage WPT+CCO*: the first stage selected the most depleted robot and applied 10 W RF charging. With that action fixed, the second stage searched the offloading and

transmit-power grid and allocated equal instantaneous bandwidth and edge-CPU shares. A 2.5% coordination-latency factor represented the serial stage boundary.

- *Fixed-SWIPT dynamic offloading*: the power-splitting ratio was fixed at $\rho = 0.35$. The RF recipient was selected from battery deficit and queue activity at 10 W, while offloading and transmit power adapted to the current task and channel state.
- *Offline Q-learning scheduler*: a 72-state table represented quantised battery, channel, workload, deadline and contention conditions. Five offloading actions were trained from 140,000 disjoint synthetic state samples using seed 12,026 and learning rate 0.06. Our evaluation on seeds 0–29 used no exploration; transmit power was selected by one-step merit minimisation and WPT used a greedy adaptive mode at 12 W.

The WPT-JCCO replay jointly scored the robot-mode pair over RF, near-field and lightwave actions at the full 15 W budget and evaluated the same offloading/power grid with its battery-fairness logic. For this stress test, the per-seed composite measure was

$$J_{SB} = \bar{E}_{\text{net}} + 18(1 - S) + 2.5 \frac{\bar{B} - B_{\min}}{B_{\max}}, \quad (73)$$

where S is the task-success fraction.

6.3. Multi-Factor Generality Protocol

To examine whether the conclusions depended on the reference 20-robot operating point, a second reduced-order experiment performed one-factor-at-a-time sweeps around the nominal configuration. Each level was evaluated for 100 supervisory epochs and seeds 0–29. When one factor was varied, every other quantity remained at the nominal value in Table 11. New task, mobility and channel traces were generated deterministically for factors that altered the exogenous process; resource-only sweeps reused matched traces. This design isolates first-order trends but does not estimate cross-factor interactions. The additional baselines used in our research were evaluated in a separate reduced-order policy stress test because the historical aggregate results in Table 12 do not contain per-epoch optimiser trajectories that can be replayed retrospectively. Table 13 reports J_{SB} normalised by the WPT-JCCO mean.

Table 11. One-factor-at-a-time generality study. The nominal level is shown in bold.

Factor	Levels	Nominal Value Held in Other Sweeps
Swarm size, N	10, 20, 40, 60 robots	20 robots in 50 m × 50 m
Task-arrival rate, λ	0.05, 0.10, 0.20, 0.30 jobs/robot/epoch	Poisson $\lambda = 0.10$
WPT power budget	0.1, 0.5, 1, 5, 15 W	15 W, one primary beam
Total bandwidth	4, 8, 12, 16 MHz	8 MHz orthogonal FDMA
Edge CPU capacity	12, 24, 36, 48 GHz	24 GHz shared pool
Mobility speed scale	0.5, 1.0, 1.5, 2.0 times nominal	Uniform speed 0.2–1.2 m/s
Channel uncertainty margin	0, 3, 6, 9 dB	No lower-bound margin

Table 12. Numerical comparison of WPT-JCCO and baseline schemes over 30 random seeds.

Scheme	Energy (J/Epoch)	Latency (s)	Success (%)	Min. Battery (%)	Charger Use (%)
Local-only, no WPT	146.2 ± 4.9	0.84	89.1	34.8	0.0
Edge-offload, no WPT	113.6 ± 4.1	0.71	91.7	44.1	0.0
CCO + static WPT	96.4 ± 3.3	0.66	93.5	52.6	51.2
Greedy WPT-aware	89.8 ± 3.0	0.64	94.1	56.4	57.6
Proposed WPT-JCCO	73.5 ± 2.6	0.59	97.3	69.2	63.9

Table 13. Matched strong-baseline policy stress test over the same 30 exogenous traces. The normalised composite cost uses Equation (73); lower is better. Values in this table belong to the released reduced-order evaluator and are not combined with Table 12.

Policy	Normalised Cost	Net Energy (J/Epoch)	Latency (s)	Success (%)	Min. Battery (%)
Two-stage WPT+CCO	1.125	9.266 ± 0.037	0.396	97.37	64.41
Fixed-SWIPT dynamic offloading	1.128	9.289 ± 0.036	0.396	97.35	64.39
Offline Q-learning scheduler	1.062	8.489 ± 0.111	0.478	96.03	65.97
Proposed WPT-JCCO	1.000	8.200 ± 0.036	0.388	97.44	66.59

The policy and energy equations were identical to the released common-trace evaluator in Section 6.2. For comparisons across different N , net depletion is reported per robot per epoch. A dimensionless within-study measure combines energy, missed deadlines and battery dispersion:

$$J_{\text{gen}} = \frac{E_{\text{net}}}{NT} + 0.90(1 - S) + 0.125 \frac{\bar{B} - B_{\min}}{B_{\max}}. \quad (74)$$

For each factor, J_{gen} was normalised by the value at that factor's nominal level. The purpose was to display relative degradation or improvement within a matched sweep, not to replace the physical units in Table 12. The configuration, evaluator, raw seed-level records and aggregated results were released as `simulation_config_v14.json`, `generality_sensitivity_v14.py` and the corresponding CSV files.

7. Results

7.1. Net Energy Depletion

Table 12 reports the main numerical results. WPT-JCCO obtained the lowest net energy depletion at 73.5 ± 2.6 J/epoch as presented in Figure 2. Compared with CCO + static WPT, this was a 23.8% reduction. Compared with local-only execution, the reduction was 49.7%. The improvement occurred because WPT-JCCO changes offloading and communication decisions according to where harvested energy has the highest mission value.




Figure 2. Net energy depletion per epoch for WPT-JCCO and baselines. Error bars indicate one standard deviation over 30 random seeds.

7.2. Latency and Task Success

Figure 3 shows that WPT-JCCO moved the operating point toward lower latency and higher task success. Mean latency fell from 0.66 s under static WPT to 0.59 s under WPT-JCCO. Task success increased from 93.5% to 97.3%. This improvement indicates that the benefit of WPT is not only an increase in residual energy; it also enables urgent tasks to receive radio and edge-computation resources at the right time.




Figure 3. Latency-success trade-off. WPT-JCCO improves both latency and task success by coupling charging with offloading and radio allocation.

7.3. Comparison with Stronger Alternatives

Table 13 and Figure 4 report the matched reduced-order stress test defined in Section 6.2. WPT-JCCO obtained the lowest composite cost. Relative to two-stage WPT+CCO, fixed-SWIPT dynamic offloading and offline Q-learning, the reduction was 11.1%, 11.3% and 5.8%, respectively. The result against the sequential baseline indicates that selecting charging before solving the communication–computation problem left useful cross-layer feedback unused. Fixing ρ produced a similar penalty because the decoding–harvesting balance could not adapt to the current link and battery state. The learned scheduler narrowed the energy and battery gap, but its quantised state and offloading-only action table yielded lower task success and higher mean latency than the joint continuous–discrete controller.

The learning result should not be interpreted as a general ranking of reinforcement learning and model-based optimisation. The Q-table intentionally uses a compact state and action space so that its training and evaluation are fully reproducible without a separate neural-network stack. A larger actor–critic policy could improve performance, but it would also require architecture, replay, exploration and hyperparameter sensitivity studies beyond the scope of the present revision.




Figure 4. Normalised composite cost for the matched strong-baseline stress test. Error bars show one standard deviation across 30 common exogenous traces; the dashed line marks the WPT-JCCO mean.

7.4. Performance Across Swarm and Resource Conditions

Figures 5 and 6 report the one-factor-at-a-time study defined in Section 6.3. The normalised cost increased smoothly rather than discontinuously as the swarm or offered workload grew. Relative to the nominal $N = 20$ case, $N = 40$ and $N = 60$ increased J_{gen} by 15.2% and 25.5%, respectively. The high-density case raised the mean latency from 0.388 to 0.647 s and lowered task success from 97.44% to 89.74%. Increasing the Poisson rate from the nominal 0.10 to 0.30 jobs/robot/epoch produced a 29.0% cost increase, with a similar fall in task success to 89.01%. These trends reflect contention for fixed bandwidth, edge CPU and charging opportunities rather than numerical non-convergence. The relative objective-change residual used in Figures 7 and 8 remains a reporting diagnostic and is not a KKT certificate.




Figure 5. Normalised WPT-JCCO cost under one-factor-at-a-time variations. The dashed line in each panel denotes the corresponding nominal operating point. Values are means over 30 deterministic seeds and 100 supervisory epochs.

The resource sweeps show where additional infrastructure was useful. Increasing total bandwidth from 8 to 16 MHz reduced the normalised cost by 3.8%, halved mean latency from 0.388 to 0.272 s and raised task success to 99.08%. Increasing edge capacity from 24 to 48 GHz reduced the cost by 2.1% and raised task success from 97.44% to 98.45%. Reducing WPT power from 15 to 0.1 W increased net depletion from 0.410 to 0.457 J/robot/epoch and lowered the minimum battery from 66.59% to 64.48%. Above approximately 5 W, the curve flattened because the adopted nonlinear harvester approached saturation; this plateau was specific to the transfer and rectifier parameters in Table 7 and is not a universal charging threshold. The Armijo backtracking halves it until the augmented-Lagrangian merit is non-increasing; consequently, a large nominal step cannot be applied merely because it appears in Table 8.

The mobility and uncertainty sweeps exposed different mechanisms. Doubling the nominal speed interval increased net depletion from 0.410 to 0.555 J/robot/epoch and the normalised cost by 32.0%; task success remained nearly unchanged because the dominant effect in this reduced-order test is propulsion energy rather than loss of alignment, as outlined in Table 14. Applying a 9 dB conservative channel margin increased mean latency to 0.433 s and lowered task success to 96.83%, while the cost rose by 1.5%. The modest energy change indicates that the controller responds by shifting the local/offload balance, but this should not be interpreted as robustness against arbitrary fading or blockage.




Figure 6. Task-success and minimum-battery responses to the seven operating-condition sweeps. Each factor was varied independently while all other factors remained at their nominal values.

Table 14. Endpoint changes in the one-factor-at-a-time generality study. Energy is net depletion per robot per epoch. Arrows follow the factor range shown in the first column and do not imply that the lower endpoint is always preferable.

Factor Range	Energy (J/Robot/Epoch)	Latency (s)	Success (%)	Min. Battery (%)
$N : 10 \rightarrow 60$	0.355 $\rightarrow$ 0.453	0.277 $\rightarrow$ 0.647	99.30 $\rightarrow$ 89.74	68.94 $\rightarrow$ 64.01
$\lambda : 0.05 \rightarrow 0.30$	0.400 $\rightarrow$ 0.464	0.282 $\rightarrow$ 0.636	99.09 $\rightarrow$ 89.01	67.06 $\rightarrow$ 65.56
$P_{WPT} : 0.1 \rightarrow 15 \text{ W}$	0.457 $\rightarrow$ 0.410	0.388 $\rightarrow$ 0.388	97.44 $\rightarrow$ 97.44	64.48 $\rightarrow$ 66.59
$B : 4 \rightarrow 16 \text{ MHz}$	0.414 $\rightarrow$ 0.408	0.547 $\rightarrow$ 0.272	93.77 $\rightarrow$ 99.08	66.56 $\rightarrow$ 66.61
$F^{\text{edge}} : 12 \rightarrow 48 \text{ GHz}$	0.411 $\rightarrow$ 0.410	0.458 $\rightarrow$ 0.344	95.51 $\rightarrow$ 98.45	66.59 $\rightarrow$ 66.59
Speed: 0.5 $\rightarrow$ 2.0 $\times$	0.374 $\rightarrow$ 0.555	0.388 $\rightarrow$ 0.388	97.54 $\rightarrow$ 97.47	66.46 $\rightarrow$ 65.91
CSI margin: 0 $\rightarrow$ 9 dB	0.410 $\rightarrow$ 0.411	0.388 $\rightarrow$ 0.433	97.44 $\rightarrow$ 96.83	66.59 $\rightarrow$ 66.58

The study supports generality over the tested first-order ranges, but it is not a factorial robustness proof. Interactions such as high mobility combined with blockage, dense swarms under reduced bandwidth, or low WPT power under severe CSI error are not estimated by a one-factor-at-a-time design. Moreover, the sensitivity evaluator uses the released reduced-order action grid rather than reconstructing unavailable per-epoch trajectories from the original headline experiment. Table 15 reports the mean, standard deviation and 95th-percentile stabilisation iteration.

7.5. Convergence Under Swarm Density and Task-Arrival Load

Figures 7 and 8 visualise the mean relative objective change defined in Equation (72). Increasing either swarm density or task-arrival rate delays numerical stabilisation because more robots compete for the same bandwidth, edge-compute and charging shares and because a larger workload strengthens the coupling between offloading, radio and charging decisions. The curves remained well behaved under the tested conditions and crossed the 10^{-3} tolerance before the 80-iteration cap.




Figure 7. Convergence residual under increasing swarm density. The fixed $50\text{ m} \times 50\text{ m}$ area yielded densities of 0.004, 0.008, 0.016 and 0.024 robots/ m^2 for $N = 10, 20, 40$ and 60 , respectively. Curves show the mean over 30 seeds at $\lambda = 0.50$ additional tasks per robot per epoch.

At the fixed arrival rate $\lambda = 0.50$, the mean stabilisation iteration increased from 26.0 for $N = 10$ to 44.7 for $N = 60$. The corresponding 95th percentile increased from 30 to 46 iterations. At fixed $N = 20$, increasing the additional task-arrival rate from $\lambda = 0.25$ to $\lambda = 1.00$ increased the mean stabilisation iteration from 24.8 to 39.6, while the rounded-up 95th percentile increased from 27 to 43 iterations. The result indicates gradual degradation rather than an abrupt loss of numerical stability over the tested range, as summarised in Table 15.

The iteration counts do not establish one-second online feasibility. The stress test is CPU-independent and records outer iterations, whereas deployment requires measured wall-clock execution together with the state-report and command-delivery delays in Section 7.6. It also isolates the continuous update map; repeated discrete WPT-mode changes, packet loss, asynchronous updates or stale channel reports may produce slower behaviour in a physical system.




Figure 8. Convergence residual under increasing Poisson task-arrival rate at $N = 20$. Each robot retained one background task, while λ denotes the mean number of additional tasks arriving per robot per epoch. Curves show the mean over 30 seeds.

Table 15. Stabilisation iteration under varying swarm size and task-arrival rate over 30 random seeds. A run was stabilised when the relative objective change remained below 10^{-3} for three consecutive iterations.

Sweep	Case	Mean	Std. Dev.	95th Percentile
Swarm size	$N = 10$	26.0	2.2	30
Swarm size	$N = 20$	30.1	1.5	33
Swarm size	$N = 40$	37.4	1.6	40
Swarm size	$N = 60$	44.7	1.2	46
Task arrivals	$\lambda = 0.25$	24.8	1.2	27
Task arrivals	$\lambda = 0.50$	30.1	1.5	33
Task arrivals	$\lambda = 0.75$	34.9	2.1	39
Task arrivals	$\lambda = 1.00$	39.6	1.9	43

7.6. Communication-Aware Online-Control Audit

Table 16 evaluates the control-plane airtime and the residual solver budget for the 1 s supervisory epoch. With 20 robots, a 256-B uplink report and a 128-B downlink command produce 61,440 control bits per epoch. At an effective serial control rate of 1 Mbps, the ideal exchange occupies 61.4 ms. After applying the 25% airtime inflation, the exchange occupies 76.8 ms, or 7.68% of the epoch. After reserving a further 15 ms for non-radio processing and actuation and a 20 ms timing guard, the target implementation may spend at most 888.2 ms on optimisation. At 0.25 Mbps, the inflated exchange occupies 307.2 ms and the admissible solver time falls to 657.8 ms.

This audit changes the interpretation of online suitability. The proposed method is feasible at a given rate only when the measured T_{opt} and the realised queueing, medium-access and retransmission delays jointly satisfy Equation (15). For example, at 1 Mbps the 888.2 ms solver allowance becomes the iteration cap in Equation (71) after T_{init} and the measured per-sweep time are inserted. A feasible online update therefore requires both enough sweeps to satisfy the four convergence/acceptance tests and enough residual time to deliver the command. If either condition fails, the fallback is used; the solver is not allowed to overrun the epoch in an attempt to reach an asymptotic limit. Contention-based

access, correlated packet loss or additional security headers can consume more than the 25% allowance; scheduled parallel uplink transmission, multicast command dissemination and report aggregation can reduce it. The table therefore gives a transparent engineering threshold rather than proof of hard real-time operation. A target-hardware experiment must measure the upper tail, not only the mean, of T_{loop} and $\bar{T}_{\text{iter}}$ before deployment.

Table 16. Analytical timing audit for one 1 s supervisory epoch. The maximum solver time is the residual after a 25% airtime inflation, 15 ms of fixed non-radio delay, and a 20 ms guard; it is not a measured execution time.

Effective Control Rate (Mbps)	Ideal Serial Airtime (ms)	Inflated Control Airtime (ms)	Maximum Solver Time (ms)
0.25	245.8	307.2	657.8
0.50	122.9	153.6	811.4
1.00	61.4	76.8	888.2
2.00	30.7	38.4	926.6

7.7. Scope Under Mobility, Estimation Error and Task Dependencies

Table 9 is essential for interpreting the numerical results. Mobility is present in the position process, but the current data do not contain measured coil offset, receiver attitude, beam-tracking error or within-epoch blockage. Likewise, the channel values used by the scheduler are not perturbed by an estimator-error model, and task latency is computed for an independently partitionable workload. Consequently, the values in Tables 12 and 17 are nominal operating points. They show the resource-allocation effect of coupling WPT, communication and computation under a common model, but they do not establish resilience to pointing loss, CSI mismatch or precedence bottlenecks.

Table 17. Ablation study for the proposed WPT-JCCO method.

Variant	Energy (J/Epoch)	Latency (s)	Success (%)	Min. Battery (%)
Full WPT-JCCO	73.5	0.59	97.3	69.2
Without adaptive WPT mode selection	85.2	0.62	95.1	61.7
Without SWIPT power splitting	79.9	0.63	95.8	64.4
Without battery-fairness dual term	77.6	0.60	96.4	60.5
Without edge offloading	101.8	0.79	90.6	58.9

The generalised model predicts the direction of these effects without assigning unsupported numerical magnitudes. Alignment loss reduces $\underline{P}_i^{\text{harv}}$, causing the controller to shorten or reject near-field/lightwave actions and, where possible, select RF WPT or defer charging. A lower communication gain reduces $\underline{R}_i$, increasing transmission time and forcing more local execution, additional bandwidth or a fallback action. DAG precedence reduces the amount of parallel local/edge execution available in Equation (22); intermediate-output transfers can further increase the critical-path latency. Quantitative sensitivity curves require hardware-calibrated error distributions and application task graphs and are therefore left to the deployment study rather than fabricated from the nominal data.

7.8. Effect of WPT Mode Selection

Figure 9 illustrates the distance-dependent harvested-power profile for the three WPT modes. Near-field transfer is best at short distances, lightwave transfer is useful for line-of-sight charging at moderate distances and RF WPT provides broad but lower-power

coverage. WPT-JCCO uses the priority index to order requests and the mode-dependent harvested-power term to distinguish physically feasible charging actions. This figure demonstrates transfer-mode diversity; it does not establish novelty for the priority score.




Figure 9. Illustrative harvested-power profiles for RF far-field, resonant near-field and directional lightwave WPT modes.

7.9. Scope of the Priority-Aware Dispatch Rule

The reviewer-relevant distinction is between the complete controller and its recipient-ranking heuristic. Equation (61) combines signals that are routinely used by priority-aware charging schemes, and no performance gain is attributed to that algebraic combination by itself. The “without adaptive WPT mode selection” ablation removes the adaptive recipient/mode block as a whole; it does not compare the proposed score with TADP, earliest-deadline-first or another priority formula. Consequently, the ablation supports the value of adapting WPT mode and recipient within the common resource loop, but it should not be interpreted as evidence that Equation (61) is intrinsically superior to established priority schedulers.

A controlled comparison of dispatch rules would require matching charger mobility, request generation, deadline semantics and feasible transfer modes. TADP and MORE were developed for mobile chargers in rechargeable sensor networks, whereas the present model uses a stationary hybrid hub, local/edge computation and SWIPT. Re-implementing those schedulers under the same swarm and hardware assumptions is therefore identified as a separate benchmark rather than inferred from the current baselines.

7.10. Ablation Analysis

Table 17 isolates the contribution of the main components. Removing adaptive WPT mode selection increases energy depletion from 73.5 to 85.2 J/epoch. Removing SWIPT splitting increased energy depletion to 79.9 J/epoch and reduced success because the energy-information trade-off was no longer adjusted to channel conditions. Removing the battery-fairness dual term reduced the minimum residual battery from 69.2% to 60.5%, even though the average energy remained competitive. These results show that battery fairness and mode-adaptive WPT are necessary for robust swarm operation.

8. Discussion

The results support three main observations. First, WPT is most effective when it is co-optimised with communication and computation. Static WPT improves average battery level but cannot decide whether harvested energy should support local computation, radio transmission or waiting for a better charging opportunity. Second, battery fairness is a mission-level property. Greedy charging can improve the weakest battery state locally but may still miss urgent tasks elsewhere in the swarm. Third, SWIPT splitting cannot be fixed independently of channel quality. The appropriate split changes with SNR, residual energy and packet-delivery constraints.

The difference from the three adjacent method classes can be stated directly. Against SWIPT-MEC, WPT-JCCO adds the discrete choice of both recipient robot and physical charging modality, and it conditions that choice on alignment, blockage, safety and dwell feasibility; the SWIPT split is only one continuous block inside the larger action. Against WPT scheduling, the charging decision is not final: it is accepted only after task placement, CPU, bandwidth, transmit-power, packet-delivery and edge-capacity variables have been co-optimised. Against swarm resource allocation, battery energy is not merely a state to conserve; replenishment is an endogenous action and the weakest robot enters the feasible set through an individual battery residual and fairness multiplier. The age-gated command and fallback then connect this cross-layer allocation to online supervisory control. These coupled mechanisms, rather than any individual energy objective or solver family, constitute the novelty claimed by the paper.

These observations should not be interpreted as establishing novelty for net-energy minimisation, joint wireless charging and offloading, SWIPT, priority-aware charging or primal–dual optimisation. All are established [6–8,10,24,25]. The evidence provided by the baselines and ablations is narrower: within the stated mobile-swarm model, performance deteriorates when adaptive mode selection, SWIPT splitting, battery-fairness feedback or edge offloading is removed from the common supervisory loop. Because the battery–backlog–deadline index was not benchmarked against alternative priority formulas, the gains are assigned to the complete coupled controller rather than to Equation (61). The published adjacent methods use different terminal, charger, channel and task abstractions, so Table 12 is not a direct numerical ranking against every MEC-CCO, SWIPT-MEC or robotic WPT algorithm. The common-trace experiment in Table 13 now provides controlled implementations of three stronger policy classes, but it remains a reduced-order policy stress test rather than a claim that every published architecture has been reproduced exactly.

The two reviewer-recommended studies clarify the boundary further. Zeng et al. [9] optimised cooperative NOMA offloading and computation in a wireless-power three-node network and therefore provided a direct WP-MEC comparison at the communication–computation layer. WPT-JCCO does not reproduce their NOMA protocol; it instead adds mobile-robot battery fairness, selectable WPT modalities and supervisory control timing. Luo et al. [17] addressed adaptive safety-critical control for cellular satellites using improved PPO. Their contribution was closer to the inner autonomy layer than to WPT-MEC resource allocation, but it highlights the need to separate a safety-critical controller from the one-second resource supervisor proposed here. Consequently, neither paper is treated as a numerically interchangeable baseline without a common system model, objective and dynamics.

8.1. Interpretation of the 1 s Online Control Loop

WPT-JCCO should be interpreted as a supervisory resource-management controller, not as the inner motion-control loop of the robots. The 1 s update selects charging mode, offloading fraction, radio allocation and processor settings. Fast stabilisation, obstacle

avoidance, collision prevention and emergency braking remain on board each robot and continue even when a base-station update is delayed or lost. This separation is essential: closing a safety-critical motion loop through the central optimiser would require substantially tighter and experimentally verified latency and reliability guarantees than those established by the present simulation.

The timing model in Section 3.3 also shows why computational complexity is only one component of online feasibility. Before the first optimisation iteration, the base station must receive a sufficiently fresh snapshot from the swarm. After optimisation, the commands must still be delivered and validated before the epoch deadline. The relevant quantity is therefore the complete sense–compute–actuate age T_{loop} , not the solver time in isolation. Table 16 indicates that the control exchange is modest at a scheduled 1 Mbps effective rate for 20 robots but that it can consume almost one-third of the epoch at 0.25 Mbps after the stated overhead allowance. The linear dependence on robot count also means that a centralised serial exchange will eventually become the dominant cost as the swarm grows.

A deployable implementation should timestamp all reports and commands, reject stale state, reserve contention-free control resources, and monitor both the mean and high-percentile loop delay. Event-triggered reports can suppress unchanged states, while aggregated or multicast downlink commands can reduce the downlink burden. For larger swarms, cluster heads or a hierarchical controller can replace the $\mathcal{O}(N)$ central exchange. The previous feasible decision provides a deterministic one-epoch fallback, but repeated misses should trigger a degraded local mode and suspend new edge offloading. These measures prevent an optimisation result from being applied after its supporting state has become obsolete.

8.2. Interpretation of the Convergence Scaling

The density and arrival-rate sweeps answer a narrower question than the main energy comparison: they show how many outer updates the continuous resource-allocation core requires before its relative objective change remains below the prescribed tolerance. The approximately monotone increase in Table 15 is consistent with stronger coupling when more robots or more tasks share fixed radio, edge and charging capacities. Across the tested range, the 95th-percentile count remained at or below 46 iterations, leaving the numerical margin below the configured cap of 80.

This result should not be converted directly into a real-time claim. The stress test used synchronous, error-free updates, fixed discrete WPT feasibility within each trial and cold-started feasible continuous shares. A deployed controller may benefit from warm starts, but asynchronous reports, recipient/mode switching, infeasibility repair and packet loss can also add iterations or terminate the loop early. The correct deployment criterion therefore remains Equation (15): the achieved iteration count must be multiplied by measured per-iteration execution time and combined with the complete communication and actuation delay. For swarms larger or busier than the tested cases, hierarchical decomposition or cluster-level charging decisions may be required.

The derivation in Sections 5.3–5.5 also narrows the mathematical claim. For a fixed binary WPT action, a sufficiently accurate augmented-Lagrangian subproblem solves and standard regularity conditions support convergence to a stationary/KKT point of the continuous subproblem. The implemented binary shortlist does not provide a mixed-integer global guarantee. The new four-part stopping test prevents objective flattening alone from being reported as convergence, while the parameter-selection subsection distinguishes the disclosed 80-sweep dual schedule from the square-summability assumption needed by an untruncated theorem. Merit-decrease acceptance prevents repeated binary cycling, and the feasibility filter prevents an infeasible or stale candidate from being transmitted.

The local tracking bound in Equation (68) explains when warm-started primal–dual updates can remain numerically stable as the swarm state changes: the fixed-binary update map must remain locally contractive, the epoch-to-epoch optimiser drift must be bounded and the timing budget must permit enough sweeps. The online contraction monitor and fallback prevent a detected non-contracting iterate from being applied. These safeguards support bounded supervisory resource actions; they are not a proof of Lyapunov stability for robot motion, formation or collision avoidance. The convergence plots therefore show empirical stabilisation of the tested continuous update map rather than proof of global convergence or safety-critical closed-loop stability.

8.3. Interpretation of the Multi-Factor Sensitivity Study

The seven sweeps separate three sources of performance loss. Swarm size and task intensity primarily consume shared radio and edge capacity; mobility primarily increases propulsion energy in the present evaluator; and a conservative channel margin primarily increases transmission and offloading delay. The bandwidth and edge sweeps show that additional infrastructure recovers service quality, but with diminishing energy benefit once the action grid has already shifted to low-energy feasible choices. Similarly, the WPT curve saturates because the adopted harvester is nonlinear. These observations are useful for system sizing, but they are conditional on the disclosed model and should not be extrapolated to a different rectifier, task graph, mobility controller or multi-beam charger without recalibration.

A one-factor-at-a-time design also hides interactions. For example, adding bandwidth may not restore success when high mobility simultaneously destroys lightwave alignment; increasing WPT power may be unavailable under a protected-zone exposure limit; and edge capacity cannot compensate for a severely aged uplink estimate. A factorial or uncertainty-propagation study with measured joint distributions is therefore still required before deployment. The released v14 evaluator is intended to make those extensions straightforward rather than to claim comprehensive robustness.

8.4. Physical Deployment and Remaining Limitations

The proposed method is a resource-management layer that can be integrated with an existing swarm planner. It does not replace collision avoidance, formation control or path planning. The additions in Sections 3.1 and 3.5 remove the implicit assumption that all commanded WPT links remain aligned and all workloads are independent, but they do not substitute for physical calibration.

Mode-specific operation under motion, misalignment and blockage: Table 3 makes explicit that the three WPT options are not interchangeable. Far-field RF is the least geometrically restrictive: a broad or electronically steered beam can continue while a robot moves, but attitude, polarisation, shadowing and protected-zone exposure reduce the admissible power and received gain. RF therefore serves as a lower-efficiency fallback rather than an assumption of uninterrupted charging. Resonant near-field transfer is a short-range coupling operation. A robot must enter a calibrated charging zone, reduce translational and angular motion, maintain coil gap and orientation and satisfy foreign-object and thermal checks for the requested dwell. Directional lightwave transfer is the most dependent on LOS and pointing. The beam tracker must keep the receiver inside the optical footprint, while any person, robot, obstacle or tracking fault in the path triggers an independent shutter or beam-stop interlock.

Equations (4)–(6) therefore use the minimum predicted clearance and gain over the whole dwell interval and add mode-specific speed and hardware-health gates. The operational transition is fail-safe: when near-field coupling or lightwave tracking is lost,

the controller may select admissible RF charging, postpone the transfer or issue the no-charge action; it does not continue the requested directional mode on stale geometry. In an implementation, the pose-error bounds, speed limits, coil and receiver alignment maps, blockage predictor, $g_m^{\min}$ and protection-state latency must be identified from the actual robot, charger and site. Because the present dataset contains no measured pose-to-efficiency, foreign-object, interlock or dynamic-blockage traces, the nominal mode-selection gains remain upper-bound operating points rather than demonstrated performance on moving hardware.

Channel and transfer-gain uncertainty: Perfect CSI can overstate the feasible uplink rate, SWIPT decoding margin and harvested energy. Equations (8) and (30) use deterministic uncertainty sets, following the conservative principle used in robust wireless-power allocation [21]. If empirical estimation errors are available, these bounds can be replaced by outage or chance constraints. The uncertainty radius should include pilot error, channel ageing during T_{loop} , blockage and model/calibration error. Larger radii reserve more bandwidth and charging margin but can make the controller conservative; this robustness–efficiency trade-off is not quantified by the current perfect-estimate Monte Carlo runs.

Safety limits and fail-safe operation: Equations (32) and (33) turn exposure into a hard feasibility condition rather than a soft energy penalty, while Equation (6) requires the corresponding hardware protection state to be healthy. For RF, the relevant quantities include electric or magnetic field and incident power density at protected locations. For resonant near-field transfer, exposure must be considered together with induced currents, foreign-object heating and receiver temperature. For lightwave transfer, accessible optical emission and the possibility of a moving beam entering an eye or skin exposure path dominate the protection design. The allowable limits depend on frequency or wavelength, averaging interval, operating environment and protected population. IEEE C95.1 covers human exposure to electric, magnetic and electromagnetic fields from 0 Hz to 300 GHz, while the ICNIRP radiofrequency guidelines cover 100 kHz to 300 GHz [26,27]. Directional lightwave hardware additionally requires classification and accessible-emission safeguards under IEC 60825-1; IEC TS 60825-19 addresses emitting apertures on moving platforms [28,29]. Consequently, the nominal 15 W resource limit is not a certification result. A deployable system needs measured field/irradiance and temperature maps, independent beam or coupler inhibition, watchdogs and fail-safe interlocks that remain effective if the optimisation process, perception system or communication link fails.

Multi-robot and multi-beam coexistence: Equation (10) separates co-channel robot traffic from WPT leakage into communication receivers, and Equation (34) adds receiver/interference and beam-count limits. This extension is required when several robots transmit concurrently, more than one charger is active, or energy and data phases overlap. The current single-beam, orthogonal, time-separated experiment sets these terms to zero. It therefore does not quantify receiver desensitisation, rectifier saturation from unintended beams, inter-hub conflicts, or coherent combining. Those quantities must be obtained from waveform-level and hardware measurements rather than selected to improve the optimiser’s outcome.

Task dependencies: Fractional parallel offloading is optimistic for perception and planning pipelines such as acquisition–preprocessing–feature extraction–classification or mapping–localisation–planning. Equations (21) and (22) provide the required DAG extension: a successor waits for all predecessors, and an intermediate output incurs communication delay when the execution location changes. This extension increases the discrete decision dimension and can reduce the energy benefit of partial offloading because fewer stages can execute concurrently. The reported task-success values therefore apply to inde-

pendently partitionable workloads and should not be transferred directly to tightly coupled pipelines without a task-graph experiment.

The present evidence remains simulation-based. The timing audit is analytical and does not measure the execution time of the solver on a particular base-station processor, the radio-driver and protocol-stack delay, packet-loss bursts, clock synchronisation error, command parsing time or the energy consumed by the control exchange. The headline energy results retain the nominal task and channel models used in the Monte Carlo study; Equation (17) identifies the additional term that must be populated from hardware measurements. Accordingly, this paper does not claim unconditional 1 s real-time feasibility or robustness to arbitrary motion and uncertainty. Hardware-in-the-loop experiments should measure the distribution of every term in Equation (14), obtain pose/alignment and CSI-error traces, repeat the optimisation under delayed and missing reports, and execute representative DAG workloads.

Other limitations remain. The priority and objective weights in Tables 7 and 8 are fixed design parameters rather than learned or optimised variables. The common-trace comparison covers sequential, fixed-SWIPT and compact Q-learning policies, and the v14 study adds seven one-factor-at-a-time sweeps; neither experiment is a full factorial uncertainty analysis or an exhaustive reproduction of TADP-, Lyapunov-, actor–critic- and multi-agent-learning schedulers. Future work should compare those recipient-selection policies under identical hardware-calibrated swarm dynamics and perform a factorial sensitivity study of the penalty, primal-step, dual-gain, stopping-tolerance and hold-length settings. The nominal FDMA model sets $I_i = 0$, uses one primary beam and time-separates WPT and data transmission; therefore, the results do not quantify the coexistence terms added in Equations (10) and (34). Mobility energy is treated as exogenous, so the current framework does not yet co-optimize trajectories and charging dwell.

The state-prediction layer is also centralised. The simulator supplies robot pose, queue and battery state to the controller, whereas Cao et al. demonstrated that federated learning can recover mobility-conditioned routing and demand information across small-cell domains without collecting all raw user records at one server [13]. A swarm-oriented extension could let robots or local clusters federatively estimate pose-transition, contact-opportunity and task-arrival models, then pass predictive distributions rather than raw trajectories to the robust WPT gate. Such an extension would introduce new costs and failure modes—model-update traffic, non-identically distributed robot data, aggregation delay, stale models and possible poisoning—which are absent from the present optimisation and numerical evaluation. The recommended study is therefore used to motivate a future distributed prediction layer, not presented as a direct baseline or as evidence already demonstrated by WPT-JCCO. These limitations delimit the claims and define the experiments required before field deployment.

9. Conclusions

This paper presented WPT-JCCO, a supervisory resource-allocation framework for swarm robots. Its novelty can be stated in one sentence: WPT-JCCO is the first model considered in this paper's comparison set to place robot-modality charging selection, SWIPT/MEC resource allocation, weakest-robot battery protection and end-to-end command freshness in one mixed discrete–continuous feasible action. Compared with SWIPT-MEC, it adds heterogeneous robot-modality selection and physical availability constraints; compared with WPT scheduling, it embeds charging inside the task/radio/CPU optimisation; compared with swarm resource allocation, it makes energy replenishment endogenous and constrains the minimum individual battery. The algorithmic contribution is the augmented-Lagrangian decomposition, integral candidate screening and feasibility/timing

acceptance sequence tailored to that model. It is not a new general-purpose primal–dual method, and no global mixed-integer optimum is claimed. Energy minimisation, wireless-power offloading, nonlinear harvesting and priority scoring remain established ingredients.

For the evaluated 20-robot scenario, WPT-JCCO reduced net energy depletion by 23.8% relative to CCO with static WPT and by 49.7% relative to local-only execution. It also improved task success and the minimum residual battery. The ablation study showed that adaptive WPT-mode selection, SWIPT splitting, battery-fairness feedback and edge offloading each contributed to the result. The convergence sweeps further showed that the continuous allocation core stabilised within a 95th-percentile limit of 46 outer iterations for $N = 10$ –60 and $\lambda = 0.25$ –1.00 in the tested setting. In the separate common-trace strong-baseline stress test, WPT-JCCO reduced the normalised composite cost by 11.1% relative to two-stage WPT+CCO, 11.3% relative to fixed-SWIPT dynamic offloading and 5.8% relative to the offline Q-learning scheduler. The seven-factor generality study showed gradual degradation under larger swarms and heavier task arrivals, service recovery with additional bandwidth and edge capacity, increasing propulsion cost with mobility, and diminishing charging returns once the adopted nonlinear harvester approached saturation. At $N = 60$ and $\lambda = 0.30$, the within-sweep normalised cost was 25.5% and 29.0% above the nominal cases, respectively; at 16 MHz and 48 GHz it was 3.8% and 2.1% below nominal. These findings support the proposed coupling within the stated evaluators; they do not imply universal superiority or factorial robustness under different hardware and joint operating conditions.

A communication-aware timing audit and the revised stability analysis further delimit the online-control claim. One-second supervisory operation is feasible only when measured state-report, queueing, per-sweep solver, command-delivery and actuation delays satisfy the complete sense–compute–actuate budget and the stationarity, feasibility, merit and binary-hold tests pass before the real-time sweep cap. Under a locally contractive fixed-binary update and bounded state-induced optimiser drift, the warm-started error obeys the stated tracking bound; repeated non-contraction triggers the revalidated fallback. This is a conditional numerical/resource-feasibility statement, not a proof of safety-critical robot-motion stability. The revised model additionally makes mode-specific RF, resonant near-field and lightwave operation explicit: motion and alignment envelopes, dynamic blockage, hardware-interlock health, imperfect communication/WPT gain estimates, protected-zone exposure, multi-beam coexistence and dependent task stages enter through conservative gain bounds, dwell and safety gates, interference limits and DAG precedence constraints. Because the current numerical data use perfect epoch-start estimates and independently partitionable tasks, the reported gains remain nominal and should not be interpreted as hardware-robust performance. Future work will therefore prioritise hardware-in-the-loop comparison with representative WP-MEC schedulers and established priority-aware charging rules, measured pose-to-efficiency, blockage, exposure, leakage and CSI-error distributions, certified lightwave interlocks, DAG-based robotic workloads, federated learning of mobility and task-demand models, priority-weight sensitivity, asynchronous and higher-density convergence tests, multi-charger coordination and joint path planning with charging-aware task allocation.

Author Contributions: Conceptualization, A.I. and H.H.; methodology, A.I.; software, A.I.; validation, A.I.; formal analysis, A.I.; investigation, A.I.; resources, J.P.; data curation, A.I.; writing—original draft preparation, A.I.; writing—review and editing, E.N.; visualization, E.N.; supervision, H.H.; project administration, J.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Simulation code and numerical data can be made available by the corresponding author upon reasonable request.

Acknowledgments: The authors acknowledge the use of institutional computing resources for numerical simulations. We are grateful to our collaborators for their contributions to this publication. Their discussions, ideas, and constructive criticism have significantly enhanced the outcomes of this research.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AoI	age of information
CCO	communication–computation optimisation
CPU	central processing unit
CSI	channel-state information
DAG	directed acyclic graph
EMF	electromagnetic field
FL	federated learning
MEC	mobile edge computing
MAC	medium access control
NF	near-field
PDR	packet delivery ratio
RF	radiofrequency
SNR	signal-to-noise ratio
SWIPT	simultaneous wireless information and power transfer
WPT	wireless-power transfer
WPT-JCCO	wireless-power-transfer-aware joint communication–computation–charging optimisation

References

1. Brambilla, M.; Ferrante, E.; Birattari, M.; Dorigo, M. Swarm Robotics: A Review from the Swarm Engineering Perspective. *Swarm Intell.* **2013**, *7*, 1–41. <https://doi.org/10.1007/s11721-012-0075-2>.
2. Mao, Y.; Zhang, J.; Letaief, K.B. Stochastic Joint Radio and Computational Resource Management for Multi-User Mobile-Edge Computing Systems. *IEEE Trans. Wirel. Commun.* **2017**, *16*, 5994–6009. <https://doi.org/10.1109/TWC.2017.2717986>.
3. Mao, Y.; You, C.; Zhang, J.; Huang, K.; Letaief, K.B. A Survey on Mobile Edge Computing: The Communication Perspective. *IEEE Commun. Surv. Tutor.* **2017**, *19*, 2322–2358. <https://doi.org/10.1109/COMST.2017.2745201>.
4. Hu, Z.; Liu, S.; Zhou, D.; Shen, C.; Wang, T. Task Offloading and Data Compression Collaboration Optimization for UAV Swarm-Enabled Mobile Edge Computing. *Drones* **2025**, *9*, 288. <https://doi.org/10.3390/drones9040288>.
5. Xu, C.; Zhang, J.; Yu, H. Control–Communication Co-Optimization for Wireless Cloud Robotic System via Multi-Agent Transfer Reinforcement Learning. *IEEE/CAA J. Autom. Sin.* **2026**, *13*, 311–326. <https://doi.org/10.1109/JAS.2025.125894>.
6. Wang, F.; Xing, H.; Xu, J. Real-Time Resource Allocation for Wireless Powered Multiuser Mobile Edge Computing with Energy and Task Causality. *IEEE Trans. Commun.* **2020**, *68*, 7140–7155. <https://doi.org/10.1109/TCOMM.2020.3011990>.
7. Malik, R.; Vu, M. Energy-Efficient Joint Wireless Charging and Computation Offloading in MEC Systems. *IEEE J. Sel. Top. Signal Process.* **2021**, *15*, 1110–1126. <https://doi.org/10.1109/JSTSP.2021.3098963>.
8. Wu, M.; Qi, W.; Park, J.; Lin, P.; Guo, L.; Lee, I. Residual Energy Maximization for Wireless Powered Mobile Edge Computing Systems with Mixed-Offloading. *IEEE Trans. Veh. Technol.* **2022**, *71*, 4523–4528. <https://doi.org/10.1109/TVT.2022.3147824>.
9. Zeng, S.; Huang, X.; Li, D. Joint Communication and Computation Cooperation in Wireless-Powered Mobile-Edge Computing Networks with NOMA. *IEEE Internet Things J.* **2023**, *10*, 9849–9862. <https://doi.org/10.1109/JIOT.2023.3236089>.
10. Wang, X.; Li, J.; Ning, Z.; Song, Q.; Guo, L.; Guo, S.; Obaidat, M.S. Wireless Powered Mobile Edge Computing Networks: A Survey. *ACM Comput. Surv.* **2023**, *55*, 263. <https://doi.org/10.1145/3579992>.

11. Chen, Y.; Kang, H.; Li, J.; Sun, G.; Wang, B.; Wang, J.; Liang, C.; Liang, S.; Niyato, D. Joint Resource Management for Energy-Efficient UAV-Assisted SWIPT-MEC: A Deep Reinforcement Learning Approach. *IEEE Internet Things J.* **2025**, *12*, 31448–31465. <https://doi.org/10.1109/JIOT.2025.3574332>.
12. Zeng, Y.; Zhang, R.; Lim, T.J. Wireless Communications with Unmanned Aerial Vehicles: Opportunities and Challenges. *IEEE Commun. Mag.* **2016**, *54*, 36–42. <https://doi.org/10.1109/MCOM.2016.7470933>.
13. Cao, Y.; Maghsudi, S.; Ohtsuki, T.; Quek, T.Q.S. Mobility-Aware Routing and Caching in Small Cell Networks Using Federated Learning. *IEEE Trans. Commun.* **2024**, *72*, 815–829. <https://doi.org/10.1109/TCOMM.2023.3327278>.
14. Ayan, O.; Kutsevol, P.; Ozkan, H.Y.; Kellerer, W. Semantics- and Task-Oriented Scheduling for Networked Control Systems in Practice. *IEEE Access* **2022**, *10*, 115673–115690. <https://doi.org/10.1109/ACCESS.2022.3218410>.
15. Ballotta, L.; Talak, R. Safe Distributed Control of Multi-Robot Systems with Communication Delays. *IEEE Trans. Veh. Technol.* **2025**, *74*, 10137–10150. <https://doi.org/10.1109/TVT.2025.3546857>.
16. European Telecommunications Standards Institute. *5G; Service Requirements for the 5G System (3GPP TS 22.261 Version 19.12.0 Release 19)*; ETSI TS 122 261 V19.12.0; ETSI: Sophia Antipolis, France, 2025.
17. Luo, X.; Liu, C.; Yue, X. Improved-PPO-Based Adaptive Safety-Critical Control for Orbit–Attitude Takeover of Failed Spacecraft with Cellular Satellites. *IEEE Trans. Autom. Sci. Eng.* **2026**, *23*, 10872–10882. <https://doi.org/10.1109/TASE.2026.3700197>.
18. Bi, S.; Ho, C.K.; Zhang, R. Wireless Powered Communication: Opportunities and Challenges. *IEEE Commun. Mag.* **2015**, *53*, 117–125. <https://doi.org/10.1109/MCOM.2015.7081084>.
19. Ojha, T.; Raptis, T.P.; Passarella, A.; Conti, M. Wireless Power Transfer with Unmanned Aerial Vehicles: State of the Art and Open Challenges. *Pervasive Mob. Comput.* **2023**, *93*, 101820. <https://doi.org/10.1016/j.pmcj.2023.101820>.
20. Alabsi, A.; Hawbani, A.; Wang, X.; Al-Dubai, A.Y.; Aziz, S.A.; Kumar, S.; Zhao, L.; Shvetsov, A.V.; Alsamhi, S.H. Wireless Power Transfer Technologies, Applications, and Future Trends: A Review. *IEEE Trans. Sustain. Comput.* **2025**, *10*, 1–17. <https://doi.org/10.1109/TSUSC.2024.3380607>.
21. Boshkovska, E.; Ng, D.W.K.; Zlatanov, N.; Koelpin, A.; Schober, R. Robust Resource Allocation for MIMO Wireless Powered Communication Networks Based on a Non-Linear EH Model. *IEEE Trans. Commun.* **2017**, *65*, 1984–1999. <https://doi.org/10.1109/TCOMM.2017.2664860>.
22. Wang, G.; Yu, X.; Xu, F.; Cai, J. Task Offloading and Resource Allocation for UAV-Assisted Mobile Edge Computing with Imperfect Channel Estimation over Rician Fading Channels. *EURASIP J. Wirel. Commun. Netw.* **2020**, *2020*, 169. <https://doi.org/10.1186/s13638-020-01780-8>.
23. Zhao, G.; Xu, H.; Zhao, Y.; Qiao, C.; Huang, L. Offloading Dependent Tasks in Mobile Edge Computing with Service Caching. In *Proceedings of the IEEE INFOCOM 2020—IEEE Conference on Computer Communications, Toronto, ON, Canada, 6–9 July 2020*; IEEE: New York, NY, USA, 2020; pp. 1997–2006. <https://doi.org/10.1109/INFOCOM41043.2020.9155396>.
24. Lin, C.; Wang, Z.; Han, D.; Wu, Y.; Yu, C.W.; Wu, G. TADP: Enabling Temporal and Distantial Priority Scheduling for On-Demand Charging Architecture in Wireless Rechargeable Sensor Networks. *J. Syst. Archit.* **2016**, *70*, 26–38. <https://doi.org/10.1016/j.sysarc.2016.04.005>.
25. Yang, P.; Wu, T.; Dai, H.; Rao, X.; Wang, X.; Wan, P.J.; He, X. MORE: Multi-Node Mobile Charging Scheduling for Deadline Constraints. *ACM Trans. Sens. Netw.* **2020**, *17*, 7. <https://doi.org/10.1145/3410454>.
26. *IEEE Std C95.1-2019*; IEEE Standard for Safety Levels with Respect to Human Exposure to Electric, Magnetic, and Electromagnetic Fields, 0 Hz to 300 GHz. Institute of Electrical and Electronics Engineers: Piscataway, NJ, USA, 2019. <https://doi.org/10.1109/IEEESTD.2019.8859679>.
27. International Commission on Non-Ionizing Radiation Protection. Guidelines for Limiting Exposure to Electromagnetic Fields (100 kHz to 300 GHz). *Health Phys.* **2020**, *118*, 483–524. <https://doi.org/10.1097/HP.0000000000001210>.
28. *IEC 60825-1:2014*; Safety of Laser Products—Part 1: Equipment Classification and Requirements, 3rd ed. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2014.
29. *IEC TS 60825-19:2024*; Safety of Laser Products—Part 19: Moving Platform Laser Products. International Electrotechnical Commission (IEC): Geneva, Switzerland, 2024.
30. Razaviyayn, M.; Hong, M.; Luo, Z.-Q. A Unified Convergence Analysis of Block Successive Minimization Methods for Nonsmooth Optimization. *SIAM J. Optim.* **2013**, *23*, 1126–1153. <https://doi.org/10.1137/120891009>.
31. Conn, A.R.; Gould, N.I.M.; Sartenauer, A.; Toint, P.L. Convergence Properties of an Augmented Lagrangian Algorithm for Optimization with a Combination of General Equality and Linear Constraints. *SIAM J. Optim.* **1996**, *6*, 674–703. <https://doi.org/10.1137/S1052623493251463>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.