# POLITECNICO DI TORINO

*Department of Electronics and Telecommunications*

*Master's Degree in ICT for Smart Societies*


Politecnico
di Torino

IMT Atlantique
Bretagne-Pays de la Loire
École Mines-Télécom

## On the impact of recent 5G/6G BS Power Models on energy efficiency studies

**Supervisors**

Prof. Michela Meo

Prof. Loutfi Nuaymi

**Candidate**

Augusta Castelli

March 2026

# Table of Contents

# List of Figures

# List of Tables

# Abstract

Energy efficiency is a major design objective in current Fifth Generation (5G) and emerging Sixth Generation (6G) radio access networks, where base stations account for a large share of operational consumption and remain far from energy-proportional under time-varying load. In this setting, the adopted base-station power-consumption model is not a neutral implementation detail: different formulations can lead to substantially different estimates of achievable energy savings and of the associated energy–performance trade-offs.

This thesis quantifies such model-driven effects through a controlled comparison between the standardized 3GPP TR 38.864 state-based model and an analytical Active Antenna Unit (AAU) formulation based on hardware-level component scaling. To isolate the impact of model choice, a unified evaluation pipeline is developed in which scenario assumptions, traffic, control knobs, and KPI definitions remain fixed while only the power model is changed. The analysis combines three complementary components: (i) static one-at-a-time sweeps over antenna activity, utilized bandwidth, and transmit-power scaling to derive energy-delay trade-off curves; (ii) a discrete-time traffic-driven simulator with FIFO queueing and multi-level sleep modes (MICRO/LIGHT/DEEP) to capture activity factors, wake-up constraints, and sleep opportunities; and (iii) a preliminary multi-small-cell extension based on tabular Q-learning to illustrate how model choice propagates into learning-based control through the energy feedback embedded in the reward.

Across the evaluated scenarios, the analytical AAU-based formulation consistently yields a more optimistic energy picture than the standardized 3GPP model, predicting larger energy savings under the same operating conditions. Overall, the results show that power-model selection is itself a methodological choice that can materially affect the interpretation of energy-efficiency analyses and learning-based energy-management studies.

# Executive Summary

Energy efficiency is increasingly regarded as a system-level requirement for current 5G and future 6G networks. Within the Radio Access Network (RAN), base stations (BSs) are a primary target for energy-saving strategies because they operate continuously and their power consumption is far from energy-proportional: a significant share persists even at low load, while the remaining part depends on transmitted activity and on the active hardware configuration. This motivates two main classes of energy-saving mechanisms: (i) radio reconfiguration during ACTIVE operation, such as bandwidth adaptation, antenna activation, and transmit-power scaling; and (ii) Advanced Sleep Modes (ASMs) during idle or low-traffic periods.

**Thesis purpose and research area.** The thesis is based on the premise that the quantitative outcome of energy-saving mechanisms depends strongly on the BS power-consumption model adopted in the evaluation. Different formulations embed different assumptions about which power terms scale with reconfiguration and which remain as non-removable floors. As a result, the selected model can significantly affect both absolute energy estimates and the relative gains measured against a common baseline. The purpose of this work is therefore to quantify these *model-driven* differences in energy metrics and energy-performance trade-offs, with specific reference to evaluation methodologies for 5G/6G RANs.

**Objective.** The study performs a controlled comparison between:

- the standardized 3GPP TR 38.864 BS power model, used as a reproducible reference for 5G-Advanced energy-efficiency studies; and
- an analytical Active Antenna Unit (AAU) model derived from hardware-level component scaling, representative of more detailed platform-aware formulations.

To isolate the effect of model selection, scenario assumptions, traffic generation, control knobs, and KPI definitions are kept unchanged across experiments, while only the BS power model is varied.

**Methodology.** BS operation is parameterized through three dimensionless scaling factors:

- **Antenna activity** $s_a$: fraction of active transmit chains.
- **Bandwidth utilization** $s_f$: fraction of system bandwidth used for downlink service.
- **Transmit-power scaling** $s_p$: multiplicative factor applied to the radiated transmit power.

The evaluation is articulated into three complementary tracks:

- **Static sensitivity analysis:** one-at-a-time sweeps over (sa, sf, sp) are used to derive trade-off curves between active-only energy per bit, reported also through a normalized energy-saving indicator (NES), and a service-time proxy represented by latency increase, under a repeatable link abstraction.
- **Traffic-driven time-domain evaluation:** a discrete-time simulator models stochastic arrivals, FIFO queueing, and multi-level sleep operation (MICRO/LIGHT/DEEP) with wake-up constraints. It reports average power, energy per served bit, the corresponding normalized gain relative to the baseline, latency statistics, and diagnostic indicators such as conditional ACTIVE power and sleep-state occupancies.
- **Preliminary learning-based extension:** a small multi-cell setup is used to illustrate how the selected BS power model propagates into the control loop through the energy term of the reward adopted by a tabular learning controller, thereby influencing the learned operating regime.

**Personal contribution.** The main personal contributions of this thesis are:

- the design and implementation of a unified evaluation pipeline that isolates the effect of BS power-model selection under fixed scenario, traffic, and KPI assumptions;

- the development of a traffic-driven simulator for BS operation with queueing dynamics and multi-level sleep-mode behavior;
- the implementation of a preliminary tabular learning-based extension to examine how model choice propagates into energy-aware control.

**Main results.** Across the evaluated scenarios, the AAU-based formulation consistently provides a more optimistic estimate of energy efficiency than the standardized 3GPP model. Under the same reconfiguration settings, it predicts larger reductions in power consumption and energy per bit. This gap reflects the different internal structure of the two formulations: in the 3GPP model, a larger share of the ACTIVE power budget remains weakly sensitive or insensitive to bandwidth and transmit-power scaling, whereas in the AAU-based model a wider fraction of the consumption is coupled to the active configuration, amplifying the benefit of reconfiguration.

The results also show that energy conclusions depend on *dynamic* system operation, not only on instantaneous ACTIVE power. In the traffic-driven simulator, actions that reduce conditional ACTIVE power may still lead to smaller net savings if they increase service time and backlog, thereby shrinking idle intervals and reducing the opportunity to enter deeper sleep states. This effect is especially relevant for aggressive bandwidth reductions: although conditional ACTIVE power decreases, the total average power can flatten or even worsen when the BS remains ACTIVE for a larger fraction of time.

In the preliminary tabular learning-based extension, the selected BS power model directly affects the energy term included in the reward signal and therefore changes the objective effectively "seen" by the learner. Even in the simplified setup considered here, the AAU-based formulation induces a more optimistic network-level energy picture, showing that power-model selection is part of the environment definition whenever learning-based energy management is assessed.

**Conclusion and future work.** Overall, the thesis shows that BS power-model selection is a methodological choice that materially affects energy-efficiency evaluation in 5G/6G RANs. Across the considered scenarios, the analytical AAU-based model consistently yields more optimistic estimates than the standardized 3GPP formulation.

The results should be interpreted within the adopted traffic assumptions, sleep-mode abstraction, and preliminary tabular learning setting. Future work may extend the framework with richer KPIs, more realistic system models, and scalable coordinated learning approaches.

# Keywords


- BS power consumption
- Third Generation Partnership Project (3GPP) TR 38.864
- AAU power model
- Power-model sensitivity
- Energy efficiency
- Energy per bit
- NES
- Traffic-driven simulation
- Queueing dynamics
- Sleep modes
- State occupancy
- Conditional ACTIVE power
- Radio Access Network (RAN)
- Learning-based energy management
- Q-Learning

# Acronyms

**3GPP** Third Generation Partnership Project. xi

**5G** Fifth Generation of Cellular Mobile Communications. 1, 3, 14, 22, 63

**6G** Sixth Generation of Cellular Mobile Communications. 1, 13, 14, 22

**AAU** Active Antenna Unit. ii, iv, xi, 1–4, 8, 10, 20, 21, 31, 33, 34, 43, 47, 54–59, 61–63

**ASM** Advanced Sleep Mode. v, 2–4, 12–17, 21

**BS** Base Station. iv, v, xi, 1–15, 17–28, 30, 31, 34, 35, 41–43, 45, 47, 49, 50, 52–54, 56–61, 63–65

**CC** Component Carrier. 14, 15, 17

**DL** Downlink. 22, 24, 30, 32–34

**EE** Energy Efficiency. 16

**ICT** Information and Communication Technology. 1, 6

**KPI** Key Performance Indicator. 15, 17, 25, 63, 65

**MIMO** Multiple-Input Multiple-Output. 2, 8, 10, 18

**NES** Normalized Energy Saving. iv, xi, 28, 29, 43, 52, 54–56

**NR** New Radio. 9, 21

**PA** Power Amplifier. 2

**QoS** Quality of Service. 2, 13–18, 21, 59, 65

**RAN** Radio Access Network. xi, 1, 3, 5, 7, 12, 22, 63

**RB** Resource Block. iv, 2, 14, 15, 17, 56

**RL** Reinforcement Learning. v, 1, 2, 4, 5, 17–22, 50, 52, 53

**SBS** Small Cell Base Station. ii, iv, v, 43–53, 58–60, 62

**SINR** Signal-to-Interference-plus-Noise Ratio. 45

**SM** Sleep Mode. 16, 22, 38, 40, 43, 54

**SNR** Signal-to-Noise Ratio. 24, 25, 36, 44, 45, 48

**UE** User Equipment. iv, 23, 37

# 1 Introduction

Improving the energy efficiency of cellular networks is a major objective for current and emerging systems, driven by the growth of mobile traffic and by the need to reduce the environmental footprint of Information and Communication Technology (ICT). Within this context, the RAN deserves particular attention because BSs account for a substantial share of operational energy consumption.

This thesis investigates how the choice of BS power-consumption model influences the estimated benefit of energy-saving mechanisms and the interpretation of energy-performance trade-offs in Fifth Generation of Cellular Mobile Communications (5G) and beyond-5G networks. To this end, it develops a controlled comparison between the standardized 3GPP TR 38.864 model and the analytical AAU-based model proposed by Piovesan et al.[18], combining static parameter sweeps, dynamic traffic-driven simulations, and a preliminary RL-based multi-BS extension.

## 1.1 Context and Motivation

The evolution from 5G towards Sixth Generation of Cellular Mobile Communications (6G) and beyond is expected to deliver major improvements in capacity, latency, and reliability while complying with increasingly stringent sustainability targets. As a result, the energy footprint of cellular infrastructures has become a central concern for operators, vendors, and policy makers. Mobile networks must support heterogeneous services and growing traffic volumes without a proportional increase in operational costs and environmental impact. Within the end-to-end system, the RAN is one of the main contributors to overall energy consumption, and BSs dominate this footprint because they operate continuously and include multiple hardware subsystems that remain partially active even at low load.

A key characteristic of mobile traffic is its strong spatio-temporal variability. During off-peak periods, the offered load may decrease significantly, yet a non-negligible fraction of BS power is still consumed by radio chains, baseband processing, and auxiliary components. This mismatch motivates energy-saving strategies that adapt the radio configuration during active operation and exploit low-power states during

idle periods. Representative mechanisms include antenna muting in massive Multiple-Input Multiple-Output (MIMO) deployments, transmit-power control, bandwidth or RB utilization adaptation, and advanced sleep modes (ASM) that progressively deactivate hardware blocks when the transmission queue is empty. The actual benefit of these mechanisms, however, depends on how BS power consumption is modeled.

For this reason, conclusions on energy efficiency are intrinsically linked to the adopted BS power-consumption model. Power models map radio-domain parameters, such as the number of active antennas, utilized bandwidth, transmit power, and sleep-state occupancy, into energy indicators that are then combined with performance metrics including latency and Quality of Service (QoS). Different modeling assumptions, for example on the fraction of power treated as configuration-dependent versus fixed, or on the granularity used to represent hardware components, can lead to substantially different conclusions for the same control action or energy-saving strategy.

This thesis focuses on two representative BS power-consumption models. The standardized 3GPP TR 38.864 (Release 18)[1] formulation is adopted as a reference model providing prescribed configurations, normalized parameters, and explicit transitions between active operation and hierarchical sleep modes. In parallel, the analytical AAU-based model proposed by Piovesan et al.[18], derived from measurements on commercial units and refined through fitting, is considered as a hardware-informed alternative. It decomposes total power into interpretable components, such as baseband processing, transceiver chains, and Power Amplifiers (PAs), and links them to scaling factors including antenna activity, bandwidth utilization, and transmit-power ratio.

The work is motivated by the observation that power-model selection is not a neutral implementation detail. It can materially alter the estimated savings of the same adaptation or sleep strategy and, consequently, the system-level conclusions drawn from the evaluation. Conservative assumptions may underestimate achievable gains, whereas optimistic scaling laws may overstate them when used outside their intended abstraction scope. A controlled comparison under aligned conditions is therefore necessary to identify when and why the predictions of different models diverge and to support more reliable energy-efficiency assessments. This is particularly important for advanced control approaches, including RL-based energy management, in which

learned decisions depend directly on the energy feedback generated by the adopted model.

## 1.2 Problem Statement and Thesis Objectives

**Problem statement.** Although energy-efficient RAN design has been widely investigated, the impact of BS power-model assumptions on the estimated energy-saving potential of control strategies remains only partially understood. Many studies evaluate radio-resource adaptations and ASM policies under a single power model and implicitly treat the resulting energy figures as broadly representative. However, standardized and analytical models differ in abstraction level, internal structure, and sensitivity to radio-domain parameters. For instance, the standardized 3GPP TR 38.864 model relies on a system-level state abstraction with hierarchical sleep modes, whereas the analytical AAU-based model proposed by Piovesan et al. decomposes total power into component contributions and scales dominant terms with configuration variables. As a consequence, the same control action may produce different estimates of power consumption, energy per bit, delay-related metrics, and sleep-state utilization depending on the selected model. Without a controlled comparison under aligned assumptions, these discrepancies risk being mistaken for genuine system-level effects, leading to biased interpretations and potentially misleading recommendations.

**Objectives and research questions.** The goal of this thesis is to quantify how BS power-model selection influences both the estimated energy-saving potential of radio adaptations and ASM operation, and the resulting energy–performance trade-offs in 5G and beyond-5G networks. To this end, a standardized reference model, namely 3GPP TR 38.864, is compared with an analytical hardware-informed alternative, namely the AAU-based model by Piovesan et al., under aligned assumptions designed to isolate model-driven effects from system-level behavior. The study is organized around the following research questions:

- **Static comparison:** how do the two models differ in the estimated energy-saving potential of antenna activity, bandwidth utilization, and transmit-power adaptations when evaluated through consistent KPIs?

- **Dynamic comparison:** under traffic-driven operation with queueing and sleep-mode transitions, how does model selection affect average power, energy per served bit, latency, and sleep-state occupancy?

- **Interpretation and guidelines:** which structural assumptions of the two models explain the observed gaps, and what practical guidelines follow for model selection and for the interpretation of reported energy-efficiency results?

- **Learning-based control under model uncertainty:** in a multi-BS scenario, how does power-model selection affect the reward signal and, consequently, the policy learned by an RL agent for coordinated sleep-mode and resource-configuration control?

## 1.3 Main Contributions

The main contributions of this work are summarized as follows:

- **Controlled comparison of two representative BS power models.** The thesis develops a systematic comparison between the standardized 3GPP TR 38.864 model and the analytical AAU-based model proposed by Piovesan et al., showing how differences in abstraction level and configuration sensitivity affect estimated energy savings.

- **Unified evaluation framework for static and dynamic analysis.** A common framework is defined to enable fair model-to-model comparisons under aligned assumptions, combining a static KPI-driven study based on radio-parameter sweeps with a dynamic time-domain simulator including traffic arrivals, queue evolution, and explicit ASM operation.

- **Quantification and interpretation of model-driven discrepancies.** The thesis quantifies the gaps between model predictions across multiple adaptation dimensions and traffic conditions, and relates these differences to specific structural assumptions of the underlying formulations.

- **Methodological guidance for energy-efficiency assessment.** The results are used to derive practical guidance on model selection and on the consistent interpretation of energy-performance trade-offs in comparative studies of RAN energy-saving strategies.

- **Preliminary RL-based extension in a multi-BS setting.** A preliminary learning-based extension is implemented to illustrate how power-model selection propagates into the reward definition and affects coordinated energy-management behavior in a simplified multi-BS scenario.

## 1.4 Thesis Structure

The remainder of this thesis is organized as follows.

Chapter 2 (*State of the Art*) reviews the main literature on energy efficiency in cellular networks, with emphasis on BS power-consumption modeling, radio-resource adaptations, advanced sleep modes, and learning-based energy management. It also highlights the methodological gap addressed by this work.

Chapter 3 (*Methodology*) presents the study scenario and the unified evaluation framework. It describes the adopted assumptions, the selected KPIs, the implementation of the two BS power models, the static analysis procedure, the dynamic traffic-driven simulator, and the preliminary multi-BS RL extension.

Chapter 4 (*Results*) reports and discusses the main findings. It compares the two models under static radio adaptations, analyzes dynamic traffic scenarios with sleep-mode operation, and examines how model selection affects the behavior observed in the learning-based extension.

Chapter 5 (*Conclusion*) summarizes the main results, discusses the limitations of the study, and outlines directions for future work.

# 2 State of the Art

Energy efficiency is a central objective in modern mobile networks, where BSs account for a large share of operational energy consumption and where the reported benefit of energy-saving mechanisms depends strongly on the adopted modeling assumptions. The literature on this topic spans several interconnected strands, including BS power-consumption models, traffic- and sleep-aware evaluation methodologies, energy-performance metrics, and learning-based control strategies. This chapter reviews these strands from a broad and narrative perspective, with particular emphasis on how modeling choices shape the interpretation of energy-saving mechanisms and the robustness of the resulting conclusions.

## 2.1 Sustainability Context and BS Focus

Sustainability has become a major design objective for communication infrastructures as mobile traffic continues to grow and networks are expected to scale in capacity without a proportional increase in energy demand. In this context, energy efficiency is not only a cost driver, but also a practical proxy for the environmental footprint of large-scale ICT systems. A representative example is the Green Meter study of the GreenTouch consortium, which frames the challenge in terms of strong traffic growth and, at the same time, highlights the long-term potential of architectural and technological advances in improving energy efficiency under idealized assumptions. [9] More broadly, recent 6G vision documents treat sustainability as a system-level requirement, with operational energy consumption, associated emissions, and the need for agreed metrics and measurement practices becoming explicit design concerns. [2, 16]

Within ICT, telecom networks are a particularly relevant case because their energy consumption depends on both infrastructure deployment and time-varying service demand. As surveyed in [12], improving the energy profile of mobile systems involves multiple layers and technologies, including more efficient radio access, virtualization and cloud-native deployments, smarter resource management, and energy-aware operational strategies. A recurring theme across these directions is that meaningful

reductions require mechanisms able to adapt to load, for example by scaling resources or transitioning equipment to low-power states, while preserving service performance. This immediately motivates evaluation methods that connect traffic dynamics and control decisions to both energy and performance outcomes. At the same time, broader sustainability discussions in networking stress that technical efficiency gains alone do not automatically translate into lower footprint unless operational practices, measurement, and control are considered jointly. [12, 15]

A natural narrowing step is to identify where energy is predominantly spent within mobile networks. The RAN is widely recognized as a major contributor, and BSs in particular dominate operational consumption because they are densely deployed, operate continuously, and include power-hungry radio-frequency and baseband components. Early quantitative assessments already attributed a large fraction of the network energy budget to BS sites, often reported around $\sim 80\%$ in the literature, making them the natural target of energy-saving strategies. [5] Complementary footprint-oriented analyses likewise identify access-network equipment and base stations as major contributors in mobile communications from both ecological and economic perspectives, reinforcing the relevance of BS-level efficiency work. [8]

The centrality of BSs, however, is not explained only by where energy is spent, but also by how consumption scales with load. A key observation is that BS power is not perfectly energy-proportional: substantial components remain even at low utilization, while additional power depends on transmitted load and active hardware. [5] For this reason, many energy-saving strategies aim either to reduce the power drawn during active operation through resource adaptation, such as bandwidth scaling, transmit-power control, and antenna activation management, or to exploit low-power states during idle or low-traffic periods. The actual benefit of these mechanisms, however, depends on two tightly coupled modeling choices: the BS power-consumption model used to translate resource and activity levels into supply power, and the traffic-driven evaluation framework used to determine when the system remains active and when it can effectively sleep.

This observation motivates the focus of the remainder of the chapter. The next sections review the evolution of BS power models, from early affine approximations

to standardized and hardware-informed formulations, and then discuss traffic- and sleep-aware evaluation methodologies, metric definitions, and learning-based control approaches. The goal is to clarify how these assumptions propagate into energy–performance conclusions and why they matter for the positioning of the present thesis. [12, 5]

## 2.2 Historical Evolution of BS Power Models

The literature on BS power-consumption modeling has evolved through a progressive refinement of three aspects: (i) *which* operating variables are treated as energy-relevant, such as load, radiated power, bandwidth, antenna activity, and component carriers; (ii) *how* BS operation is abstracted into states, from simple active/idle views to explicit multi-level sleep representations; and (iii) *how* model parameters are obtained, ranging from empirical fits and standardized parameter sets to measurement-driven hardware decomposition. This evolution reflects a recurring trade-off between portability and analytical simplicity on the one hand, and hardware fidelity and configuration sensitivity on the other. In practice, the selected abstraction level directly affects the estimated benefit of radio adaptations and sleep strategies, and therefore the robustness of the resulting energy-performance conclusions. [5, 1, 18]

**Legacy affine models: EARTH (2011).** A widely used early reference is the EARTH model, introduced to provide a simple yet practical mapping from radio operation to supply power. In its most common form, it captures the empirical observation that BS consumption is not energy-proportional: a non-negligible baseline persists even at low load, while an additional component grows with radiated power and with the number of active transceiver chains. The model is therefore typically expressed as an affine relation in the load/output-power domain, with parameters selected by BS class and, in some variants, complemented by a small number of coarse low-power states. This abstraction proved effective for large-scale system studies in 4G and early 5G, where the main objective was to compare deployments and estimate first-order savings rather than to isolate the contribution of bandwidth scaling, carrier shutdown, or massive-MIMO hardware blocks. Its main limitation is precisely this simplicity: modern AAUs expose

additional energy knobs whose effect is difficult to represent through a purely affine view. [5]

**Enhanced lumped parametrizations: GreenTouch (2015).** As sustainability-oriented studies increasingly targeted network-wide transformation, models were enriched to support large-scale "what-if" evaluations while remaining computationally lightweight. The Green Meter study of the GreenTouch consortium is representative of this phase: it combines ambitious efficiency targets with scenario-based modeling that aggregates per-site and network-level effects across traffic growth, architectural evolution, and energy-saving techniques. From the perspective of BS modeling, the emphasis is on richer parametrization and multiple operating points, such as traffic regimes and equipment classes, rather than on direct hardware-block decomposition. This preserves portability and enables end-to-end accounting, but it also keeps the model relatively black-box in nature: inputs remain high-level operating conditions, and calibration is intended for macro evaluations rather than for short time-scale control analysis. [9]

**Standardized NR models: 3GPP TR 38.864 (Release 18, 2024).** Recent standardization efforts moved beyond simple "active + idle" abstractions by introducing a structured BS state machine and a standardized split between static and dynamic power components for New Radio (NR). In TR 38.864, BS operation is described through explicit active downlink/uplink states together with hierarchical sleep modes, such as micro, light, and deep sleep, entered during idle intervals and associated with steady-state power levels as well as transition penalties, including wake-up delay and optionally extra energy. Within the active state, the model separates load-independent contributions, such as baseline circuitry and overheads, from a dynamic component that scales with the active configuration and radio usage. This dynamic part explicitly accounts for antenna activity and for joint scaling effects related to bandwidth usage and transmit-power density through an efficiency term. The main strength of this family is comparability: by fixing the functional form and providing reference parameter sets for representative equipment categories, 3GPP models enable reproducible sensitivity analyses and more consistent cross-study benchmarking. [1]

**Hardware-based and measurement-driven models: AAU models (2022).** In parallel, the literature increasingly recognized that the sensitivity of energy outcomes to radio adaptations depends on whether the model explicitly represents the dominant hardware blocks of modern AAUs. Piovesan et al. [18] provide a representative approach that combines a measurement-driven learning model trained on a large dataset of commercial AAUs with a compact analytical formulation derived from it in order to preserve interpretability and tractability. The resulting analytical model decomposes active power into baseline, baseband processing, transceiver-chain terms, amplifier overheads, and a load-dependent output-power contribution scaled by efficiency. This yields a more direct mapping between operating decisions and physical consumption drivers. Importantly, the same structure also supports the representation of low-power operation as partial component deactivation, for example through symbol, channel, or carrier shutdown and deep dormancy, thereby aligning the model with the granularity required by modern energy-management policies. The main limitation of this family is that calibration and generalization become critical issues: although measurement-driven fitting improves fidelity for the targeted equipment class, parameters may need re-fitting across vendors, radio configurations, and implementation generations. [18]

**Implications and emerging requirements.** Two trends clearly emerge from recent 5G/6G energy-efficiency studies. First, standardized models such as TR 38.864 provide a common baseline for comparing energy knobs across time, power, antenna, and frequency domains under shared assumptions, which is particularly valuable when the goal is benchmarking or ranking mechanisms. Second, hardware-based and measurement-grounded models are increasingly preferred when fine-grained control or realistic quantification of AAU savings is required, especially in massive-MIMO and multi-carrier settings where transceiver and amplifier overheads are structurally important. At the same time, recent analytical frameworks emphasize that advanced sleep modes cannot be evaluated in isolation from traffic, queueing, and state-transition constraints, which strengthens the coupling between the adopted power model and the surrounding evaluation methodology. [10, 6, 7, 18]

Overall, the historical trajectory of BS power models reflects a shift from portable

affine approximations to standardized state-based formulations and, more recently, to hardware-informed and measurement-grounded decompositions. These model families are not interchangeable: they encode different assumptions on which power terms scale with radio adaptations and which remain as non-removable floors. As a consequence, model choice can materially alter the interpretation of energy-saving mechanisms and their energy-performance trade-offs. This motivates the next step of the state of the art, which turns to traffic-driven and sleep-aware evaluation methodologies, since the benefits of both radio adaptations and multi-level sleep critically depend on how activity periods, idle opportunities, and transition constraints are generated by the traffic process and by the system-level control logic. [1, 7]

A compact comparison of these model families, in terms of main inputs, sleep/state support, and calibration or validation intent, is summarized in Table 1.

| **Family** | **Main inputs** | **Sleep / states** | **Calibration & typical use** |
|---|---|---|---|
| Legacy affine (EARTH) [5] | Load, radiated/output power, active TRX chains | Few coarse idle/sleep states | Empirical fit; simple and portable; suitable for network-level comparison. |
| Enhanced lumped (GreenTouch) [9] | Scenario operating points, traffic regime, equipment class | Aggregate duty-cycling / occupancy | Scenario-level accounting; roadmap and upper-bound studies; not intended for fine-grained control. |
| Standardized NR (3GPP TR 38.864) [1] | Antenna activity, bandwidth use, TX-power density, efficiency | Multi-level sleep (micro/light/deep) with transitions | Reference parameterization; reproducible baseline; benchmarking and comparative evaluation. |
| Hardware-based AAU (Piovesan) [18] | BB/RF/PA component terms, carriers, TX power, RF chains | Partial deactivation and deep dormancy | Measurement-driven calibration; higher fidelity; may require recalibration across platforms. |

*Table 1: Representative families of BS power-consumption models and the main assumptions that differentiate them.*

## 2.3 Traffic-Driven Evaluation: Sleep Modes, Scenarios, and Variable Resources

Energy-saving mechanisms in the RAN are intrinsically time-dependent: both the *opportunity* to enter low-power states and the *cost* of exiting them depend on how traffic arrives, how queues evolve, and how quickly the BS can adapt its operating point. Static comparisons at fixed load can therefore be misleading whenever the considered mechanism affects service rate or introduces transition overheads. For example, bandwidth or antenna reductions may decrease instantaneous active-state power while simultaneously increasing transmission time and backlog, thereby shrinking idle intervals and reducing the fraction of time spent in sleep. This motivates traffic-driven evaluation frameworks in which BS operation is represented as a time-evolving process with explicit activity, queueing, and sleep transitions.

**From load snapshots to time-domain evaluation.** A time-domain view is typically built by combining: (i) a traffic model, including the arrival process and packet-size assumptions; (ii) a service model, mapping the instantaneous radio configuration and channel conditions to a departure rate; and (iii) an operational policy deciding when and how to transition between active and low-power states. Within this perspective, the key quantity is not only active-state power, but also the *active-time fraction*, or more generally the state occupancies induced by traffic burstiness and control decisions. Studies differ in how realistic the traffic process is and in the time scale at which sleep transitions are represented, but they share the same methodological point: energy outcomes depend on temporal dynamics, not only on static operating points.

**Sleep modes and ASM: taxonomy, decisions, and transition constraints.** A substantial body of work identifies sleep as a dominant lever for reducing the baseline, or load-independent, consumption of BSs. Recent surveys organize this space by distinguishing (i) *sleep depth*, that is, which hardware blocks are deactivated and which remain ready; (ii) the *decision logic*, ranging from reactive timers and thresholds to predictive policies; and (iii) the *time scale* of transitions, since deeper sleep generally implies longer wake-up delays and tighter operational constraints. [20] In this thesis,

the term ASM refers to the advanced end of this spectrum, namely multi-level sleep with non-negligible wake-up delays and potentially hierarchical shutdown, for example at component or carrier level, whose benefit must be assessed jointly with QoS impact. Representative recent studies highlight that, in the transition towards 6G, ASM design must explicitly account for both transition penalties and traffic burstiness. [6]

**Traffic and scenario choices: what is "load" in a dynamic setting?** Dynamic evaluations typically rely on stochastic arrivals, such as Poisson-like packet processes, aggregated flows, or bursty ON/OFF variants, and emulate diurnal or scenario-specific load variations by changing the arrival intensity over time. The scenario definition then specifies user distributions, channel models, and scheduling assumptions that determine how traffic translates into service rate. Even when simplified traffic models are adopted, the time-domain coupling remains essential: queueing and state transitions introduce memory, so average-load snapshots cannot capture how often the BS can exploit idle gaps. This is also why many studies report not only average power, but also state occupancy and delay indicators, in order to distinguish whether savings arise from lower active-state power or from increased sleep time. [20, 11]

**Policies for entering and exiting sleep.** Early, and still widely used, sleep-control policies are reactive: inactivity timers and threshold rules trigger a transition when the queue is empty, or sufficiently small, for a given duration, while wake-up is activated by new arrivals or by anticipated demand. [11] More advanced approaches incorporate prediction or context-awareness, for example through traffic forecasting or mobility hints, but the same fundamental constraint remains: if the wake-up time is not negligible with respect to the traffic time scale, overly aggressive sleeping can increase delay and induce oscillatory behavior. For this reason, evaluations often distinguish micro, light, and deep sleep behaviors, or more generally shallow versus deep ASM, and explicitly account for wake-up delays when computing QoS penalties. [20, 6]

**Short time-scale evaluation: the role of queues.** When sleep transitions occur on the order of milliseconds to seconds, queueing becomes the natural abstraction for coupling arrival burstiness, service variability, and sleep overhead. Recent analytical

approaches propose queueing-based models for short time-scale sleep evaluation in 5G BSs, explicitly incorporating wake-up times and the impact of state-dependent service. [7] The methodological value of this line of work lies in making the energy–delay coupling explicit: deeper sleep reduces idle power, but it also introduces effective service interruptions during wake-up, which may amplify tail delays under bursty arrivals. This reinforces a broader point that is central to the present thesis: mechanisms that appear attractive under long-term averaging may become less favorable, or even counterproductive, once transition overheads are resolved at short time scales.

**Variable resources beyond sleep: multi-domain adaptations.** In parallel with sleep-oriented studies, a related literature investigates *resource adaptation* as an energy lever during active periods by adjusting the radio operating point according to traffic and channel conditions. Typical knobs include bandwidth usage, for example through active bandwidth parts or the number of allocated RBs; antenna or RF-chain activity; transmit-power settings; and, in multi-carrier systems, Component Carrier (CC) activation. [10] Recent 6G-oriented studies explicitly analyze the combined potential of multi-domain adaptations and advanced sleep, showing that the resulting gains are not automatically additive: reducing active resources may lower instantaneous power while increasing time-to-serve traffic, thereby shrinking the idle windows that enable deeper ASM. [10, 6] This interaction motivates evaluation frameworks that jointly represent active-power scaling and the time-domain creation of sleep opportunities.

**Time horizons and reporting choices.** A recurring theme across the above strands is the coexistence of at least two relevant time scales. At *short time scales* (ms–s), queueing and transition constraints dominate the QoS impact of ASM, making metrics such as tail delay or wake-up frequency particularly informative. At *longer time scales* (minutes–hours), the main question becomes how energy-saving mechanisms reshape energy accounting over traffic profiles, such as diurnal cycles, and how robust the resulting conclusions are across scenario assumptions. A compact summary of the main ingredients of traffic-driven evaluation and of their typical reporting implications is provided in Table 2. The next section builds on these notions by reviewing the most common energy–performance metrics used in the literature.

| Evaluation ingredient | Typical choices | What it captures | Representative refs. |
|---|---|---|---|
| Traffic model & scenario | Stochastic arrivals (Poisson/ON–OFF), time-varying rate, user/channel model | Idle gaps, queue build-up, sensitivity to burstiness and load profile | [20, 11] |
| Sleep/ASM model | Multi-level sleep (micro/light/deep), wake-up time, optional transition energy | State occupancy and energy–delay trade-off due to transitions | [20, 6] |
| Policy class | Reactive timers/thresholds vs. predictive rules, hysteresis/cooldown | Sleep timing, stability, and oscillations under bursty traffic | [11, 20] |
| Resource adaptations | Bandwidth/RBs, antenna/RF-chain activity, TX power, CCs on/off | Active-power scaling and impact on time-to-serve and idle windows | [10] |
| Time horizon & KPIs | ms–s: queue/tail delay/wake-ups; min–h: average power, energy/bit, occupancy | Short-term QoS effects vs. long-term energy accounting across traffic profiles | [7, 20] |

*Table 2: Key ingredients of traffic-driven evaluations for sleep/ASM and resource adaptations, and the main aspects they reveal.*

## 2.4 Metrics in the Literature: Definitions and Common Choices

Energy-efficiency studies in cellular networks typically report a relatively small set of energy- and QoS-related Key Performance Indicators (KPIs), yet identical metric names often conceal different normalization, aggregation, and baseline choices. This is methodologically important because the interpretation of an energy-saving mechanism depends not only on the adopted power model, but also on how its outcome is measured and reported.

At the BS or cell level, the most common energy-side metrics are (i) average power, or average energy per unit time, (ii) energy per delivered bit (J/bit), and (iii) energy efficiency (bit/Joule). Although closely related, these indicators emphasize different aspects depending on what is held fixed, for example offered load versus carried throughput, or fixed-rate service versus adaptive service. At the network level,

broader scenario-wide accounting is also common, especially in roadmap-style studies that aggregate across sites and time-varying traffic conditions. [5, 9]

On the performance side, latency is reported in forms that depend strongly on the evaluation scope. In sleep-mode and ASM studies, it is often interpreted as packet delay driven by queueing and service dynamics, with additional waiting introduced by wake-up time and by state-dependent service interruptions. Other works adopt a broader end-to-end view and treat radio delay as one component of the overall QoS budget. To capture the effect of burstiness and transitions, the literature frequently reports not only mean delay, but also tail statistics, such as p95 or p99, and violation indicators, such as the fraction of packets exceeding a given threshold. [11, 19]

Beyond the metric name itself, several recurring choices affect comparability and may alter qualitative conclusions:

- **Energy normalization:** averaging per unit time (average power/energy) versus per delivered data (energy/bit, Energy Efficiency (EE)) can produce different rankings when a mechanism changes service time or carried throughput.
- **Throughput definition:** offered load, scheduled rate, and delivered goodput are not equivalent; energy per bit depends directly on which quantity is used in the denominator.
- **Time horizon and aggregation:** long-term averages may hide short time-scale effects such as wake-up bursts or tail-delay penalties, whereas short windows may overemphasize transient behavior.
- **Transition accounting:** including or excluding wake-up time and, when modeled, transition energy can materially change the perceived benefit of deeper Sleep Mode (SM)/ASM.
- **Baseline dependence:** many studies report relative metrics, such as ratios or savings, against a reference configuration; the conclusions depend on how that baseline is defined and on the scenario under which it is measured.
- **Hidden state composition:** reporting only average power can obscure whether savings come from lower active-state power or from increased sleep time; state

occupancy is often needed as a diagnostic indicator.

Recent standardization efforts and evaluation guidelines increasingly emphasize consistent KPI definitions, comparable baselines, and state-aware reporting conventions, especially in the presence of sleep modes. [1] In this thesis, energy and delay are therefore interpreted jointly, and the exact definitions, averaging horizons, and normalization rules adopted in the experimental pipeline are specified in Section 3.2.

## 2.5 Reinforcement Learning for Energy Management

The growing configuration space of modern BSs, including multi-level sleep, antenna or RF-chain activation, multi-carrier operation, and traffic-adaptive resource allocation, has motivated learning-based control as an alternative to purely rule-based policies. In this context, RL is typically formulated as a sequential decision problem in which a controller observes the network state, selects an action, and receives a reward that trades off energy consumption and QoS. Compared with static optimization, RL is naturally suited to time-domain effects such as traffic burstiness, queue memory, and wake-up constraints, and can adapt policies to non-stationary operating conditions. [19, 4]

**MDP abstraction: state, action, reward.** Most approaches cast energy management as a Markov decision process (MDP), where the state provides a compact description of traffic and radio operation, for example through load indicators, queue or backlog proxies, channel-quality summaries, and the current sleep or active mode. Actions typically include at least one sleep-related decision, such as shallow versus deep sleep, ASM level, or wake/sleep commands, often combined with radio-configuration knobs including antenna or RF-chain activity, transmit-power settings, bandwidth or RBs allocation, and, in multi-carrier systems, CC activation. The reward is commonly defined as a weighted combination of energy and QoS terms, where QoS penalties account for delay, drops, or violation events caused by insufficient service or by wake-up interruptions. This formulation makes the control objective explicit, but it also shows that learned behavior depends on normalization choices and on the specific latency indicator included in the reward. [20, 11]

**Single-agent RL for multi-sleep control.** A substantial line of work considers a single agent controlling the sleep behavior of a BS, or of a small cell, under dynamic traffic. Early contributions employ tabular methods, such as Q-learning, with discretized states and actions to learn when to enter deeper sleep levels while keeping delay within acceptable bounds. This approach is particularly suitable when the action space is limited and interpretability or stability is prioritized. More recent works enrich the state representation and sleep logic while still focusing on the central energy-delay coupling introduced by wake-up constraints. [4, 3]

**Deep RL and risk/constraint handling.** As the number of controllable knobs grows, deep value-based methods such as Deep Q-Networks (DQN) are often adopted to generalize across larger state spaces. In this regime, constraint handling becomes especially important: aggressive sleeping or overly restrictive radio adaptations can produce unstable behavior under bursty traffic or rare high-load events. Representative approaches introduce risk-aware reward shaping and auxiliary mechanisms to mitigate unsafe actions, for example by relying on digital-twin-inspired components to anticipate outcomes and penalize high-risk decisions. [14] Complementary deep RL studies also investigate traffic-adaptive sleep policies and radio-resource adaptation, showing that learning can outperform static strategies when the best operating point depends jointly on queue state and radio conditions. [13, 17]

**Multi-agent settings: interaction, offloading, and coordination.** In multi-BS deployments, decisions are coupled through interference, user association, and offloading. This motivates multi-agent RL formulations in which each BS, or sector, acts as an agent and learning is coordinated through paradigms such as centralized training with decentralized execution (CTDE). Recent work in massive-MIMO scenarios combines sleep-mode selection with antenna configuration and user offloading, showing that cooperation can improve both energy outcomes and QoS under dynamic traffic, while also making learning more sensitive to non-stationarity and partial observability. [22]

**Model sensitivity: the reward depends on the power model.** Across the above families, the energy term in the reward is typically computed through a BS power-

consumption model that maps the chosen operating point, including sleep level, active configuration, and load, into supply power. This makes the learned policy inherently *model-sensitive*: if the power model changes, for example from a standardized state-based formulation to a measurement-driven hardware-based one, the relative attractiveness of actions may change as well, potentially leading to different policies under identical traffic conditions. This observation motivates careful alignment between control design, the assumed power model, and the validation methodology, especially when conclusions are transferred across equipment classes. [1, 18]

A compact summary of the main families of RL-based approaches for BS energy management, together with their control scope and main methodological issues, is provided in Table 3.

| **RL family** | **Typical control scope** | **Main methodological issue** | **Representative refs.** |
|---|---|---|---|
| Tabular single-agent RL | Sleep-mode selection for a single BS or small cell, with limited action space and explicit delay constraints | Good interpretability and stable benchmarking, but limited scalability and coarse state/action discretization | [19, 4, 3] |
| Deep single-agent RL | Joint sleep and radio-resource adaptation with richer state representation and larger action space | Constraint handling, reward shaping, and robustness under bursty traffic or rare high-load events | [13, 14, 17] |
| Multi-agent RL | Coordinated control across multiple BSs or sectors, including offloading and interference-coupled decisions | Non-stationarity, partial observability, and coordination complexity | [22] |
| Model-sensitive reward design | Energy-aware control under a chosen BS power-consumption model | The learned policy depends on how the reward computes the energy term, including power-model assumptions and KPI normalization | [1, 18] |

*Table 3: Main families of RL-based approaches for BS energy management and their main methodological implications.*

Overall, and consistently with the summary reported in Table 3, RL provides a flexible framework for handling time-varying traffic and multi-knob energy management, ranging from tabular single-agent control to deep and multi-agent approaches. At the same time, important open issues remain, including reward and metric consistency, constraint satisfaction, robustness to scenario shifts, and the sim-to-real gap. These issues help position the present thesis, where RL is considered not as the main contribution, but as a preliminary extension used to illustrate how power-model assumptions propagate into learning-based control. [21]

## 2.6 Open Issues and Thesis Positioning

The state of the art reviewed in Sections 2.2–2.5 shows substantial progress in both BS power modeling and energy-management strategies. At the same time, several methodological gaps and sources of uncertainty remain, affecting the reliability, comparability, and interpretation of reported energy gains:

- **Model alignment and calibration across equipment classes.** Standardized models provide portable parameter sets and a shared functional form, whereas hardware-grounded models can better reflect internal consumption drivers but may require platform-specific calibration. This creates an inherent tension between comparability and fidelity, and raises the question of how conclusions generalize across vendors, AAUs, and technology generations when the same control knob maps to different internal power budgets. [1, 18]

- **Limited cross-model comparisons under controlled inputs.** Many studies adopt a single reference power model, or a single parameter set, and treat model choice as a fixed background assumption. As a result, it is often unclear whether reported gains are optimistic or conservative, and which qualitative trends remain robust when the underlying power model is changed.

- **Dynamic realism and time-scale dependence.** Energy outcomes depend on traffic burstiness, queue build-up, and the creation of idle gaps that enable sleep. Long-horizon averages can hide short time-scale effects, such as wake-up bursts

and tail-delay penalties, whereas short-horizon analyses may overemphasize transients. This motivates evaluation pipelines that explicitly couple traffic, service dynamics, and state occupancy when assessing energy–QoS trade-offs. [20, 7]

- **Advanced sleep complexity and transition constraints.** Multi-level sleep (ASM) introduces non-negligible wake-up delays and, potentially, transition-energy costs. The benefit of deeper sleep therefore depends on how often it can be entered and on the resulting delay distribution; ignoring transition constraints can substantially overestimate achievable savings. [6]

- **Learning-based control sensitivity to modeling assumptions.** In RL-based energy management, the reward includes an energy term computed through the adopted power model. This makes the model part of the environment definition: changing the power model changes the learning signal and can lead to different learned policies and operating regimes, especially in multi-cell settings where actions are coupled through interference and offloading. [14, 22]

This thesis addresses a subset of these gaps by focusing on *controlled comparability*. Specifically, it performs a side-by-side evaluation of a standardized NR power model and a measurement-grounded AAU-oriented model within a unified pipeline in which scenario assumptions, traffic generation, and KPI definitions are kept fixed while only the BS power model is changed. The analysis combines (i) a static sensitivity study to isolate model-dependent scaling behavior under repeatable configurations, (ii) a traffic-driven time-domain evaluation with multi-level sleep to quantify the interaction between instantaneous power, service time, and sleep opportunities, and (iii) a learning-based extension to illustrate how model choice propagates into the control loop through model-dependent energy feedback.

Within this scope, the thesis positions power-model selection as a methodological variable in its own right. Its goal is not only to compare two formulations, but to clarify which reported energy-saving trends remain robust across modeling granularities and which are primarily induced by the selected power-consumption abstraction.

# 3 Methodology

This chapter describes the methodological framework used to quantify how the choice of a BS power-consumption model affects energy-efficiency estimates and energy–performance trade-offs in 5G/6G RANs. The central design principle is *controlled comparability*: all experiments share the same scenario assumptions, traffic generation, and KPI definitions and normalization conventions, while the BS power-consumption model is the only component that changes. As a result, deviations in the outcomes can be attributed to differences in the power-model formulation rather than to mismatched inputs, traffic realizations, or aggregation conventions. Both models are evaluated through a common input interface to isolate formulation effects.

Under this unified setup, the study combines three complementary evaluation pipelines. A *static* pipeline performs one-at-a-time sweeps of radio-adaptation variables to produce KPI-driven trade-off curves under a repeatable link abstraction. A *dynamic* pipeline relies on a discrete-time simulator in which stochastic arrivals and queueing interact with a multi-level SM state machine, enabling the estimation of long-run averages (e.g., dynamic energy per served bit) as well as transient effects (e.g., wake-up delays and sleep-state occupancy). Finally, the framework is instantiated in a multi-BS RL pipeline to assess how model-dependent energy feedback affects learned control policies and the resulting energy–performance outcomes.

## 3.1 System Model and Assumptions

This section defines the reference single-BS Downlink (DL) scenario and the assumptions shared across the methodological pipelines in this thesis. The objective is to fix a common system description (topology, large-scale channel/noise abstraction, and controllable radio variables) so that differences in energy-related outcomes can be attributed to the selected BS power-consumption model rather than to inconsistencies in the simulated setting. Unless otherwise stated (Section 3.6), all experiments refer to the single-cell configuration illustrated in Figure 1. The same control tuple $\mathbf{s} = (s_a, s_f, s_p)$ is applied to all compared power models; the static and dynamic pipelines differ only in how link performance and system operation are evaluated (Sections 3.4 and 3.5).

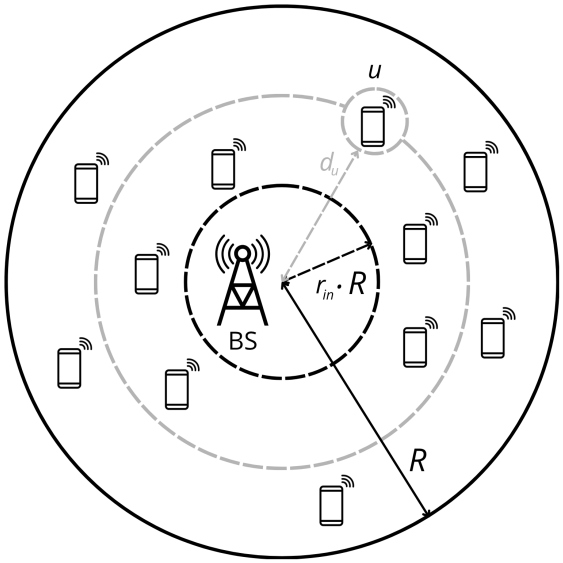


*Figure 1: Single-BS single-cell scenario. UEs are placed in an annulus centered at the BS, with outer radius R and inner radius $R_{\text{in}} = r_{\text{in}}R$. The distance of user u from the BS is denoted by $d_u$.*

### 3.1.1 Topology and user distribution (Fig. 1)

A single BS is placed at the origin and serves $U$ UEs located within a circular cell of radius $R$. To avoid unrealistically short links, users are placed in an annulus with inner radius $R_{\text{in}} = r_{\text{in}}R$ and outer radius $R$. In the implementation adopted here, the polar angle is sampled uniformly and the radial coordinate is sampled uniformly in $[R_{\text{in}}, R]$:

$$\theta_u \sim \mathcal{U}(0, 2\pi), \qquad r_u = \big(r_{\text{in}} + (1 - r_{\text{in}})\xi_u\big)R, \;\; \xi_u \sim \mathcal{U}(0, 1), \tag{1}$$

and mapped to Cartesian coordinates $\mathbf{x}_u = [r_u \cos\theta_u, \; r_u \sin\theta_u]^\top$. The distance between user $u$ and the BS is $d_u = \|\mathbf{x}_u\|$.

### 3.1.2 Channel and noise modeling

Time-domain simulations rely on a large-scale channel abstraction derived from a distance-dependent pathloss model with an additive shadowing term. Distances are lower-bounded to 1 m for numerical stability, i.e., $D_u = \max(d_u, 1\text{ m})$. The pathloss

(in dB) at carrier frequency $f_c$ (in GHz) is modeled as

$$\mathrm{PL}_{\mathrm{dB}}(D_u) = 16.9 \log_{10}(D_u) + 32.8 + 20 \log_{10}(f_c) + X_{\mathrm{sh}}, \tag{2}$$

where $X_{\mathrm{sh}}$ denotes the shadowing term in dB. In the baseline configuration, the shadowing dispersion is parameterized by the value reported in Table 4, while a fixed realization is used within each simulation run. The corresponding linear channel gain is then

$$h_u = 10^{-\mathrm{PL}_{\mathrm{dB}}(D_u)/10}. \tag{3}$$

Thermal noise power is computed from a noise spectral density $N_0$ (in dBm/Hz), receiver noise figure NF (in dB), and bandwidth $B$ (in Hz) as

$$N_{\mathrm{dBm}}(B) = N_0 + 10 \log_{10}(B) + \mathrm{NF}, \qquad N(B) = 10^{N_{\mathrm{dBm}}(B)/10} \cdot 10^{-3} \ \ [\mathrm{W}]. \tag{4}$$

In the static link proxy, noise is computed equivalently via a $kTB$ formulation using the configured temperature $T$ and NF, while the dynamic simulator uses the $(N_0, B, \mathrm{NF})$ expression above; both yield a consistent thermal-noise scaling with bandwidth.

Inter-cell interference is not explicitly modeled (single-cell abstraction). Accordingly, the link-quality proxy used throughout the study is effectively an Signal-to-Noise Ratio (SNR), with model-dependent energy differences originating from the BS power-consumption mapping rather than from interference dynamics.

### 3.1.3 Radio adaptation variables

The BS operation is parameterized by three dimensionless scaling factors that represent the energy-saving adaptations considered in this work:

- **Antenna activity** $s_a \in (0, 1]$: fraction of active transmit chains.
- **Bandwidth utilization** $s_f \in (0, 1]$: fraction of the system bandwidth allocated to the DL.
- **Transmit-power scaling** $s_p \in (0, 1]$: multiplicative factor applied to the DL transmit power.

Let $N_{\text{tx}}$ be the total number of transmit chains. The number of active chains is obtained by rounding to the nearest integer and enforcing at least one active chain:

$$N_{\text{act}}(s_a) = \max(1, \text{round}(s_a N_{\text{tx}})) \,. \tag{5}$$

The effective bandwidth is $B(s_f) = s_f B_{\text{sys}}$, where $B_{\text{sys}}$ denotes the system bandwidth. The total radiated transmit power in the static proxy is

$$P_{\text{tx}}(s_a, s_p) = N_{\text{act}}(s_a)\, P_{\text{tx,ant}}\, s_p, \tag{6}$$

with $P_{\text{tx,ant}}$ the per-antenna reference transmit power. In the static link abstraction, spatial multiplexing is represented through an effective number of layers,

$$L(s_a) = \min(L_{\max},\, N_{\text{act}}(s_a)) \,, \tag{7}$$

where $L_{\max}$ is the configured maximum number of layers.

In the dynamic simulator, the large-scale SNR is computed per user from the baseline (full-bandwidth, full-activity) reference and then scaled consistently with the adaptation tuple: reducing bandwidth decreases noise proportionally to $s_f$, while antenna activity and transmit-power scaling increase the available radiated power proportionally to $s_a s_p$. This yields an effective per-user SNR scaling factor proportional to $(s_a s_p)/s_f$, used to drive the per-slot service rates in Section 3.5.

### 3.1.4 Reference parameters

Tables 4 and 5 summarize the baseline parameters used throughout the single-BS evaluations, unless otherwise stated. The discrete sweep sets for $(s_a, s_f, s_p)$ follow the configuration used in the experimental campaign (Section 3.4).

## 3.2 KPIs and Normalization

This section defines the KPIs used throughout the thesis to quantify energy consumption and service performance, together with the normalization conventions adopted for fair comparisons across BS power-consumption models. For any experiment, the same

*Table 4: Baseline scenario, channel, and time-domain parameters (single-BS).*

| **Parameter** | **Value** |
|---|---|
| Number of users $U$ | 50 |
| Cell radius $R$ | 500 m |
| Inner radius fraction $r_{\text{in}}$ | 0.125 |
| Carrier frequency $f_c$ | 3.5 GHz |
| Shadowing standard deviation | 8 dB |
| Noise density $N_0$ | $-174$ dBm/Hz |
| Noise figure NF | 7 dB |
| Temperature $T$ | 300 K |
| Antenna gain $G_{\text{AE}}$ / UE gain $G_{\text{UE}}$ | 8 dBi / 5 dBi |
| Total TX chains $N_{\text{tx}}$ | 64 |
| Maximum layers $L_{\max}$ | 8 |
| System bandwidth $B_{\text{sys}}$ | 100 MHz |
| Per-antenna TX power $P_{\text{tx,ant}}$ | 5 W |
| Packet size | 50 kB |
| Slot duration $\Delta t$ | 0.5 ms |
| Simulation horizon | 100 s |
| Traffic rate per user $\lambda$ | 20 pkt/s |

*Table 5: Baseline sleep-mode thresholds (single-BS).*

| **Parameter** | **Value** |
|---|---|
| $\tau_{\text{micro}}$ | 2 ms |
| $\tau_{\text{light}}$ | 10 ms |
| $\tau_{\text{deep}}$ | 50 ms |

radio configuration $\mathbf{s} = (s_a, s_f, s_p)$ and (where applicable) the same traffic realization are used across models; the only model-dependent component is the mapping from common inputs (and operational state) to power consumption.

### 3.2.1 Baseline and normalization conventions

Let $\mathbf{s} = (s_a, s_f, s_p)$ denote the radio adaptation vector defined in Section 3.1.3. The baseline configuration is

$$\mathbf{s}_0 = (1, 1, 1), \tag{8}$$

corresponding to full antenna activity, full bandwidth, and nominal transmit power. Given any scalar metric $X(\mathbf{s})$, its ratio and relative change with respect to $\mathbf{s}_0$ are

$$\widetilde{X}(\mathbf{s}) \triangleq \frac{X(\mathbf{s})}{X(\mathbf{s}_0)}, \qquad \Delta X(\mathbf{s})[\%] \triangleq 100\big(\widetilde{X}(\mathbf{s}) - 1\big). \tag{9}$$

For metrics where "lower is better" (e.g., latency and energy per bit), improvements are reported as a *gain* (saving):

$$G_X(\mathbf{s})[\%] \triangleq 100\left(1 - \frac{X(\mathbf{s})}{X(\mathbf{s}_0)}\right). \tag{10}$$

For metrics where "higher is better" (e.g., rate), degradations are reported as a *loss*:

$$L_X(\mathbf{s})[\%] \triangleq 100\left(1 - \frac{X(\mathbf{s})}{X(\mathbf{s}_0)}\right). \tag{11}$$

These conventions are used consistently across static sweeps and time-domain simulations.

### 3.2.2 Static evaluation KPIs

Static experiments construct energy-performance trade-off curves by sweeping one adaptation variable at a time while keeping the others fixed (Section 3.4). For each configuration $\mathbf{s}$, service performance is summarized through a deterministic link-level proxy: the rate proxy $R_{\text{norm}}(\mathbf{s})$ and the corresponding transmission-time proxy $T_{\text{tx}}(\mathbf{s})$ are defined by the static abstraction of Section 3.4 (with $s_t = 1$ in static sweeps).

Let $P_{\text{active}}(\mathbf{s})$ denote the BS `ACTIVE` downlink power returned by the selected power model under configuration $\mathbf{s}$ (Section 3.3). The active-only energy per transmitted bit is

$$E_{\text{bit}}^{\text{act}}(\mathbf{s}) \triangleq \frac{P_{\text{active}}(\mathbf{s})\, T_{\text{tx}}(\mathbf{s})}{N_{\text{bits}}} \quad [\text{J/bit}], \tag{12}$$

where $N_{\text{bits}}$ is the fixed packet size (in bits) used in the proxy. Thus, model differences in $E_{\text{bit}}^{\text{act}}(\mathbf{s})$ originate solely from $P_{\text{active}}(\mathbf{s})$.

The primary static indicators used to build trade-off curves are:

- **Rate loss**:

$$L_R(\mathbf{s})[\%] = 100\left(1 - \frac{R_{\text{norm}}(\mathbf{s})}{R_{\text{norm}}(\mathbf{s}_0)}\right). \tag{13}$$

- **Latency increase**:

$$\Delta T_{\text{tx}}(\mathbf{s})[\%] = 100\left(\frac{T_{\text{tx}}(\mathbf{s})}{T_{\text{tx}}(\mathbf{s}_0)} - 1\right). \tag{14}$$

- **NES gain (static)**:

$$G_{NES}^{\text{static}}(\mathbf{s})[\%] \triangleq 100\left(1 - \frac{E_{\text{bit}}^{\text{act}}(\mathbf{s})}{E_{\text{bit}}^{\text{act}}(\mathbf{s}_0)}\right). \tag{15}$$

### 3.2.3 Dynamic (time-domain) KPIs and diagnostics

Dynamic simulations (Section 3.5) produce time traces of instantaneous BS power, served traffic, packet completion times, and sleep-state evolution. Let $T_{\text{sim}}$ be the simulation horizon and $\Delta t$ the slot duration. Denote by $P(t)$ the slot-level power consumption (including sleep states) in slot $t$, and let $\{E_{\text{wake},k}\}$ be the set of one-off wake-up energies incurred when transitioning from a sleep level to `ACTIVE`. The total consumed energy is

$$E_{\text{tot}} \triangleq \sum_{t=1}^{T_{\text{sim}}/\Delta t} P(t)\,\Delta t \;+\; \sum_{k} E_{\text{wake},k} \quad [\text{J}]. \tag{16}$$

Let $B_{\text{srv}}$ be the total number of served bits during the simulation (including partially served packets). The dynamic energy per served bit is defined as

$$E_{\text{bit}}^{\text{dyn}}(\mathbf{s}) \triangleq \frac{E_{\text{tot}}}{\max(B_{\text{srv}}, \epsilon)} \quad [\text{J/bit}], \tag{17}$$

where $\epsilon$ is a small positive constant used to avoid division by zero in degenerate cases.

**Energy-related KPIs.** The main energy KPIs reported from dynamic simulations are:

- **Average power**:

$$\overline{P}(\mathbf{s}) \triangleq \frac{E_{\text{tot}}}{T_{\text{sim}}} \quad [\text{W}]. \tag{18}$$

- **NES gain (dynamic)** relative to $\mathbf{s}_0$:

$$G_{NES}^{\text{dyn}}(\mathbf{s})[\%] \triangleq 100\left(1 - \frac{E_{\text{bit}}^{\text{dyn}}(\mathbf{s})}{E_{\text{bit}}^{\text{dyn}}(\mathbf{s}_0)}\right). \tag{19}$$

**Delay-related KPIs.** Let $\{\ell_i\}$ be the set of packet sojourn times (arrival-to-completion) for completed packets during the simulation, and let $N_{\text{pkts}}$ be the number of completed packets. Service performance is summarized through

$$\overline{\ell}(\mathbf{s}) \triangleq \frac{1}{N_{\text{pkts}}} \sum_{i=1}^{N_{\text{pkts}}} \ell_i, \qquad \ell_{95}(\mathbf{s}) \triangleq \text{percentile}_{95}\big(\{\ell_i\}\big), \tag{20}$$

i.e., mean latency and 95th percentile latency. Relative degradations with respect to $\mathbf{s}_0$ are reported via $\Delta\overline{\ell}[\%]$ and $\Delta\ell_{95}[\%]$ using the convention of Section 3.2.1.

**Activity and sleep-mode diagnostics.** To interpret energy and delay results, diagnostic indicators are extracted from the sleep-mode trace:

$$\phi_{\text{state}} \triangleq \frac{T_{\text{state}}}{T_{\text{sim}}}, \qquad \text{state} \in \{\texttt{ACTIVE}, \texttt{MICRO}, \texttt{LIGHT}, \texttt{DEEP}, \texttt{WAKING}\}, \tag{21}$$

where $T_{\text{state}}$ is the cumulative time spent in the given state. In particular, the **active-time fraction** is

$$s_t^{\text{dyn}} \triangleq \phi_{\texttt{ACTIVE}}. \tag{22}$$

Wake-up counters (e.g., number of activations from each sleep level) can be reported as supporting diagnostics when discussing transient behavior.

### 3.2.4 Aggregation rules and fairness across models

Static metrics are computed deterministically for each swept configuration $\mathbf{s}$ under the same link proxy and packet size; for a given $\mathbf{s}$, the only model-dependent quantity is $P_{\text{active}}(\mathbf{s})$. Dynamic metrics are obtained by running the same discrete-time simu-

lator under identical scenario parameters, user layout, and traffic generation settings; model dependence enters only through the power values returned by the selected model (in `ACTIVE` and, if model-specific, in sleep states). All KPI transformations (ratios, relative changes, gains, and losses) use the same baseline $\mathbf{s}_0$ and the conventions of Section 3.2.1.

## 3.3 Base Station Power-Consumption Models

This section presents the two BS power-consumption models compared in this thesis and describes their integration into a common evaluation pipeline. For any experiment, the scenario, radio configuration $\mathbf{s} = (s_a, s_f, s_p)$, and evaluation procedure are kept fixed, while the power model is swapped in a controlled manner. In particular, both models expose the `ACTIVE` downlink power $P_{\text{active}}(s_a, s_f, s_p)$ and the sleep-state parameters (steady powers and wake-up times) required by the shared sleep-mode controller used in the time-domain simulator (Section 3.5). The pipeline also supports a model-specific pre-processing hook, invoked once per configuration, before evaluating `ACTIVE` power.

### 3.3.1 3GPP TR 38.864 state-based model

The 3GPP TR 38.864 BS power model represents downlink consumption through a static contribution in `ACTIVE` mode and a dynamic contribution that scales with radio adaptations and load proxies. In this thesis, the `ACTIVE` DL power is modeled as

$$P_{\text{active}}^{\text{3GPP}}(s_a, s_f, s_p) = P_{\text{static}} + s_a \left( P_{\text{dyn,ante}} + \frac{s_f s_p}{\eta} P_{\text{dyn,joint}} \right), \tag{23}$$

where $P_{\text{static}}$ denotes the configured static `ACTIVE` component, $s_a$ is the antenna-activity factor, $s_f$ the bandwidth-utilization factor, $s_p$ the transmit-power scaling factor, and $\eta$ is an efficiency factor (kept constant in the considered configuration). The dynamic contribution is parameterized by the total active DL power at nominal operation, $P_{\text{active,DL}}$, and by a split factor $A \in [0, 1]$ that partitions the dynamic component into: (i) an antenna-dependent term and (ii) a joint term affected by bandwidth and power scaling. Letting

$$\Delta \triangleq \max\big(P_{\text{active,DL}} - P_{\text{static}},\, 0\big), \tag{24}$$

the dynamic parts are defined as

$$P_{\text{dyn,ante}} = A\Delta, \qquad P_{\text{dyn,joint}} = (1 - A)\Delta. \tag{25}$$

Equations (23)–(25) match the implementation used in the evaluation pipeline.

**Sleep-state parameters.** In addition to `ACTIVE` power, the model exposes steady-state sleep powers and wake-up times for the `MICRO`, `LIGHT`, and `DEEP` states. These parameters are consumed by the shared sleep-mode controller of the time-domain simulator (Section 3.5) to determine the instantaneous power $P(t)$ in non-active states.

### 3.3.2 AAU analytical hardware-driven model (Piovesan)

The AAU analytical model describes BS power as the sum of measurement-inspired hardware blocks plus a load-dependent RF-output contribution weighted by the PA efficiency. As illustrated in Figure 2, the `ACTIVE` power is decomposed into baseline/control and baseband contributions, per-chain transport and auxiliary terms, and an RF output contribution, as discussed by Piovesan et al. [18]

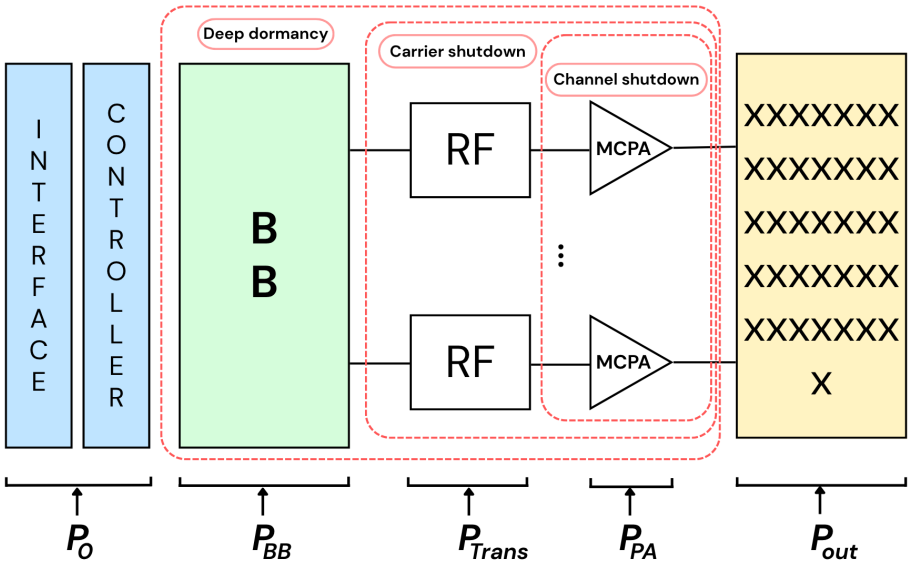


*Figure 2: AAU power-component decomposition underlying the hardware-driven model. The total `ACTIVE` power is expressed as the sum of (i) baseline/control contributions ($P_0$), (ii) baseband processing ($P_{\text{BB}}$), (iii) transport/RF-chain related terms ($P_{\text{Tran}}$), (iv) auxiliary PA-chain terms ($P_{\text{PA,aux}}$), and (v) a load-dependent RF output contribution ($P_{\text{out}}/\eta$). The dashed boxes conceptually relate these blocks to increasingly aggressive energy-saving actions (e.g., channel shutdown, carrier shutdown, and deep dormancy).*

In the formulation adopted here, the `ACTIVE` DL power is written as

$$P_{\text{active}}^{\text{AAU}}(s_a, s_f, s_p) = P_{\text{static}}^{\text{AAU}}(s_a) + P_{\text{dyn}}^{\text{AAU}}(s_a, s_f, s_p), \tag{26}$$

where the static (hardware) contribution is

$$P_{\text{static}}^{\text{AAU}}(s_a) = k_{\text{static}} \left( P_0 + P_{\text{BB}} + P_{\text{Tran}}(m_{\text{act}}) + P_{\text{PA,aux}}(m_{\text{act}}) \right). \tag{27}$$

Here $P_0$ and $P_{\text{BB}}$ account for baseline/control and baseband processing, respectively, while the remaining terms scale with the number of active RF chains (or equivalent active hardware units) $m_{\text{act}}$. Consistently with the implementation, the number of active chains is modeled as

$$m_{\text{act}}(s_a) = \max(1,\, \text{round}(s_a\, M_{\text{av}})), \tag{28}$$

where $M_{\text{av}}$ is the maximum available chains per transceiver. In the considered setup, $M_{\text{av}}$ matches the system-level number of transmit chains used to define antenna activity. The per-chain transport and auxiliary contributions are linear in $m_{\text{act}}$:

$$P_{\text{Tran}}(m_{\text{act}}) = T_{\text{bands}}\, m_{\text{act}}\, D_{\text{Tran}}, \qquad P_{\text{PA,aux}}(m_{\text{act}}) = m_{\text{act}}\, D_{\text{PA}}, \tag{29}$$

with $T_{\text{bands}}$ the number of aggregated bands (set to one in the considered scenario) and $D_{\text{Tran}}$, $D_{\text{PA}}$ per-chain coefficients.

The dynamic term models the dependence on radiated power. In the adopted pipeline, the total RF output power under radio adaptations is computed from the scenario-level per-antenna output power and the active antenna count:

$$P_{\text{out,tot}}(s_a, s_p) = P_{\text{tx,ant}}\, N_{\text{act}}(s_a)\; s_p, \tag{30}$$

where $N_{\text{act}}(s_a)$ is defined in Section 3.1.3. Bandwidth utilization is mapped to RF output through the continuous policy used in this thesis:

$$P_{\text{out}}(s_a, s_f, s_p) = s_f\, P_{\text{out,tot}}(s_a, s_p). \tag{31}$$

The dynamic consumption is then expressed as an efficiency-weighted term with an additional calibration factor:

$$P_{\text{dyn}}^{\text{AAU}}(s_a, s_f, s_p) = k_{\text{dyn}} \frac{P_{\text{out}}(s_a, s_f, s_p)}{\eta}. \tag{32}$$

Equations (26)–(32) match the implementation used to produce the results in Chapter 4.

**Sleep-state parameters.** As for the 3GPP model, the AAU model provides steady-state sleep powers and wake-up times for `MICRO`, `LIGHT`, and `DEEP`. These parameters are consumed by the shared sleep-mode controller of the time-domain simulator (Section 3.5) and are treated as configuration inputs in the evaluation pipeline.

### 3.3.3 Implementation mapping and alignment notes

Both models are exposed through a common interface that supports: (i) a one-time pre-processing step per configuration (used to bind scenario-level parameters to the model implementation), (ii) `ACTIVE` DL power $P_{\text{active}}(s_a, s_f, s_p)$, and (iii) sleep-state parameters (steady powers and wake-up times) consumed by the shared sleep-mode controller. This enables a controlled substitution of the power model while preserving the same radio adaptation variables and the same sleep-controller logic.

**Mapping of inputs $(s_a, s_f, s_p)$.** The adaptation variables introduced in Section 3.1.3 are passed to both models:

- $s_a$ modulates antenna activity and affects `ACTIVE` power in both models (directly in (23) and through $m_{\text{act}}$ in (28)).
- $s_f$ represents bandwidth utilization and affects the 3GPP joint dynamic term via $s_f s_p$ in (23); in the AAU model it scales the RF output power via (31).
- $s_p$ scales the radiated power, entering (23) and (30).

**Component-carrier handling (simplified assumption).** In the AAU implementation, a simplified uniform allocation over component carriers is assumed when distributing the RF output power across $C_{\text{cc}}$ carriers. This choice keeps the power pipeline

consistent with the abstracted bandwidth utilization factor $s_f$ while avoiding additional carrier-level scheduling assumptions.

**Shared sleep parameters.** Unless otherwise stated, the sleep-state steady powers ($P_{\text{micro}}, P_{\text{light}}, P_{\text{deep}}$) and wake-up times are configured identically across the compared models. Therefore, differences in energy outcomes primarily arise from the `ACTIVE` power formulation. Time-domain energy accounting and KPI computation follow Section 3.2.

**Parameter summary.** Table 6 summarizes the main parameters and their role in the two formulations. Values used in the experiments are reported in the configuration tables introduced in Section 3.1.4.

*Table 6: Main parameters of the compared BS power models and their role in the formulation.*

| Symbol / key | Model | Meaning / role |
|---|---|---|
| $P_{\text{active,DL}}$ (`P_active_DL`) | 3GPP | Active DL power at nominal operation |
| $P_{\text{static}}$ (`P_static_active`) | 3GPP | Static `ACTIVE` component |
| $A$ (`A`) | 3GPP | Dynamic split factor between antenna-dependent and joint part |
| $\eta$ (`eta`) | 3GPP/AAU | Efficiency factor (constant in this study) |
| $P_0$ (`P0`) | AAU | Baseline/control hardware power |
| $P_{\text{BB}}$ (`PBB`) | AAU | Baseband processing contribution |
| $D_{\text{Tran}}$ (`DTran_per_chain`) | AAU | Per-chain transport coefficient |
| $D_{\text{PA}}$ (`DPA_per_chain`) | AAU | Per-chain auxiliary PA-chain coefficient |
| $M_{\text{av}}$ (`M_av_per_trx`) | AAU | Maximum available chains per transceiver |
| $C_{\text{cc}}$ (`C_cc`) | AAU | Number of component carriers (uniform allocation assumed) |
| $T_{\text{bands}}$ (`T_bands`) | AAU | Number of aggregated bands (set to one) |
| $k_{\text{static}}$ (`k_static`) | AAU | Calibration factor applied to static blocks |
| $k_{\text{dyn}}$ (`k_dyn`) | AAU | Calibration factor applied to dynamic term |
| $P_{\text{micro}}, P_{\text{light}}, P_{\text{deep}}$ | both | Steady power levels of sleep states |
| $t_{\text{wake,micro}}, t_{\text{wake,light}}, t_{\text{wake,deep}}$ | both | Wake-up times from sleep states |

## 3.4 Static Evaluation Methodology

This section describes the static analysis procedure used to evaluate energy consumption under controlled radio-parameter adaptations. The objective is to isolate the structural scaling behavior of each BS power-consumption model by comparing their predicted energy savings under identical radio configurations, without confounding effects due to traffic dynamics, queueing, or sleep-mode transitions. The static study therefore provides a clean baseline for interpreting the time-domain results presented later in Section 3.5.

### 3.4.1 One-at-a-time sweep design

The static evaluation is performed through one-at-a-time sweeps over the radio adaptation variables defined in Section 3.1.3. For each experiment, a single control dimension is varied while the remaining variables are kept fixed at their baseline values. The baseline configuration is $\mathbf{s}_0 = (1, 1, 1)$ as defined in Section 3.2.1.

Let $\theta \in \{s_a, s_f, s_p\}$ denote the swept variable and let $\mathcal{V}_\theta$ be its discrete sweep set. Unless stated otherwise, the sweep sets follow the experimental configuration:

$$\begin{aligned} \mathcal{V}_{s_a} &= \{1.0,\ 0.75,\ 0.5,\ 0.25,\ 0.125\}, \\ \mathcal{V}_{s_p} &= \{1.0,\ 0.75,\ 0.5,\ 0.25,\ 0.125\}, \\ \mathcal{V}_{s_f} &= \{1.0,\ 0.95,\ 0.90,\ \ldots,\ 0.125\}. \end{aligned} \tag{33}$$

For each value $v \in \mathcal{V}_\theta$, the configuration vector is built as

$$\mathbf{s}(v) = \begin{cases} (v,\ 1,\ 1), & \theta = s_a, \\ (1,\ v,\ 1), & \theta = s_f, \\ (1,\ 1,\ v), & \theta = s_p. \end{cases} \tag{34}$$

### 3.4.2 Static link abstraction and performance proxy

For each configuration point, service performance is quantified through a deterministic link-level *proxy*. In this context, a proxy is a simplified, repeatable surrogate metric

that is used in place of a full physical-layer and scheduling simulation: it captures the intended monotonic dependence on link quality and radio adaptations, while avoiding additional assumptions (e.g., MCS selection, HARQ, fading dynamics, and scheduler details) that would otherwise confound controlled comparisons.

In the static setting, the active-time factor is fixed to $s_t = 1$, and the achievable rate proxy is computed as

$$R_{\text{norm}}(\mathbf{s}) = s_f \, L(s_a) \, \log_2\big(1 + \gamma(\mathbf{s})\big), \tag{35}$$

where $L(s_a)$ is the effective number of spatial layers (Section 3.1.3) and $\gamma(\mathbf{s})$ is the effective SNR obtained under controlled assumptions. In the static proxy, the pathloss is treated as a scenario-level fixed value (configured in the experiments) to generate fully repeatable trade-off curves; distance-dependent large-scale variations are instead captured in the time-domain simulator (Section 3.5). The corresponding latency proxy is defined as the transmission time of a packet of size $N_{\text{bits}}$:

$$T_{\text{tx}}(\mathbf{s}) = \frac{N_{\text{bits}}}{R_{\text{norm}}(\mathbf{s})}. \tag{36}$$

Here $R_{\text{norm}}(\mathbf{s})$ is a comparative indicator (used to quantify relative rate and latency changes across configurations) rather than a fully physical throughput model.

For a fixed $\mathbf{s}$, $R_{\text{norm}}(\mathbf{s})$ and $T_{\text{tx}}(\mathbf{s})$ are shared across all compared power models; only the mapping from $\mathbf{s}$ to `ACTIVE` power differs (Section 3.3).

### 3.4.3 Energy evaluation and KPI construction

For each configuration $\mathbf{s}$, the compared power models are queried using the same input tuple $(s_a, s_f, s_p)$ to obtain the `ACTIVE` downlink power $P_{\text{active}}(\mathbf{s})$ (Section 3.3). A model-specific pre-processing hook is invoked once per configuration before evaluating `ACTIVE` power, to bind scenario parameters to the implementation. In the static analysis, only `ACTIVE` operation is considered; sleep-state powers and wake-up dynamics are not used.

Energy efficiency in the static setting is measured through the active-only energy per transmitted bit $E_{\text{bit}}^{\text{act}}(\mathbf{s})$, and trade-off curves are constructed using the indicators

**Algorithm 1** Static one-at-a-time sweep and KPI computation.

**Input:** Sweep variable $\theta \in \{s_a, s_f, s_p\}$ and discrete set $\mathcal{V}_\theta$
**Input:** Baseline $\mathbf{s}_0 = (1, 1, 1)$ and packet size $N_{\text{bits}}$
**for** $v \in \mathcal{V}_\theta$ **do**
  Build configuration $\mathbf{s}(v)$ by setting $\theta \leftarrow v$ and fixing the others to 1
  Compute $R_{\text{norm}}(\mathbf{s}(v))$ and $T_{\text{tx}}(\mathbf{s}(v))$ via the static link proxy
  **for each** power model $m \in \{3\text{GPP}, \text{AAU}\}$ **do**
    Invoke model pre-processing hook for $\mathbf{s}(v)$
    Query $P^{(m)}_{\text{active}}(\mathbf{s}(v))$
    Compute $E^{\text{act},(m)}_{\text{bit}}(\mathbf{s}(v))$
    Compute normalized KPIs w.r.t. $\mathbf{s}_0$ (Section 3.2.2)
  **end for**
**end for**
Aggregate results to construct trade-off curves for the swept dimension $\theta$

defined in Section 3.2.2. The packet size in bits is

$$N_{\text{bits}} = 8 \cdot 1024 \cdot \texttt{packet\_size\_kB}, \tag{37}$$

matching the implementation. Operationally, for each $\mathbf{s}$ the procedure is: (i) compute $(R_{\text{norm}}, T_{\text{tx}})$ from the link proxy above; (ii) obtain $P_{\text{active}}(\mathbf{s})$ from the selected model; (iii) compute $E^{\text{act}}_{\text{bit}}(\mathbf{s})$ and the corresponding normalized KPIs with respect to $\mathbf{s}_0$ following Section 3.2.

### 3.4.4 Workflow and reproducibility hooks

Algorithm 1 summarizes the static sweep workflow used to generate per-configuration KPIs for each power model. Numerical results are exported in tabular form for subsequent plotting and comparison; implementation details (file organization and scripts) are summarized in Section 3.7.

## 3.5 Dynamic Simulator and Sleep-Mode Operation

This section describes the discrete-time simulator used for traffic-driven evaluations. The simulator operates with slot duration $\Delta t$ and models: (i) stochastic downlink packet arrivals at the UEs, (ii) FIFO queue evolution and packet sojourn times, (iii) a service process driven by a fixed radio configuration $\mathbf{s} = (s_a, s_f, s_p)$, and (iv) a multi-level

sleep-mode (SM) controller that governs transitions between `ACTIVE` and low-power modes, including wake-up transients. The simulator produces raw traces (queues, served bits, state timeline, instantaneous power) that are post-processed into KPIs and diagnostics according to Section 3.2. For a fixed scenario and traffic realization, the only model-dependent component is the mapping from ($\mathbf{s}$, state) to power consumption (Section 3.3).

### 3.5.1 Discrete-time timeline and simulator loop

The simulation horizon $T_{\text{sim}}$ is divided into $N_{\text{slots}} = \lceil T_{\text{sim}}/\Delta t \rceil$ slots indexed by $t$, with slot $t$ covering the interval $[t\Delta t, (t+1)\Delta t)$. Within each slot, the simulator applies the following update order:

1. **Arrivals:** generate new packets and append them to per-user FIFO queues.
2. **Backlog detection:** identify the set of active users (non-empty queues).
3. **Sleep-controller step:** update the SM state machine (Algorithm 2) and determine whether service is allowed.
4. **Power/energy accounting:** compute instantaneous power from the current state and accumulate slot energy.
5. **Service (if enabled):** serve queued bits according to the bandwidth-sharing rule and per-user achievable rates.

This structure ensures that delay and energy consumption are jointly determined by the interaction between stochastic arrivals, sleep transitions, and the chosen radio configuration.

### 3.5.2 Traffic model and FIFO queueing

Each of the $U$ users generates downlink packets independently according to a Poisson process with rate $\lambda$ packets/s. In the slotted simulator, arrivals are implemented through a Bernoulli trial per user and per slot:

$$p_{\text{arr}} = \min(\lambda\Delta t, 1), \tag{38}$$

so that at most one packet per user can arrive within a slot. All packets have fixed size $N_{\text{bits}}$ and are enqueued upon arrival. Each user queue is served in FIFO order and may be partially served within a slot; a packet is marked as completed when its remaining bits reach zero.

Packet delay is measured as sojourn time (arrival-to-completion). Let $t_{\text{arr}}^{(i)}$ be the arrival time of packet $i$ and let $t_{\text{cmp}}^{(i)}$ be its completion time (recorded at the slot boundary when the packet finishes service). The corresponding latency is

$$\ell_i = t_{\text{cmp}}^{(i)} - t_{\text{arr}}^{(i)}, \tag{39}$$

which accounts for both waiting time in the queue and transmission time under the available service rate.

### 3.5.3 Channel abstraction and SNR scaling under radio adaptations

Users are placed according to the annulus topology described in Section 3.1.1. For each user $u$, the simulator computes a large-scale channel gain from a distance-dependent pathloss model with an additive shadowing offset/realization. Distances are lower-bounded to 1 m for numerical stability, i.e., $D_u = \max(d_u, 1\ \text{m})$. At carrier frequency $f_c$ (in GHz), the pathloss (in dB) is modeled as

$$\text{PL}_{\text{dB}}(D_u) = 16.9 \log_{10}(D_u) + 32.8 + 20 \log_{10}(f_c) + X_{\text{sh}}, \tag{40}$$

where $X_{\text{sh}}$ denotes the configured shadowing offset (fixed in the baseline configuration). The linear channel gain is

$$h_u = 10^{-\text{PL}_{\text{dB}}(D_u)/10}. \tag{41}$$

Thermal noise power is computed from the noise spectral density $N_0$ (dBm/Hz), receiver noise figure NF (dB), and bandwidth $B$ (Hz) as in Section 3.1.2. The dynamic simulator adopts an **SNR-based** abstraction (single cell, no explicit inter-cell interference).

For a given run, the radio configuration $\mathbf{s} = (s_a, s_f, s_p)$ is fixed. Let $B_{\text{sys}}$ be the system bandwidth. The simulator first computes a reference per-user SNR at full

bandwidth and nominal scaling:

$$\mathrm{SNR}_{u,\mathrm{ref}} \triangleq \frac{P_{\mathrm{tx,ref}}\, h_u}{N(B_{\mathrm{sys}})}, \qquad P_{\mathrm{tx,ref}} = P_{\mathrm{tx,ant}}\, N_{\mathrm{tx}}, \tag{42}$$

and then applies a multiplicative scaling consistent with the adaptation logic. Transmit power is scaled proportionally to $s_a s_p$, while noise power scales proportionally to $B(s_f) = s_f B_{\mathrm{sys}}$, hence to $s_f$. This yields the effective per-user SNR

$$\mathrm{SNR}_{u,\mathrm{eff}}(s_a, s_f, s_p) = \mathrm{SNR}_{u,\mathrm{ref}} \cdot \frac{s_a s_p}{\max(s_f, \epsilon_f)}, \tag{43}$$

where $\epsilon_f > 0$ is a small implementation constant used to avoid division by zero.

### 3.5.4 Service process and equal bandwidth sharing

At each slot, the set of active users is defined as those with non-empty queues, and $U_{\mathrm{act}}(t)$ denotes their count in slot $t$. When service is enabled (i.e., when the SM controller is in `ACTIVE`), the available downlink bandwidth is shared uniformly across active users:

$$B_u(t) = \frac{B(s_f)}{U_{\mathrm{act}}(t)}. \tag{44}$$

The instantaneous rate of user $u$ is then modeled as

$$R_u(t) = B_u(t)\, \log_2\Big(1 + \mathrm{SNR}_{u,\mathrm{eff}}(s_a, s_f, s_p)\Big), \tag{45}$$

and the service budget within a slot is $R_u(t)\Delta t$ bits. These bits are removed from the head of the user FIFO queue; when a packet completes, its sojourn time is recorded. When service is not enabled (sleep states or `WAKING`), no bits are served during the slot and queues can accumulate backlog due to arrivals. Spatial multiplexing is not explicitly modeled in this dynamic service-rate abstraction.

### 3.5.5 Sleep-mode controller and state occupancy

Sleep-mode dynamics are governed by a multi-level state machine with states

$$\mathcal{S} = \{\texttt{ACTIVE}, \texttt{MICRO}, \texttt{LIGHT}, \texttt{DEEP}, \texttt{WAKING}\},$$

executed once per slot. The controller tracks an idle timer that increases when there is no backlog and triggers progressive transitions into `MICRO`, `LIGHT`, and `DEEP` after thresholds ($\tau_{\text{micro}}, \tau_{\text{light}}, \tau_{\text{deep}}$). When backlog becomes non-zero while the BS is in a sleep state, a wake-up transient is initiated: the system enters `WAKING` for a state-dependent duration and does not provide service during this interval. The per-slot controller logic is summarized in Algorithm 2.

**Algorithm 2** Sleep-mode controller (per-slot update, compact rule-based form).

**Input:** Slot duration $\Delta t$, backlog flag `has_traffic`
**Input:** Thresholds ($\tau_{\text{micro}}, \tau_{\text{light}}, \tau_{\text{deep}}$)
**Input:** Sleep parameters: steady powers $P_{\texttt{MICRO}}, P_{\texttt{LIGHT}}, P_{\texttt{DEEP}}$ and wake-up times $t_{\texttt{MICRO}}, t_{\texttt{LIGHT}}, t_{\texttt{DEEP}}$
**Input:** Current state `state` $\in$ {`ACTIVE`, `MICRO`, `LIGHT`, `DEEP`, `WAKING`}, idle timer $t_{\text{idle}}$, wake timer $t_{\text{wake}}$, origin sleep level `src`
**Output:** Instantaneous power $P(t)$ and flag `service_ok` (True iff `state=ACTIVE`)

**if** `state=WAKING` **then**
  Decrease wake timer: $t_{\text{wake}} \leftarrow \max(0,\ t_{\text{wake}} - \Delta t)$
  Set power to the sleep level being exited: $P(t) \leftarrow P_{\texttt{src}}$; `service_ok`←False
  **if** $t_{\text{wake}} = 0$ **then**
    End transient: `state`←`ACTIVE`, $t_{\text{idle}} \leftarrow 0$
  **end if**
  **return** $(P(t), \texttt{service\_ok})$
**end if**
**if** `has_traffic` **then**
  **if** `state=ACTIVE` **then**
    Reset idle timer: $t_{\text{idle}} \leftarrow 0$
    $P(t) \leftarrow P_{\text{active}}(s_a, s_f, s_p)$; `service_ok`←True
  **else**
    Trigger wake-up from current sleep level:
    `src`←`state`, $t_{\text{wake}} \leftarrow t_{\texttt{src}}$, `state`←`WAKING`
    $P(t) \leftarrow P_{\texttt{src}}$; `service_ok`←False
  **end if**
  **return** $(P(t), \texttt{service\_ok})$
**end if**
**No backlog:** increase idle timer $t_{\text{idle}} \leftarrow t_{\text{idle}} + \Delta t$
Select target state from thresholds:
  `ACTIVE` if $t_{\text{idle}} < \tau_{\text{micro}}$; `MICRO` if $\tau_{\text{micro}} \le t_{\text{idle}} < \tau_{\text{light}}$; `LIGHT` if $\tau_{\text{light}} \le t_{\text{idle}} < \tau_{\text{deep}}$; `DEEP` if $t_{\text{idle}} \ge \tau_{\text{deep}}$
Update `state`←`target`
Set power: $P(t) \leftarrow P_{\text{active}}(s_a, s_f, s_p)$ if `state=ACTIVE`, else $P(t) \leftarrow P_{\texttt{state}}$
`service_ok`← (`state=ACTIVE`)
**return** $(P(t), \texttt{service\_ok})$

The simulator logs the cumulative time spent in each state, $T_{\text{state}}$, and reports the corresponding occupancy fractions

$$\phi_{\text{state}} = \frac{T_{\text{state}}}{T_{\text{sim}}}, \qquad \text{state} \in \mathcal{S}. \tag{46}$$

The active-time fraction used as a diagnostic KPI in Section 3.2.3 is

$$s_t^{\text{dyn}} \triangleq \phi_{\texttt{ACTIVE}}. \tag{47}$$

Wake-up event counters (number of activations from `MICRO`, `LIGHT`, and `DEEP`) are also recorded.

### 3.5.6 Power computation and energy accounting

At each slot, the instantaneous BS power $P(t)$ is obtained from the current operational state and the selected power model:

- in `ACTIVE`: $P(t) = P_{\text{active}}(s_a, s_f, s_p)$, where $P_{\text{active}}(\cdot)$ is given by either model in Section 3.3;
- in `MICRO/LIGHT/DEEP`: $P(t)$ equals the corresponding steady-state sleep power;
- in `WAKING`: $P(t)$ equals the steady-state sleep power of the sleep level from which the BS is waking up.

Slot energy is accumulated as the steady contribution plus any (optional) one-off wake-up energy terms:

$$E_{\text{tot}} \leftarrow E_{\text{tot}} + P(t)\,\Delta t + E_{\text{wake}}(t), \tag{48}$$

where $E_{\text{wake}}(t)$ is non-zero only in slots where a wake-up is triggered (and is set to zero in the baseline configuration).

### 3.5.7 Simulator outputs and KPI post-processing

Over the full simulation horizon, the simulator outputs the total served bits $B_{\text{srv}}$ (including partially served packets), total consumed energy $E_{\text{tot}}$, and the set of packet latencies

$\{\ell_i\}$ (computed over completed packets), as well as the sleep-mode state trace and wake-up counters. These raw outputs are converted into the reported energy/performance KPIs and diagnostics using the definitions and normalization rules in Section 3.2.3 (e.g., average power, dynamic energy per served bit, NES gain, mean/p95 delay, and state occupancies). Optionally, the simulator records time series traces of $P(t)$, cumulative energy, served bits, and the SM state, enabling time-resolved inspection of transient behavior and validation of steady-state averages.

## 3.6 Multi-BS RL Extension

This section extends the single-BS time-domain framework (Section 3.5) to a small multi-BS deployment, where multiple small base stations (SBSs) operate simultaneously and their decisions interact through co-channel interference. The objective is to quantify how the choice of the BS power model (3GPP TR 38.864 vs. AAU-based) affects the *learned* energy-management policy and the resulting energy-performance trade-offs. The extension is intentionally lightweight and follows the simulator logic summarized in Algorithms 3–4, adopting simplifying assumptions to keep the learning loop tractable. In particular, the controller uses a tabular action-value table with an immediate-utility update rather than a full bootstrapped Q-learning rule; this choice is revisited in Section 3.6.8. Unless otherwise stated, the multi-SBS baseline parameters are those reported in Table 7.

### 3.6.1 Multi-SBS scenario and shared stochastic dataset

The considered topology includes $N_{\mathrm{SBS}} = 4$ SBSs placed at fixed planar coordinates (as specified in the configuration). Time is slotted and all SBSs take one control decision per slot (synchronized operation). Per-slot arrivals are modeled independently per SBS as a Poisson random variable; each newly arrived user is assigned a geometrically distributed holding time (in slots) and remains associated to the serving SBS for the entire holding time (no offloading/handover). A conceptual depiction of the multi-SBS scenario is shown in Fig. 3.

To ensure fairness across power-model comparisons, all runs reuse the *same* pre-computed stochastic dataset, comprising (for each episode and slot): (i) the number of

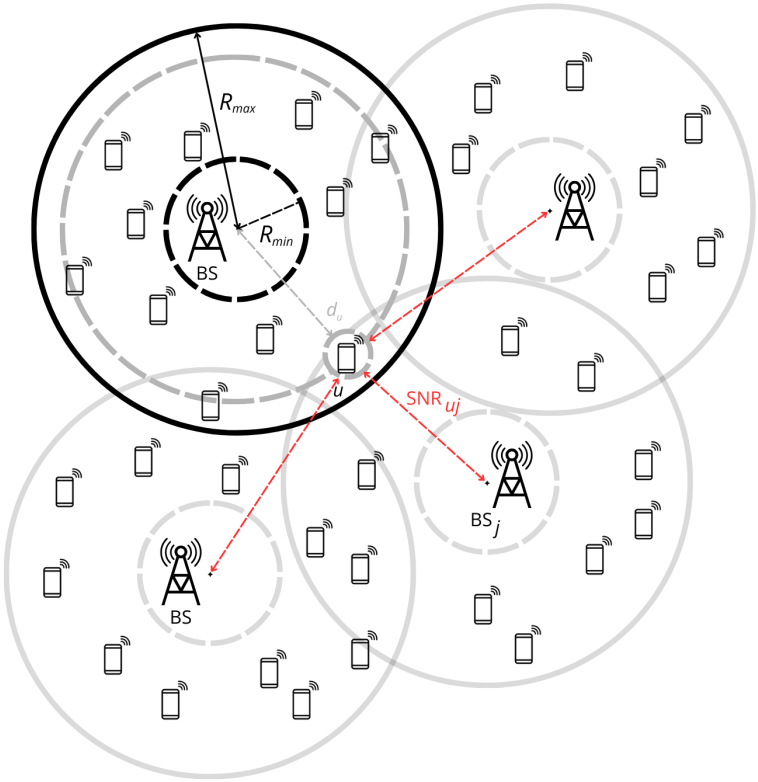


*Figure 3: Multi-SBS scenario: users are generated around a serving SBS, but their channel quality is evaluated w.r.t. all neighboring SBSs, enabling interference-aware SINR and rate computation.*

new users per SBS, (ii) their holding times, and (iii) their large-scale SNR vectors to all SBSs. Therefore, the traffic and propagation realizations are identical across compared runs; only the per-slot energy feedback depends on the chosen power model.

### 3.6.2 Interference-aware SINR and rate computation (simplified coupling)

Within a slot, each SBS $s$ applies a single radio configuration shared by all its active users: a bandwidth fraction $f^{(s)} \in [0, 1]$ and an antenna-activity factor $s_a^{(s)} \in [0, 1]$, selected from discrete sets (Table 7). The link evaluation uses a simplified interference coupling:

- only *ACTIVE* SBSs contribute interference (sleeping SBSs do not);

- the desired signal is scaled by the serving antenna-activity factor $s_a^{(s)}$;
- the interference from an interfering SBS $j$ is scaled by both its antenna activity $s_a^{(j)}$ and its bandwidth fraction $f^{(j)}$.

Let $\mathrm{SNR}_{u\to j}$ denote the precomputed large-scale SNR of user $u$ with respect to SBS $j$ (obtained from pathloss, shadowing, and a nominal noise reference; Table 7). The effective Signal-to-Interference-plus-Noise Ratio (SINR) used for a user $u$ served by SBS $s$ is:

$$\mathrm{SINR}_u^{(s)} = \frac{s_a^{(s)}\,\mathrm{SNR}_{u\to s}}{1 + \sum\limits_{j\neq s:\, j\ \mathrm{ACTIVE}} s_a^{(j)}\, f^{(j)}\,\mathrm{SNR}_{u\to j}}. \tag{49}$$

When SBS $s$ is ACTIVE and has $U_{\mathrm{act}}^{(s)}(t)$ active users, the available downlink bandwidth is shared uniformly:

$$B_u^{(s)}(t) = \frac{B_{\mathrm{SBS}}\, f^{(s)}}{U_{\mathrm{act}}^{(s)}(t)}. \tag{50}$$

The instantaneous user rate proxy is computed via a Shannon-type expression:

$$R_u^{(s)}(t) = B_u^{(s)}(t)\,\log_2\!\big(1 + \mathrm{SINR}_u^{(s)}\big), \tag{51}$$

and the per-SBS service indicator used by the learning signal is the average across its active users:

$$\bar{R}^{(s)}(t) = \frac{1}{U_{\mathrm{act}}^{(s)}(t)} \sum_{u=1}^{U_{\mathrm{act}}^{(s)}(t)} R_u^{(s)}(t). \tag{52}$$

**Note on modeling scope.** This extension aims to introduce interference coupling while keeping the simulator lightweight. In particular, the precomputed $\mathrm{SNR}_{u\to j}$ uses a nominal noise reference, and the bandwidth fraction $f^{(s)}$ affects the rate through bandwidth sharing and affects interference through (49). This differs from the single-BS scaling rule in Section 3.5.3 and is treated as a simplifying assumption.

### 3.6.3 Control problem: discrete state and action spaces

Each SBS is controlled by an independent tabular learner. The observed state is a compact discrete tuple:

$$x \triangleq (\texttt{Mode}, \texttt{Ubin}, \texttt{Tbin}, \texttt{RBbin}),$$

where $\texttt{Mode} \in \{0, 1, 2, 3\}$ encodes ACTIVE/MICRO/LIGHT/DEEP, Ubin is a binned user-load indicator, Tbin is a binned transition-timer indicator derived from an internal timer $\tau$ (in slots), and RBbin is a binned memory of the previously applied bandwidth fraction. The bin thresholds used in this thesis are reported in Table 7.

Feasibility constraints are enforced in state indexing:

- if Mode is ACTIVE, then Tbin is forced to 0;
- if Mode is a sleep mode, then RBbin is forced to 0.

**Action space.** At each slot, the controller selects one action

$$a \triangleq (\texttt{mode\_cmd}, \texttt{rb\_idx}, \texttt{sa\_idx}),$$

where $\texttt{mode\_cmd} \in \{\texttt{ACTIVE}, \texttt{MICRO}, \texttt{LIGHT}, \texttt{DEEP}\}$ sets the operating mode. If mode_cmd is ACTIVE, then rb_idx and sa_idx select $(f^{(s)}, s_a^{(s)})$ from predefined discrete sets; if mode_cmd is a sleep mode, bandwidth and antenna indices are ignored. A special NOOP action is also included to keep the previous configuration unchanged.

**Transition timer and derived WAKING condition.** When switching from a previous mode to a sleep mode, the simulator sets a transition timer $\tau$ (in slots) based on mode-specific transition durations (Table 7). While $\tau > 0$, the only admissible action is NOOP, so the configuration is frozen until the timer expires. For reporting purposes, when $\texttt{Mode} \neq \texttt{ACTIVE}$ and $\tau > 0$, the SBS is labeled as WAKING; this label is not a selectable mode but a derived transient condition used in traces and state-occupancy diagnostics.

### 3.6.4 Learning signal and role of the power model

The learning signal combines normalized energy feedback with a rate utility term and QoS penalties. Since actions are selected greedily through $\arg\max_a Q(x,a)$ (with $\epsilon$-greedy exploration during training), the scalar signal is formulated as a per-slot *utility* to be maximized:

$$\mathcal{U}^{(s)}(t) = -w \cdot \frac{E_{\mathrm{BS}}^{(s)}(t)}{E_{\max}} + (1-w) \cdot \frac{\bar{R}^{(s)}(t)}{R_{\mathrm{norm}}} - \lambda_s \cdot \mathrm{Shortfall}^{(s)}(t) - \lambda_o \cdot \mathrm{OverProvision}^{(s)}(t), \tag{53}$$

where $E_{\mathrm{BS}}^{(s)}(t)$ is the per-slot energy (Wh/slot) returned by the selected power model given the current mode and radio controls, $E_{\max}$ is a fixed reference energy corresponding to `ACTIVE` at full bandwidth and full antenna activity, $\bar{R}^{(s)}(t)$ is the average served rate (Mbps), and $R_{\mathrm{norm}}$ is a constant rate normalizer (Table 7). The weight $w \in [0,1]$ tunes the energy-throughput emphasis, while $(\lambda_s, \lambda_o)$ weight shortfall and over-provision penalties derived from per-user rate requirements. If users exist but the SBS is not ACTIVE, the simulator enforces a strong QoS penalty by setting $\mathrm{Shortfall}^{(s)}(t) = 1$.

**Role of the power model.** The **only** component that changes across the two compared settings is the mapping from the current operating mode and radio controls to $E_{\mathrm{BS}}^{(s)}(t)$, i.e., the selected BS power model (3GPP TR 38.864 vs. AAU-based). The dataset, interference/SINR computation, and rate-related terms remain identical.

### 3.6.5 Tabular learning algorithm and update rule

The learning mechanism follows a tabular action-value structure with $\epsilon$-greedy exploration and a one-step immediate-utility update. Although the implementation maintains a Q-table notation, the discount factor is not used in the current update rule, and no bootstrap term based on the next state is included. Therefore, the update should be interpreted as a simplified non-bootstrapped action-value update rather than as full Q-learning:

$$Q^{(s)}(x,a) \leftarrow Q^{(s)}(x,a) + \alpha\Big(\mathcal{U}^{(s)}(t) - Q^{(s)}(x,a)\Big), \tag{54}$$

where $\alpha$ is the learning rate (Table 7) and $\mathcal{U}^{(s)}(t)$ is given by (53).

Because SBSs are interference-coupled, each slot is evaluated in two passes: first all SBSs select and apply actions, then utilities are computed under the resulting joint configuration and the Q-tables are updated. Algorithms 3–4 summarize the procedure.

**Algorithm 3** Multi-SBS tabular learning: training and evaluation loop.

**Input:** $N_{\text{SBS}}$ SBSs; state space $\mathcal{X}$; action space $\mathcal{A}$ (incl. `NOOP`); learning rate $\alpha$; exploration schedule $\epsilon(e)$; episode length $T$; episodes $E$.

**Input:** Shared dataset: per-slot arrivals, holding times, and per-user SNR vectors to all SBSs.

Initialize per-SBS Q-tables: $Q^{(s)}(x, a) \leftarrow 0, \forall s, x, a$.
**for** episode $e = 1, \ldots, E$ **do**
    Set $\epsilon \leftarrow \epsilon(e)$ and reset simulator state (users, modes, timers, and previous controls).
    **for** slot $t = 1, \ldots, T$ **do**
        Update user lists: remove expired users; add new arrivals from the shared dataset.
        Execute **SlotStep** (Alg. 4) and update Q-tables.
    **end for**
**end for**
Run a greedy test episode ($\epsilon = 0$) and log traces/KPIs.

**Algorithm 4** SlotStep: per-slot multi-SBS decision and interference-coupled evaluation.

**Input:** Current slot $t$; exploration $\epsilon$; Q-tables $\{Q^{(s)}\}$; timers $\{\tau^{(s)}\}$; current modes/controls; current user lists; selected power model.

**Output:** Updated modes/controls, timers, and Q-tables.

**PASS A (decide actions for all SBSs):**
**for** each SBS $s$ **do**
    Build state $x_t^{(s)}$; if $\tau^{(s)} > 0$ select `NOOP`, else select via $\epsilon$-greedy.
    Apply the action: update mode; if ACTIVE set $(f^{(s)}, s_a^{(s)})$; if entering sleep set transition timer $\tau^{(s)}$.
**end for**
**PASS B (evaluate and update tables):**
**for** each SBS $s$ **do**
    Compute per-slot energy $E_{\text{BS}}^{(s)}(t)$ from the selected power model.
    If ACTIVE with users: compute $\bar{R}^{(s)}(t)$ using (49) and derive QoS penalties; otherwise set $\bar{R}^{(s)}(t) = 0$ and apply penalties.
    Compute $\mathcal{U}^{(s)}(t)$ via (53) and update $Q^{(s)}$ via (54).
**end for**
Decrement timers: $\tau^{(s)} \leftarrow \max(\tau^{(s)} - 1, 0)$ for all $s$.

### 3.6.6 Training and evaluation protocol

For each power model, training is executed for a fixed number of episodes and slots per episode (Table 7), using the shared dataset to ensure identical arrivals/holding times/channel realizations. After training, evaluation simulates the learned greedy policy ($\epsilon = 0$) and logs:

- **Per-slot traces:** per-SBS power/energy, throughput, mode/transition condition, and user load.
- **Aggregated KPIs:** average network power, normalized network power (w.r.t. the `ACTIVE` full-load reference), average network rate, and energy efficiency (Mbps/W).
- **State occupancy:** time shares in `ACTIVE`, the derived `WAKING` condition, and `MICRO/LIGHT/DEEP`.

### 3.6.7 Outputs

The multi-SBS extension produces: (i) convergence diagnostics from training, (ii) state time shares vs. $w$, and (iii) network-level traces over time (power, rate, and energy efficiency), including best-vs-best overlays across power-model choices. In the grid-based evaluation, the best operating points are selected according to maximum achieved energy efficiency (Mbps/W).

### 3.6.8 Limitations and future refinements

While the multi-SBS extension is more realistic than the single-BS case, it still relies on simplifying assumptions that keep the simulator and learning loop tractable. Key limitations include the simplified interference coupling (ACTIVE-only interference with linear scaling by $(f, s_a)$), the lack of user mobility and association dynamics, and the use of a non-bootstrapped tabular action-value update in (54) instead of full Q-learning. These aspects are suitable for a first controlled extension aimed at isolating the impact of the power-model energy feedback on the learned policy, and they motivate further refinements that can be addressed in subsequent work.

**Reference parameters (multi-SBS RL baseline)**

Table 7 summarizes the baseline configuration and discretization parameters adopted for the multi-SBS RL experiments described in this section.

## 3.7 Reproducibility Notes and Limitations

This section collects the implementation- and reporting-level details needed to reproduce the results in Chapter 4 and delineates the validity scope of the proposed methodology. The evaluation relies on two pipelines: (i) a single-BS discrete-time simulator (Section 3.5), and (ii) a multi-BS extension where multiple SBSs are controlled via tabular RL (Section 3.6). In both cases, results are obtained by executing the reference implementation with the configuration settings described below, without introducing mechanisms beyond those explicitly implemented.

### 3.7.1 Pipelines, configuration sources, and horizons

Reproducibility requires keeping time discretization, simulation horizons, and configuration sources consistent across runs:

- **Single-BS pipeline:** the simulator evolves on a slot-based timeline with slot duration $\Delta t$ and total horizon $T_{\text{sim}}$ specified in `config.json`. The same file provides scenario, channel, traffic, and sleep-controller parameters used throughout the run.

- **Multi-BS RL pipeline:** the simulator evolves on a slot-based timeline with slot duration and horizon specified in `cfg.py` (episode length in slots, repeated over a fixed number of episodes). The same file defines topology, channel/noise parameters, traffic and holding-time generation, and the discretized state/action spaces used by the controller.

Unless explicitly stated in Chapter 4, all metrics are computed over the full simulated horizons (i.e., no burn-in removal) for all compared power models.

*Table 7: Baseline configuration for the multi-SBS RL extension.*

| Group | Parameter | Value |
|---|---|---|
| Training | Episodes $E$ | 200 |
| | Slots per episode $T$ | 100 |
| | Slot duration | 0.640 s |
| Learning / cost | Learning rate $\alpha$ | 0.05 |
| | Discount factor $\gamma$ | 0.95 (configured, not used in update) |
| | Weight sweep $w$ | $\{0.0, 0.1, \ldots, 1.0\}$ (grid); examples: $\{0.1, 0.3, 0.5, 0.7, 0.9\}$ |
| | Shortfall penalty $\lambda_s$ | 1.0 (baseline; may be tuned in sweeps) |
| | Over-provision penalty $\lambda_o$ | 0.1 (baseline; may be tuned in sweeps) |
| | Rate normalizer $R_{\text{norm}}$ | 50 |
| Topology | Number of SBSs $N_{\text{SBS}}$ | 4 |
| | SBS radius | 40 m |
| | SBS coordinates | fixed planar layout (configuration-defined) |
| | Per-SBS TX power $P_{\text{tx,SBS}}$ | 32 W |
| Channel | Bandwidth $B_{\text{SBS}}$ | 20 MHz |
| | Carrier frequency $f_c$ | 3.5 GHz |
| | Shadowing (std.) | 3 dB |
| | Noise density $N_0$ | $-174$ dBm/Hz |
| | Noise figure NF | 9 dB |
| Traffic | Arrival rate (per SBS) $\lambda$ | 0.5 users/s |
| | Mean holding time | 20 slots (geometric) |
| | Per-user rate requirement | 20 Mbps |
| State discretization | User-load bins (`Ubin`) | thresholds: 0, 4, 10 users |
| | Transition bins (`Tbin`) | thresholds: 0, 7 slots |
| | Previous-band bins (`RBbin`) | thresholds: $0.0, 0.1, \ldots, 0.9$ |
| | Mode encoding | 0 = `ACTIVE`, 1 = `MICRO`, 2 = `LIGHT`, 3 = `DEEP` |
| Action levels | Bandwidth fractions $f$ | 11 levels in $[0.1, 1.0]$ |
| | Antenna activity $s_a$ | $\{0.1, 0.4, 0.6, 0.8, 1.0\}$ |
| | Actions | `ACTIVE`$(f, s_a)$, sleep modes, and `NOOP` |
| Transition durations | `MICRO` | 0 ms (0 slots) |
| | `LIGHT` | 640 ms (1 slot) |
| | `DEEP` | 10,000 ms ($\lceil 10{,}000/(0.640 \cdot 1000) \rceil = 16$ slots) |

### 3.7.2 Randomization and seed policy

Both pipelines include stochastic components; therefore, controlled seeding is required:

- **Single-BS:** a single seed specified in `config.json` governs all pseudo-random draws in the simulation (user placement, traffic arrivals, and any other stochastic components), enabling repeated runs under identical conditions.
- **Multi-BS RL:** two seeds are used to decouple *environment randomness* from *learning randomness*. A `seed_data` drives the generation of the shared stochastic dataset (arrivals, holding times, and per-user channel/SNR quantities used by the environment), while a `seed_train` governs exploration and other stochastic choices in the learning loop. This separation enables fair model-to-model comparisons (shared environment realizations) while allowing reproducible training trajectories.

### 3.7.3 Metric computation and aggregation conventions

Metrics are computed according to the KPI definitions and normalization rules in Section 3.2. For clarity, the reporting conventions adopted in Chapter 4 are summarized as follows:

- **Single-BS:** total energy is accumulated by summing per-slot energy, and the dynamic energy per served bit is computed as consumed energy divided by served bits. The reported NES gain is then obtained by normalizing this quantity with respect to the baseline configuration. Latency statistics (e.g., mean and 95th percentile) are computed over completed packets.
- **Multi-BS RL:** per-SBS energy is accumulated per slot and aggregated at network level by summation across SBSs. Average network power and rate are obtained by time-averaging over slots and episodes. Network energy efficiency is computed as the ratio between average rate and average power. State occupancies are derived from the logged mode timeline (including waking/transition phases) and summarized as time shares.

Across both pipelines, comparisons are controlled by applying the same scenario assumptions, controller settings, and (when applicable) the same stochastic realizations; the only model-dependent component is the mapping from controls/operational state to power.

### 3.7.4 Validity scope and methodological limitations

The methodology is designed for controlled and explainable comparisons of power-model impact; consequently, it adopts simplifying abstractions that bound the external validity of absolute numerical values:

- **Traffic/service abstraction:** arrivals and service are modeled at slot level using simplified rate expressions; protocol-level effects (e.g., scheduling granularity, retransmissions, control signaling) are not represented.
- **Sleep-mode abstraction:** sleep operation is represented through a timer-driven state machine with fixed wake-up times and steady-state sleep powers; implementation-specific constraints and vendor-dependent state hierarchies are not modeled.
- **Multi-BS dynamics:** inter-SBS coupling is captured through a simplified interference model and fixed association during a user holding time; mobility, handover, and offloading dynamics are not considered.
- **RL scope:** the controller is tabular and uses discretized states/actions, improving interpretability and reproducibility but limiting scalability and policy expressiveness.

Accordingly, Chapter 4 emphasizes *relative* differences induced by the power-model structure and by its feedback within the control loop, rather than absolute performance claims.

# 4 Results

This chapter presents the outcomes of the comparative evaluation between the standardized 3GPP TR 38.864 BS power model and the analytical AAU-based model. Following the methodology introduced in Chapter 3, the analysis is organized into three steps. First, static radio-parameter sweeps are used to isolate model-dependent scaling behavior under controlled configurations. Second, dynamic time-domain simulations quantify how traffic-driven operation and hierarchical SM utilization shape both energy- and delay-related outcomes. Third, a preliminary multi-BS learning-based extension illustrates how power-model selection affects the energy feedback available to the controller and, consequently, the learned operating regime. The chapter concludes with a consolidated synthesis of the main findings across all scenarios.

## 4.1 Static Results

This subsection presents the outcomes of the static evaluation introduced in Section 3.4. The two BS power models are compared under controlled, traffic-free configurations by sweeping one radio control dimension at a time while keeping the remaining adaptation variables fixed. At each operating point, results are summarized through the normalized KPIs defined in Section 3.2. In particular, the trade-off curves are reported in terms of latency increase $\Delta T_{\mathrm{tx}}(\mathbf{s})$ (14) and static NES gain $G_{NES}^{\mathrm{static}}(\mathbf{s})$ (15), both measured with respect to the baseline configuration $\mathbf{s}_0$.

**Antenna and transmit-power scaling ($s_a$ and $s_p$).** Fig. 4 compares the trade-off curves obtained by scaling the number of active antennas ($s_a$) and the transmit power ($s_p$). For both control knobs, the curves are monotone: larger energy savings are obtained at the price of increased transmission time. Across the full sweep, the analytical AAU-based model (Piovesan) systematically predicts higher NES gains than the 3GPP TR 38.864 model at comparable latency increases. The gap is moderate for antenna scaling, for example around $\approx 25\%$ latency increase, where the AAU-based model predicts $\approx 70\%$ NES gain versus $\approx 64\%$ for 3GPP. The discrepancy becomes much more pronounced under transmit-power scaling: at a similar latency increase, the

AAU-based model reaches ≈ 57% NES gain, whereas 3GPP remains around ≈ 28%.

This difference already reveals an important structural feature of the two formulations. In the 3GPP model, transmit-power scaling affects only the joint dynamic portion of the `ACTIVE` power budget, whereas in the AAU formulation it directly drives the RF-output-related contribution. As a result, the AAU-based model reacts more strongly to reductions in $s_p$, while both models remain closer under antenna scaling, where the relevant terms are more evenly distributed across the two formulations.

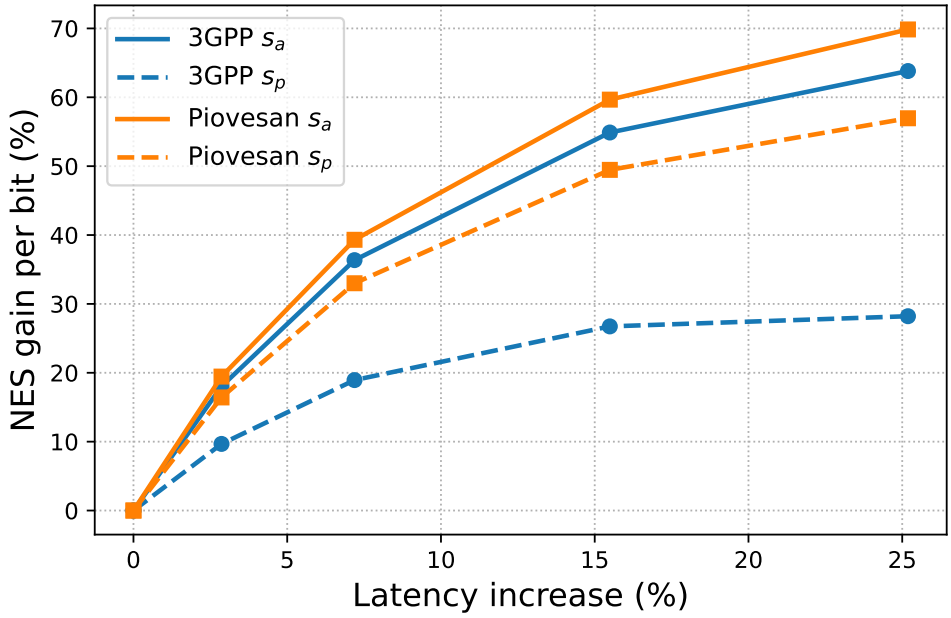


*Figure 4: Static radio-parameter sweeps: NES gain versus latency increase for antenna scaling ($s_a$) and transmit-power scaling ($s_p$). Each point corresponds to one scaling value and is normalized to the baseline configuration.*

**Bandwidth scaling ($s_f$).** Fig. 5 reports the trade-off induced by bandwidth scaling ($s_f$). Unlike antenna and transmit-power scaling, reducing the utilized bandwidth can lead to negative NES gains, that is, to a deterioration of energy efficiency even when the estimated `ACTIVE` power decreases. This follows directly from the structure of the static KPI: active-only energy per bit combines the model-dependent power term with the transmission-time proxy, and under bandwidth reduction the increase in transmission time can dominate the power reduction, thereby increasing energy per bit rather than decreasing it.

The curve makes this mechanism visible. Latency penalties become very large, reaching several hundreds of percent, while the NES trend turns increasingly negative. In this regime, the 3GPP model is more severely penalized than the AAU-based model: around $\approx$ 300% latency increase, the reported NES gain is approximately $-150\%$ for 3GPP and $-60\%$ for the AAU-based model; at the highest-latency point, the degradation reaches approximately $-280\%$ and $-130\%$, respectively. The reason is again structural: in the 3GPP formulation, only part of the dynamic `ACTIVE` power responds to $s_f$, whereas in the AAU model bandwidth scaling directly affects the RF-output-related term. Even in the more responsive AAU formulation, however, the rate penalty remains large enough to make aggressive bandwidth reduction unattractive in the static setting.

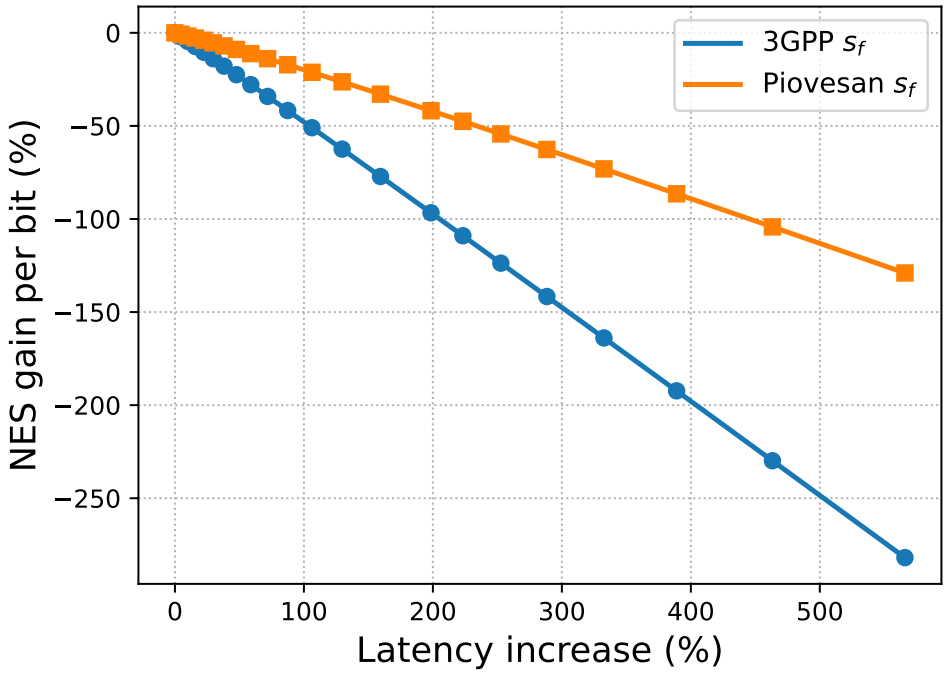


*Figure 5: Static bandwidth scaling: NES gain versus latency increase for $s_f$ (utilized bandwidth / RB scaling).*

## 4.2 Dynamic Results

This subsection presents the dynamic results obtained with the time-domain simulator introduced in Section 3.5. Traffic arrivals and queue evolution drive BS operation over time, triggering transitions between `ACTIVE` transmission and hierarchical low-power states (`MICRO`, `LIGHT`, and `DEEP`) as described in Chapter 3. For each configuration,

the comparison focuses on how the two power models translate identical traffic traces and the resulting state trajectories into different energy estimates. The main reported indicator is average power $\overline{P}(\mathbf{s})$ defined in (18). To better interpret the effect of bandwidth scaling, results are complemented by the average power observed while the BS is in `ACTIVE`, defined as

$$\overline{P}_{\texttt{ACTIVE}}(\mathbf{s}) \triangleq \frac{1}{T_{\texttt{ACTIVE}}} \sum_{t:\,\text{state}(t)=\texttt{ACTIVE}} P(t)\,\Delta t \quad [\text{W}], \tag{55}$$

where $T_{\texttt{ACTIVE}}$ is the cumulative time spent in `ACTIVE` during the simulation.

**Antenna and transmit-power scaling ($s_a$ and $s_p$).** Fig. 6 reports $\overline{P}(\mathbf{s})$ as a function of antenna activity and transmit-power scaling. For completeness, Fig. 6(b) also reports the conditional `ACTIVE` average $\overline{P}_{\texttt{ACTIVE}}(\mathbf{s})$. The two trends are closely aligned, which indicates that, for these two control knobs, differences in total average power are driven mainly by the scaling of `ACTIVE` consumption rather than by major changes in state occupancy.

As expected, decreasing either scaling factor reduces average consumption, and both models converge to the same value at the baseline point ($\approx$ 4.7 W at scaling factor 1). Across the sweep, however, the AAU-based model (Piovesan) consistently yields lower average power than 3GPP for the same scaling. The gap remains limited under antenna scaling, but becomes clearly visible under aggressive transmit-power reduction. For instance, at the smallest reported scaling value, the average power under $s_p$ is approximately 3.30 W for the AAU-based model versus 3.81 W for 3GPP, whereas the corresponding values under $s_a$ are approximately 3.07 W and 3.17 W. This confirms, also in the dynamic setting, that transmit-power adaptation is the dimension along which the two formulations diverge most strongly.

**Bandwidth scaling ($s_f$).** Fig. 7 reports the effect of bandwidth scaling on both $\overline{P}(\mathbf{s})$ and the conditional `ACTIVE` average $\overline{P}_{\texttt{ACTIVE}}(\mathbf{s})$. Here the contrast between the two quantities becomes central. The `ACTIVE`-only curve decreases monotonically with $s_f$ for both models, and the reduction is markedly stronger under the AAU-based model at low $s_f$ (for example, at $s_f \approx 0.2$, $\overline{P}_{\texttt{ACTIVE}} \approx 12$ W for Piovesan versus $\approx 18$ W for

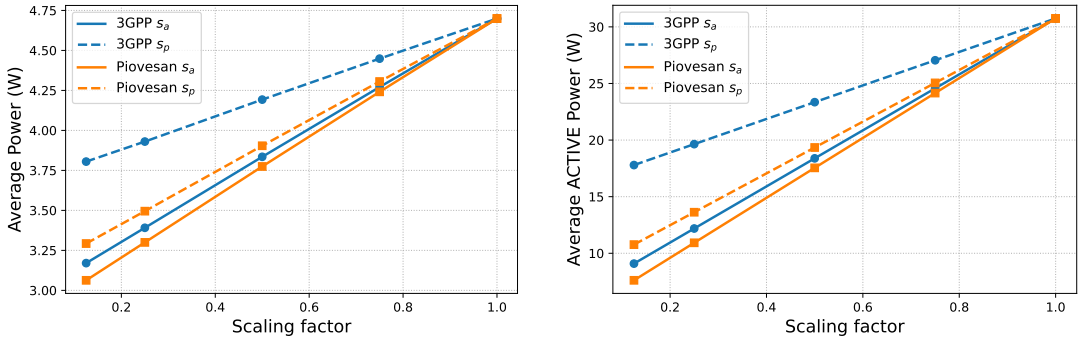


**(a)** Total average power $\overline{P}(\mathbf{s})$.

**(b)** Conditional `ACTIVE` average power $\overline{P}_{\texttt{ACTIVE}}(\mathbf{s})$.

*Figure 6: Dynamic antenna and transmit-power scaling ($s_a$ and $s_p$): (a) total average power; (b) conditional `ACTIVE` average power. Results compare the 3GPP TR 38.864 and AAU-based (Piovesan) models.*

3GPP, while both coincide around $\approx 27.5$ W at $s_f = 1$). If one looked only at `ACTIVE` power, bandwidth reduction would therefore appear uniformly beneficial.

The total average power tells a different story. Under moderate bandwidth reductions, $\overline{P}(\mathbf{s})$ decreases for both models, reaching approximately 3.8 W for the AAU-based model and 4.3 W for 3GPP. Under more aggressive reductions, however, the 3GPP curve rises again and can exceed the baseline, reaching approximately 4.85 W at the smallest $s_f$. This divergence between $\overline{P}_{\texttt{ACTIVE}}$ and $\overline{P}$ reflects the interaction between two effects: lower instantaneous `ACTIVE` power, but longer service time and therefore less opportunity to exploit sleep states. In other words, aggressive bandwidth reduction keeps the BS active for longer, shrinks idle intervals, and reduces the energy benefit that would otherwise come from deeper sleep. The AAU-based model compensates this effect more effectively because it predicts a stronger reduction in `ACTIVE` power, whereas in the 3GPP formulation the loss of sleep opportunity dominates in the low-$s_f$ regime.

## 4.3 Multi-BS RL Results

This subsection reports preliminary outcomes of the multi-BS learning-based extension introduced in Section 3.6. In this setting, $N_{\text{SBS}} = 4$ SBSs are controlled by independent tabular learners that select, at each slot, a sleep-mode command and, when `ACTIVE`, a shared radio configuration $(f^{(s)}, s_a^{(s)})$. Because the learning signal is computed using

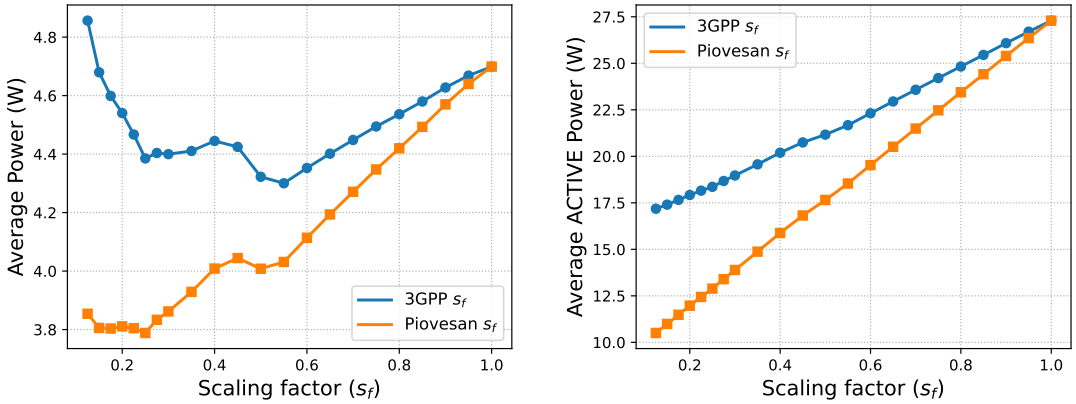


**(a)** Total average power $\overline{P}(\mathbf{s})$.

**(b)** Conditional `ACTIVE` average power $\overline{P}_{\texttt{ACTIVE}}(\mathbf{s})$.

*Figure 7: Dynamic bandwidth scaling ($s_f$): (a) total average power; (b) conditional `ACTIVE` average power. Results compare the 3GPP TR 38.864 and AAU-based (Piovesan) models.*

model-dependent energy feedback, the selected BS power model affects the per-slot utility and, consequently, the learned operating regime. In particular, the tabular update is driven by a utility signal that combines normalized energy feedback, rate utility, and QoS penalties, as defined in Section 3.6.4.

**Best operating-point comparison: network power over time.** Fig. 8(a) shows the network power trace, summed over the four SBSs and averaged over episodes, for the best operating point selected within the explored weight grid. After the initial transient, both curves settle into a quasi-stationary regime, but the AAU-based model (Piovesan) remains systematically below the 3GPP-based one. This indicates that the more optimistic energy tendency already observed in the single-BS analyses carries over to the interference-coupled learning setting, yielding lower network-level power throughout the steady portion of the trace and slightly lower peaks as well.

**Effect of the energy-rate weight.** Fig. 8(b) reports the average network power as a function of the weight $w$ that balances the energy term and the throughput utility in the per-slot signal (cf. (53)). As $w$ increases, the learned behavior becomes more energy-oriented and the average network power generally decreases for both models. Across most of the sweep, the Piovesan-based runs yield lower average power than

the 3GPP-based ones, with the largest separations typically appearing at intermediate values of $w$. For large $w$, where the objective becomes strongly energy-dominant, the curves tend to move closer, suggesting that the learned policies approach similar low-power operating regimes under a heavily energy-driven objective.

**Interpretation and scope.** The main value of these results is methodological rather than performance-oriented. They show that the choice of power model affects not only the post-processed energy KPIs, but also the utility signal used during learning. As a consequence, the controller is effectively trained in different energy environments depending on the adopted model. At the same time, these outcomes remain exploratory: they are intended to illustrate the propagation of model assumptions into the learning loop, not to provide a definitive benchmark of scalable multi-BS control.

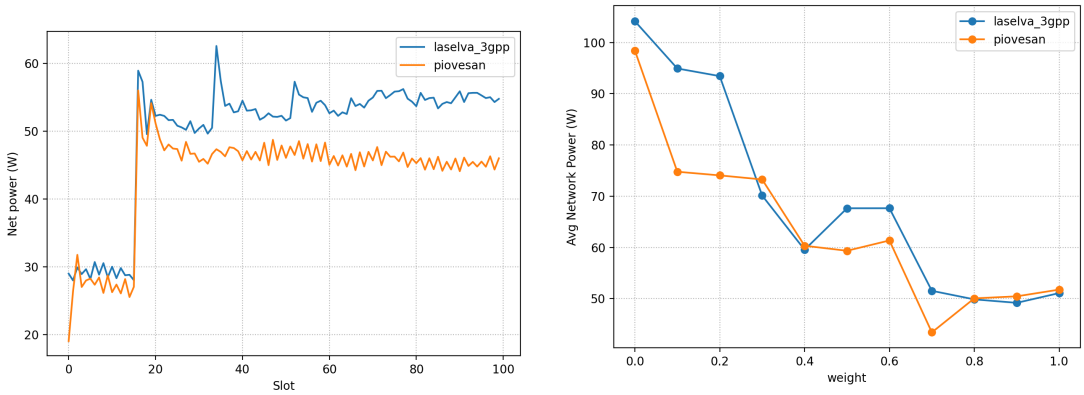


**(a)** Network power over time (mean over episodes).

**(b)** Average network power versus weight $w$.

*Figure 8: Preliminary multi-SBS learning-based extension: network-level power outcomes comparing the 3GPP-based and AAU-based models.*

## 4.4 Key Findings

This subsection consolidates the main outcomes across the static sweeps, the traffic-driven simulations with sleep states, and the multi-SBS learning-based extension. The goal is to separate systematic model-driven effects from scenario-specific details and to clarify how the two BS power formulations should be interpreted in comparative energy-efficiency studies.

**A structural explanation for the model gap.** A consistent gap between the two models emerges from how `ACTIVE` power scales with the radio adaptation variables. In the 3GPP formulation, (23)–(25) partition the dynamic component into an antenna-dependent term and a joint term; only the latter is scaled by the product $s_f s_p$, while the static contribution remains and the residual dynamic share is not reduced by bandwidth or transmit-power scaling. In contrast, the AAU formulation explicitly separates hardware blocks and ties them more directly to the active configuration: per-chain contributions depend on $m_{\text{act}}(s_a)$ via (28)–(29), and the RF-output-related term is driven by $P_{\text{out}}(s_a, s_f, s_p)$ through (30)–(32). As a result, reductions in $(s_a, s_f, s_p)$ affect a wider portion of the AAU power budget, which systematically amplifies the estimated savings.

**Which adaptation levers remain most effective.** Across the examined scenarios, antenna-activity and transmit-power adaptations provide the clearest energy-performance trade-offs. The largest discrepancy between models is observed under transmit-power scaling, because in the 3GPP model the dependence on $s_p$ is limited to the joint dynamic fraction, whereas in the AAU formulation it directly drives the RF-output-related term. By comparison, bandwidth scaling is consistently the least attractive lever. Although it can reduce `ACTIVE` power, it also reduces service capacity almost proportionally and therefore increases transmission time. This makes bandwidth reduction particularly sensitive to how strongly the selected model responds to $s_f$, and explains why it becomes unattractive much earlier under the 3GPP formulation.

**Why dynamic operation changes the interpretation.** The time-domain simulations show that average power cannot be inferred from `ACTIVE` power alone. Any adaptation that reduces service rate can increase queue occupancy and keep the BS active for longer, thereby shrinking the time available for low-power states. This mechanism is especially visible for bandwidth scaling: the conditional `ACTIVE` power decreases, but the total average power may flatten or even worsen because the loss of sleep opportunity offsets the instantaneous gain. Dynamic evaluation is therefore essential whenever traffic burstiness, queueing, and sleep transitions are part of the system behavior.

**Implications for learning-based control and model selection.** In the multi-SBS learning-based extension, the energy term in the utility signal is computed through the selected power model. Model choice therefore affects not only post-processing, but also the objective perceived by the controller during training. Even in this simplified setting, the same qualitative tendency observed in the static and dynamic analyses persists: the AAU-based formulation yields a more optimistic energy picture at network level, while 3GPP behaves as a more conservative reference. Taken together, these results suggest a practical interpretation of the two formulations: the 3GPP model is particularly suitable as a standardized baseline for reproducible and conservative comparison, whereas the AAU-based model is more informative when the goal is to reason about hardware-sensitive adaptation effects and potentially more aggressive savings under the assumed platform behavior.

# 5 Conclusion

This thesis investigated how the choice of the BS power-consumption model affects the estimated energy-saving potential of radio adaptations and sleep-mode operation, as well as the resulting energy-performance trade-offs in 5G and beyond-5G RANs. To address this question, it developed a controlled comparison between the standardized 3GPP TR 38.864 model and an analytical AAU-based formulation within a unified evaluation pipeline combining static sensitivity analysis, traffic-driven time-domain simulation, and a preliminary multi-BS learning-based extension.

**Main conclusions and answers to the research questions.** The results consistently show that energy-efficiency conclusions are strongly *model-dependent*. Across the evaluated scenarios, the AAU-based formulation provides a systematically more optimistic estimate of achievable savings, whereas 3GPP TR 38.864 yields more conservative figures under the same operating conditions. This answers the first research question by showing that even under aligned inputs and common KPIs, the selected power model can materially change the estimated benefit of antenna activity, bandwidth adaptation, and transmit-power scaling.

The dynamic analysis further shows that the relevant comparison cannot be limited to instantaneous `ACTIVE` power. Under traffic-driven operation, reconfiguration actions also affect service time, backlog evolution, and therefore the opportunity to enter deeper sleep states. As a result, the net energy outcome depends on the interaction between active-state scaling and sleep-state occupancy. This answers the second research question and confirms that dynamic operation is essential for interpreting energy-saving strategies.

The observed gap between the two models can be traced back to their internal structure. The 3GPP formulation retains a larger configuration-insensitive share of the `ACTIVE` power budget, whereas the AAU-based model couples a wider fraction of consumption to the active radio configuration. This structural difference explains why the latter amplifies the apparent benefit of reconfiguration and sleep-enabled operation. In this sense, the thesis also answers the third research question by showing that the

divergence between models is not incidental, but rooted in different assumptions about which hardware contributions scale and which remain as non-removable floors.

The preliminary multi-BS learning-based extension suggests that the same issue propagates into control. Since the energy term in the learning signal is computed through the adopted power model, changing the model changes the objective effectively perceived by the controller. The qualitative gap observed in the static and dynamic analyses therefore persists at network level, although the learning-based results remain exploratory and should be interpreted accordingly. This addresses the fourth research question and supports the view that the power model is part of the environment definition in learning-based energy management.

**Methodological implications.** Several methodological implications follow from these results. First, the BS power model should be treated as a first-order modeling assumption and reported explicitly in comparative studies. Second, when sleep modes are involved, average-power results should be complemented by diagnostic indicators such as conditional `ACTIVE` power, state occupancies, and wake-up behavior, since similar instantaneous savings may lead to different long-run outcomes. Third, radio reconfiguration levers should be interpreted carefully: actions that reduce active-state power may still reduce idle opportunities and therefore weaken the net benefit once traffic and queueing effects are taken into account. Finally, in learning-based control, the adopted power model should be recognized as part of the reward and environment specification rather than as a neutral implementation detail.

**Limitations.** The proposed framework was designed to isolate model-driven effects under controlled conditions and therefore relies on simplified abstractions. Traffic and service are modeled at a reduced level, the sleep-mode controller is intentionally lightweight, and PHY/MAC details, vendor-specific constraints, and full deployment realism are outside the present scope. In the multi-BS extension, the use of tabular learning and simplified interference coupling further limits scalability and realism. For these reasons, the reported results should be interpreted primarily as *relative* differences induced by power-model selection rather than as absolute performance predictions.

**Future work.** Several extensions can build on this framework. A first direction is to enrich the evaluation with additional KPIs, including broader energy-delay trade-offs, stricter QoS indicators, and penalties related to mode-switching frequency. A second direction is to move from the present lightweight learning extension towards scalable coordinated multi-agent methods and to study robustness under model mismatch. A third direction is to increase realism through heterogeneous hardware configurations, multi-carrier operation, richer traffic models including traces, and calibration or validation against measurements or higher-fidelity simulators.

Overall, the thesis shows that BS power-model selection is not a secondary implementation choice, but a methodological variable that can materially affect the interpretation of energy-efficiency studies and the behavior of energy-aware controllers. Within the adopted scope, the work provides a controlled framework for identifying which trends are robust across model formulations and which are primarily induced by the selected power-consumption abstraction.

# References

[1] 3rd Generation Partnership Project (3GPP). TR 38.864 V18.0.0: Study on Energy Efficiency Enhancements for 5G Advanced (Release 18). Technical report, 3GPP, June 2024. Available at: https://www.3gpp.org/DynaReport/38864.htm.

[2] 6G-IA Vision and Societal Challenges Working Group. Sustainability of 6G: Ways to Reduce Energy Consumption. White paper, 6G-IA, 2024. Version 1.2 (Nov. 2023–Dec. 2024). doi:10.5281/zenodo.13986789.

[3] A. El Amine, J.-P. Chaiban, H. Al Haj Hassan, P. Dini, L. Nuaymi, and R. Achkar. Energy Optimization with Multi-Sleeping Control in 5G Heterogeneous Networks using Reinforcement Learning. *IEEE Transactions on Network and Service Management*, 19(4):4310–4322, 2022. doi:10.1109/TNSM.2022.3157650.

[4] Ali El Amine, Paolo Dini, and Loutfi Nuaymi. Reinforcement Learning for Delay-Constrained Energy-Aware Small Cells with Multi-Sleeping Control. In *IEEE Int. Conf. on Communications Workshops (ICC Workshops)*, pages 1–6, 2020. doi:10.1109/ICCWorkshops49005.2020.9145431.

[5] G. Auer et al. How Much Energy is Needed to Run a Wireless Network? *IEEE Wireless Communications*, 18(5):40–49, October 2011. doi:10.1109/MWC.2011.6056691.

[6] Juan Guillermo Borja, Philipp Bruhn, and Marina Petrova. Towards 6G: Leveraging Advanced Sleep Modes to Improve Energy Performance. In *IEEE Int. Conf. on 6G Networking (6GNet)*, pages 112–116, 2024. doi:10.1109/6GNET63182.2024.10765786.

[7] Sana Elayoubi, Julien Monteil, and Anastasios Aravanis. A Queuing Model for Short Time-Scale Sleep Mode Evaluation in 5G Base Stations. In *IEEE Mediterranean Communication and Computer Networking Conference (MeditCom)*, 2025.

[8] A. Fehske, G. Fettweis, J. Malmodin, and G. Biczók. The global footprint of mobile communications: The ecological and economic perspective. *IEEE Communications Magazine*, 49(8):55–62, August 2011. doi:10.1109/MCOM.2011.5978416.

[9] GreenTouch Consortium. GreenTouch Final Results: The Green Meter Research Study. Technical report, GreenTouch Foundation, June 2015. White Paper V1.0. Available online [accessed Oct. 2025].

[10] D. Laselva et al. On the Potential of Radio Adaptations for 6G Network Energy Saving. In *European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, 2024. doi:10.1109/EuCNC/6GSummit60053.2024.10597015.

[11] P. Lähdekorpi, M. Hronec, P. Jolma, and J. Moilanen. Energy Efficiency of 5G Mobile Networks with Base Station Sleep Modes. In *IEEE Conf. on Standards for Communications and Networking (CSCN)*, 2017.

[12] M. Masoudi et al. Green Mobile Networks for 5G and Beyond. *IEEE Access*, 7, 2019. doi:10.1109/ACCESS.2019.2932777.

[13] M. Masoudi, M. G. Khafagy, E. Soroush, D. Giacomelli, S. Morosi, and C. Cavdar. Reinforcement Learning for Traffic-Adaptive Sleep Mode Management in 5G Networks. In *Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 1–6, 2020. doi:10.1109/PIMRC48278.2020.9217276.

[14] M. Masoudi, E. Soroush, J. Zander, and C. Cavdar. Digital Twin Assisted Risk-Aware Sleep Mode Management Using Deep Q-Networks. *IEEE Transactions on Vehicular Technology*, 72(1):1224–1239, January 2023. doi:10.1109/TVT.2022.3206498.

[15] M. Meo. On the sustainability of communications and networking. In *Proc. 31st International Telecommunication Networks and Applications Conference*

*(ITNAC)*, pages 147–150, Sydney, Australia, November 2021. doi:10.1109/ITNAC53136.2021.9652139.

[16] Orange. Orange's vision for 6G. White paper, Orange, March 2022. White Paper.

[17] S. K. G. Peesapati, M. Olsson, M. Masoudi, S. Andersson, and C. Cavdar. Q-learning based Radio Resource Adaptation for Improved Energy Performance of 5G Base Stations. In *Proc. IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC)*, pages 979–984, 2021. doi:10.1109/PIMRC50174.2021.9569420.

[18] N. Piovesan et al. Machine Learning and Analytical Power Consumption Models for 5G Base Stations. *IEEE Communications Magazine*, 60(10):56–62, October 2022. doi:10.1109/MCOM.001.2200023.

[19] F. E. Salem et al. Reinforcement Learning Approach for Advanced Sleep Modes Management in 5G Networks. In *IEEE Vehicular Technology Conference (VTC-Fall)*, 2018.

[20] Jann Camilo Sánchez Huertas, Greta Vallero, Loutfi Nuaymi, and Anne-Cécile Orgerie. Review and classification of the use of SMs for energy saving in next-generation radio access networks. *Computer Networks*, 273, 2025. doi:10.1016/j.comnet.2025.111745.

[21] G. Vallero, A.-C. Orgerie, and L. Nuaymi. Energy Efficiency in 5G and 6G Radio Access Networks: A Comprehensive Survey. *Computer Communications*, 2025. Available online: https://univ-rennes.hal.science/IRISA/hal-05310147v1.

[22] Shuai Zhang, Tianzhang Cai, Özlem Tuğfe Demir, and Cicek Cavdar. Multi-Agent RL for Sleep Mode and Antenna Configuration With User Offloading Under Dynamic Traffic in Massive MIMO Networks. *IEEE Transactions on Vehicular Technology*, 74(6):9734–9749, June 2025. doi:10.1109/TVT.2025.3541136.