

Constrained Discrete Bayesian Optimization for Base Station Sleep Mode in Cellular Networks

Ryad Madi, Salah-Eddine Elayoubi, and Michel Kieffer
L2S, CNRS

Centralesupélec, Université Paris-Saclay

Gif-sur-Yvette, France

{ryad.madi, salah-eddine.elayoubi, michel.kieffer}@centralesupelec.fr

Abstract—Sleep mode strategies, which deactivate under-utilized base stations (BSs), are a promising solution for reducing the energy consumption of cellular networks. However, selecting which BSs to deactivate is a challenging combinatorial optimization problem under Quality of Service (QoS) constraints, such as cell-edge throughput and coverage. Moreover, evaluating the impact of a sleep mode decision on the QoS is costly, since it requires either digital twin-like system-level simulation or testing on the live network. In this paper, we propose a discrete constrained Bayesian Optimization (BO) framework for BS deactivation. We use Random Forest (RF) surrogates for throughput and coverage to build a low-complexity acquisition function (AF) that combines active-BS reduction with joint constraint satisfaction. Numerical results show that the proposed method outperforms greedy heuristics and random search, and achieves competitive energy-QoS trade-offs with respect to a constrained discrete BO baseline, while reducing the computational complexity of the BO decision step.

Index Terms—Constrained Bayesian Optimization, Mobile Network Optimization, Energy Saving.

I. INTRODUCTION

Base Stations (BSs) account for a major share of the energy consumption of cellular networks, yet their power usage does not scale with traffic load variations [1]. A significant portion of this consumption is load-independent, due to components such as cooling systems that remain active even under low traffic conditions. This mismatch between traffic demand and energy consumption motivates the development of adaptive *sleep modes*, where BSs or some of their components are switched off during periods of low traffic demand [2]. This mechanism has been extensively studied as a key enabler for improving the energy efficiency of cellular networks [3], [4].

From an operational perspective, sleep mode management consists of dynamically adapting the network configuration by selectively deactivating BSs while preserving Quality of Service (QoS), such as coverage and throughput. If the objective is to derive a **global, network-level optimal sleep mode**, the classical distributed approaches that rely on local decision-making or neighbor-to-neighbor coordination, *e.g.*, [3], [4] are to be replaced by a centralized global approach, operating on a large time scale (*e.g.*, minutes), coupled with

local decisions that adapt the capacity to the traffic on a smaller time scale. However, this global approach leads to a combinatorial optimization problem over binary decision variables, and evaluating a candidate configuration is heavy as it involves testing on a live network with a risk of performance degradation, or costly digital-twin simulation.

In this paper, we propose a Bayesian Optimization (BO) approach for discrete and constrained BS deactivation. The method relies on a structural property of the power model: because the load-independent term dominates per-BS consumption, the energy objective is essentially monotone in the number of active BSs and thus easily evaluated. We exploit this to depart from standard constrained BO: instead of fitting an objective surrogate at each iteration, we treat the objective as given and learn only the feasibility of the constraints. This lowers the per-iteration complexity of the acquisition step.

Contributions. Our main contribution is a novel constrained discrete BO formulation for BS sleep-mode selection, featuring i) A low-complexity acquisition function (AF) combining objective improvement with a learned QoS feasibility probability from Random Forest (RF) throughput and coverage surrogates, and ii) an empirical and complexity comparison with random search, greedy heuristics, and constrained discrete BO, showing comparable or better energy-QoS trade-offs at equal evaluation budget and lower acquisition-step complexity.

The remainder of the paper is organized as follows. Section II reviews related work. Section III formulates BS sleep mode as a discrete constrained BO problem. Section IV describes the evaluation setup and simulation results. Finally, Section V concludes the paper.

II. RELATED WORK

We review approaches to BS sleep mode management and recent advances in BO for constrained and discrete problems.

A. Sleep Mode Strategies in Cellular Networks

A wide range of sleep mode enabling strategies have been proposed in the literature, aiming to reduce energy consumption while maintaining acceptable QoS. These approaches differ not only in their modeling assumptions, but also in the level at which decisions are made, ranging from local, distributed policies to more coordinated or network-level strategies. As local solutions struggle to ensure global coverage, we focus

This work has been partially supported by the Sustainable 6G chair, funded by Orange and held by CentraleSupélec. It was also supported by the French government, in the framework of the France 2030 program (INTENTION-6G project).

on global-level and network coordinated strategies. These coordinated approaches are generally categorized as follows.

Structural or random decisions. Specific BSs are deactivated based on spatial or deployment considerations. For instance, cluster-based approaches switch off central cells assuming that neighboring cells can compensate for coverage [4]. In [5], a sleep mode candidate list is built and selects at random one of its elements for sleep mode. While simple to implement, these methods may lead to suboptimal decisions.

Cell Breathing and Dynamic Clustering. An important set of works adaptively adjusts the physical coverage area of active BSs to fill "coverage holes" left by sleeping sites [3]. Furthermore, the remaining BSs may make use of Coordinated Multi-Point (CoMP) transmission to cooperatively serve users in coverage holes [4]. Advanced algorithms, such as Sleep Mode with Dynamic Clustering (SMDC), maximize energy savings by selecting an optimized set of points for coordination to ensure that savings are not offset by increased transmission power in neighboring cells [3], [4]. These methods improve performance but they increase system complexity and rely on strong coordination or modeling assumptions.

Heterogeneous Network (HetNet) Coordination. In multi-tier deployments, macro-cells provide umbrella coverage, allowing small cells (pico, femto) to remain in deep sleep until traffic mandates activation [3], [4]. Coordination between tiers significantly improves coverage probability and energy efficiency compared to random sleeping policies [6]. However, multi-tier macro/small cell deployments are still restricted to highly dense networks, limiting their practical gains.

O-RAN-based sleep mode. Recent architectures enable finer-grained energy optimization by applying sleep decisions at the level of Radio Units (RU) and Distributed Units (DU), rather than entire BSs. This is made possible by the disaggregation of the RAN and the introduction of programmable controllers such as the RAN Intelligent Controller (RIC). Learning-based approaches have been proposed to dynamically control RU activation based on network conditions [7], [8]. While O-RAN architecture increase flexibility by allowing a decision at the DU level, controlling multiple RUs, the sleep mode decision algorithm remains an open issue.

In summary, despite their diversity, existing approaches suffer from inherent drawbacks. Most of them rely on complex coordination between BSs, inducing overhead without convergence guarantees. Centralized approaches that rely on Key Performance Indicator (KPI) collection have the advantage of providing a global view, but they need to accurately anticipate the impact of the sleep mode decisions. Network system level simulations are used in [5], while a recent trend is to mobilize network digital twins for a more accurate and dynamic evaluation [9]. Digital twins, however, are not always available and require regular synchronization from live networks to remain accurate. Relying on feedback from the network is thus inevitable for reaching the optimal setting, but operators cannot afford many test iterations on their live network.

In contrast to classical approaches, we adopt a centralized approach that relies on BO to minimize the number of

trials. Our method can be regarded as an experiment design approach, relying on surrogate models continuously enhanced by carefully selected additional measurements.

B. BO for Constrained and Discrete Problems

BO is an effective framework for optimizing expensive black-box functions [10]. It relies on a surrogate model, typically a Gaussian Process (GP), and an AF that balances exploration and exploitation [11].

Continuous Constrained BO. Extensions of BO to constrained settings have been widely studied. For example, [12] jointly accounts for objective improvement and constraint satisfaction. In [13], both objective and constraints are modeled with GPs, leading to the expected constrained improvement criterion. While effective, these approaches are designed for continuous domains and rely on GP-based surrogates.

Discrete BO. A strategy for discrete variables is to relax the search space to a continuous domain and to round solutions, introducing mismatches and redundant evaluations [14]. To address this issue, [15] incorporates discretization directly into the kernel via input transformations. This approach, however, still relies on a continuous representation of the input space, and distance-based kernels may fail to capture similarities in binary settings such as BS deactivation, where small input changes can induce large variations in the objective. Sequential Model-Based Optimization (SMBO) [16] addresses this limitation by using surrogate models such as RF, which do not rely on distance-based similarity and can handle discrete input spaces. In standard SMBO, constraints, however, are not explicitly considered.

Constrained Discrete BO. Recently, d-MALIBOO [17] extends BO to constrained discrete settings. Models for constraint feasibility are incorporated into the AF using simple predictors (e.g., ridge regression or classification). Such approach may struggle in combinatorial settings due to limited modeling capacity. Some variants evaluate the AF over the entire discrete domain, which is not feasible in large-scale settings.

We propose a constrained discrete BO method to address the BS sleep-mode problem. Unlike d-MALIBOO, used as our baseline, which keeps a GP surrogate on the objective, our proposed approach exploits the problem structure to directly evaluate the objective without surrogate, concentrate the modeling effort on RF surrogates of the QoS constraints, and optimize over sampled candidates rather than exhaustively.

III. PROBLEM FORMULATION

We consider the problem of minimizing the energy consumption of a cellular network by selecting a subset of BSs to deactivate as illustrated in Fig. 1, subject to QoS constraints.

A. Optimization Problem

Consider N BSs and the binary vector $\mathbf{x} \in \{0, 1\}^N$ representing a network configuration, where $x_i = 1$ indicates that BS i is active, and $x_i = 0$ that it is switched off.




Fig. 1. Voronoi tessellations with (left) all BSs active (green) and (right) a subset of BSs switched off (red).

Let $T(\mathbf{x})$ denote the cell-edge throughput, achieved by the worst 5% of users, and $C(\mathbf{x})$ the coverage under configuration $\mathbf{x}$, defined as the proportion of users whose Signal-to-Interference-plus-Noise Ratio (SINR) exceeds a given threshold. The BS deactivation problem is formulated as

$$\begin{aligned} \min_{\mathbf{x} \in \{0,1\}^N} \quad & f(\mathbf{x}) \\ \text{s.t.} \quad & T(\mathbf{x}) \geq T_{\min}, \\ & C(\mathbf{x}) \geq C_{\min}, \end{aligned} \quad (1)$$

$f(\mathbf{x}) = \sum_{i=1}^N x_i (P_0 + \rho_i(\mathbf{x})P_{\text{var}})$, where P_0 is the static power consumption of an active BS and P_{var} the power factor multiplying the load $\rho_i(\mathbf{x})$ of BS i under configuration $\mathbf{x}$.

The problem is a discrete constrained optimization problem over a combinatorial space of size 2^N with costly evaluations. Indeed, evaluating $\rho_i(\mathbf{x})$, $i = 1, \dots, N$, $T(\mathbf{x})$, and $C(\mathbf{x})$ requires live tests or digital-twin simulations. We propose a BO methodology for solving (1) under a limited evaluation budget. The QoS constraints are modeled with RF surrogates from which feasibility probabilities are derived. They are then combined with an objective-improvement term in the AF.

B. Surrogate Modeling of Constraints

Feasibility is determined by the constraint functions $T(\mathbf{x})$ and $C(\mathbf{x})$. Owing to their combinatorial and highly non-linear dependence on $\mathbf{x}$, and to the binary nature of the inputs (for which distance-based kernels are ill-suited), we employ RF as surrogates. Given a dataset $\mathcal{D}_t = \{(\mathbf{x}_i, T_i, C_i)\}_{i=1}^{n_t}$ collected from previous simulator evaluations or real-world measurements, two RF models of B trees are trained to approximate $T(\mathbf{x})$ and $C(\mathbf{x})$ respectively. For a configuration $\mathbf{x}$, each RF provides the ensemble mean and variance, $\mu_T(\mathbf{x})$, $\sigma_T^2(\mathbf{x})$ and $\mu_C(\mathbf{x})$, $\sigma_C^2(\mathbf{x})$, estimated across the trees [16], as well as the individual per-tree predictions $\hat{T}_b(\mathbf{x})$ and $\hat{C}_b(\mathbf{x})$, $b = 1, \dots, B$.

C. Feasibility Probability

A configuration is feasible when both QoS constraints hold, encoded by the binary label $Y(\mathbf{x}) = \mathbb{1}[T(\mathbf{x}) \geq T_{\min}, C(\mathbf{x}) \geq C_{\min}]$. The feasibility probability $P_{\text{feas}}(\mathbf{x}) = \mathbb{P}(Y = 1 \mid \mathbf{x})$ is estimated by a binary classifier trained on $\mathcal{D}_t$, where each label Y_i follows from (T_i, C_i) .

We consider two variants, differing in the features fed to the classifier:

Direct ($P_{\text{feas}}^{\text{clf}}$): the classifier (logistic regression) predicts feasibility from the raw configuration $\mathbf{x}$, learning the joint feasible region of T and C without an independence assumption.

Hybrid ($P_{\text{feas}}^{\text{hybrid}}$): to account for the uncertainty captured by the constraint surrogates, the classifier (logistic regression) exploits

$$\psi(\mathbf{x}) = [s_T, s_C, \mu_T - T_{\min}, \mu_C - C_{\min}, \sigma_T, \sigma_C, \sum_i x_i],$$

where $s_T(\mathbf{x}) = \frac{1}{B} \sum_{b=1}^B \mathbb{1}[\hat{T}_b(\mathbf{x}) \geq T_{\min}]$ is the fraction of trees voting that the throughput constraint holds, and s_C is defined analogously for coverage. This maps the RF ensemble statistics to a feasibility probability while retaining their calibrated uncertainty.

D. Acquisition Function

The improvement associated with a candidate configuration is $I(\mathbf{x}) = \max(0, f^* - f(\mathbf{x}))$, where f^* denotes the best feasible value observed so far. As $f(\mathbf{x})$ depends on $\rho_i(\mathbf{x})$, $i = 1, \dots, N$, it is costly to evaluate. Instead, it is natural to consider its expectation under the predictive distribution of a surrogate, leading to $\text{EI}(\mathbf{x}) = \mathbb{E}[\max(0, f^* - f(\mathbf{x})) \mid \mathcal{D}_t]$. Under the assumption that the predictive distribution of $f(\mathbf{x})$ is a Gaussian distribution with mean $\mu(\mathbf{x})$ and standard deviation $\sigma(\mathbf{x})$, the EI admits the closed-form expression [18]

$$\text{EI}(\mathbf{x}) = (f^* - \mu(\mathbf{x}))\Phi(z) + \sigma(\mathbf{x})\phi(z)$$

where $z = \frac{f^* - \mu(\mathbf{x})}{\sigma(\mathbf{x})}$, and $\Phi(\cdot)$ and $\phi(\cdot)$ denote the cumulative distribution function and probability density function of the standard normal distribution, respectively.

In the BS sleep-mode problem, we do not consider a surrogate for the objective. The realistic energy $f(\mathbf{x})$ is considered for evaluation and reporting, but the search is driven by the active-BS count $g(\mathbf{x}) = \sum_{i=1}^N x_i$. This is motivated by the dominance of the static load-independent term in the BS power model: among feasible configurations, reducing the number of active BSs is the main driver of energy reduction. Unlike $f(\mathbf{x})$, the value of $g(\mathbf{x})$ is known for any candidate and does not require a surrogate model, reducing the complexity of the optimization.

Let n^* be the number of active BSs in the best feasible configuration found so far. The improvement in active-BS count is then computed exactly as $I_g(\mathbf{x}) = \max(0, n^* - g(\mathbf{x}))$.

The proposed AF combines this exact improvement with the learned probability of joint QoS feasibility:

$$\alpha(\mathbf{x}) = I_g(\mathbf{x})P_{\text{feas}}(\mathbf{x}). \quad (2)$$

E. Optimization Procedure

At each iteration, the next configuration is selected by maximizing the AF over the discrete space

$$\mathbf{x}_t = \arg \max_{\mathbf{x} \in \{0,1\}^N} \alpha(\mathbf{x}). \quad (3)$$

Because of the combinatorial size of the space, the maximization is approximated by evaluating $\alpha(\mathbf{x})$ over a set of S distinct

candidates sampled uniformly, providing a scalable alternative to exhaustive search.

Algorithm 1 summarizes the proposed approach. To bias the search toward energy-efficient solutions, the candidate set is further restricted to configurations satisfying $\sum_i x_i \leq n^*$, which focuses the search on increasingly sparse configurations and avoids revisiting suboptimal regions with denser configurations.

Algorithm 1: Constrained BO for BS deactivation.

Input: Total evaluation budget N_{eval} , number of candidates S , initial dataset $\mathcal{D}_0$ of size n_0

Output: Best feasible configuration $\mathbf{x}^*$

- 1: Initialize $\mathcal{D}_0 = \{(\mathbf{x}_i, T_i, C_i)\}_{i=1}^{n_0}$ with $\lfloor n_0/2 \rfloor$ random and $\lceil n_0/2 \rceil$ feasible configurations
 - 2: Initialize best feasible $\mathbf{x}^*$ and $n^* = \sum_i x_i^*$, breaking ties by higher constraint margins
 - 3: **for** $t = 1$ **to** $N_{\text{eval}} - n_0$ **do**
 - 4: Train RF surrogates for $T(\mathbf{x})$ and $C(\mathbf{x})$ on $\mathcal{D}_{t-1}$, and fit the feasibility classifier ($P_{\text{feas}}^{\text{clf}}$ or $P_{\text{feas}}^{\text{hybrid}}$)
 - 5: Sample S distinct candidates $\mathcal{X}_{\text{cand}}$ such that $\sum_i x_i \leq n^*$
 - 6: Select $\mathbf{x}_t$ using (3) over $\mathcal{X}_{\text{cand}}$
 - 7: Evaluate $T(\mathbf{x}_t)$, $C(\mathbf{x}_t)$, $\rho(\mathbf{x}_t)$ with the simulator (or real-world measurements) and compute $f(\mathbf{x}_t)$
 - 8: $\mathcal{D}_t = \mathcal{D}_{t-1} \cup \{(\mathbf{x}_t, T(\mathbf{x}_t), C(\mathbf{x}_t))\}$
 - 9: Update best feasible $\mathbf{x}^*$ and n^*
 - 10: **end for**
-

F. Computational Complexity

Since all methods are compared under the same budget for BS configuration evaluation (Step 7 of Algorithm 1), the following complexity analysis focuses on the BO decision step. At iteration t , let n_t be the dataset size, S the number of candidate configurations scored by the AF, B the number of trees, m the features tried per split, d the average tree depth, and recall the N binary inputs.

d-MALIBOO. Following the standard GP implementation [19], the posterior mean and variance are obtained through a Cholesky factorization of the $n_t \times n_t$ kernel matrix and solving the resulting triangular systems. The factorization requires $\mathcal{O}(n_t^3)$ operations, and the posterior standard deviation over the S candidates $\mathcal{O}(Sn_t^2)$. A degree-2 polynomial expansion of the binary inputs adds $q = 1 + N + \frac{N(N-1)}{2} = \mathcal{O}(N^2)$ features; the squared terms vanish since $x_i^2 = x_i$ for binary inputs, and the ridge constraint models require $\mathcal{O}(n_t q^2 + q^3)$ operations. Scoring the S candidates against the n_t observations in the q -dimensional expanded feature space has a complexity $\mathcal{O}(Sn_t q)$. The dominant per-iteration complexity is

$$\mathcal{O}(n_t^3 + Sn_t^2 + Sn_t q + n_t q^2 + q^3), \quad q = \mathcal{O}(N^2), \quad (4)$$

(without GP hyperparameter tuning factor).

Proposed method. No surrogate is placed on the objective and no linear system is solved. Each iteration trains two RF regressors in $\mathcal{O}(B m n_t \log n_t)$ operations and a feasibility

classifier on a fixed p -dimensional input ($p = 7$ for the hybrid variant) in $\mathcal{O}(n_t p)$. Scoring the S candidates requires $\mathcal{O}(S B d)$ operations for the surrogates, $\mathcal{O}(S p)$ for the classifier, and $\mathcal{O}(S)$ for the exact improvement. The dominant complexity is

$$\mathcal{O}(B m n_t \log n_t + S B d). \quad (5)$$

Hence the proposed approach scales near-linearly in n_t (versus cubic), avoids the $\mathcal{O}(N^2)$ feature expansion.

IV. EXPERIMENTS

In this section, we describe the experimental setup used to evaluate the proposed method. We detail the QoS constraints, the surrogate model configuration, and the baseline methods used for comparison, and then present the simulation results.

A. Simulation Environment

We consider a network of $N = 21$ BSs with an average inter-site distance $\text{ISD} = 450$ m, over an area of 1.8 km^2 . The physical layer operates at a carrier frequency of 3.5 GHz with a system bandwidth of 90 MHz . Each BS is equipped with a 4×8 antenna array, while user terminals have a single receive antenna. Propagation includes path-loss with exponent $\beta = 3$, log-normal shadowing with standard deviation 4 dB , and small-scale fading modeled using a multi-path channel with 50 rays in an urban micro-cell (UMi) scenario. Users are spatially distributed according to a homogeneous Poisson Point Process (PPP), with density λ_{UE} . Traffic demand is modeled through a flow-level model, where user arrivals follow a Poisson process, and file sizes are exponentially distributed with mean μ_d . We consider two traffic scenarios corresponding to day and night conditions. In the **day scenario**, the user density is $\lambda_{\text{UE}} = 60 \text{ users/km}^2$ with an average file size $\mu_d = 20 \text{ Mbits}$. In the **night scenario**, the user density $\lambda_{\text{UE}} = 30 \text{ users/km}^2$ with $\mu_d = 20 \text{ Mbits}$.

Performance metrics are averaged over $N_{\text{real}} = 100$ independent UE PPP realizations. For each UE realization, we perform $N_{\text{MC}} = 500$ Monte Carlo channel realizations. For a given configuration $\mathbf{x}$, the simulator [20] provides the cell-edge throughput $T(\mathbf{x})$, the network coverage $C(\mathbf{x})$ and the per-BS loads $\rho_i(\mathbf{x})$.

B. QoS Constraints

The QoS constraints are defined relative to the fully active day configuration: $T_{\min} = 0.85, T_{\text{full}}$ and $C_{\min} = 0.95, C_{\text{full}}$, where T_{full} and C_{full} denote the corresponding cell-edge throughput and coverage. This ensures that energy savings are obtained without significant degradation of QoS.

C. Baseline Methods

We compare the proposed approach against several baselines under a fixed trial budget $N_{\text{eval}} = 40$, defined as the number of configuration evaluations on the network.

Random search. Configurations to test are sampled uniformly. **Greedy heuristics.** Starting from the fully active network, we consider different greedy variants:

Standard greedy: BSs are iteratively deactivated one at a time. At each step, all configurations obtained by switching off one

BS are evaluated on the network. If feasible configurations exist, the one with the largest constraint margin is chosen.

Aggressive greedy: To accelerate the search, 2 BSs are simultaneously deactivated during the first quarter of the evaluation budget, enabling larger steps in the search space. The algorithm then switches to the standard greedy procedure.

Greedy with backtracking: To improve robustness against early suboptimal decisions, we store the top- K configurations at each iteration. The search proceeds with the best candidate and keeps the others as alternatives. In case of no feasible configuration, the algorithm backtracks to the next best candidate and resumes the search. This algorithm is deterministic and we use it for exploring the outcome of a good greedy scheme, letting it converge without imposing an evaluation budget.

D. BO algorithms

d-MALIBOO. The d-MALIBOO algorithm [17] is our basic discrete constrained BO. It keeps a surrogate on the objective and models feasibility using ridge regression with second-degree polynomial feature expansion. We use an ε -greedy strategy with $\varepsilon = 0.1$ to avoid repeated evaluations [17]. While d-MALIBOO can, in principle, evaluate the AF over the entire discrete search space, we fix the sample set of distinct candidates to $S = 10\,000$, following [16].

Our BO proposal. RF models are used as surrogates for cell-edge throughput and coverage. The RF hyperparameters follow standard configurations used in [16], with $B = 10$ trees, a feature subsampling ratio of $5/6$, and a minimum of 10 samples per split; bootstrap sampling is enabled. Combining the two classifiers of Section III-C (direct, hybrid) with the optional log-transform (it has been shown in [16] to improve surrogate modeling performance) of the throughput yields four variants; we report log-direct, hybrid, and log-hybrid, omitting the plain direct variant, which was consistently dominated.

E. Initial Dataset Construction

Initially, 20 configurations are generated. To guarantee feasibility, the fully active configuration is always included. Additional configurations with a high number of active BSs (typically 19 or 20 active BSs) are sampled, as they are more likely to satisfy the QoS constraints. The remaining configurations are randomly sampled across the search space, with the number of active BSs ranging from 1 to 18. This strategy ensures covering a diversity of configurations with a balance between feasible and infeasible ones. Overall, the initial dataset contains half of feasible configurations.

F. Results

Table I summarizes the performance of each method over 20 independent runs. Each run uses a total budget of 40 network evaluations (20 initial and 20 guided). Using multiple runs assesses the robustness of the methods in operational scenarios. For each run we record the objective f of the best feasible configuration found: $\bar{f}$ and $\text{std}(f)$ are its mean and standard deviation across runs, $\text{saving} = (f_{\text{ref}} - \bar{f})/f_{\text{ref}}$ is the relative energy reduction, mean BS is the average number of active

TABLE I
COMPARISON OF OPTIMIZATION METHODS. THE REFERENCE (21 ACTIVE BSs) CORRESPONDS TO A GLOBAL ENERGY CONSUMPTION OF 7271 W AND 6018 W, IN THE DAY AND NIGHT SCENARIOS, RESPECTIVELY.

Method	saving (%)	$\bar{f}$ (W)	$\text{std}(f)$ (W)	mean BS	best (%)	2nd best (%)
Day scenario						
BO hybrid	11.415	6442	42	15.90	0.375	0.475
BO log-hybrid	10.707	6493	60	16.30	0.225	0.100
BO log-direct	10.234	6528	63	16.75	0.075	0.050
BO d-MALIBOO	<u>11.139</u>	<u>6462</u>	73	<u>16.05</u>	<u>0.325</u>	0.175
Greedy aggressive	9.224	6601	138	17.40	0.000	0.150
Random search	4.749	6927	178	18.75	0.000	0.050
Night scenario						
BO hybrid	13.862	5184	51	14.65	0.000	0.000
BO log-hybrid	<u>15.518</u>	<u>5084</u>	54	<u>14.05</u>	<u>0.100</u>	0.750
BO log-direct	15.362	5094	64	13.75	0.050	0.250
BO d-MALIBOO	16.942	4999	65	<u>14.05</u>	0.850	0.000
Greedy aggressive	11.168	5346	30	17.00	0.000	0.000
Random search	8.579	5502	130	16.80	0.000	0.000

BSs, and the best/2nd-best columns indicate the proportion of runs for which each method identifies the best and second-best configuration (in terms of power consumption).

In the **day scenario**, the proposed BO hybrid variant performs best. It achieves the highest energy saving (11.4%), the lowest mean objective ($\bar{f} = 6442$ W), and the fewest active BSs on average (15.9). It is also the most robust, with the smallest run-to-run variance ($\text{std}(f) = 42$ W). d-MALIBOO is the closest competitor (11.1% saving, 16.1 BSs).

In the **night scenario**, d-MALIBOO achieves the highest saving (16.9%), the lowest mean objective ($\bar{f} = 4999$ W), and the best configuration in 85% of the runs. Among the proposed variants, the log-scaled ones take the lead: BO log-hybrid reaches 15.5% saving (and the second-best configuration in 75% of runs), while BO log-direct reaches 15.4% with the fewest active BSs on average (13.75).

The night scenario illustrates a limitation of using the active-BS count as the objective. Configurations with the same number of active BSs may exhibit different load-dependent energy consumption. Since d-MALIBOO models the realized objective, it can distinguish such configurations, whereas the proposed AF cannot.

The comparison with baselines highlights the benefit of BO-based candidate selection. The standard greedy method converges to a configuration with 19 active BSs (not shown in the table as it is deterministic). Random search generally finds feasible configurations with a larger number of active BSs, while the aggressive greedy heuristic is sensitive to early deactivation choices and may become trapped in suboptimal regions of the search space. Finally, Greedy backtracking deterministically reaches 15 (day) and 13 (night) active BSs, but only after 122 and 134 network evaluations respectively, whereas the BO-based methods reach often this target within a fixed budget of 40.

Fig. 2 provides additional insight into the learned feasibility models. For one representative run, it shows the maximum predicted feasibility probability among the sampled candidate

configurations as a function of the number of active BSs, at different BO iterations. These curves should be interpreted as a diagnostic of the feasibility model rather than as the selected configurations. Indeed, the AF is obtained by weighting the feasibility probability by the active-BS improvement term; therefore, the candidate with the highest predicted feasibility is not necessarily selected if it contains too many active BSs and offers little or no improvement.




Fig. 2. Maximum predicted feasibility probability among sampled candidate configurations as a function of the number of active BSs in the day scenario for different BO iterations. The curves are shown for run 7.

As expected, the maximum predicted feasibility generally increases with the number of active BSs, since denser configurations are more likely to satisfy coverage and throughput constraints. As the BO loop progresses and better feasible configurations are found, the candidate set is restricted to configurations with at most n^* active BSs. This explains why the feasibility curves at later iterations are shown only over a reduced range of active-BS counts.

V. CONCLUSION

This paper addresses the problem of BS sleep mode management under QoS constraints and formulates it as a constrained optimization problem over a combinatorial search space, with each configuration costly to evaluate.

We propose a BO solution approach, tailored to discrete domains with expensive evaluations. It relies on RF surrogate models to approximate the QoS metrics and to estimate feasibility probabilities through a learned classifier, enabling the design of an AF that jointly accounts for energy efficiency and constraint satisfaction. Because the energy objective reduces to the number of active BSs, no surrogate is placed on it and no Gaussian assumption is required, which also yields a lower per-iteration complexity than GP-based baselines. A sampling-based optimization strategy is introduced to efficiently explore the large combinatorial search space.

Simulation results show that the proposed method identifies feasible, energy-efficient configurations with fewer evaluations than greedy heuristics, and remains competitive with a state-of-the-art discrete BO baseline at equal evaluation budget while lowering the decision-step complexity.

Future work will focus on extending the proposed framework towards active and adaptive data acquisition strategies, where the objective is not only to identify optimal configurations, but also to select the most informative experiments.

REFERENCES

- [1] O. Arnold, F. Richter, G. Fettweis, and O. Blume, "Power consumption modeling of different base station types in heterogeneous cellular networks," in *Proc. Fut. Netw. & Mob. Summit*, 2010, pp. 1–8.
- [2] L. Saker and S. E. Elayoubi, "Sleep mode implementation issues in green base stations," in *IEEE PIMRC*, 2010, pp. 1683–1688.
- [3] J. Wu, Y. Zhang, M. Zukerman, and E. K.-N. Yung, "Energy-efficient base-stations sleep-mode techniques in green cellular networks: A survey," *IEEE Com. Surv. & Tut.*, vol. 17, no. 2, pp. 803–826, 2015.
- [4] F. Salahdine, J. Opadere, Q. Liu, T. Han, N. Zhang, and S. Wu, "A survey on sleep mode techniques for ultra-dense networks in 5G and beyond," *Comput. Com.*, vol. 150, pp. 559–578, 2021.
- [5] G. Micallef, L. Saker, S.-E. Elayoubi, and H. O. Scheck, "Realistic energy saving potential of sleep mode for existing and future mobile networks," *J. Commun.*, vol. 7, no. 10, pp. 740–748, 2012.
- [6] S.-E. Elayoubi, L. Saker, and T. Chahed, "Optimal control for base station sleep mode in energy efficient radio access networks," in *Proc. IEEE INFOCOM*, 2011, pp. 106–110.
- [7] X. Liang, A. Al-Tahmeesschi, S. Chetty, and H. Ahmadi, "Green O-RAN operation: a modern ML-driven network energy consumption optimisation," in *Proc. IEEE GLOBECOM*, 2025, pp. 4475–4480.
- [8] M. Bordin, A. Lacava, M. Polese, S. Satish, M. A. Nittoor, R. Sivaraj, F. Cuomo, and T. Melodia, "Design and evaluation of deep reinforcement learning for energy saving in open RAN," in *IEEE CCNC*, 2025, pp. 1–6.
- [9] M. Masoudi, E. Soroush, J. Zander, and C. Cavdar, "Digital twin assisted risk-aware sleep mode management using deep q-networks," *IEEE Trans. Vehicul. Techn.*, vol. 72, no. 1, pp. 1224–1239, 2022.
- [10] J. Snoek, H. Larochelle, and R. P. Adams, "Practical bayesian optimization of machine learning algorithms," *Advances in Neural Information Processing Systems*, vol. 25, 2012.
- [11] P. I. Frazier, "A tutorial on bayesian optimization," *arXiv preprint*, 2018. [Online]. Available: <https://arxiv.org/abs/1807.02811>
- [12] R. B. Gramacy and H. K. H. Lee, "Optimization under unknown constraints," *Bayesian Statistics 9*, pp. 229–256, 2011.
- [13] J. R. Gardner, M. J. Kusner, Z. E. Xu, K. Q. Weinberger, and J. P. Cunningham, "Bayesian optimization with inequality constraints," in *Proc. ICML*, vol. 2014, 2014, pp. 937–945.
- [14] P. H. Luong, S. Gupta, S. Rana, and S. Venkatesh, "Bayesian optimization for discrete and categorical variables," in *Proc. Australasian Joint Conference on AI*. Springer, 2019, pp. 473–484.
- [15] E. C. Garrido-Merchán and D. Hernández-Lobato, "Dealing with categorical and integer-valued variables in bayesian optimization with gaussian processes," *Neurocomputing*, vol. 380, pp. 20–35, 2020.
- [16] F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Sequential model-based optimization for general algorithm configuration," *Journal of Artificial Intelligence Research*, vol. 40, pp. 507–547, 2011.
- [17] R. Sala, B. Guindani, D. Ardagna, and A. Guglielmi, "d-maliboo: a bayesian optimization framework for dealing with discrete variables," in *Proc. IEEE MASCOTS*, 2024, pp. 1–8.
- [18] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive black-box functions," *Journal of Global Optimization*, vol. 13, no. 4, pp. 455–492, 1998.
- [19] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [20] I. Merah, M. Kieffer, S.-E. Elayoubi, T. Clessienne, and B. Sayrac, "Towards network design-to-impact: RAN pruning for sustainability," in *Proc. IEEE International Conference on Communications (ICC)*, 2026.