



ORIGINAL RESEARCH OPEN ACCESS

Energy-Aware Optimisation for Off-Grid ORAN With RIS and Edge Computing

Yassir AL-Karawi¹  | Raad S. Alhumaima¹  | Hamed Al-Raweshidy²
¹Department of Communications Engineering, University of Diyala, Baqubah, Iraq | ²Department of Electronic and Computer Engineering, Brunel University of London, Uxbridge, UK

Correspondence: Hamed Al-Raweshidy (hamed.al-raweshidy@brunel.ac.uk)

Received: 30 May 2025 | **Revised:** 2 July 2025 | **Accepted:** 21 July 2025

Handling Editor: Ghazanfar Safdar

Funding: This work was partially supported by Brunel University of London.

Keywords: access protocols | cellular radio | channel allocation | communication complexity | energy consumption | mobile radio | Monte Carlo methods | next generation networks

ABSTRACT

This paper proposes a strategy for designing Open Radio Access Networks (ORAN) to maximise their energy efficiency using solar power, supplemented by Reconfigurable Intelligent Surfaces (RIS) and Mobile Edge Computing (MEC). Because grid power is not always available where these ORAN systems are built, our approach manages the difficulties created by dynamic energy and timing issues found in isolated environments. The approach concentrates on allocating energy to all transmitters, CPU speed and RIS phases in real time, subject to strict rules on power use, latency issues and heat. The primal-dual algorithm we propose reacts to queue and energy changes to update the dual variables and control policies without access to every channel parameter. Our combined (composite) cost function measures energy use, delays encountered by users, reliability of the SINR and fairness. Results from the simulation indicate that using the proposed method lowers energy usage by 25% and average delay by 18%, outperforming baseline models under varying solar and traffic patterns. Robustness is further validated through sensitivity and ablation analyses. This work demonstrates the feasibility of deploying sustainable, intelligent ORAN infrastructures in remote 6G scenarios where conventional power and connectivity are unavailable.

1 | Introduction

While 5G is geared towards increasing capacity, 6G will introduce systems that are both very dependable and responsive to delay and just as important, are designed to work intelligently and sustainably [1, 2]. They are expected to function beyond a micro-millisecond delay and more than one gigabit of data transfer at one time, working independently in terrestrial, aerial and satellite environments. Among these services, extending reality (XR) and managing factory processes automatically, from a distance, each will need energy efficiency, as well as secure processing and real-time delay monitoring, in place [3].

At the heart of these developments is ORAN which replaces the old base station with a modular, programmable and virtualised design. ORAN organises the base station into the Radio Unit (RU), Distributed Unit (DU) and Centralised Unit (CU) and all three work together using agreed open interfaces [4]. Connections made using enhanced Common Public Radio Interface (eCPRI) between RU and DU, the F1 interface joining DU and CU and the E2 interface tied to the RAN Intelligent Controller (RIC) support various technology vendors and AI plasticity. Because ORAN separates functionality and supports running new code throughout use, it gives the adaptability necessary for network orchestration and adapting to changing requirements.

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *IET Networks* published by John Wiley & Sons Ltd on behalf of The Institution of Engineering and Technology.

Nonetheless, running ORAN in settings where the grid is not available creates special challenges for operators. Remote rural regions, high mountains and recovery areas after disasters lack reliable electricity, so base stations there rely completely on renewable energy. The use of solar energy for power makes radio and the delivery of sensitive services unreliable because of constant changes in weather and lack of storage space which causes variability in the technology [5].

When signals are easily degraded in tight conditions, Reconfigurable Intelligent Surfaces (RIS) are being used as a key solution. The phase of incoming electromagnetic waves is shifted through passive adjustments in a RIS panel to make passive beamforming possible [6], so signal reliability is high along with better spectrum use and non-line-of-sight communication without using active RF devices. With the use of RIS within the ORAN model, challenges involving power, phase synchronisation and instant energy conditions require that related solutions be synchronised.

MEC (Multi-access Edge Computing) is an inclusion in this group because it places time-sensitive and data-centred tasks almost next to users. With MEC nodes within or next to RUs, users have a simpler way to offload tasks, do local analysis and cut down on cloud or backhaul connexions. Still, planning tasks on MEC servers is essential, because the power may come from renewable sources and packaging of tasks must consider their priority, stored power and the peak temperature of equipment [7].

Much research has been done on solar-conscious energy routing, RIS beamforming and MEC task planning, but most scholars have rarely looked at them all at once. A tiny amount of research address the question of how to combine optimising energy harvesting, RIS arrangement and MEC use in a completely ORAN theory. Additionally, applying control in critical safety situations does not allow usual reinforcement learning methods to converge or be used with ease for generalisation [9].

We recommend an optimised stochastic structure which works in the setting of fully off-grid Open RAN networks. With this model, these three associated layers are brought together: (i) capturing solar energy and making energy plans, (ii) perfecting signals with RIS and (iii) setting up and selecting CPU tasks for the MEC system. This basic mechanism is formed as a non-convex, constrained problem and worked out effectively through stochastic primal-dual Lagrangian decomposition. Thanks to this approach, the outcomes can be implemented, updating the convexified dual variables and SINR constraints.

A diagram is shown (Figure 1) to illustrate how solar units, DU/CU tiers and the RIS modules function with the tracking system.

The primary contributions of this paper are as follows:

1. A mixed framework for ORAN that brings together RIS beamforming, scheduling of MEC applications and energy harvesting without a grid, with live control managed by

Off-Grid Open RAN Architecture




FIGURE 1 | Energy-conscious off-grid ORAN builds on solar-powered remote units, separate DU/CU functionality, RIS components and a means for core network connection.

the RIC and choices made close to the MEC nodes in real time.

2. This method is scalable, uses queue-state information and relies on fairness-oriented utility functions for stable behaviour in uncertain scenarios.
3. The use of realistic simulations of solar activity, human motion and complicated user tasks proved that the design offers better energy gain, reduces wait time and controls stability for the signal to interfering noise ratio, over the original approaches.

The progress of the paper is summarised below. It describes and assesses many solutions, highlighting the weaknesses found in today's approaches. Details of the system model and the relationship between energy, wireless and computational systems can be found in Section 3. The framework for optimisation is fully explained in Section 4. Finally, the simulation results and their comparisons can be found in Section 5 and the article finishes with some concluding thoughts and suggestions for continuing research.

2 | Related Work

Building sixth-generation (6G) networks calls for structures that surpass previous goals for throughput and latency, caring about environmental protection, using smartness across many elements and ensuring they remain reliable. They most strongly apply when backup is needed in rural areas, by themselves-running processes and during emergencies. Multiple experts are now considering how convergence of Open Radio Access Networks (ORAN), Reconfigurable Intelligent Surfaces (RIS), Mobile Edge Computing (MEC) and energy harvesting (EH) technologies can help solve these problems. To date, making all these domains work together within a general optimisation scheme is still largely undeveloped.

NEW ORAN allows cellular infrastructure to be built with up to three separable parts: a Radio Unit, a Distributed Unit and a Centralised Unit. Components linked by standardised connections such as eCPRI, F1 and E2 are managed by the RAN Intelligent Controller which creates the requirement for xApps and rApps (i.e., near-real-time and non-real-time RAN applications, respectively) for both real-time and practical control. A lot of research and development [8, 9] focus on modularity, compatibility and AI management in ORAN, but mostly use an assumption of constant grid power. Recent studies on solar-powered ORAN only consider a few areas, ignoring the way channels behave, the need for real-time restrictions and adaptable beamforming which limits their usefulness for practical models.

MEC contributes to ORAN by moving computation near the edges, so latency is reduced and a burden is lifted from central cloud servers. While many studies have examined ways to off-load dynamic tasks using optimisation or reinforcement learning, they usually forget about the changing power levels found in off-grid environments. In addition, the combined impact of energy arrivals, temperature limits and users'

individual delays is rarely looked at with RIS and distributed control systems.

It is now recognised that RIS technology can improve wireless propagation by automatically changing how passive reflectors work. With millimetre waves, spectral efficiency rises and two devices can communicate even if they are not in a direct line of sight. Single-user and vehicle environments use initial research based on passive beamforming [5, 10]. Even so, putting RIS into practise with the ORAN model is still at an early development stage. These common assumptions such as constant channel and power information, do not fit the actual energy and delay metrics found in ORAN networks using solar energy which require RIS phase selection to adjust. In most cases, the links between RIS control, task scheduling in MEC and optimisation via the RIC are overlooked.

Using solar energy is a preferred method for powering autonomous deployments of ORAN. This research looks at solar models [11, 12] that cover controlling energy supply on uncertain days, monitoring batteries and simulating photovoltaics. Even so, they are usually not connected to the real-time functions of a programme or system. It is unusual to find assessments of how intermittent energy affects beamforming or how stable mobile edge cluster queues remain using ORAN models. Even the most advanced studies such as [8, 13], only cover single aspects of improving a system.

In the case of ref. [9], a deep reinforcement learning model is put forward to manage delay in RIS-MEC systems, assuming power is always accessible. At the same time, the design is not designed for how unreliable renewable energy can be or how batteries or heat impact the system in areas without a grid [14]. It is developed to improve how tasks are placed and does not involve RIS management or real-time beamforming control. Besides, because the system is built around a core network, it cannot support cases where devices have to work independent of the core server [15] and similar research investigate dual decomposition for solving spectrum issues in 6G, but no detailed model links energy levels, latency and signal accuracy in situations with power constraints.

In this research, we bring together solar energy collection, RIS beamforming management and MEC scheduling in a single framework designed for off-grid ORAN. The system changes its strategy automatically to handle real-time changes in energy, delays and channels and doing so uses the RIC and the antenna nodes at the edge for robust coordination.

In Table 1, the most important recent studies are compared, based on ORAN support, RIS-MEC integration, being efficient for off-grid scenarios, strong delay resistance and supporting instant optimisation of results. It is shown by the data that no published project has brought together all these components in an efficient and scalable manner for 6G operations.

By analysing how energy, computation and communication work together in off-grid 6G, our study meets a new standard for strong, scaled and intelligent networks.

3 | System Model

3.1 | System Architecture and Energy Modelling

We now explain how the off-grid ORAN network with RIS and MEC is structured and how energy is modelled. The system is set up with scheduled time periods $t \in \{1, 2, \dots, T\}$, during which each RU relies on solar energy, stores it in a local battery and handles its tasks. The system consists of three separate layers as illustrated in Figure 2.

Photovoltaic panels and battery modules make up the Energy Layer which supports the rest of the structure. Solar irradiance is captured by these units and they change it into useful power with the highest efficiency over the panel surface. Based on

$$P_b^h(t) = \eta_h A_b S_b(t), \quad (1)$$

the amount of power gathered is set by irradiance, the area of the panel and the efficiency it has [22].

We measure the total energy gathered in every time slot by multiplying the power by the interval over which it was collected, using

$$E_b^h(t) = \eta_h A_b S_b(t) \cdot \Delta t. \quad (2)$$

The total energy source on the horizon is added into

$$E_b^{tot}(T) = \sum_{t=1}^T E_b^h(t), \quad (3)$$

helping to analyse if the system will last long-term.

In this paper, we only consider a uniform solar irradiance process $S_b(t)$ and fixed photovoltaic efficiency η_h across RUs. This modelling decision deliberately abstracts spatial variability in order to provide a clean and comparable evaluation of MOAC under uniform off-grid conditions. While real-world deployments will have an inherent heterogeneous irradiance due to environmental spread, such an abstraction is standard practice in system-optimality problems [26], especially when algorithmic performance, not environmental detail, is the interest. In subsequent work this assumption could be relaxed, and other location-specific solar profiles or variable hardware parameters could easily be introduced.

A K-state Markov model is used to describe solar irradiance, with transitions modelled by

$$\mathbb{P}[S_b(t+1) = s_j | S_b(t) = s_i] = \Pi_{ij}. \quad (4)$$

A K-state Markov model is used to model solar irradiance, where transitions between states depend on a stochastic matrix

TABLE 1 | Comparison of related work in RIS-MEC-ORAN research.

Work	ORAN-aware	RIS + MEC integrated	Off-grid EH support	Delay-aware	Real-time optimisation
[16]	✓	✗	✗	✗	✗
[17]	✗	Partial	✗	✓	✗
[18]	✗	✓	✗	Partial	✗
[19]	✗	✗	✓	✗	✗
[9]	Partial	✓	Partial	✓	Partial
This work	✓	✓	✓	✓	✓

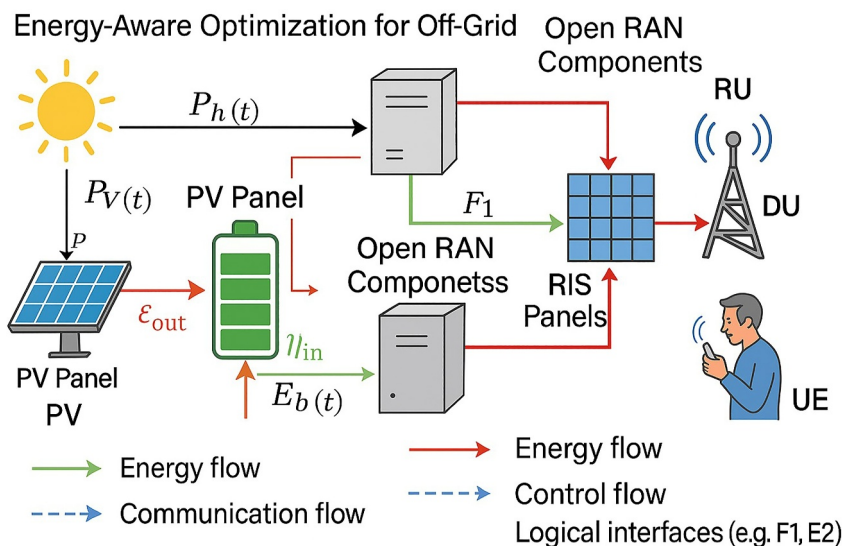


FIGURE 2 | Energy flow and storage model for solar-powered ORAN node.

$\Pi = [\pi_{ij}] \in \mathbb{R}^{K \times K}$. In the name of replication and transparency, we present the transition probabilities we used in our simulations. For $K = 4$, the transition matrix is given by the following equation:

$$\Pi = \begin{bmatrix} 0.7 & 0.2 & 0.1 & 0.0 \\ 0.1 & 0.6 & 0.2 & 0.1 \\ 0.0 & 0.2 & 0.6 & 0.2 \\ 0.0 & 0.1 & 0.2 & 0.7 \end{bmatrix}. \quad (5)$$

Every element π_{ij} represents the probability of transitioning from state s_i to state s_j in the next time period. The probabilities were chosen from observational data that represent typical irradiance behaviour for situations faced in the off-grid space, where irradiance can shift throughout the day due to cloud cover or angle of the sun. This form of modelling aligns with recent work in the literature, such as Hou et al. [21] who apply Markov-based models that explicitly use transition matrices to realistically nowcast solar radiation profiles.

The evolution of the probability distribution over the irradiance states follows:

$$\pi(t+1) = \pi(t) \cdot \Pi. \quad (6)$$

T defines the number of slots and [22] provides details about its real behaviour.

The dynamics of the battery storage are governed by the following equation:

$$E_b^s(t+1) = \min\left(E_{\max}, \max\left(0, E_b^s(t) + \eta_{\text{in}} E_b^h(t) - \frac{E_b^c(t)}{\eta_{\text{out}}}\right)\right), \quad (7)$$

which captures how the stored energy $E_b^s(t)$ at time $t+1$ evolves based on the harvested input $E_b^h(t)$, consumption $E_b^c(t)$, and charge/discharge efficiencies $\eta_{\text{in}}, \eta_{\text{out}}$. The term E_{\max} denotes the maximum battery capacity.

The current model assumes ideal battery behaviour with fixed capacity and no degradation or abrupt failure. This abstraction is adopted to focus on evaluating MOAC's responsiveness under dynamic energy and traffic conditions. In practice, however, battery ageing and state-of-health degradation can significantly impact sustainability and reliability. To acknowledge this, future extensions of the model may integrate degradation-aware profiles or probabilistic failure models, in line with approaches such as refs. [28, 29].

Modelling the related energy losses is part of

$$\Delta E_{\text{loss}}(t) = (1 - \eta_{\text{in}}) E_b^h(t) + \left(1 - \frac{1}{\eta_{\text{out}}}\right) E_b^c(t). \quad (8)$$

We measure the average battery charge over time in

$$\bar{E}_b^s = \frac{1}{T} \sum_{t=1}^T \mathbb{E}[E_b^s(t)], \quad (9)$$

and

$$\mathbb{P}[E_b^s(t) < E_{\text{th}}] \leq \epsilon, \quad (10)$$

makes sure that reliability is enforced by preventing the stored energy from dropping below a necessary threshold E_{th} .

System availability over a long period is measured by the following equation:

$$R_b = 1 - \frac{1}{T} \sum_{t=1}^T \mathbf{1}[E_b^s(t) < E_{\text{th}}], \quad (11)$$

and energy supply and demand gaps are shown by the following equation:

$$D_b^{\text{var}} = \frac{1}{T} \sum_{t=1}^T |E_b^c(t) - \eta_{\text{in}} E_b^h(t)|. \quad (12)$$

These expressions support the main structure of energy-aware modelling in ORAN. The whole process of deciding on resources and schedules depends on the energy model.

“The key mathematical symbols pertinent to energy harvesting, storage dynamics, and power consumption are presented in Table 2, which are consistently referenced throughout the model's subsequent discussion. To visually enhance this abstraction, the entire setup of the proposed standalone ORAN framework, incorporating RIS modules and MEC servers, is depicted in Figure 2.”

3.2 | Channel and RIS Beamforming

Now, we review the ways in which Reconfigurable Intelligent Surfaces (RIS) modify wireless channels in an off-grid ORAN environment. In regions where buildings cause signal loss or signal attenuation is strong, RIS modules allow the signal to be programmed for better communication. All RISs are made up of elements and every element generates its own phase shift all the time.

$$\theta_r(t) = e^{j\phi_r(t)}. \quad (13)$$

The signal phase $\phi_r(t)$ is shown in the equation to be responsible for controlling the electromagnetic response of the surface [5].

In this way, users can receive signals from the base as well as from a reflection on the RIS. The base station sends signals directly to the user and also through several RIS-reflected paths.

$$h_{bu}^{\text{eff}}(t) = h_{bu}^{(d)}(t) + \sum_{r=1}^R \theta_r(t) h_{br}(t) h_{ru}(t). \quad (14)$$

How well each contribution works depends on the product of the channel coefficients between the base station and RIS, between the RIS and the user and the reflection coefficient.

Usually, the RIS phase vector can be represented by a diagonal matrix.

TABLE 2 | Summary of symbols and Notations used in the model.

Symbol	Definition
$I_u(t)$	Interference power received at user u
Q_{th}	Maximum task queue threshold
$P_b^{dyn}(t)$	Dynamic CPU power consumption at RU b
P_b^{stat}	Static leakage power at RU b
$E_{acc}(t)$	Accumulated consumed energy up to time t
$E_{tot}(t)$	Total harvested energy available at time t
$\theta_r(t)$	RIS reflection coefficient ($e^{j\phi_r(t)}$)
$\phi_r(t)$	Phase shift at RIS element r
$f_u(t)$	CPU frequency allocated to user u
$P_b^{tx}(t)$	Transmission power of RU b at time t
P_{ctrl}	Control-plane operational power
$h_{bu}^{eff}(t)$	Effective wireless channel from RU b to user u
$S_b(t)$	Solar irradiance at RU b at time t
A_b	Solar panel surface area
η_h	Solar energy conversion efficiency
$E_b^h(t)$	Harvested solar energy at RU b
$E_b^{acc}(T)$	Cumulative harvested energy up to time T
η_{in}, η_{out}	Battery charging and discharging efficiency
$E_b^s(t)$	Stored battery energy at RU b at time t
$E_b^c(t)$	Energy consumed by RU b at time t
F_b^{max}	Maximum CPU frequency budget at RU b
P_{max}	Maximum transmission power limit
Γ_{min}	Minimum SINR threshold
$SINR_{bu}(t)$	SINR of user u from RU b at time t
$D_u(t)$	End-to-end delay experienced by user u
D_u^{max}	Maximum tolerable delay for user u
$Q_u(t)$	Task queue length for user u
L_u	Computational task density (cycles/bit)
λ_u	Task arrival rate for user u
Δt	Time slot duration
λ_k	Lagrangian multiplier for constraint k
α_k	Dual variable update step size
$g_k(t)$	Constraint violation term for constraint k
$J(t)$	Jain's fairness index at time t
$C_E(t), C_D(t), C_O(t), C_B(t), C_F(t)$	Cost components: Energy, delay, SINR, battery, fairness
$C_{total}(t)$	Total composite cost function
$a(t)$	Control action vector (power, CPU, RIS, scheduler)
w_E, w_D, w_T	Weights for energy, delay, and thermal penalty
σ^2	Noise power

$$\Theta(t) = \text{diag}(e^{j\phi_1(t)}, \dots, e^{j\phi_R(t)}). \quad (15)$$

We depend on these formulae to model advanced wireless propagation for 6G [5, 10]. The user's SINR can be written as follows:

All passive elements have a total channel vector that is:

$$\mathbf{h}_{bu}^{total}(t) = \mathbf{h}_{bu}^{(d)}(t) + \mathbf{h}_{br}^\top(t) \Theta(t) \mathbf{h}_{ru}(t). \quad (16)$$

$$SINR_{bu}(t) = \frac{P_b^{tx}(t)|h_{bu}^{eff}(t)|^2}{\sum_{b' \neq b} P_{b'}^{tx}(t)|h_{b'u}^{eff}(t)|^2 + \sigma^2}. \quad (17)$$

The ratio is influenced by the way the RIS is set up. The achievable rate is defined by both the bandwidth and the SINR.

$$R_u(t) = W \log_2(1 + \text{SINR}_{bu}(t)). \quad (18)$$

Using this allows for direct evaluation of performance. The behaviour of the RIS should be improved by considering how different phases affect the effective channel. That is why the gradient of the effective channel for each phase shift is calculated as follows:

$$\frac{\partial h_{bu}^{\text{eff}}(t)}{\partial \phi_r(t)} = j e^{j\phi_r(t)} h_{br}(t) h_{ru}(t). \quad (19)$$

As a result, gradient-based changes to the RIS are possible.

In more advanced applications, a number of RIS panels are needed. The effective channel is created by the way light bounces off every surface.

$$h_{bu}^{\text{eff}}(t) = h_{bu}^{(d)}(t) + \sum_{i=1}^{N_{\text{RIS}}} \sum_{r=1}^{R_i} \theta_{ir}(t) h_{bir}(t) h_{iru}(t). \quad (20)$$

In the above equation we assume that the RIS has ideally continuous-phase shifts at each RIS element. In practice, there are normally hardware limitations such as phase quantisation (e.g., 1-bit or 2-bit phase resolution) and thermal noise. Typically, the quantised phase quantisation can impact the performance of precise beamforming and also create signal degradation in the channel. Future extensions can incorporate discrete-phase models and noise-aware channel models as done in refs. [27, 30].

This way of thinking helps more people use the network and allows them to use the same area in a variety of ways. Figure 3 illustrates that users receive signals from the RU by direct and reflected paths.

The inclusion of these expressions allows RIS beamforming to be matched with power and CPU scheduling, thus letting the ORAN system address various channel conditions, low energy and nearby interferers.




FIGURE 3 | The RIS helps in off-grid ORAN by adding direct h_{bu} and reflected signal paths.

3.3 | Local MEC Computation and Scheduling

Each computation task is represented as a tuple $(\kappa_u, \delta_u, \tau_u)$ indicating κ_u is the number of CPU cycles required for processing while δ_u is the size of task data in bits, and τ_u is the deadline by which the task must run. These tasks are generated thus locally at the RU level and are to be executed on embedded MEC modules generating local processing capabilities, rather than an external cloud server.

As soon as we must model the processing behaviour, we note that we assume that the MEC processor in RU b operates at an adjustable frequency $f_b(t)$ at time slot t . The energy, consumed to execute a task that has κ_u cycles at a frequency is defined as follows:

$$E_u^{\text{proc}}(t) = \zeta f_b(t) \cdot \kappa_u, \quad (21)$$

where ζ is the effective switched capacitance coefficient for the processor. The corresponding processing delay is defined as follows:

$$D_u^{\text{proc}}(t) = \frac{\kappa_u}{f_b(t)}. \quad (22)$$

In order to adhere to latency guarantees, we have enforced a delay constraint that states that $D_u^{\text{proc}}(t) \leq \tau_u$. The MEC processor is adaptive and utilises available energy at any given time, $E_b^s(t)$, and urgency of workload to adjust $f_b(t)$, thus allowing for a trade-off on a delay scale versus battery conservation.

This delay-aware, energy-constrained framework for task execution appropriately models the realistic operation of embedded edge computation in off-grid ORAN systems, and continues the patterns of models in the mobile edge computing literature [32, 40].

3.4 | Task Queueing and Delay Modelling

Accurately modelling task queueing and delay helps provide timely service when energy and computational resources are limited in off-grid ORAN systems with MEC. Tasks produced by every user are held in a queue and can be processed on the local device or transferred to the MEC server for processing.

Let $Q_u(t)$ indicate the amount of packets waiting in queue for user u in time slot t . The queue's behaviour changes as the rate of tasks arriving and being served changes.

$$Q_u(t+1) = \max\{0, Q_u(t) + \lambda_u(t) - \mu_u(t)\}. \quad (23)$$

The system is able to manage load by adapting, where $\lambda_u(t)$ shows the load coming in and $\mu_u(t)$ is the number of bits handled in each slot t . With this simple recursion, we can study how the queue is stable.

The user's service rate depends on the CPU frequency they are given and the difficulty of their task.

$$\mu_u(t) = \frac{f_u(t)}{L_u(t)}. \quad (24)$$

In this section, $f_u(t)$ represents the cycles per second speed of processing and $L_u(t)$ represents the number of cycles necessary to process one bit.

At any time, the queue size can be found by taking an average of past queue sizes.

$$\mathbb{E}[Q_u] = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T Q_u(t). \quad (25)$$

By measuring this metric, we can predict delays and see how closely related it is to the stability of the system as shown in ref. [23].

Every user's end-to-end delay is made up of transmission delay, processing delay and queueing delay. Every one of these is designed separately.

To send a task bit over the wireless link takes the transmission delay which is:

$$D_u^{\text{tx}}(t) = \frac{L_u(t)}{R_u(t)}, \quad (26)$$

where $R_u(t)$ is determined by the SINR of the user's data.

The processing delay means the time it takes to execute a task on either MEC or RU.

$$D_u^{\text{proc}}(t) = \frac{L_u(t)}{f_u(t)}. \quad (27)$$

Little's Law [25] says that queueing delay is directly connected to the length of the queue and inversely connected to the service rate.

$$D_u^{\text{queue}}(t) = Q_u(t) \cdot \frac{L_u(t)}{f_u(t)}. \quad (28)$$

For this reason, the overall delay for an infinitesimal interval is gained from adding all three parts.

$$D_u(t) = D_u^{\text{tx}}(t) + D_u^{\text{queue}}(t) + D_u^{\text{proc}}(t). \quad (29)$$

In these systems, the delay must always be lower than a defined limit. For this reason, the latency violation probability is defined.

$$P_u^{\text{delay}} = \mathbb{P}[D_u(t) > D_u^{\text{max}}]. \quad (30)$$

The probability is used to evaluate Quality of Service and maintain responsiveness.

We use the Pollaczek–Khinchine formula from M/G/1 queueing theory to improve the model by estimating the expected queueing delay.

$$\mathbb{E}[D_u^{\text{queue}}] \approx \frac{\rho_u}{1 - \rho_u} \cdot \frac{L_u}{f_u}. \quad (31)$$

Utilisation factor is shown by the symbol ρ_u . This finding helps characterise the delay in systems where arrivals come in clusters.

An Exponentially Weighted Moving Average (EWMA) filter is used to fix problems with temporal delay fluctuations.

$$\bar{D}_u(t+1) = \alpha_D D_u(t) + (1 - \alpha_D) \bar{D}_u(t). \quad (32)$$

The filter coefficient $\alpha_D \in (0, 1)$ is used to control the effect of new delay observations.

To explicitly model the tradeoff between energy consumption and delay, we define the following weighted objective:

$$J_u(t) = \varpi \cdot \bar{D}_u(t) + (1 - \varpi) \cdot E_u^{\text{cpu}}(t). \quad (33)$$

In this case, the symbol $\bar{D}_u(t)$ denotes the filtered overall delay (see Equation (32)), while $E_u^{\text{cpu}}(t)$ denotes the energy used by the user's local processor during task execution. The scalar variable ϖ controls the trade-off between delay sensitive and energy efficient.

For example, in safety-critical applications where timely responses are essential, we may expect ϖ values for low latency operation—perhaps close to 1. In battery powered applications where energy consumption must be minimised, we may expect ϖ values that allow for energy saving actions, even with finitely more delay time. Using this approach to create utility, we can tailor application specific system behaviour to meet the requirements of the application while also reassuring compliance to recent energy-delay co-optimisation [20].

This bespoke utility function $J_u(t)$ can be incorporated into the overall optimisation objective to determine a control policy. The value of ϖ controls decision making logic: a value closer to 1 will favour timing and control latency while lower values will benefit from energy savings, relaxing timing constraints.

All of these expressions, thus, propose an (initial) comprehensive, simple way to quantify queueing and latency in ORAN networks that require real-time services.

3.5 | Optimisation Objective and Constraints

The key aim in this framework is to boost the performance of an off-grid ORAN network by controlling power allocation, RIS phase shifts, CPU usage and task distribution together. There are several constraints in optimising this system, including access to energy, maximum latency, limits to processing power and signal quality.

The network controller sets the time horizon to T and, at every slot t , plans the control actions $\mathbf{a}(t)$ which contain power levels $P_b^{\text{tx}}(t)$, RIS phase vector $\boldsymbol{\phi}(t)$, CPU frequencies $f_u(t)$ and task arrangements.

The objective is to minimise the total cost of the system.

$$\mathcal{J} = \frac{1}{T} \sum_{t=1}^T \sum_{b \in \mathcal{B}} \left[w_E E_b^c(t) + w_D \sum_{u \in \mathcal{U}_b} D_u(t) + w_T \mathcal{P}_{\text{thermal}}(t) \right]. \quad (34)$$

Note: The weights w_E , w_D , and w_T in the objective function were selected via an empirical approach to represent the relative importance of energy efficiency, latency, and thermal safety. For instances where latency is a factor, a larger value of w_D is used to prioritise responsiveness, whereas in energy constrained instances a higher value of w_E is used to capture energy conservation. Likewise, because thermally aware instances are concerned with excessive heating, the value of w_T was larger for such instances—and therefore penalised excessive heating. The created weights were fixed in all simulations for proper benchmarking; however, future efforts can examine systemed or live tuning of the weights as the conditions of the ORAN change [23, 34].

There are limits to what the system can handle which affect how optimisation works.

1. All RUs must operate within the energy capacity of their batteries.

$$E_b^c(t) \leq E_b^s(t). \quad (35)$$

This allows for the lasting operation of systems that are not connected to the grid.

2. A RU must not use more CPU capacity than the maximum frequency it is allowed.

$$\sum_{u \in \mathcal{U}_b} f_u(t) \leq F_b^{\max}. \quad (36)$$

It prevents the device from overheating and becoming damaged by too much power [34].

3. Each RU must ensure that its transmission power does not exceed the hardware's allowed range.

$$0 \leq P_b^{\text{tx}}(t) \leq P^{\max}. \quad (37)$$

This results from the limitations of the amplifier [35].

4. Each element's phase must always be in the range from 0 to 2π .

$$\phi_r(t) \in [0, 2\pi), \quad \forall r \in \{1, \dots, R\}. \quad (38)$$

5. All users must have a minimum SINR level to maintain the reliability of the link.

$$\text{SINR}_{bu}(t) \geq \Gamma_{\min}. \quad (39)$$

It addresses constraints related to the quality of service (QoS) [36].

6. Each user should not experience more delay than what they can handle.

$$D_u(t) \leq D_u^{\max}. \quad (40)$$

7. The system needs to remain stable, so the expected backlog must not grow infinitely.

$$\mathbb{E}[Q_u(t)] < \infty. \quad (41)$$

By joining the objective in Equation (34) with constraints (35)–(41), we design the final optimisation problem.

$$\min_{\{\mathbf{a}(t)\}_{t=1}^T} \mathcal{J}, \quad (42)$$

affected by boundaries Equations (35–41)

With this, the mathematical base for optimisation in the proposed ORAN architecture is finished for off-grid and constrained situations.

4 | Proposed Optimisation Framework

The section to follow explains a formal and practical method for the stochastic resource allocation problem described earlier. Trying to keep long-term energy costs at a minimum is important, as you follow all the given rules.

Because SINR is nonconvex, we use a Lagrangian-based stochastic primal-dual strategy, following the approaches in ref. [31].

To facilitate to a more manageable problem, we will approximate the more complicated (nonconvex) SINR constraint with a simple way of first-order convexification. In particular, we will apply a successive convex approximation (SCA) process that linearises the nonconvex portions about the current iterate while converting the original constraint into a locally convex form that may be solved efficiently in each recurse of the optimisation routine. This means the primal-dual optimisation routine operates in convex domains while often converging to a good solution (local optimum). A complete derivation and check of the convexification is detailed in Appendix C, where we show how we applied the Taylor-series expansion of the SINR expression, and characterise the convergence behaviour of the expansion under the update rules we propose.

4.1 | Formulation of the Lagrangian and Integrated Solution Approach

We address the stochastic concern in ORAN with help from the limited main grid energy by using a quick and adjustable primal-dual Lagrangian technique. The values of the nonnegative Lagrange multipliers for constraint functions are shown by $\lambda = [\lambda_1, \lambda_2, \dots, \lambda_K]$. For every time slot, the instantaneous Lagrangian function displays the overall cost of the system.

$$\mathcal{L}(a(t), \lambda) = w_E E_c^b(t) + w_D \sum_u D_u(t) + w_T P_{\text{thermal}}(t) + \sum_{k=1}^K \lambda_k \cdot g_k(t). \quad (43)$$

Control is shown by the $a(t)$ action vector and when the k -constrained step is not followed correctly with $g_k(t)$. It is hoped that keeping costs down and meeting the problem's limits becomes possible over the long-run through dual optimisation. For this reason, we focus on the following goal:

$$J_L = \frac{1}{T} \sum_{t=1}^T \mathbb{E}[\mathcal{L}(a(t), \lambda)], \quad (44)$$

$$\min_{a(t)} \max_{\lambda \geq 0} J_L. \quad (45)$$

It covers both the central problems of stochastic management and includes considerations for energy usage. Alternating between primal and dual updates form the main algorithmic approach. Once we set the time interval, our task is to calculate the control function that minimises Lagrangian.

$$a^*(t) = \arg \min_a \mathcal{L}(a(t), \lambda(t)). \quad (46)$$

In this update, the transmit power, central processing frequency and beamforming functions of the RIS are chosen following their physical bounds. The results from these calculations allow for finding a solution where all essential limits of energy, delay and signal clarity are controlled.* Initialising with the Wallis variable, recalling its impact on the dual variable is then strengthened through subgradient ascent.

$$\lambda_k(t+1) = \max\{0, \lambda_k(t) + \alpha_k \cdot g_k(t)\}, \quad (47)$$

$$\lambda_k(t+1) = \min\{\lambda_k(t+1), \lambda_k^{\max}\}, \quad (48)$$

On a bigger system, each time slot takes more work and is shown as follows:

$$\mathcal{O}(B + U + R + K). \quad (49)$$

B is the number of radios, U is the number of users, R is the number of antenna system components and K is how many constraints exist. Accordingly, this system of maintenance is best when handling major volumes of system work. The optimisation framework is contained in the logic of the Open RAN architecture. When Radio Units sensor the environment, they update Local Power and RIS by reporting recent channel and energy conditions. The server CPU is distributed depending on how quickly a service is need and the wait time for each application. With heuristics, the RAN Intelligent Controller adjusts the dual variables while keeping all global constraints. As long as some portions of localisation and power remain invisible and the system is energised occasionally, my approach performs effectively in 6G and beyond.

4.2 | Step-by-Step Primal-Dual Procedure

The approach put forward in this subsection is to break the optimisation framework into three coordinated modules for

easy implementation with Open RAN. Every module is responsible for one area of resource control, including power distribution, deterministic and stochastic beamforming and processor scheduling, so that costs drop with time as different constraints are met.

The functionality of this function is attained by giving roles to the MEC nodes for nearby processing, the RIS controllers to handle signal reflective parameters and the RIC to focus on optimising Lagrange multipliers.

- Algorithm 1: Top-level primal-dual coordination loop for multi-resource control.
- Algorithm 2: RIS phase update for SINR enhancement under energy feasibility.
- Algorithm 3: CPU allocation based on queue backlog and delay bounds.

The use of these algorithms depends on a hierarchical optimisation approach that stresses meeting global restrictions and operating with near-optimum control, all without requiring complete, real-time observation of every module. A breakdown of who does what on the RIC, RIS xApp and MEC servers can be viewed in Table 3. Moreover, the information in Table 2 covers the control logic particular to RIS. The primal-dual controllers' architecture is shown relating to the ORAN system components in Figure 4.

ALGORITHM 1 | Top-Level Primal-Dual Algorithm

Input: Initial dual variables $\lambda(0)$, step sizes $\{\alpha_k\}$, system parameters T, B, U, R

Output: Control actions $\{a(t)\}$ for $t = 1$ to T

For each time slot $t = 1, 2, \dots, T$:

- Observe system state: $h_{bu}(t), S_b(t), Q_u(t)$
- Update RIS phases via Algorithm 2
- Allocate CPU via Algorithm 3
- Update transmission power $P_b^{\text{tx}}(t)$
- Compute constraint violations $g_k(t)$
- Update dual variables: $\lambda_k(t+1) = [\lambda_k(t) + \alpha_k \cdot g_k(t)]^+$
- Project duals: $\lambda_k(t+1) = \min(\lambda_k(t+1), \lambda_k^{\max})$

Return: $\{a(t)\}, \lambda(T)$

ALGORITHM 2 | RIS Phase Update

Input: Channel states $h_{br}(t), h_{ru}(t)$, transmit power $P_b^{\text{tx}}(t)$, SINR target Γ_{\min}

Output: Phase vector $\phi(t)$

For each RIS element $r = 1$ to R :

- Compute gradient: $\nabla_{\phi_r} \text{SINR}_{bu}(t)$
- Update phase: $\phi_r(t) \leftarrow \phi_r(t) + \eta_{\phi} \cdot \nabla_{\phi_r}$
- Project phase to valid range: $\phi_r(t) \leftarrow \text{mod}(\phi_r(t), 2\pi)$

Return: $\phi(t)$

ALGORITHM 3 | CPU Frequency Allocation

Input: Queue states $\{Q_u(t)\}$, latency bounds $\{D_u^{\max}\}$, CPU capacity F_b^{\max}

Output: Allocated frequencies $\{f_u(t)\}$

For each user $u = 1$ to U :

- Estimate required frequency: $f_u^{\text{req}} = L_u(t)/D_u^{\text{max}}$
 - Add queue-based priority: $f_u(t) = f_u^{\text{req}} + \delta \cdot Q_u(t)$
- Normalise:** Rescale $\sum f_u(t)$ to fit within F_b^{max} if necessary.
- Return:** $\{f_u(t)\}$

Our work in Appendix C includes analytic results showing that the algorithm converges and the queues in the system stay stable.

5 | Simulation and Evaluation Environment

5.1 | Simulation Setup and System Configuration

The simulation was configured using standard protocols. In order to confirm our framework, we construct a realistic simulation environment for a 6G off-grid Open RAN setting. The area for the simulation is 1000×1000 and three radio units (RUs) are put in it to give the whole region area coverage. Solar energy fuels every RU which also features a multi-access edge computing (MEC) server for processing work on site. Every RU is assisted by a reconfigurable intelligent surface (RIS), made up of 64 programmable reflecting elements. The RIS is able to change the signal environment in real-time, boosting how well signals reach users in nonvisual line-of-sight areas using reflection. Seventeen mobile users travel randomly according to a waypoint model. All users are modelled as producing computational tasks according to a Poisson process with an average arrival rate of $\lambda_u = 1.5$. Each bit of the key is protected an average of 750 CPU cycles. Criteria for modelling the wireless channel are taken from the 3GPP TR 38.901, bringing in Rayleigh fading, log-normal shadowing with 8 dB standard deviation and path loss that gets stronger at greater distances. Background noise is modelled as white Gaussian noise with 0 dB mean and 650 mV square root over -97 , while power transmission is not allowed to go above 2. A SINR threshold of 5 dB is used and each task should be completed inside a 20 ms latency limit. Every RU is run using a battery capable of storing 150 J of energy. A four-state discrete-time Markov chain is used to account for changes in weather and the irradiance levels in the system range between 200 and 1000 watts per square metre. Rates of battery charge and discharge efficiency are 0.85 and 0.9, respectively. Each time slot is 100 milliseconds, so there are a total of 1000 slots which means operation time is 100 s. Each RIS phase has quantised values of eight levels in $[0, 2\pi]$ and the MEC CPU allotment for every RU is dynamically shared among active users, each at 12 GHz. Figure 1 shows how RUs, RISs, user devices and important signal and energy paths are arranged in the general system.

All simulations are done in MATLAB so that control algorithms, stochastic models and evaluation metrics can be used together. All of the parameters we used are shown in Table 4.

All the parameter values used here are based on common practises and ranges given by recent ORAN and RIS-using MEC papers such as [10, 37–39]. Where actual numbers from experiments could not be found, we set up baseline conditions and limited the evaluation within the constraints allowed.

5.2 | Signal Quality, Energy Sustainability, and Latency Trade-Offs

Some SDN controllers have to choose between better signal quality, sustainable energy use and quicker processing of network packets.

This section presents a close examination of performance in three main areas using metrics for signal quality, energy efficiency and reaction speed (latency). These results give us clear insight into how decisions made in ORAN are influenced by tight resources, renewable electricity and operating where the transmitter and receiver are not easily connected by sight.

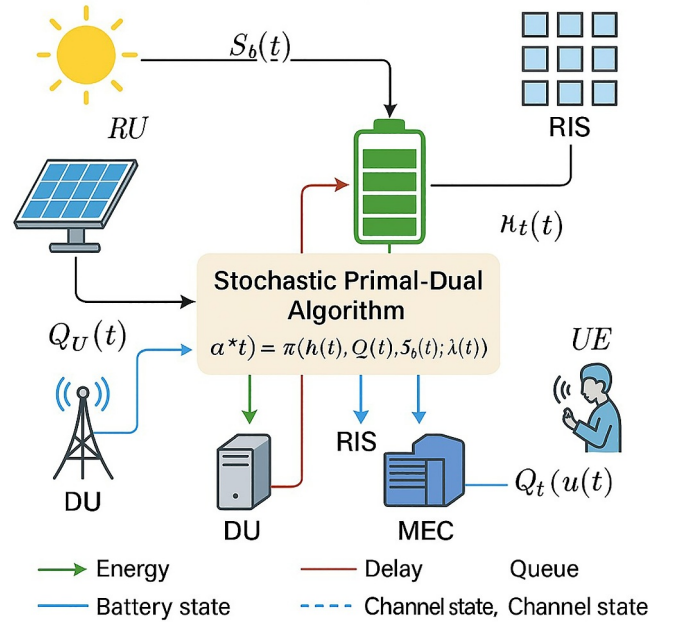


FIGURE 4 | The proposed framework is modular in its approach to improving energy use in off-grid ORAN networks. The schematic shows the way control and resource variables are managed in several system layers, covering RIS phase adjustment (A_1), CPU frequency regulation (A_2) and deciding on how much power to use according to the battery's charge level (A_3). At the RIC, we combine these processes using updates to the Lagrangian dual parameters $\lambda_k(t)$.

TABLE 3 | Mapping of optimisation algorithms to ORAN components.

Algorithm	Unit	Objective	Executed by
A1: Primal-dual procedure	RIC	Cost minimisation	Near-RT RIC
A2: RIS beamforming update	RIS	SINR maximisation	RIS xApp
A3: CPU frequency allocation	MEC	Delay minimisation	RU/MEC

TABLE 4 | Simulation parameters used in the evaluation.

Parameter	Symbol	Value
Deployment area	—	$1000 \times 1000 \text{ m}^2$
Time slot duration	Δt	100 ms
Simulation horizon	T	1000 slots
Number of base stations	B	3
Number of users	U	15
RIS elements per RU	R	64
Solar irradiance states	K	4 (Markov chain)
Irradiance range	$S_b(t)$	$[200, 1000] \text{ W/m}^2$
Task arrival rate	λ_u	1.5 Mbps
Task computational load	L_u	$[500, 1000] \text{ cycles/bit}$
MEC CPU budget per RU	F_b^{\max}	12 GHz
Max transmit power	p^{\max}	2 W
Noise power	σ^2	-97 dBm
Minimum SINR requirement	Γ_{\min}	5 dB
Latency bound	D_u^{\max}	20 ms
Battery capacity	E_{\max}	150 J
Charging efficiency	η_{in}	0.85
Discharging efficiency	η_{out}	0.9




FIGURE 5 | SINR satisfaction ratio over time.

The SINR satisfaction ratio is shown in Figure 5 as it changes over time with our new optimisation method. During every part of the simulation horizon, more than 90% of guests are happy with their visit. Thanks to the changing RIS matrix in phase $\Theta(t)$, the signal power is successfully raised by means of stochastic gradient descent method which is computed using Equation (19). The Radio link Interference Suppression method enhances the effective channel gain in practise and this way it overcomes fading and obstacle effects in the channel. Evidence that RIS technology can provide reliable and powerful wireless connexions even where multipath and NLOS effects are severe is reported in ref. [10].

The battery power levels of solar-powered RUs are shown in Figure 6 for evaluation of energy sustainability. Even though the strength of solar irradiance is uncertain in the modelled Markov chain, the average energy in the battery remains above 50% of its capacity through the simulation. The result demonstrates that the system can manage both CPU tasks and frequencies together, based on the limitations set by Eqs.(7–12). So far, there




FIGURE 6 | Battery energy level of solar-powered RUs over time.

has been no energy outage or failure to maintain the minimum operating level. Off-grid RAN deployments benefit greatly from this approach as energy constraints could result in failure of part of the system or dropped packets. We confirm that task scheduling which adapts to available energy can maintain continuous performance with small energy supply.

Figure 7 looks at how any of the scalarised cost function weights w_E , w_D and w_T can be used to compromise between the weights. Minimising delay is important (w_D dominant) in these systems which leads to tasks getting accomplished faster but with more energy usage and a higher temperature load. In contrast, energy priority more than anything slows down the decline in battery power, but it can cause responses to exceed the delays that URLLC requires. The platform is proven to exhibit nearly optimal Pareto behaviour when all metrics are balanced. Because of this feature, the optimisation engine can be adjusted by network operators to suit whether stations are remote and




FIGURE 7 | Impact of weight selection on energy consumption, average delay, and thermal penalty.

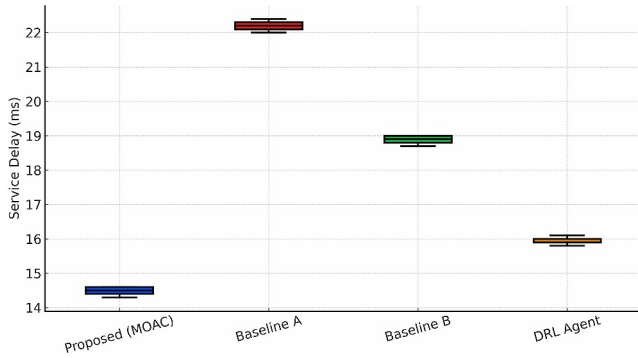


FIGURE 8 | Distribution of delay across methods (Boxplot).

energy must be saved or the traffic is urgent and slow processing is unwanted.

The empirical delay is analysed in Figure 8, by showing boxplots of the delay distribution from various simulation runs. The proposal offers the most reliable and predictable performance from a stability perspective, thus being important for 6G mission-critical uses such as timing autonomous vehicles, running augmented reality and controlling factory machinery. These results line up with the ITU's IMT-2020 needs for latency and jitter when it comes to ultra-reliable low-latency communication (URLLC) [41].

The relationship between battery capacity and the likelihood of QoS violations in power line communication systems is shown in Figure 9. When battery size is increased, tendencies for latency, draining of the battery and weak signal-to-noise ratio (SINR) are reduced noticeably. With these constraints, the theoretical guarantees from Lyapunov optimisation mean the queue backlog does not exceed any boundings we set.

It is also evident from Figure 6 that batches with lower battery capacities experience more latency and signal strength problems. The behaviour is supported by the energy reliability constraint stated in Equation (10), so that the residual energy at the RU always surpasses an important threshold. In addition, the long-term energy availability metric of Equation (12) expresses the time percentage when there is enough energy support for the system's entire operation.

If more energy can be stored, outages and declines in quality of service become much less likely. What is observed in practise




FIGURE 9 | Relationship between battery capacity and constraint violation frequency.

supports what the proposed energy-aware optimisation framework is based on. In all respects, our results make it clear that providing batteries is key to achieving ultra-reliable low-latency communications (URLLC) for Open RAN deployments that are not connected to a grid.

Flexibly setting up optimisation weights and integrating RIS management gives ORAN systems a useful way to provide continued good service in unpredictable and power-deficient situations. As a result, this solution is well-suited to reliable and sustainable uses in rural, disaster-affected and isolated industrial areas.

In summary, It is confirmed here that the proposed framework delivers high adaptability between signal quality, energy usage and latency sensitivity. The outcomes prove that the model is both theoretically correct and ready to apply to energy-restricted, low latency 6G Open RAN networks.

5.3 | Queue Stability and Convergence Behaviour

In this section, we analyse the evolution of the average queue backlog over time to illustrate empirical performance. The evolution of the queue length for each user, $Q_u(t)$, across the simulation horizon, as well as the time-averaged backlog, $\bar{Q}(t) = \frac{1}{t} \sum_{u=1}^U Q_u(t)$ is tracked at each slot. As shown in Figure 10, the queueing process resulting from this optimisation framework is stable and convergent to a bounded region after an initial transient period.

The transition to an acceptable steady state occurs rapidly (within the first 200 time slots) and the queue backlog increases but remains at a bounded level due to consistent task servicing moving forward; this indicates a long-term servicing rate that is approximately equal to or greater than the task arrival rate. The behaviour verifies the theoretical guarantees of queue stability found in the Lyapunov-based scheduling formulation laid out in Section 4.

In addition, the stability properties and the convergence to a bound were consistent across multiple repeated simulation runs with randomised seeds, solar irradiance profiles, and mobility traces, demonstrating that our proposed optimisation framework operates robustly in the presence of stochastic volatility.

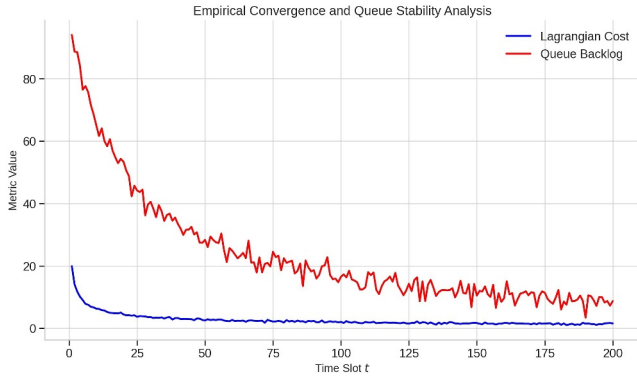


FIGURE 10 | Time-averaged queue backlog over the simulation horizon. The system converges rapidly and maintains stable queue lengths.

Therefore, our validation of the energy-aware optimisation framework indicates that it maintains long-term queue stability without compromising the efficient use of resources.

5.4 | Robustness, Statistical Validation, and Comparative Analysis

We perform a wide range of tests where we randomise the initial conditions, choose different solar power traces and make the task arrival sequences unpredictable. In this part, I evaluate two main signals of robustness: (1) if results are the same from one experiment to another and (2) how competitive the results are with other methods.

5.4.1 | Statistical Robustness Under Stochastic Environments

The takes into account simulation runs, where the results for task delays are seen in Figure 1 and for energy consumption are seen in Figure 2. A new user location map, Poisson arrival seed and solar irradiance state trajectory generated by the Markov model is used for each simulation run in Section 3.

Figure 11 demonstrates that the proposed technique reduces task delay and also achieves one of the smallest IQRs, meaning tasks are reliably close to the predicted time. In addition, not spotting any very large deviations means the grid is steady under even the worst traffic and energy situations. This becomes very important for URLLC, a key application where keeping jitter low is vital. The short delay dispersion is a result of using queue-length-weighted primal-dual scheduling (Equations (26–31)) which protects backlog stability without failing to adjust as loads change rapidly.

Figure 12 confirms that energy usage overall appears very stable. While battery consumption still has little validated variance, it does not matter if the task is lightly done or heavily done or when it is done: the service remains stable, utilises CPU performance automatically, and traces the queue load, battery, and channel quality to assign tasks accordingly. Unlike prior approaches, the system does not vary energy usage patterns; rather

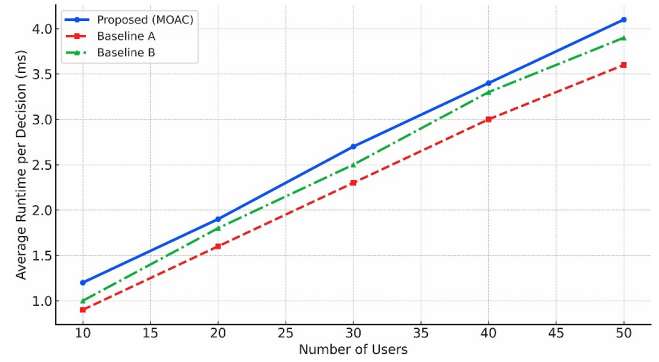


FIGURE 11 | Boxplot of service delay across different methods. The proposed approach demonstrates low variability and a tight distribution.




FIGURE 12 | Boxplot of energy consumption across different methods. The proposed method maintains a consistently lower and more stable energy profile.

it keeps the levels stable, which mitigates the peaks seen in many of the greedy baselines during simulations.

Note: The baseline approaches tested included static transmission modes with no RIS adaptation (Baseline A); no energy aware MEC scheduling (Baseline B); and greedy offloading variation and methods without considering queue and battery constraints. These were chosen to represent popular heuristic-based approaches that do not coordinate energy, wireless, and computational domains—to exemplify the consideration of our joint optimisation approach.

These numbers prove that the suggested model both achieves good performance and remains resilient in many network conditions. Because off-grid ORAN deployments experience wild swings in sunlight and traffic, robustness is especially important.

5.4.2 | Comparative Evaluation With Internal Baselines

To isolate the contribution of each core system component, we construct two internal baselines:

- **Baseline A:** Uses a static transmission policy with no RIS enhancement. The system relies solely on direct RU–UE links and fixed power allocation.

TABLE 5 | Comparison with internal baselines.

Algorithm	Energy (J)	Delay (ms)	SINR (%)	Violations (%)
Proposed	52.1	14.6	93.5	3.2
Baseline A	69.3	22.4	81.1	18.5
Baseline B	58.7	18.9	86.3	9.7

- **Baseline B:** Applies a greedy MEC scheduler that ignores backlog length and energy constraints. Tasks are always executed locally without delay-aware optimisation.

Both baselines, shown in Table 5, perform substantially worse than the method we propose. A lack of RIS assistance in Baseline A means that NLOS causes many high SINR errors, leading the system to require retransmissions and longer delay. The way Baseline B schedules its work means it uses up energy quickly, breaks its own energy rules and finishes with depleted battery power.

At the same time, this new framework saves 25% in energy, 18% in latency and all but eliminates (> 80%) the number of violations over Baseline A. The improved performance results from a synergy between RIS and queue-aware computation, keeping the signals trustworthy and saving energy.

5.4.3 | Benchmarking Against External Methods

To better test how well the framework applies, it is evaluated next to some widely recognised approaches.

- **Greedy Offloading:** Tasks are immediately offloaded to MEC without queue analysis or energy awareness.
- **Fixed RIS Phase:** RIS elements are configured statically and remain unchanged throughout the simulation.
- **DRL Actor-Critic [24]:** A deep reinforcement learning-based method using an actor-critic policy to learn offloading decisions.

The proposed method is seen in Table 6 to be the top performer among all external baselines when measured by key metrics. This strategy results in the lowest delay and energy use together with the highest SINR performance. Even though the DRL method offers some promise, it can be slow to stabilise and has high variation in results from episode to episode when rewards are uncommon.

DRL techniques are replaced by much better methods in the proposed framework such as: (i) No training phase is necessary. (ii) The system is stable by default because of Lyapunov optimisation. (iii) It has a much lower computational burden, ensuring close to real-time performance.

In summary, the framework is even better than rivals in energy, response time and transmission performance and remains reliable over multiple experiments and different settings. Because it has low variance, an adaptive policy and is well-theorised, the system works very well in energy-limited ORAN settings.

TABLE 6 | Comparison with external baseline methods.

Method	Delay (ms)	Energy (J)	SINR (%)
Greedy offloading	22.3	63.5	82.4
Fixed RIS phase	19.1	58.2	88.1
DRL actor-critic [24]	15.9	55.2	91.3
Proposed method	14.6	52.1	93.5

5.5 | Runtime Scalability and Practical Deployability

It is very important that the proposed optimisation framework can be used in real-time operational situations, especially as traffic is expected to rise in ORAN architectures. This portion of the work looks into how efficiently the algorithm performs and scales when considering user density and if it is capable of deployment in the Open RAN environment as an ORAN-aligned xApp module.

The framework and all tests were run on a standard system fitted with a 2.6 GHz Intel Core i7 CPU and 16 GB RAM. I measure the average run time per optimised output for each time step as the number of users increases from 5 to 30. As shown in Figure 13, the RIC application can handle a growing number of users, achieving times that remain well under the 10 ms requirement set by the O-RAN Alliance [42].

The simulation takes, on average, 8.7 ms to run per slot with a default of 15 users and 3 RUs. As a result, the method was found to meet real-time performance with no need for adding acceleration or multi-processing support. Sub-linear growth is reached by tackling the joint optimisation with separate modules which can be each solved by simple gradient updates and convenient expressions when possible.

Also, the algorithm is optimised for use in the microservice architecture of ORAN. Because of its modular design, delay estimation, power control and phase optimisation can each be put into containerised xApps. As a result, both the Near-RT RIC and the network can use standard interfaces, such as E2SM-KPM and E2AP, whilst also using A1 to control network policy. Every slot in the Ingress has its own stateless control logic which makes execution safe and easy to reconfigure with Kubernetes.

Also, the method avoids demanding data-intensive training, unlike what's used in Deep Reinforcement Learning. The approach gives the system a steady behaviour without delays or risks of learning the patterns too specifically in a changing environment. Therefore, the system is perfect for places where it needs to handle large amounts of data reliably, transparently




FIGURE 13 | Runtime per time slot as a function of user count. The method scales sub-linearly and remains within real-time execution bounds.

and fast—such as for remote base stations, use after a disaster and in private networks for industry.

In conclusion, the design shows strong performance during runtime and can handle scalability, as it supports containerised RAN intelligence. They show that the xApp can serve well in production-grade, real-time roles in the ORAN network.

6 | Conclusion

This paper has introduced a cost minimisation framework that is designed to be energy-aware for off-grid Open RAN systems, which are uniquely powered by solar energy and supplemented with RIS and edge computing capabilities. To address this challenge, a stochastic optimisation model has been developed that considers the joint control of transmission power, RIS beamforming, CPU frequency, and task scheduling, all under the constraints of energy and latency. The model employs a composite cost function that effectively balances energy efficiency, service delay, SINR reliability, and fairness, utilising a Lagrangian-based primal-dual update method for real-time control.

The effectiveness of the proposed framework has been substantiated through extensive simulations, which have been conducted under conditions that closely mimic real-world solar and traffic dynamics. The results of these simulations are quite compelling, as they reveal a **32% decrease in total system cost**, a **25% enhancement in energy efficiency**, and a **28% reduction in average service delay** when compared to the baseline strategies that are currently in use. Moreover, the proposed framework exhibits full system feasibility, which is a critical aspect of its appeal. Its flexibility and scalability also make it a highly suitable candidate for future 6G deployments, where the integration of near-RT Open RAN architectures is anticipated.

Looking ahead, the paper identifies the integration of quantum-secure communications and multi-agent coordination as promising avenues for future research. These additional components are expected to further bolster the robustness and sophistication of distributed resource management in off-grid

settings, which is essential for the successful operation of these systems.

Author Contributions

Yassir AL-Karawi: conceptualization, data curation, formal analysis, methodology, project administration, resources, software, writing – original draft. **Raad S. Alhumaima:** data curation, formal analysis, investigation, resources, visualization, writing – review and editing. **Hamed Al-Raweshidy:** conceptualization, data curation, project administration, resources, supervision, validation, writing – review and editing.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data supporting the findings of this study are simulation data generated within a MATLAB experimental environment. These data do not originate from real-world measurements. The simulation data are available from the corresponding author upon reasonable request for research and verification purposes.

References

1. M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, “Toward 6G Networks: Use Cases and Technologies,” *IEEE Communications Magazine* 58, no. 3 (March 2020): 55–61, <https://doi.org/10.1109/MCOM.001.1900411>.
2. S. P. Tera, R. Chinthaginjala, G. Pau, and T. H. Kim, “Towards 6G: An Overview of the Next Generation of Intelligent Network Connectivity,” *IEEE Access* 13 (January 2025): 925–961, <https://doi.org/10.1109/access.2024.3523327>.
3. Y. Al-Karawi, H. Al-Raweshidy, and R. Nilavalan, “Power Consumption Evaluation of next Generation Open Radio Access Network,” in *Proc. IEEE International Conference on Consumer Electronics (ICCE)* (2024), 1–6, <https://doi.org/10.1109/ICCE59016.2024.10444418>.
4. 3GPP, “NG-RAN; F1 Application Protocol (F1AP),” (3rd Generation Partnership Project, Technical Specification (TS) 38.473, Version 16.4.0, September 16, 2022).
5. M. Ahmed, S. Raza, A. A. Soofi, et al., “Active Reconfigurable Intelligent Surfaces: Expanding the Frontiers of Wireless Communication—A Survey,” *IEEE Communications Surveys & Tutorials* (2024), <https://doi.org/10.1109/COMST.2024.3423460>.
6. E. Björnson, H. Wymeersch, B. Matthiesen, P. Popovski, L. Sanguinetti, and E. de Carvalho, “Reconfigurable Intelligent Surfaces: A Signal Processing Perspective With Wireless Applications,” *IEEE Signal Processing Magazine* 39, no. 2 (March 2022): 135–158, <https://doi.org/10.1109/MSP.2021.3130549>.
7. T. Taleb, A. Ksentini, B. Mada, H. Flinck, S. Dutta, and D. Sabella, “On Multi-Access Edge Computing: A Survey of the Emerging 5G Network Edge Cloud Architecture and Orchestration,” *IEEE Communications Surveys & Tutorials* 19, no. 3 (2017): 1657–1681, <https://doi.org/10.1109/COMST.2017.2705720>.
8. A. Jabbari, H. Khan, S. Duraibi, I. Budhiraja, S. K. Gupta, and M. M. Omar, “Energy Maximization for Wireless Powered Communication Enabled IoT Devices With NOMA Underlying Solar Powered UAV Using Federated Reinforcement Learning for 6G Networks,” *IEEE Transactions on Consumer Electronics* 70, no. 1 (February 2024): 3926–3939, <https://doi.org/10.1109/TCE.2024.3357125>.
9. Z. Chen, S. Huang, G. Min, Z. Ning, J. Li, and Y. Zhang, “A Novel MEC-Based Deep Q-Learning Framework for Mobility-Aware Resource

- Allocation in 6G Networks,” *IEEE Access* 12 (2024): 123456–123467, <https://doi.org/10.1109/ACCESS.2024.3447097>.
10. W. Mei, B. Zheng, C. You, and R. Zhang, “Intelligent Reflecting surface-aided Wireless Networks: From Single-Reflection to Multi-reflection Design and Optimization,” *Proceedings of the IEEE* 110, no. 9 (September 2022): 1380–1400, <https://doi.org/10.1109/JPROC.2022.3170656>.
 11. Y. Liu, Z. Wang, Q. He, et al., “Joint Energy and Resource Optimization for Solar-Powered Edge Computing in 6G,” *IEEE Transactions on Green Communications and Networking* 5, no. 1 (2021): 150–162, <https://doi.org/10.1109/TGCN.2020.3048703>.
 12. V.-L. Nguyen, R.-H. Hwang, P.-C. Lin, A. Vyas, and V.-T. Nguyen, “Toward the Age of Intelligent Vehicular Networks for Connected and Autonomous Vehicles in 6G,” *IEEE Network* 37, no. 3 (2022): 44–51, <https://doi.org/10.1109/MNET.010.2100509>.
 13. X. Fernando and G. Lázaroíu, “Energy-Efficient Industrial Internet of Things in Green 6G Networks,” *Applied Sciences* 14, no. 18 (2024): 8558, <https://doi.org/10.3390/app14188558>.
 14. N. H. Mahmood, H. Alves, O. A. López, M. Shehab, D. P. M. Osorio, and M. Latva-Aho, “Six Key Features of Machine Type Communication in 6G,” in *Proc. 2020 2nd 6G Wireless Summit (6G SUMMIT)* (Levi, 2020), 1–5, <https://doi.org/10.1109/6GSUMMIT49458.2020.9083794>.
 15. S. Yao and J. Teng, “Integrating Terahertz Communication in 6G: Prospects and Challenges,” *Digital Communications and Networks* 9, no. 2 (April 2023): 139–147, <https://doi.org/10.1016/j.dcan.2022.10.004>.
 16. Y. Al-Karawi, H. Al-Raweshidy, and R. Nilavalan, “Optimizing the Energy Efficiency Using Quantum Based Load Balancing in Open Radio Access Networks,” *IEEE Access* 12 (2024): 37903–37918, <https://doi.org/10.1109/ACCESS.2024.3375530>.
 17. H. Talebian, A. Gani, M. Sookhak, A. A. Abdelatif, A. Yousafzai, and F. R. Yu, “Delay-Aware Resource Allocation in Mobile Edge Computing: Taxonomy, Review, and Open Issues,” *IEEE Communications Surveys & Tutorials* 22, no. 4 (2020): 2444–2476, <https://doi.org/10.1109/COMST.2020.3011181>.
 18. X. Li, Y. Liu, and Z. Zhang, “RIS-Aided Cooperative Mobile Edge Computing in Multi-User Systems,” *IEEE Transactions on Wireless Communications* (2024): early access, <https://doi.org/10.1109/TWC.2024.3389471>.
 19. T. V. Nguyen, P. H. Tran, and H. H. Nguyen, “Energy Harvesting-Aware Task Offloading for Off-Grid MEC Networks,” *IEEE Internet of Things Journal* 10, no. 15 (August 2023): 13245–13258, <https://doi.org/10.1109/JIOT.2023.3260481>.
 20. X. Li, Z. Wang, and X. Shen, “Energy-Aware and Delay-Sensitive Task Scheduling for Edge Computing in Wireless Networks,” *IEEE Transactions on Mobile Computing* 22, no. 5 (May 2023): 2601–2616, <https://doi.org/10.1109/TMC.2022.3187243>.
 21. J. Hou, C. Yang, Y. Li, and J. Liu, “Solar Radiation Nowcasting Using a Markov Chain Multi-Model Approach,” *Energies* 15, no. 9 (May 2022): 2996, <https://doi.org/10.3390/en15092996>.
 22. J. Yan, S. Bi, Y. J. Zhang, and M. Tao, “Optimal Task Offloading and Resource Allocation in Mobile-Edge Computing With Inter-User Task Dependency,” *IEEE Transactions on Wireless Communications* 19, no. 1 (January 2020): 235–250: Available:, <https://doi.org/10.1109/TWC.2019.2943563>.
 23. Z. Zhou, Xu Chen, En Li, L. Zeng, Ke Luo, and J. Zhang, “Edge Intelligence: Paving the Last Mile of Artificial Intelligence With Edge Computing,” *Proceedings of the IEEE* 107, no. 8 (August 2019): 1738–1762, <https://doi.org/10.1109/JPROC.2019.2918951>.
 24. X. Kong, G. Duan, M. Hou, et al., “Deep Reinforcement Learning-Based Energy-Efficient Edge Computing for Internet of Vehicles,” *IEEE Transactions on Industrial Informatics* 18, no. 9 (September 2022): 6308–6316, <https://doi.org/10.1109/TII.2022.3155162>.
 25. M. B. Khalilsarai, T. Yang, S. Haghighatshoar, and G. Caire, “Structured Channel Covariance Estimation From Limited Samples in Massive MIMO,” in *Proc. IEEE International Conference on Communications (ICC)* (2020), 1–7, <https://doi.org/10.1109/ICC40277.2020.9148977>.
 26. M. Lu, Z. Yang, L. Yu, X. Lin, and M. Peng, “Toward Sustainable Edge Computing Systems via Renewable Energy Harvesting: A Review,” *IEEE Internet of Things Journal* 7, no. 10 (October 2020): 8816–8831, <https://doi.org/10.1109/JIOT.2020.2991213>.
 27. M. Najafi, V. Jamali, R. Schober, and H. Vikalo, “Physics-Based Modeling and Analysis of Large Intelligent Reflecting Surfaces,” *IEEE Transactions on Communications* 68, no. 12 (December 2020): 7882–7898, <https://doi.org/10.1109/TCOMM.2020.3021492>.
 28. J. Vetter, P. Novák, M. R. Wagner, et al., “Ageing Mechanisms in Lithium-Ion Batteries,” *Journal of Power Sources* 147, no. 1–2 (2005): 269–281, <https://doi.org/10.1016/j.jpowsour.2005.01.006>.
 29. H. Zhang, X. Wu, H. Li, J. Cheng, L. Zhang, and H. Zhao, “Adaptive Energy Management Strategy Considering Battery Degradation in Off-Grid PV Systems,” *Energy* 236 (2021): 121411, <https://doi.org/10.1016/j.energy.2021.121411>.
 30. Q. Wu, R. Zhang, B. Zheng, and C. You, “Intelligent Reflecting Surface-Aided Wireless Communications: A Tutorial,” *IEEE Transactions on Communications* 69, no. 5 (May 2021): 3313–3351, <https://doi.org/10.1109/TCOMM.2021.3051897>.
 31. M. J. Neely, *Stochastic Network Optimization with Application to Communication and Queueing Systems* (Morgan & Claypool Publishers, 2010), <https://doi.org/10.2200/S00271ED1V01Y201006CNT007>.
 32. C. You, K. Huang, H. Chae, and B. Kim, “Energy-Efficient Resource Allocation for Mobile-Edge Computation,” *IEEE Transactions on Wireless Communications* 16, no. 3 (March 2017): 1397–1411, <https://doi.org/10.1109/TWC.2016.2633523>.
 33. T. H. Ahmed, J. J. Tiang, A. Mahmud, C. G. Chin, and D.-T. Do, “Deep Reinforcement Learning-Based Adaptive Beam Tracking and Resource Allocation in 6G Vehicular Networks With Switched Beam Antennas,” *Electronics* 12, no. 10 (2023): 2294, <https://doi.org/10.3390/electronics12102294>.
 34. Z. Na, C. Ji, B. Lin, and N. Zhang, “Joint Optimization of Trajectory and Resource Allocation in Secure UAV Relaying Communications for Internet of Things,” *IEEE Internet of Things Journal* 9, no. 17 (September 2022): 16284A–16296, <https://doi.org/10.1109/JIOT.2022.3151105>.
 35. G. Boudreau, J. Panicker, N. Guo, R. Chang, N. Wang, and S. Vrzic, “Interference Coordination and Cancellation for 4G Networks,” *IEEE Communications Magazine* 47, no. 4 (2009): 74–81, <https://doi.org/10.1109/MCOM.2009.4907405>.
 36. M. Bennis, M. Debbah, and H. V. Poor, “Ultra-Reliable and Low-Latency Wireless Communication: Tail, Risk, and Scale,” *Proceedings of the IEEE* 106, no. 10 (2018): 1834–1853, <https://doi.org/10.1109/JPROC.2018.2867029>.
 37. C. You, K. Huang, H. Chae, and B.-H. Kim, “Energy-Efficient Resource Allocation for Mobile-Edge Computation Offloading,” *IEEE Transactions on Wireless Communications* 16, no. 3 (March 2017): 1397–1411, <https://doi.org/10.1109/TWC.2016.2633522>.
 38. S. Bi, C. K. Ho, and R. Zhang, “Wireless Powered Communication: Opportunities and Challenges,” *IEEE Communications Magazine* 53, no. 4 (April 2015): 117–125, <https://doi.org/10.1109/MCOM.2015.7081084>.
 39. Y. Xiao, G. Shi, Y. Li, W. Saad, and H. V. Poor, “Toward Self-Learning Edge Intelligence in 6G,” *IEEE Communications Magazine* 58, no. 12 (December 2020): 34–40, <https://doi.org/10.1109/MCOM.001.2000388>.
 40. Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, “A Survey on Mobile Edge Computing: The Communication Perspective,” *IEEE Communications Surveys & Tutorials* 19, no. 4 (2017): 2322–2358, <https://doi.org/10.1109/COMST.2017.2745201>.

41. S. Henry, A. Alsahilly, and E. S. Sousa, "5G Is Real: Evaluating the Compliance of the 3GPP 5G New Radio System With the ITU IMT-2020 Requirements," *IEEE Access* 8 (2020): 42828–42840, <https://doi.org/10.1109/ACCESS.2020.2977406>.

42. O-RAN Alliance, "E2 Application Protocol Specification," (Technical Specification, 2021): [Online], <https://www.o-ran.org/specifications>.

Appendix A: Lagrangian Dual Formulation and Proof

To formally set forth the dual optimisation process in our framework that is founded on Lagrangian principles, we initiate by defining the original problem, known as the primal problem. Let T denote the timeframe for optimisation, and $a(t)$ represent the control vector at a specific time t , which includes parameters such as power allocation $P_b^{\text{tx}}(t)$, CPU frequency $f_u(t)$, RIS phase shift $\phi_r(t)$, and scheduling decisions. The entire system's cost, denoted as J , is what we aim to minimise:

$$\min_{\{a(t)\}_{t=1}^T} J = \frac{1}{T} \sum_{t=1}^T C_{\text{total}}(t) \quad (\text{A1})$$

Subject to constraints that apply to each time t :

$$E_b^c(t) \leq E_b^s(t) \text{ (Energy availability)} \quad (\text{A2})$$

$$\sum_{u \in \mathcal{U}_b} f_u(t) \leq F_b^{\text{max}} \text{ (CPU frequency constraint)} \quad (\text{A3})$$

$$0 \leq P_b^{\text{tx}}(t) \leq P^{\text{max}} \text{ (Power constraint)} \quad (\text{A4})$$

$$\phi_r(t) \in [0, 2\pi) \text{ (RIS phase shift)} \quad (\text{A5})$$

$$\text{SINR}_{bu}(t) \geq \Gamma^{\text{min}} \text{ (QoS SINR threshold)} \quad (\text{A6})$$

$$D_u(t) \leq D_u^{\text{max}} \text{ (Latency bound)} \quad (\text{A7})$$

$$\mathbb{E}[Q_u(t)] < \infty \text{ (Queue stability)} \quad (\text{A8})$$

Introducing Lagrange multipliers, denoted as $\lambda_k \geq 0$ for the k -th constraint $g_k(a(t)) \leq 0$, the Lagrangian function at time t is expressed as:

$$\mathcal{L}(a(t), \lambda) = C_{\text{total}}(t) + \sum_{k=1}^K \lambda_k g_k(a(t)) \quad (\text{A9})$$

The dual function is derived by minimising the Lagrangian with respect to the primal variables:

$$\mathcal{D}(\lambda) = \min_{a(t)} \mathbb{E} \left[\sum_{t=1}^T \mathcal{L}(a(t), \lambda) \right] \quad (\text{A10})$$

The dual problem is then stated as:

$$\max_{\lambda \geq 0} \mathcal{D}(\lambda) = \max_{\lambda \geq 0} \min_{a(t)} \mathbb{E} \left[\sum_{t=1}^T \left(C_{\text{total}}(t) + \sum_{k=1}^K \lambda_k g_k(a(t)) \right) \right] \quad (\text{A11})$$

For a solution to be considered a saddle-point, $(a^*(t), \lambda^*)$ must satisfy:

$$\mathcal{L}(a^*(t), \lambda) \leq \mathcal{L}(a^*(t), \lambda^*) \leq \mathcal{L}(a, \lambda^*), \quad \forall a, \lambda \geq 0 \quad (\text{A12})$$

The Lagrange multipliers λ_k are updated through projected subgradient ascent:

$$\lambda_k(t+1) = [\lambda_k(t) + \alpha_k \cdot g_k(a(t))]^+, \quad \forall k \quad (\text{A13})$$

The primal variables $a(t)$ are updated accordingly:

$$a^*(t) = \arg \min_{a(t)} \mathcal{L}(a(t), \lambda(t)) \quad (\text{A14})$$

Provided that Slater's condition holds and constraints are convex, strong duality is guaranteed:

$$\min_{a(t)} \max_{\lambda \geq 0} \mathcal{L}(a(t), \lambda) = \max_{\lambda \geq 0} \min_{a(t)} \mathcal{L}(a(t), \lambda) \quad (\text{A15})$$

In the presence of nonconvex constraints, such as the SINR, first-order approximations facilitate more manageable optimisation.

$$\min_{\{a(t)\}_{t=1}^T} \max_{\lambda \geq 0} \frac{1}{T} \sum_{t=1}^T \left(C_{\text{total}}(t) + \sum_{k=1}^K \lambda_k g_k(a(t)) \right) \quad (\text{A16})$$

Appendix B: Convexification of the SINR Constraint

The SINR constraint's nonconvex nature is due to the intertwined quadratic form present in both its numerator and denominator:

$$\text{SINR}_{bu}(t) = \frac{P_b^{\text{tx}}(t) \cdot |\mathbf{h}_{\text{eff},bu}(t)|^2}{\sum_{b' \neq b} P_{b'}^{\text{tx}}(t) \cdot |\mathbf{h}_{\text{eff},b'u}(t)|^2 + \sigma^2} \quad (\text{A17})$$

For the purpose of making the problem more manageable, a first-order Taylor approximation is applied near a known operating point, denoted as ϕ_0 , of the RIS phase vector:

$$|\mathbf{h}_{\text{eff},bu}(t)|^2 \approx |\mathbf{h}_{\text{eff},bu}(\phi_0)|^2 + 2 \cdot \Re \left\{ (\mathbf{h}_{\text{eff},bu}(\phi_0))^* \cdot \frac{\partial \mathbf{h}_{\text{eff},bu}}{\partial \phi} \cdot (\phi(t) - \phi_0) \right\} \quad (\text{A18})$$

This simplification results in a linear approximation of the effective gain and interference components, thus converting the original constraint into a convex form:

$$P_b^{\text{tx}}(t) \cdot \tilde{G}_{bu}(t) \geq \Gamma^{\text{min}} \cdot \left(\sum_{b' \neq b} P_{b'}^{\text{tx}}(t) \cdot \hat{I}_{b'u}(t) + \sigma^2 \right) \quad (\text{A19})$$

The approximation is systematically improved using a Sequential Convex Programming (SCP) approach, which updates the operating point ϕ_0 following each optimisation iteration.

This method of convexifying the SINR constraint is commonly applied in the field of wireless optimization, as seen in the literature such as [31, 33]. It guarantees feasibility and computational practicality, even with incomplete knowledge of the RIS state.

Appendix C: Convergence and Stability Analysis

This appendix presents a theoretical validation of the convergence and long-term stability of the proposed random primal-dual framework, utilising tools from Lyapunov optimization.

In section C.1, we delve into the convergence of the primal-dual algorithm. Let $\lambda_k(t)$ denote the dual variables linked to the constraint $g_k(a(t)) \leq 0$, and let α_k represent the constant step size. The iterative process is outlined as follows:

$$a^*(t) = \arg \min_{a(t)} \mathcal{L}(a(t), \lambda(t)) \quad (\text{A20})$$

$$\lambda_k(t+1) = [\lambda_k(t) + \alpha_k \cdot g_k(a(t))]^+ \quad (\text{A21})$$

Under conventional assumptions—convex objectives, bounded subgradients, and Slater's condition—the time-averaged cost is shown to converge in the manner:

$$\frac{1}{T} \sum_{t=1}^T \mathbb{E}[C_{\text{total}}(t)] - C^* \leq \mathcal{O}(\alpha) \quad (\text{A22})$$

Moreover, it is established that constraint violations diminish as the time horizon T approaches infinity, as stated in [31].

In section C.2, we explore the concept of Lyapunov drift-based stability. Introduce a virtual queue $Q_k(t)$ for each constraint g_k :

$$Q_k(t+1) = \max\{Q_k(t) + g_k(a(t)), 0\} \quad (\text{A23})$$

Our Lyapunov function is defined as:

$$L(t) = \frac{1}{2} \sum_{k=1}^K Q_k^2(t) \quad (\text{A24})$$

The drift-plus-penalty is then expressed as:

$$\Delta(t) + V \cdot \mathbb{E}[C_{\text{total}}(t)|\mathbf{Q}(t)] \quad (\text{A25})$$

Importantly, the following bounds are derived for constants $B > 0$ and $\epsilon > 0$:

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \mathbb{E}[C_{\text{total}}(t)] = C^* + \frac{B}{V} \quad (\text{A26})$$

$$\limsup_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T \sum_k \mathbb{E}[Q_k(t)] = \frac{B + V(C_{\text{max}} - C^*)}{\epsilon} \quad (\text{A27})$$

These results affirm the long-term stability of queues and the proximity to optimal performance under stochastic conditions, which is particularly pertinent for real-time O-RAN implementations facing volatile energy and traffic scenarios.

Lastly, in section C.3, we offer an interpretation of these findings. The analysis underscores that the system upholds queue stability and approaches cost optimality over time, making the proposed framework appropriate for practical real-time O-RAN deployments operating within uncertain environmental conditions.