

Alqahtani, Talal Marei Hussain (2026) *Blockchain for spectrum management, trading, and settlement in next-generation wireless networks.* PhD thesis.

https://theses.gla.ac.uk/86343/

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten: Theses
https://theses.gla.ac.uk
research-enlighten@glasgow.ac.uk

Blockchain For Spectrum Management, Trading, and Settlement in Next-Generation Wireless Networks

Talal Marei Hussain Alqahtani

Submitted in fulfilment of the requirements for the
Degree of Doctor of Philosophy

James Watt School of Engineering
College of Science and Engineering
University of Glasgow



June 2026

Abstract

The rapid growth of mobile traffic and localised capacity pressure in 5G and beyond networks has exposed the limitations of conventional licensed spectrum management. Long-term exclusive spectrum assignment provides regulatory order and interference protection, but it is poorly suited to short-term, location-dependent variations in demand. In practice, spectrum may remain underutilised in some areas while others experience congestion, leading to inefficient use of a scarce resource. Existing dynamic spectrum trading approaches address part of this problem. Still, many remain limited at the execution layer: they assume trusted intermediaries, offer incomplete settlement models, and lack auditable coordination among independent stakeholders. Meanwhile, in shared indoor environments where multiple operators use neutral-host infrastructure, settlement between operators is typically governed by static fees that do not reflect the radio performance actually delivered to users. This thesis investigates how blockchain, smart contracts, and learning-assisted demand analysis can be combined to support flexible, accountable, and operationally credible spectrum coordination and settlement in 5G and beyond wireless networks. The work is developed around three connected contributions.

The first contribution presents a blockchain-based decentralised spectrum trading marketplace implemented on the Ethereum Sepolia test network. The framework defines the roles of three main entities. The Telecommunications Regulatory Authority, spectrum holders, and spectrum buyers are defined, and a market mechanism comprising a dynamic posted-price function with a scarcity multiplier, a rule-based batch matching algorithm, and interference-aware feasibility constraints that ensure spatial, temporal, and frequency-domain coexistence among allocations is developed. Four smart contract functions cover the complete lifecycle of a spectrum trade, from listing through matching to payment settlement and ownership transfer. Security evaluation using static analysis across one hundred vulnerability detectors confirmed robustness against replay, payment manipulation, and front-running attacks. In contrast, a queueing-based scalability analysis demonstrated that the marketplace can support meaningful trading

workloads within blockchain throughput constraints.

The second contribution extends this foundation into a demand-aware and learning-assisted spectrum coordination framework. Analysis of cellular traffic data from Milan offers concrete evidence of underutilised spectrum across space and time. By applying a normalised index of utilisation together with a two-tier k-means clustering approach, the results show clear demand heterogeneity at both the city-wide and sub-regional level. An artificial neural network (ANN)-based traffic prediction model is introduced to enable proactive spectrum trading by forecasting low-demand intervals. To validate the framework in a realistic setting, it was extended and connected to a 5G testbed built on OpenAirInterface. The results show that once a spectrum trade is finalised on-chain, the gNB and UE can be reconfigured automatically, with no changes required to the PHY or MAC layers. This confirms, end-to-end, that blockchain-based spectrum trading can operate alongside real cellular deployments.

The third contribution extends blockchain-based coordination into indoor neutral-host and Open RAN environments, addressing the problem of performance aware inter-operator settlement. A system model based on 3GPP channel models and a Monte Carlo system-level simulation quantifies the distributional improvement in received power, SINR, and achievable downlink throughput for neutral-host indoor service compared to a macro outdoor-to-indoor baseline. A throughput-centric QoS metric and a reward-penalty pricing mechanism are developed to adjust the wholesale price paid by the MNO to the neutral-host operator based on the measured QoS gain and the per-user degradation probability. Settlement is enforced on-chain via a smart contract that forwards payments with immutable audit trails. Sensitivity analysis under Worst and Random admission policies demonstrates that the risk-aware pricing mechanism appropriately rewards beneficial offloading and discounts unstable regimes, aligning neutral-host revenue with delivered performance.

Together, the three contributions demonstrate that blockchain-based coordination can address real and complementary limitations in wireless spectrum management and shared infrastructure settlement. The results show that blockchain and smart contracts provide a credible foundation for auditable, multi-party telecom coordination, that learning-assisted demand analysis can improve responsiveness to localised resource imbalance. Performance-aware pricing can introduce accountability into neutral-host operation. The thesis identifies limitations of the present work and several directions for future investigation, including consortium blockchain architectures, tighter integration with O-RAN intelligent controllers, and richer multi-metric settlement models.

University of Glasgow

College of Science & Engineering

Statement of Originality

Name: Talal Marei Alqahtani
Registration Number: xxxxxxxx

I certify that the thesis presented here for examination for a PhD degree of the University of Glasgow is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it) and that the thesis has not been edited by a third party beyond what is permitted by the University's PGR Code of Practice.

The copyright of this thesis rests with the author. No quotation from it is permitted without full acknowledgement.

I declare that the thesis does not include work forming part of a thesis presented successfully for another degree.

I declare that this thesis has been produced in accordance with the University of Glasgow's Code of Good Practice in Research.

I acknowledge that if any issues are raised regarding good research practice based on review of the thesis, the examination may be postponed pending the outcome of any investigation of the issues.

Signature: ***Talal Marei H. Alqahtani***
Date: 30/07/2026

List of Publications

The following publications are the outcomes of the research done for this PhD.

- **T. Alqahtani**, S. Ansari, Y. Sambo and M. A. Imran, "Blockchain Based Decentralized Spectrum Trading Using Smart Contracts," 2024 7th International Conference on Advanced Communication Technologies and Networking (CommNet), Rabat, Morocco, 2024, pp. 1-5, doi: 10.1109/CommNet63022.2024.10793377.

- **T. Alqahtani**, M. Ozturk, Y. Chen, A. Hameed, S. Ansari, M. A. Imran, and Y. Sambo, "Blockchain-Enabled Dynamic Spectrum Trading for Next-Generation Wireless Networks," *IEEE Open Journal of the Communications Society*, Accepted for publication, 2026.

- **T. Alqahtani**, S. Ansari, O. Onireti, M. A. Imran, and Y. Sambo, "Blockchain-Based Settlement for Indoor Neutral-Host Offloading," *Telecom MDPI*, Submitted, 2026.

Contents

Abstract i

Statement of Originality iii

List of Publications iv

List of Tables x

List of Figures xi

List of Acronyms xv

List of Symbols xviii

Acknowledgements xxiii

1 Introduction 1

1.1 Background and Motivation 1
1.2 Problem Statement 2
1.3 Aim and Objectives of the Thesis 3
1.4 Research Questions 4
1.5 Thesis Contributions 4
1.6 Scope of work 5
1.7 Research Methodology Overview 6
1.8 Significance of the Study 6
1.9 Thesis Organisation 8

2 Literature Review 11

2.1 Introduction 11
2.2 Traditional Spectrum Management and Its Limitations 12
2.2.1 Conventional Licensed Spectrum Allocation 12
2.2.2 Inefficiencies of Static Spectrum Assignment 12
2.2.3 Implications for 5G and Beyond 13

2.3 Dynamic Spectrum Access and Spectrum Trading 14
2.3.1 Dynamic Spectrum Access 14
2.3.2 Shared Access Spectrum Frameworks 14
2.3.3 Spectrum Trading and Secondary Spectrum Markets . . . 17
2.3.4 Auction-Based and Market-Oriented Spectrum Allocation Approaches 18
2.4 Blockchain and Smart Contracts in Telecommunications 19
2.4.1 Blockchain as a Coordination Infrastructure 19
2.4.2 Smart Contracts as Programmable Execution Logic 20
2.4.3 Permissioning, Access Control, and Auditability 20
2.4.4 Blockchain Limitations in Telecom Systems 21
2.5 Blockchain-Based Spectrum Management and Spectrum Trading . 22
2.5.1 Early Blockchain-Based Spectrum Management Concepts 22
2.5.2 Blockchain for Spectrum Sharing and Trading 22
2.5.3 Smart-Contract-Enabled Market Mechanisms 23
2.5.4 Consortium, Hybrid, and Public Blockchain Models 23
2.5.5 Critical Comparison of Existing Blockchain-Based Spectrum Studies 25
2.6 Learning-Assisted Demand Analysis and Spectrum Prediction . . 26
2.6.1 Learning for Traffic Prediction and Spectrum Awareness . 26
2.6.2 Spatial Heterogeneity and Clustering-Based Analysis . . . 27
2.6.3 Prediction for Proactive Spectrum Trading 27
2.6.4 Limits of ML-Only Approaches 28
2.7 Interference, Coexistence, and Regulatory Constraints 28
2.7.1 Interference-Aware Sharing and Reuse 28
2.7.2 Policy and Regulatory Feasibility 29
2.7.3 Feasibility Constraints 29
2.8 Indoor Offloading and Neutral-Host Settlement 30
2.8.1 Indoor Traffic Growth and Outdoor-to-Indoor Limitations 30
2.8.2 WiFi-Based Traffic Offloading 30
2.8.3 Mobile and Cellular Data Offloading 31
2.8.4 HetNet and Neutral-Host Offloading Models 31
2.8.5 Blockchain-Enabled Offloading 32
2.9 Research Gaps and Thesis Positioning 33

3 Blockchain-Based Spectrum Trading Marketplace 36
3.1 Introduction 36
3.2 Positioning Relative to Existing Work 38

3.3 System Architecture . . . 38
3.3.1 Architectural Overview . . . 39
3.3.2 Participating Entities . . . 39
3.3.3 Smart Contracts as the Coordination Mechanism . . . 41
3.3.4 Proof-of-Stake Validators and Blockchain Operation . . . 42
3.3.5 Interaction Flow . . . 42
3.4 Market Mechanism Design . . . 43
3.4.1 Request and Offer Model . . . 43
3.4.2 Dynamic Pricing Function . . . 44
3.4.3 Feasibility Conditions . . . 46
3.4.4 Rule-Based Batch Matching Heuristic . . . 47
3.4.5 Approximation, Clearing, and Stability Properties . . . 48
3.5 Interference and Coexistence Handling . . . 49
3.5.1 Spatial and Temporal Conflict Detection . . . 50
3.5.2 Adjacent Channel and Guard-Band Constraints . . . 50
3.5.3 Interference-Aware Feasibility Conditions . . . 51
3.5.4 Interference-Aware Rule-Based Matching . . . 51
3.5.5 Post-Trade Compliance Monitoring . . . 52
3.6 Ethereum and the Choice of Proof-of-Stake . . . 53
3.6.1 Layered Blockchain Architecture . . . 53
3.7 Implementation . . . 54
3.7.1 Marketplace Overview . . . 55
3.7.2 Ethereum Smart Contract . . . 56
3.7.3 Python Automation Scripts . . . 59
3.8 Results and Evaluation . . . 61
3.8.1 Smart Contract Deployment . . . 61
3.8.2 Spectrum Trading Transaction . . . 63
3.8.3 Performance Metrics . . . 65
3.8.4 Security and Robustness Evaluation . . . 68
3.8.5 Scalability Analysis . . . 71
3.9 Summary . . . 76

4 Blockchain-Enabled Dynamic Spectrum Trading for Next-Generation Wireless Networks 77

4.1 Introduction . . . 77
4.2 Motivation: Empirical Evidence for Spectrum Trading . . . 79
4.2.1 Spatial Traffic Variations . . . 80
4.2.2 Temporal Traffic Variations . . . 83

4.2.3 Quantifying Temporal and Spatial Trading Opportunities 85
4.3 ANN-Based Traffic Prediction for Proactive Trading 87
4.3.1 Role and Structure of the ANN 87
4.3.2 Model Configuration 88
4.3.3 Full and Ablated Models 89
4.3.4 From Reactive to Proactive Trading 91
4.4 Extended Marketplace for Next-Generation Networks 92
4.4.1 Deployment and Consensus Considerations 92
4.4.2 Marketplace Objectives in the 5G and Beyond Context . . 93
4.4.3 Interference-Aware and Policy-Driven Constraints 93
4.5 Integration with an OAI-Based 5G Testbed 93
4.5.1 OAI Architecture and Integration Point 94
4.5.2 Dynamic Configuration of Testbed Components 94
4.6 Discussion 97
4.7 Reproducibility and Experimental Setup 99
4.7.1 Released Artefacts and Configurable Parameters 99
4.7.2 Blockchain Environment 99
4.7.3 5G Testbed Hardware and Software Stack 100
4.8 Summary 101

5 Blockchain-Based Settlement for Indoor Neutral-Host Offloading 102
5.1 Introduction 102
5.2 System Model 104
5.2.1 Achievable Rate Proxy from SINR 108
5.3 QoS Metric and Settlement Mechanism 109
5.3.1 Throughput-Centric QoS 109
5.3.2 Serving-Condition Aggregation and QoS Gain 110
5.3.3 Usage Accounting 110
5.3.4 Reward–Penalty Pricing (Risk-Aware) 111
5.3.5 On-Chain Settlement via Smart Contract 112
5.4 System Architecture and Settlement Workflow 113
5.4.1 Framework Overview 113
5.4.2 Radio Simulation Model 114
5.4.3 Middleware QoS Processing and Usage Accounting 114
5.4.4 Blockchain-Based Settlement Layer and Implementation . 114
5.5 Performance Evaluation and Results 118
5.5.1 On-Chain Settlement Validation 120

5.5.2 Sensitivity Results under Worst vs Random Admission Policies . . . 122
5.6 Summary . . . 132

6 Conclusions and Future Directions **133**
6.1 Summary of the Thesis . . . 133
6.1.1 Blockchain-Based Spectrum Trading Marketplace (Chapter 3) . . . 134
6.1.2 Blockchain-Enabled Dynamic Spectrum Trading (Chapter 4) . . . 134
6.1.3 Blockchain-Based Settlement for Indoor Neutral-Host Offloading (Chapter 5) . . . 135
6.2 Revisiting the Research Questions . . . 136
6.3 Limitations . . . 138
6.4 Future Directions . . . 139

bibliography **141**

List of Tables

2.1 Comparison of representative blockchain-based spectrum studies against key dimensions addressed in this thesis. 25

3.1 Pricing and matching parameters and their admissible ranges. . . 46
3.2 Core Smart Contract Functions 58
3.3 Measured Block-Level Parameters (Sepolia, Block #6536159) . . 72
3.4 Measured Transaction Parameters for `makeOffer()` 72

4.1 Columns of the Milan telecommunications dataset (one record per grid square, 10-minute interval, and country code). 80
4.2 Dimensions ($m \times n$) of the data used at each stage of the analysis. 81
4.3 Colour coding of the clustering maps in Figs. 4.1 and 4.2. 83
4.4 Regression accuracy of the ANN forecaster on the independent 7-day test week (Grid 3089, $N = 1{,}008$ ten-minute slots). Activity is in the dimensionless units of the Milan dataset. Only the MSE was recorded for the ablated model; its RMSE follows from the MSE. 91

5.1 System-level simulation parameters and model sources. 106
5.2 Smart-contract interface for risk-aware settlement and payment forwarding. 116
5.3 End-to-end offloading and settlement phases (message-level view). 117

List of Figures

1.1 Thesis contributions in blockchain-enabled spectrum coordination and settlement. 5
1.2 Thesis outline. 10

2.1 Spectrum access and sharing frameworks, from exclusive licensed access to licence-exempt operation. 15
2.2 Thesis positioning across the reviewed literature. Three gaps identified in the literature are addressed by the contributions in Chapters 3, 4, and 5, respectively. 35

3.1 Blockchain-based spectrum trading architecture. 40
3.2 Layered architecture of the blockchain used by the spectrum trading marketplace. 54
3.3 Blockchain-based spectrum trading workflow. 57
3.4 Smart contract interface in the Remix IDE with MetaMask injected provider on the Sepolia network. 58
3.5 Random spectrum parameters generation output from the Python script. 60
3.6 Smart contract deployment transaction details in MetaMask. . . . 62
3.7 Available spectrum listing before the transaction. 63
3.8 Spectrum purchase transaction interface in MetaMask. 64
3.9 Sold spectrum listing after the transaction. 65
3.10 Gas usage of core spectrum trading transactions on the Ethereum Sepolia network. 66
3.11 Transaction cost of spectrum trading operations derived from measured gas usage. 67
3.12 Throughput characteristics of the proposed spectrum trading marketplace based on measured blockchain parameters. 68
3.13 Transaction revert when attempting to repurchase an already sold spectrum listing. 69
3.14 Transaction revert caused by submitting an incorrect payment value. 70

3.15 Slither static analysis output for the spectrum trading smart contract with Solidity v0.8.26, showing zero findings across 100 vulnerability detectors. 71
3.16 Achieved transaction throughput as a function of arrival rate λ. Throughput increases linearly under low load and saturates at the gas-limited service capacity $\mu \approx 10.42$ tx/s. 73
3.17 End-to-end latency as a function of arrival rate λ for confirmation depths $k = 1$ and $k = 2$. 74

4.1 First-layer clustering results for Milan city grids, grouped by average traffic intensity, revealing spatial demand heterogeneity across the metropolitan area. 82
4.2 Second-layer clustering results for the selected grids, showing micro-level spatial variation in traffic demand within regions of similar average traffic. 83
4.3 Daily fluctuations in traffic demand within a selected grid, showing how temporal variation affects spectrum usage over the course of the day. 84
4.4 Empirical CDF of $U_g(t)$ for a representative high-demand grid, drawn from the top 10% of grids by average traffic demand. The curve traces the proportion of time the grid spends operating below its peak demand. 86
4.5 Spatial heterogeneity in $U_g(t)$ across representative high-demand grids. The boxplots show that utilisation distributions differ considerably by location, pointing to localised spectrum trading opportunities even within high-demand regions. 86
4.6 Workflow of the ANN forecaster: time-of-day and day-of-week encodings feed a single hidden layer of 20 `tansig` neurons and a linear output neuron, whose normalised output is rescaled to activity units. The ablated model omits the day-of-week inputs. . . 88
4.7 ANN-based traffic load prediction for a representative Milan grid (Grid 3089), trained on 14 days of data and evaluated over an independent 7-day test week. Upper panel: observed and forecast activity. The horizontal axis is the index of the 10-minute slot in the test week (1–1,008), and the vertical axis is activity in the dimensionless units of the Milan dataset. Lower panel: mean squared error over the test week, in squared activity units, for the full model (20 neurons) and the ablated model without day-of-week features (10 neurons). 90

4.8 OAI 5G testbed architecture and the integration point of the proposed framework (red). A completed on-chain trade updates the carrier frequency and bandwidth of the gNB and UE without modifying OAI. 95
4.9 Network running with a newly allocated frequency following a completed on-chain spectrum trade. 96
4.10 Network status following spectrum integration, showing the updated parameters and confirming healthy operation across all components. 96
4.11 gNB configuration logs confirming the application of the newly traded spectrum parameters. 97
4.12 End-user equipment configuration logs confirming operation on the newly allocated frequency. 97
4.13 Network topology of the testbed: the UE, gNB, and 5G core network run as containers on a single host, and the Python/Web3 scripts link the host to the Ethereum Sepolia network. 100

5.1 Deployment scenario showing indoor users served by either an outdoor macro base station (O2I path with penetration loss L_{pen}) or an indoor neutral-host small cell (InH path). The macro load factor ρ controls inter-cell interference, while the leakage factor η captures cross-tier coupling into the indoor layer. 105
5.2 Throughput-aware settlement architecture for an indoor neutral-host Open RAN deployment. Throughput KPIs are collected via O1 and processed off-chain by the QoS and pricing engine to compute reward/penalty adjustments. Settlement is executed on-chain via a smart contract maintained by the NHO and verified by lightweight audit clients at participating MNO cores. 113
5.3 Solidity excerpt illustrating configuration and pricing functions (e.g., neutral-host payee update, parameter update, and the read-only pricing function `quotePriceWeiPerGB` used by the middleware). 116
5.4 Empirical CDF of received signal strength (RSRP proxy) for indoor users under macro O2I and neutral-host indoor service. . . . 118
5.5 Empirical CDF of downlink SINR for indoor users under macro and neutral-host operation. 119
5.6 Empirical CDF of achievable downlink throughput derived from SINR-based spectral efficiency mapping. 120
5.7 Penalty-regime settlement on Sepolia: internal transfer of 9.3482795×10^{-5} ETH from the contract to the NHO address (Tx: `0x5cb6...efd9`). 121

5.8 Reward-regime settlement on Sepolia: internal transfer of $2.281673339 \times 10^{-3}$ ETH from the contract to the NHO address (Tx: `0x0c6e...b6e`). 121
5.9 Worst admission: heatmap of mean gain ΔQ versus macro load ρ and leakage η. 122
5.10 Worst admission: mean gain ΔQ versus macro load ρ at fixed leakage $\eta = 0.70$. 123
5.11 Worst admission: penalty probability $p_{\text{pen}} = \Pr(\Delta Q_u < 0)$ versus (ρ, η). 124
5.12 Worst admission: penalty probability versus leakage η at fixed macro load $\rho = 0.70$. 125
5.13 Worst admission: price per GB versus mean gain ΔQ under risk-aware settlement using $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$ with $\lambda = 0.8$ and penalty factor $\kappa = 1.5$ (base price $P_0 = 0.001$ ETH/GB). 126
5.14 Random admission: heatmap of mean gain ΔQ versus macro load ρ and leakage η. 127
5.15 Random admission: mean gain ΔQ versus macro load ρ at fixed leakage $\eta = 0.70$. 128
5.16 Random admission: penalty probability $p_{\text{pen}} = \Pr(\Delta Q_u < 0)$ versus (ρ, η). 129
5.17 Random admission: penalty probability versus leakage η at fixed macro load $\rho = 0.70$. 130
5.18 Random admission: price per GB versus mean gain ΔQ under risk-aware settlement. 131

List of Acronyms

Third Generation Partnership Project	3GPP
Fifth Generation Mobile Network	5G
Sixth Generation Mobile Network	6G
Artificial Intelligence	AI
Artificial Neural Network	ANN
Application Programming Interface	API
Augmented Reality	AR
Citizens Broadband Radio Service	CBRS
Cumulative Distribution Function	CDF
Directed Acyclic Graph	DAG
Distributed Ledger Technology	DLT
Ether (Ethereum Cryptocurrency)	ETH
Ethereum Virtual Machine	EVM
Federal Communications Commission	FCC
Gigabyte	GB
Next Generation Node B	gNB
GPRS Tunnelling Protocol – User Plane	GTP-U
Heterogeneous Network	HetNet
Internet Engineering Task Force	IETF

Indoor Hotspot (3GPP Channel Model)	InH
Internet Protocol	IP
Internet Protocol Security	IPsec
Internet of Things	IoT
International Telecommunications Society	ITS
International Telecommunication Union Radiocommunication Sector	ITU-R
Key Performance Indicator	KPI
Local IP Access	LIPA
Medium Access Control	MAC
Mobile Network Operator	MNO
Next Generation Application Protocol	NG-AP
Neutral-Host Operator	NHO
New Radio	NR
Open Radio Access Network	O-RAN
Open Central Unit – Control Plane	O-CU-CP
Open Central Unit – User Plane	O-CU-UP
Open Distributed Unit	O-DU
Open Radio Unit	O-RU
Outdoor-to-Indoor	O2I
Physical Layer	PHY
Proof of Stake	PoS
Proof of Work	PoW
Physical Resource Block	PRB
Quality of Service	QoS
Radio Access Network	RAN

RAN Intelligent Controller	RIC
Reference Signal Received Power	RSRP
Spectrum Buyer	SB
Software-Defined Networking	SDN
Spectrum Holder	SH
Selected IP Traffic Offload	SIPTO
Signal-to-Interference-plus-Noise Ratio	SINR
Service Level Agreement	SLA
Service Management and Orchestration	SMO
Short Message Service	SMS
Telecommunications Regulatory Authority	TRA
User Equipment	UE
Urban Macro (3GPP Channel Model)	UMa

List of Symbols

$A_g(t)$	Aggregated traffic activity observed in grid g at time interval t.
$A_{i,t}(L, F)$	Available spectrum capacity offered by holder i at time t for location L and band F.
A_j	Available bandwidth of spectrum resource S_j.
B	System bandwidth.
$b^\star$	Serving macro base station index.
B_k	Bandwidth requested by buyer k.
C	Total settlement cost.
d_{3D}	Three-dimensional transmitter–receiver distance.
$D_t(L, F)$	Aggregated demand at location L, band F, and time t.
f_c	Carrier frequency.
F	Frequency band.
F_k	Frequency band requested by buyer k.
$G^{(b)}$	Antenna gain of transmitter b (dB).
$G_{\min}$	Minimum guard band.
I_u	Aggregate interference power for user u.
I_{macro}	Inter-cell macro interference power.
I_{leak}	Leaked macro interference into the neutral-host indoor layer.

$I(k, j)$	Interference indicator between allocations k and j.
K_{nho}	Neutral-host service capacity (users per accounting interval).
L	Geographic location.
L_k	Requested location of buyer k.
L_{pen}	Building penetration loss (dB).
N_0	Noise power in linear scale.
NF	Receiver noise figure (dB).
N_{off}	Number of offloaded users admitted to the neutral-host.
P	Adjusted price per GB after reward–penalty multiplier.
P_0	Base price per GB.
p_{pen}	Penalty probability: fraction of offloaded users experiencing QoS degradation.
$P_t(L, F, B, T)$	Posted market price for location L, band F, bandwidth B, and duration T.
$P_t^{(b)}$	Transmit power of transmitter b (dBm).
$P_{r,u}^{(b)}$	Received power at user u from transmitter b.
$P_{s,u}$	Serving signal power for user u.
PL_{InH}	Indoor Hotspot pathloss (3GPP InH model).
PL_{UMa}	Urban Macro pathloss (3GPP UMa model).
Q_{eff}	Risk-aware effective gain: $\Delta QoS - \lambda\, p_{\text{pen}}$.
QoS_u	Normalised QoS of user u.
$\overline{QoS}_s$	Mean QoS under serving condition s.
ΔQoS	QoS gain achieved by neutral-host service over macro service.
R_i	Revenue generated by spectrum holder i.

R_k	Spectrum request submitted by buyer k.
R_{int}	Policy-defined interference radius.
S_j	Spectrum resource or spectrum listing indexed by j.
$S_t(L, F)$	Aggregated supply at location L, band F, and time t.
T	Trading duration.
T_k	Requested duration of buyer k.
T_u	Achievable throughput of user u.
T_u^{peak}	Peak achievable-rate proxy for user u (before contention).
$T_{u,\text{nho}}^{\text{del}}$	Delivered NHO throughput for user u (after contention scaling).
$\mathcal{U}$	Set of all indoor users within an accounting interval.
$U_g(t)$	Normalised demand-based utilisation index.
U_{GB}	Metered traffic usage in gigabytes.
U_{mGB}	Metered traffic usage in milli-gigabytes.
W	Set of accepted spectrum requests.
X_σ	Log-normal shadowing term.
α	Reward–penalty pricing multiplier.
β	Neutral-host load-coupling factor.
Δ	Market clearing interval duration.
ϵ	Small positive constant used for numerical stabilisation.
η	Cross-tier leakage factor (Chapter 5) or pricing sensitivity exponent (Chapter 3).
κ	Penalty amplification factor.
λ	Risk-sensitivity weight (Chapter 5) or arrival rate (Chapter 3).

μ	Service rate of blockchain transaction processing.
ρ	Macro-layer load factor.
ρ_{nho}	Neutral-host load factor.
σ_{In}	Indoor shadowing standard deviation (dB).
τ	Accounting interval duration.
$\phi_t(L, F)$	Scarcity multiplier in dynamic pricing.

وَآخِرُ دَعْوَاهُمْ أَنِ الْحَمْدُ لِلَّهِ رَبِّ الْعَالَمِينَ

"The last of their call will be, Praise to Allah, Lord of the worlds" Quran, Yunus, Verse [10].

Acknowledgements

All praise and thanks are due to God, who granted me the wisdom to begin this journey and guided me through every step of it.

I would like to express my sincere gratitude to my supervisors, Dr. Yusuf Sambo, Prof. Muhammad Ali Imran, and Dr. Shuja Ansari, for their invaluable support, guidance, and patience throughout the years of my PhD.

I am also deeply grateful to King Khalid University for awarding me a scholarship and for providing additional support toward my travel expenses during this PhD journey.

I would also like to acknowledge the Saudi Arabian Cultural Bureau in London, whose support was instrumental to my success.

I would like to thank my colleagues for their encouragement and support. The productive and enjoyable discussions we shared during my PhD made this journey both enriching and memorable.

To my Saudi friends in Glasgow, I am truly thankful for your companionship and support, which helped ease the stresses of doctoral life.

Finally, I wish to express my heartfelt gratitude to my family, especially my father and mother, for their endless encouragement and prayers. Above all, I am deeply grateful to my wife for her unwavering love, patience, understanding, and support throughout my PhD journey.

Chapter 1

Introduction

1.1 Background and Motivation

Wireless communication networks are experiencing sustained growth in traffic volume, service diversity, and performance demands. The evolution toward 5G and beyond has introduced denser deployments, more heterogeneous applications, stricter latency requirements, and stronger expectations for localised service quality [1]. At the same time, radio spectrum remains a finite and tightly regulated resource [2, 3]. Although licensed spectrum allocation has historically enabled orderly governance and interference protection, it is still commonly based on long-term and relatively static assignment structures that do not always reflect short-term, location-dependent variations in actual network demand [1, 4].

This mismatch creates an important efficiency problem. In practice, operator demand is not constant across time and space. Traffic intensity varies by hour, by location, by user concentration, and by service type. As a result, valuable licensed spectrum may remain underutilised in some regions or periods, while other operators or locations experience congestion and reduced service quality. This problem becomes more significant in 5G and beyond, where network conditions are increasingly dynamic and where localised demand peaks may emerge faster than conventional coordination mechanisms can respond [1, 5].

To address these limitations, research has increasingly explored more flexible approaches to spectrum coordination. Dynamic spectrum access and secondary spectrum market models aim to improve utilisation by enabling temporary or conditional reassignment of spectrum resources [5, 6]. However, while these approaches improve flexibility in theory, many existing models still rely on assumptions of trusted intermediaries, simplified settlement, or coordination mechanisms that are difficult to audit and automate in practice [1, 7].

In parallel, blockchain and smart contracts have emerged as promising tools

for multi-party coordination in telecommunications [8, 9]. Blockchain provides tamper-evident record-keeping and distributed state consistency, while smart contracts provide programmable execution of rules such as authorisation, exchange conditions, payment handling, and settlement logging [7, 10]. These characteristics make blockchain-based coordination particularly valuable in settings where several operators or stakeholders need to work together without relying entirely on a single trusted central party [11].

At the same time, learning-assisted traffic analysis offers another important capability. Machine learning techniques can help identify recurring traffic patterns, characterise spatio-temporal heterogeneity, and predict likely periods of underutilisation or congestion [12, 13]. In the context of dynamic spectrum coordination, such predictive insight can support more proactive decision-making by identifying likely trading windows before demand imbalance becomes operationally visible [14].

A related coordination problem also appears in indoor and shared infrastructure deployments. As more mobile traffic originates indoors and shared-access models such as neutral-host and O-RAN architectures become more relevant, issues of accountability, service attribution, and settlement are increasingly important. In such settings, it is not sufficient to share infrastructure. Mechanisms are also needed to verify delivered performance and support fair compensation among participating operators or providers [11, 15].

Against this background, this thesis investigates how blockchain, smart contracts, and learning-assisted demand analysis can be combined to support more flexible, trusted, and accountable spectrum coordination and settlement mechanisms.

1.2 Problem Statement

Traditional licensed spectrum management remains too rigid to efficiently support short-term and localised resource variation in modern wireless networks. Although dynamic spectrum access and spectrum trading models have been proposed to improve utilisation, many existing approaches do not fully address trusted authorisation, multi-party coordination, auditable settlement, and implementation-level accountability [1, 7]. Furthermore, while machine learning can improve the prediction of traffic demand and potential underutilisation, prediction alone cannot enforce access rights, execute transactions, or provide a trustworthy audit trail [14].

As wireless systems evolve toward denser, more heterogeneous, and more

shared operational environments, there is a need for a coordination framework that can do more than merely identify spectrum inefficiency. Such a framework should support short-term and location-aware trading opportunities, verify feasibility constraints, automate transaction execution, and maintain trustworthy settlement records across multiple independent participants. A similar challenge also arises in indoor neutral-host and shared RAN environments, where infrastructure sharing introduces the need for performance-aware accountability and settlement based on delivered service rather than static assumptions [11, 15].

This thesis addresses these problems by investigating blockchain and smart contract coordination mechanisms for spectrum trading and by extending the same coordination logic toward QoS-aware settlement in shared indoor telecom environments.

1.3 Aim and Objectives of the Thesis

The aim of this thesis is to design and evaluate blockchain-enabled coordination frameworks that improve trust, accountability, and operational flexibility in dynamic spectrum trading and related multi-operator telecom resource-sharing scenarios.

To achieve this aim, the thesis pursues the following objectives:

1. To examine the limitations of conventional licensed spectrum allocation in the context of short-term and localised demand variation.

2. To analyse dynamic spectrum access and spectrum trading mechanisms as flexible alternatives to static assignment.

3. To design a blockchain-based decentralised spectrum trading framework capable of supporting trusted multi-party coordination.

4. To incorporate smart-contract logic for authorisation, transaction execution, and settlement recording in spectrum exchange.

5. To investigate how spatio-temporal traffic analysis and learning-assisted demand characterisation can support proactive spectrum trading decisions.

6. To consider interference-aware and policy-aware feasibility constraints in blockchain-based spectrum coordination.

7. To extend blockchain-based accountability mechanisms to indoor neutral-host and O-RAN settings through performance-aware settlement based on delivered QoS.

1.4 Research Questions

The thesis is guided by the following research questions:

- How can licensed spectrum coordination be made more flexible and responsive to short-term, location-dependent demand variation?
- How can blockchain and smart contracts be used to support trusted authorisation, transaction execution, and auditable settlement in dynamic spectrum trading?
- How can spatio-temporal traffic analysis and learning-assisted prediction improve the identification of proactive spectrum trading opportunities?
- How can interference and regulatory feasibility constraints be integrated into dynamic spectrum coordination mechanisms?
- How can blockchain-based coordination be extended to indoor neutral-host and O-RAN environments to support performance-aware settlement?

1.5 Thesis Contributions

This thesis makes three connected contributions as demonstrated in Figure 1.1 and can be summarized as follows:

- **A blockchain-based spectrum trading marketplace (Chapter 3).** The original contribution is a marketplace in which spectrum holders list spectrum-use rights and spectrum buyers purchase them through an Ethereum smart contract. The listing, purchase, ownership transfer, and payment settlement were implemented as a Solidity smart contract and tested on the Ethereum Sepolia test network, including security tests and static analysis. The dynamic pricing function, the batch-matching algorithm, and the interference-aware constraints are proposed at the design level, and the scalability results are analytical.
- **A demand-aware and testbed-integrated extension (Chapter 4).** The original contribution is linking real traffic data and a 5G testbed to the marketplace. The Milan traffic analysis, the clustering, and the ANN forecasting were implemented and evaluated, and the integration with an OpenAirInterface 5G testbed in RF-simulator mode was demonstrated. The ANN forecast is illustrative and is not connected to the smart contract in a closed loop.

- **A performance-aware settlement framework for indoor neutral-host offloading (Chapter 5).** The original contribution is a settlement mechanism that links the price paid by an MNO to the quality of service delivered by a neutral-host operator. The radio simulation, the QoS and pricing computation, and the on-chain settlement contract were implemented and tested on Sepolia. The integration with O-RAN management interfaces is proposed as a target architecture.

Across all three contributions, operations that require trust between competing parties, such as ownership transfer, payment, and settlement records, are executed on-chain. Computationally heavy or data-intensive operations, such as matching, traffic analysis, prediction, radio simulation, and QoS computation, are executed off-chain, and only their outcomes are recorded on the blockchain.

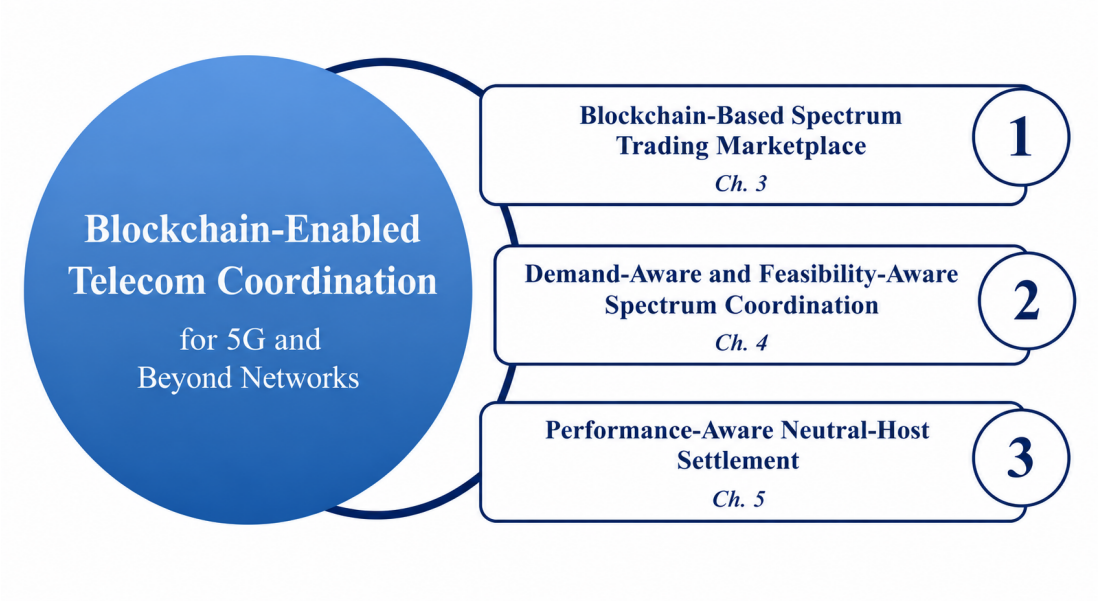


Figure 1.1: Thesis contributions in blockchain-enabled spectrum coordination and settlement.

1.6 Scope of work

This thesis focuses on dynamic coordination of licensed or shared telecom resources in multi-party environments. Its primary emphasis is on short-term and localised spectrum trading and coordination rather than broad national spectrum policy reform. The work does not attempt to replace regulator authority or redefine the legal foundations of licensed spectrum ownership. Instead, it investigates coordination mechanisms that could operate within controlled or regulated sharing contexts.

The learning component of the thesis is used to support demand characterisation and prediction, not to replace rule-based authorisation or smart-contract execution. Similarly, blockchain is used as a coordination and settlement infrastructure rather than as a universal substitute for all network control functions. The indoor extension of the thesis focuses on neutral-host and O-RAN accountability and settlement, rather than attempting a complete redesign of indoor network architecture.

1.7 Research Methodology Overview

The thesis adopts a design-oriented and evaluation-driven research methodology. It begins with a structured review of the literature on conventional spectrum management, dynamic spectrum access, spectrum trading, blockchain in telecommunications, learning-assisted demand analysis, interference-aware feasibility, and indoor shared infrastructure [16]. Based on the identified gaps, blockchain-based coordination frameworks are designed to address trusted execution and settlement limitations in existing approaches.

Following the research onion of Saunders *et al.* [17], this research adopts a pragmatic philosophy, combining deductive and inductive approaches. The strategy is design science, with a mixed-methods choice. The time horizon is cross-sectional, focusing on specific implementations and evaluations. Techniques include literature review, smart contract development and deployment on the Ethereum Sepolia test network, simulation, integration with an OpenAirInterface 5G testbed, and empirical analysis of traffic data.

The research then combines conceptual framework development with implementation-oriented validation. Smart-contract and blockchain prototypes are used to examine execution logic, transaction behaviour, and practical coordination workflows. In addition, traffic analysis and learning-assisted methods are incorporated to study how proactive trading opportunities can be identified from spatio-temporal demand variation. The resulting frameworks are analysed with attention to feasibility, execution cost, latency, and settlement accountability.

1.8 Significance of the Study

This thesis is significant in both academic and practical terms. Academically, this thesis connects several streams of literature that are often treated separately: spectrum trading, blockchain-based telecom coordination, traffic prediction, interference-aware feasibility, and neutral-host settlement. By bringing these

areas together, the thesis contributes a more integrated perspective on how flexible telecom resource coordination can be achieved in practice.

Practically, the work is relevant to operators, infrastructure providers, and telecom system designers seeking more responsive and accountable coordination mechanisms. In spectrum trading, the thesis contributes toward operational frameworks that can support trusted transactions rather than merely theoretical allocation models. In indoor neutral-host environments, it contributes toward settlement approaches that are more closely linked to actual service delivery and measurable performance.

1.9 Thesis Organisation

The remainder of this thesis is organised as follows.

Chapter 2 presents the literature review and establishes the research context of the thesis. It begins by examining conventional licensed spectrum management and the limitations of static allocation in wireless networks. It then reviews dynamic spectrum access, spectrum trading, and secondary spectrum markets as more flexible alternatives to rigid allocation. The chapter further discusses blockchain and smart contracts in telecommunications, blockchain-based spectrum management and trading, learning-assisted demand analysis and spectrum prediction, interference and regulatory feasibility constraints, and indoor neutral-host settlement. The chapter concludes by identifying the research gaps that motivate the thesis.

Chapter 3 introduces the first core contribution of the thesis: a blockchain-based decentralised spectrum trading framework. The chapter presents the system architecture and identifies the three principal entities involved in the trading process: the Telecommunications Regulatory Authority, spectrum holders, and spectrum buyers. It develops a market mechanism comprising a dynamic posted-price function, a rule-based batch matching algorithm, and interference-aware feasibility constraints that ensure spatial, temporal, and frequency-domain coexistence among allocations. The framework is implemented as a set of Ethereum smart contracts deployed on the Sepolia test network, and evaluated in terms of transaction execution, security robustness, and scalability.

Chapter 4 extends the initial trading framework by incorporating demand-aware and learning-assisted spectrum coordination. The chapter provides empirical evidence for spectrum trading opportunities through analysis of the Milan city dataset, using a normalised utilisation index and a two-layer k-means clustering procedure to expose spatio-temporal demand heterogeneity. An ANN-based traffic prediction model is introduced to enable proactive spectrum trading by forecasting low-demand intervals. The extended framework is integrated with an OpenAirInterface-based 5G testbed. This shows that a completed on-chain spectrum trade can trigger the automated reconfiguration of the gNB and UE without altering the underlying PHY/MAC mechanisms, thereby providing end-to-end validation of the proposed approach.

Chapter 5 develops the final major contribution of the thesis by extending blockchain-based coordination into indoor neutral-host environments and performance-aware inter-operator settlement. The chapter introduces a system model based on 3GPP channel models and a Monte Carlo system-level simu-

lation that quantifies the distributional improvement in throughput of neutral-host indoor service over macro outdoor-to-indoor coverage. A throughput-centric QoS metric and a reward-penalty pricing mechanism are developed to adjust the wholesale price paid by the MNO to the neutral-host operator based on measured QoS gains and per-user degradation risk. Settlement is enforced on-chain via an Ethereum smart contract deployed on the Sepolia test network. Running the sensitivity analysis under Worst and Random admission policies shows something useful: the risk-aware pricing mechanism reacts to changing macro load and cross-tier leakage conditions, it discounts regimes that turn out unstable, and it rewards the ones that deliver consistent performance improvements.

Chapter 6 presents the concluding remarks of the thesis and provides a comprehensive synthesis of the research as a whole. The chapter first summarises the thesis's motivation and the broader research problem arising from the limitations of static spectrum management, the execution-layer weaknesses of existing spectrum trading approaches, and the lack of accountable settlement mechanisms in shared telecom infrastructures. It then consolidates the principal contributions of the thesis by revisiting the blockchain-based decentralised spectrum trading framework introduced in Chapter 3, the extended demand-aware and feasibility-aware spectrum coordination framework developed in Chapter 4, and the blockchain-enabled performance-aware settlement framework for indoor neutral-host Open RAN environments presented in Chapter 5. The chapter also revisits the research questions set out in Chapter 1 and discusses how they have been answered through the theoretical development, implementation, and evaluation carried out in the thesis. Finally, it reflects on the limitations of the present work and identifies several promising avenues for future investigation, including more advanced interference-aware coordination, heterogeneous spectrum and multi-RAT resource trading, learning-integrated market mechanisms, scalable consortium blockchain architectures, tighter integration with programmable network control and orchestration platforms, and richer multi-metric settlement models that incorporate broader QoS and SLA considerations.

The main sections of the chapters in this thesis are indicated in Figure 1.2.

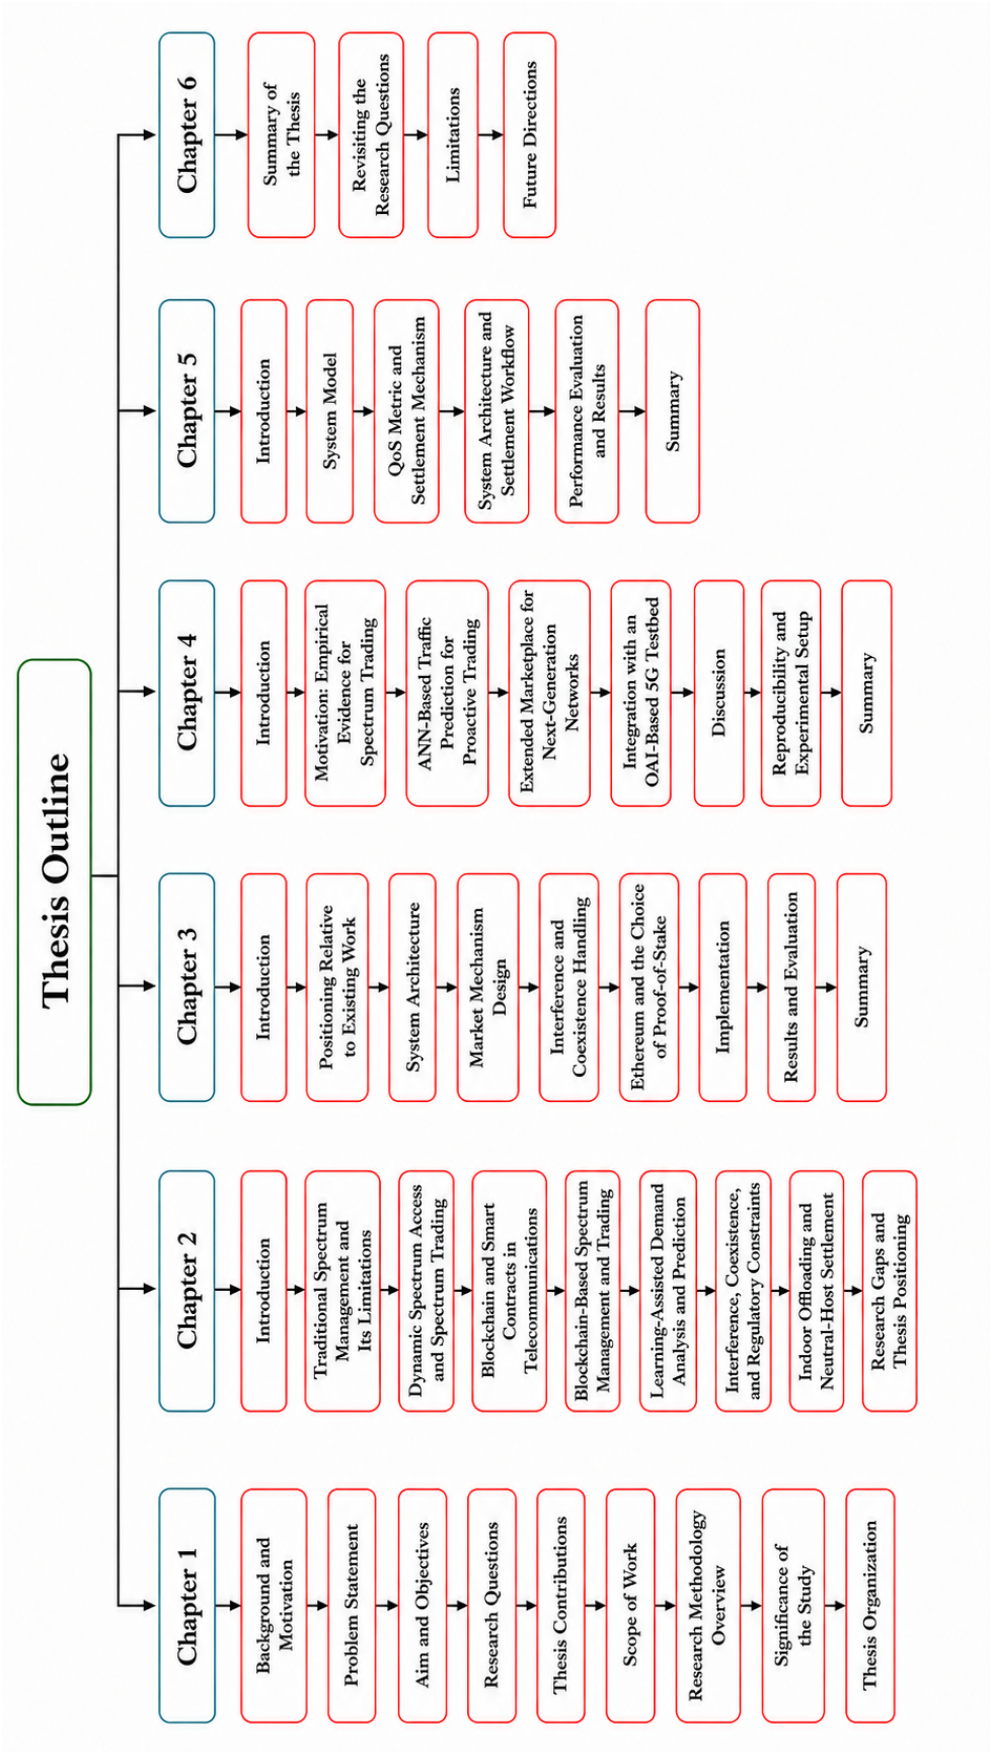


Figure 1.2: Thesis outline.

Chapter 2

Literature Review

2.1 Introduction

Wireless spectrum management has become increasingly challenging due to the rapid growth of mobile traffic, service heterogeneity, and the dynamic requirements of modern and future cellular networks [1, 18]. Conventional regulatory and operator-led spectrum assignment approaches have historically relied on long-term, static allocation mechanisms that often fail to reflect short-term, location-dependent demand. As a result, valuable spectrum resources may remain underutilised in certain regions and time periods, while other locations experience congestion and access delays [4, 6]. These inefficiencies motivate the search for more flexible, transparent, and responsive coordination mechanisms. Recent research has explored several possible approaches to address this problem, including dynamic spectrum access, secondary spectrum markets, blockchain-based coordination, and learning-assisted demand analysis [4, 5, 7, 12]. Dynamic access frameworks aim to improve utilisation by enabling spectrum reassignment across time, location, and stakeholders [18, 19].

Blockchain and smart contracts have emerged as promising tools for automating and securing multi-party coordination, while machine learning techniques provide new ways to identify and anticipate traffic variation that could justify more adaptive spectrum use. In parallel, related developments in indoor neutral-host deployment and Open RAN have created new forms of inter-operator coordination problems, particularly regarding service accountability and settlement. This chapter reviews the literature across these themes and positions the contributions of this thesis within the broader state of the art. It begins with conventional spectrum management and its limitations, then reviews dynamic spectrum access and spectrum trading approaches, and finally examines blockchain and smart contracts as coordination mechanisms. The chapter subsequently discusses

learning-assisted demand analysis, interference and feasibility considerations, and finally, indoor offloading and neutral-host settlement. The chapter concludes by synthesising the main research gaps that motivate the contributions presented in Chapters 3, 4, and 5.

2.2 Traditional Spectrum Management and Its Limitations

2.2.1 Conventional Licensed Spectrum Allocation

Spectrum has traditionally been managed through centralised regulatory frameworks in which national authorities award operators high-priced, long-term, and geographically broad exclusive licenses, a model that Mustonen characterises as the conventional basis of modern spectrum policy and that Sheharyar et al. describe as fixed spectrum access issued for long periods under static conditions [6, 20]. Massaro notes that this arrangement has historically been valued because it provides regulatory certainty, supports interference control, and gives operators clear rights and responsibilities over licensed resources, while Parvini et al. similarly present conventional mobile spectrum management as built around exclusive operator-held bands and clearly defined usage rights [3,21]. At the same time, however, Mustonen emphasises that such licensing structures are typically rigid and poorly matched to short-term or local variations in demand, while Webb et al. argue that legacy exclusive-rights frameworks often require substantial regulatory intervention to modify, making them administratively cumbersome and slow to adapt [20, 22].

2.2.2 Inefficiencies of Static Spectrum Assignment

A central criticism of traditional spectrum management is that it often equates exclusive assignment with efficient utilisation, even though Buddhikot argues that command-and-control style allocation can leave spectrum underused across time and geography, and Ren et al. state explicitly that static spectrum allocation leads to severe underutilisation alongside apparent scarcity [19, 23]. Sheharyar et al. argue that fixed and long-term spectrum assignment can lead to inefficient utilisation and underuse of licensed spectrum, while dynamic and shared-access models are better able to account for variation across time, place, and user demand, and Parvini et al. identify this inefficient use of radio spectrum as one of the main reasons that spectrum sharing frameworks have attracted growing in-

terest [6,21]. In the same direction, Gupta et al. note that newer sharing schemes are specifically motivated by the existence of unused or underutilised spectrum in conventional centralised allocation systems, reinforcing the point that static assignment struggles to reflect real fluctuations in daily, regional, and application-specific demand [1]. The scale of this inefficiency has been documented empirically. The Federal Communications Commission has reported that licensed spectrum utilisation rates range between 15% and 85% [24], while millimetre-wave spectrum holdings remain substantially underutilised even as operators continue to seek additional resources [25].

2.2.3 Implications for 5G and Beyond

The limitations of conventional spectrum management become more severe in 5G and beyond networks. 5G networks are characterised by significantly denser deployments than previous generations, with small cells, micro cells, and indoor access points operating alongside traditional macro base stations. Static spectrum allocation assigns fixed frequency bands to each operator across wide geographic areas, but in a dense heterogeneous deployment, traffic demand varies sharply across cells and over time. A macro cell in a suburban area may be lightly loaded at night while a nearby indoor small cell in a shopping centre is congested during the day. Because static allocation cannot redistribute spectrum between these deployments in response to demand, capacity is wasted in one location while users experience degraded service in another [18]. 5G must simultaneously support services with very different performance requirements: enhanced Mobile Broadband (eMBB) demands high throughput, Ultra-Reliable Low-Latency Communication (URLLC) requires ultra-low latency and high reliability, and massive Machine-Type Communication (mMTC) involves a large number of low-data-rate devices. Static allocation treats all spectrum uniformly regardless of the service being delivered. It cannot dynamically prioritise a URLLC user who needs immediate access to spectrum over an mMTC device that can tolerate delay, nor can it temporarily reallocate underused eMBB spectrum to meet a sudden surge in URLLC demand. This rigidity makes static allocation fundamentally misaligned with the service-diverse requirements of 5G [26]. 5G resource management is further complicated by ultra-low latency requirements, heterogeneous quality-of-service demands, and more complex allocation conditions than earlier generations. In 4G networks, spectrum allocation decisions could be made on relatively long timescales because the service mix was predominantly mobile broadband. In 5G, the coexistence of eMBB, URLLC, and mMTC on the same network means that allocation decisions must account for multiple competing objectives simultane-

ously, which static assignment structures are not designed to handle [27, 28].

For these reasons, more adaptive spectrum coordination frameworks are increasingly necessary. Coordinated dynamic sharing, rather than static hard splitting, can respond to localised demand variations in real time [29]. Temporally and spatially exploiting underutilised spectrum in 5G and beyond can improve overall system efficiency without requiring additional spectrum releases [1]. Dynamic allocation approaches can identify and use underused spectrum while still preserving accountability, coexistence, and regulatory control [30].

2.3 Dynamic Spectrum Access and Spectrum Trading

2.3.1 Dynamic Spectrum Access

Dynamic spectrum access (DSA) has emerged as a response to the inefficiencies of static spectrum assignment. Instead of treating spectrum rights as fixed and immutable, dynamic access approaches seek to enable more flexible use of available resources according to prevailing demand, availability, and coexistence conditions. Early DSA literature frames this shift as a move away from rigid exclusive-use models toward access arrangements that better reflect temporal and spatial variability in actual spectrum use [31, 32]. In mobile and 5G-oriented settings, licensed sharing surveys further show that dynamic access gained prominence because conventional assignment alone could not efficiently accommodate heterogeneous traffic, localised congestion, and growing demand for flexible use of licensed bands [5, 33]. In this sense, DSA is best understood as a broader coordination paradigm for matching spectrum use to changing network conditions rather than opportunistic access.

2.3.2 Shared Access Spectrum Frameworks

In parallel with the academic literature on dynamic spectrum access, several national regulators have implemented practical shared spectrum frameworks that allow multiple users to access the same frequency bands under defined rules. Figure 2.1 provides an overview of the spectrum access and sharing landscape, ranging from exclusive licensed access on the left, through the shared and local licensing frameworks discussed in this section, to licence-exempt operation on the right. As flexibility increases from left to right, interference protection decreases, and the need for trusted coordination and accountable settlement becomes more

pressing.

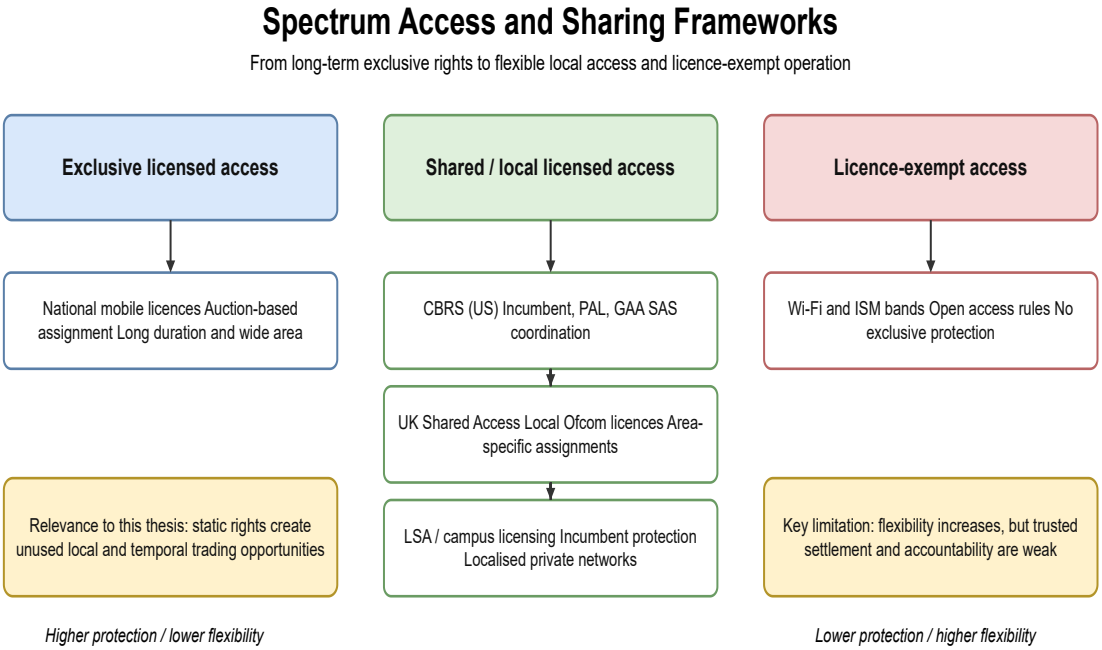


Figure 2.1: Spectrum access and sharing frameworks, from exclusive licensed access to licence-exempt operation.

United States: Citizens Broadband Radio Service (CBRS)

The most developed shared spectrum framework is the Citizens Broadband Radio Service (CBRS), established by the Federal Communications Commission (FCC) in 2015 for the 3550–3700 MHz band (150 MHz) [34]. CBRS introduces a three-tier hierarchical access model. At the top tier, Incumbent Access users, primarily U.S. Navy shipborne radar systems, satellite earth stations, and grandfathered wireless broadband licensees, receive full interference protection. The second tier comprises Priority Access Licensees (PALs), who acquire 10 MHz channels within the 3550–3650 MHz portion of the band through competitive auction on a county-by-county basis. PALs are 10-year renewable licences and receive interference protection from lower-tier users. The third and lowest tier is General Authorised Access (GAA), which is licensed-by-rule and permits open, flexible access to any portion of the band not occupied by higher-tier users, without requiring an individual licence and with no expectation of interference protection.

Spectrum coordination among the three tiers is managed by an automated frequency coordinator known as the Spectrum Access System (SAS), which dynamically authorises transmit parameters for Citizens Broadband Radio Service Devices (CBSDs) based on real-time information about incumbent activity, PAL assignments, and device location. An Environmental Sensing Capability (ESC)

sensor network detects incumbent radar transmissions and communicates this information to the SAS, enabling dynamic protection of military users. By mid-2025, over 420,000 active CBRS base station radios had been deployed across the United States, supporting private enterprise networks, neutral-host in-building systems, and fixed wireless access for rural broadband [35].

United Kingdom: Ofcom Shared Access Licence

In the United Kingdom, Ofcom introduced a Shared Access licence framework in 2019 to enable local wireless connectivity for businesses and enterprises [36]. The framework makes spectrum available in four bands: 1800 MHz (1781.7–1785 MHz paired with 1876.7–1880 MHz), 2300 MHz (2320–2340 MHz and 2390–2400 MHz), 3.8–4.2 GHz, and 26 GHz (the latter for indoor low-power use only).

Unlike the CBRS model, Ofcom does not employ a real-time automated spectrum coordinator. Instead, spectrum is allocated on a first-come, first-served basis through an application and coordination process managed by Ofcom. Two licence types are offered: Low Power licences, which authorise deployment of multiple base stations within a 50-metre radius area, and Medium Power licences, which authorise a single base station with higher transmit power, generally restricted to rural areas. Annual licence fees are modest, at £80 per 10 MHz in rural areas for the 3.8–4.2 GHz band, making the framework accessible to small and medium enterprises. As of late 2024, Ofcom had issued approximately 988 Shared Access licences, with the 1800 MHz and 3.8–4.2 GHz bands being the most popular. Recent updates to the framework have improved spectrum availability through updated coordination rules, additional indoor-use channels in the 2300 MHz band, and greater flexibility for medium-power urban deployments [37]. However, the administrative overhead of the current process remains a practical barrier: obtaining a Shared Access licence through Ofcom can take up to six weeks [38], which limits the ability of operators to respond rapidly to localised or short-term demand.

European Union: Licensed Shared Access (LSA)

The European regulatory approach to shared spectrum is centred on the Licensed Shared Access (LSA) model, developed through ETSI and CEPT standardisation activities [39]. Under LSA, an incumbent user and one or more additional users share a frequency band under defined conditions set by the national regulatory authority. Each user receives an individual (but not exclusive) licence that guarantees a predictable quality of service. The LSA architecture comprises an LSA

Repository, which stores incumbent protection requirements, and an LSA Controller, which manages the mobile network's operating parameters to prevent harmful interference.

LSA was initially defined for the 2.3–2.4 GHz band (3GPP Band 40) and has since been extended through the evolved LSA (eLSA) framework to support local, temporary wireless networks for use cases including industrial automation, programme making and special events, and public safety. The European Commission has recommended that EU member states consider sharing the 3.8–4.2 GHz band for local vertical industry applications using low and medium power networks [40].

Germany: Local 5G Campus Network Licences

Germany has been a leader in local spectrum licensing for industrial applications. In November 2019, the Federal Network Agency (BNetzA) reserved 100 MHz of spectrum in the 3.7–3.8 GHz band specifically for local 5G campus network deployments [41]. Applications are processed on a rolling, first-come, first-served basis, and licences are location-specific. By November 2024, BNetzA had awarded over 430 local licences to organisations spanning research and development, manufacturing, automotive, aviation, and logistics. Notable licensees include Bosch, Siemens, BMW, Volkswagen, and Lufthansa. BNetzA has also opened the 26 GHz band for local private network deployments since 2021, with spectrum assigned in blocks of 200 MHz.

2.3.3 Spectrum Trading and Secondary Spectrum Markets

Spectrum trading extends the logic of dynamic access by introducing market-based mechanisms through which spectrum holders can lease or sell temporarily unused resources to spectrum buyers. In the United Kingdom, Ofcom's spectrum trading framework permits licensees to transfer all or part of their licence rights to another party, or to lease spectrum use to third parties under a contractual arrangement, subject to Ofcom's consent and, for mobile spectrum, a competition assessment [36]. In such settings, spectrum becomes a tradable asset, and transfer decisions may depend on parameters such as location, bandwidth, time duration, and service requirements. Surveys on spectrum trading in DSA networks explain that these markets allow secondary users to obtain access to licensed bands through structured transactions rather than purely technical opportunism [42]. Game-theoretic and pricing studies further model spectrum markets as interactions among multiple sellers and buyers, showing how underutilised licensed

resources can be converted into exchangeable access rights under appropriate economic incentives [43, 44]. Accordingly, secondary spectrum markets represent the economic operationalisation of DSA, turning flexible access into a governed exchange process.

2.3.4 Auction-Based and Market-Oriented Spectrum Allocation Approaches

Spectrum auctions have become the dominant mechanism through which national regulators assign licensed spectrum to mobile operators. In a typical primary spectrum auction, the regulator (such as the FCC in the United States or Ofcom in the United Kingdom) divides a designated frequency band into discrete lots defined by frequency range, geographic coverage area, and licence duration. Eligible bidders, usually established MNOs or new entrants meeting financial and technical qualification criteria, submit bids over multiple rounds. The winning bidders receive exclusive national or regional spectrum licences, typically lasting 15 to 25 years, and assume ownership of the usage rights for the duration of the licence. Common auction formats include the Simultaneous Multiple Round Ascending (SMRA) auction used by the FCC and the Combinatorial Clock Auction (CCA) used by Ofcom, both designed to handle interdependencies between spectrum lots across different bands and regions [31, 45].

Beyond primary regulatory auctions, a parallel body of research has explored auction-based mechanisms for secondary and dynamic spectrum allocation, where spectrum is traded between operators or between primary and secondary users on shorter timescales. These secondary auctions differ from primary regulatory auctions in several important ways: the seller is typically a spectrum holder rather than the regulator, the allocation duration is minutes to hours rather than years, and the geographic scope is localised rather than nationwide. Truthful auction designs, based on the Vickrey-Clarke-Groves (VCG) mechanism, have been proposed to improve incentive compatibility in such dynamic spectrum markets, ensuring that bidders reveal their true valuations rather than engaging in strategic underbidding [46, 47]. Zhou and Zheng proposed the TRUST framework for truthful double spectrum auctions, addressing the specific challenge of matching multiple sellers and multiple buyers in a spectrum market while maintaining truthfulness and individual rationality [47]. Subsequent work adapted auction mechanisms to Licensed Shared Access systems, where incumbent and secondary users share the same band under regulatory constraints [48], and to 5G and beyond secondary allocation settings where concurrent auctions across

multiple bands must be coordinated to avoid fragmentation and ensure fair allocation [49, 50].

Game-theoretic and pricing studies have further modelled spectrum markets as strategic interactions among multiple sellers and buyers, demonstrating how underutilised licensed spectrum can be converted into exchangeable access rights under appropriate economic incentives [43, 44]. These models treat spectrum allocation not only as a technical assignment problem but as a market-design problem in which pricing, bidding rules, and allocation logic shape both utilisation efficiency and strategic behaviour among competing operators.

Auction-based and market-oriented approaches therefore remain one of the most widely studied mechanisms in the spectrum coordination literature. However, as discussed in the following subsection, many of these models assume the existence of a trusted auctioneer or centralised broker to manage bids, enforce outcomes, and settle payments, an assumption that becomes problematic in multi-stakeholder settings where no single party is universally trusted.

2.4 Blockchain and Smart Contracts in Telecommunications

2.4.1 Blockchain as a Coordination Infrastructure

The literature on dynamic access and spectrum trading establishes a convincing case for flexibility, but often leaves execution-layer questions unresolved, especially trusted authorisation, auditable settlement, and accountable multi-party coordination. Regulatory studies emphasise that dynamic spectrum management requires not only access logic but also institutional mechanisms governing brokerage, exchange, and compliance [51]. Likewise, recent agenda-setting work on automated spectrum sharing argues that dynamic market contracts require stronger support for authorisation, enforcement, and machine-executable agreements than conventional mechanisms typically provide [52, 53]. As a result, while these trading models are strong on economic logic and allocation theory, they are weaker on trusted operational execution across multiple actors.

Blockchain has attracted increasing attention in telecommunications because it offers properties that directly address multi-party coordination problems, including immutability, distributed record-keeping, transparency, and tamper resistance. Telecom surveys present blockchain as a coordination layer for settings where operators, infrastructure providers, service platforms, and regulators must interact without relying on a single fully trusted intermediary [8,9]. In these stud-

ies, blockchain is not treated merely as a financial ledger, but as an infrastructure for recording transactions, enforcing shared state, and improving accountability across distributed telecom processes [10, 54]. Resource-management-oriented work extends this view by showing that blockchain can support orchestration, policy enforcement, and trusted logging in networked environments involving heterogeneous stakeholders [55, 56]. Thus, the main appeal of blockchain in telecom lies in its ability to support shared coordination under trust. Beyond telecommunications, blockchain has been successfully applied across a range of domains, including healthcare [57], vehicular networks [58], electric vehicle charging [59], and the Internet of Things [60], demonstrating its versatility as a coordination technology for multi-stakeholder environments.

2.4.2 Smart Contracts as Programmable Execution Logic

A major reason blockchain is relevant to this thesis is the role of smart contracts. Smart contracts allow predefined rules to be encoded and executed automatically once triggering conditions are met, enabling functions such as access authorisation, resource listing, price retrieval, ownership update, payment transfer, or settlement logging [8, 61]. Telecom-focused studies repeatedly emphasise that smart contracts are what transform blockchain from a passive record system into an operational coordination mechanism [8, 62]. They have been proposed for automating telecom service transactions, network control logic, and spectrum-related workflows, including access control and settlement [10, 63]. Broader 5G/B5G blockchain surveys also describe smart contracts as a key mechanism for embedding policy and operational rules directly into distributed systems [54, 64, 65]. In practical terms, smart contracts matter because they make rule-governed multi-party execution programmable rather than manually administered.

2.4.3 Permissioning, Access Control, and Auditability

Although blockchains are often described as decentralised, telecom coordination typically requires controlled participation. Blockchain-based telecom systems therefore often combine ledger-level decentralisation with application-level or smart-contract-level access control. Reviews of communication-network blockchain systems note that deployment choices directly affect latency, privacy, governance, and scalability, and that permissioned blockchains are often preferred because they avoid the open participation and mining-style delay associated with public chains [62]. Telecom marketplace studies likewise report that private or consortium architectures are usually better suited to operator environments be-

cause they provide stronger governance and more predictable control over participants [10]. Similar conclusions appear in 5G and B5G surveys, which argue that permissioned designs fit telecom requirements better when accountability, privacy, and authorised participation matter [66]. What results isn't permissionless openness, but governed decentralisation with auditable execution.

2.4.4 Blockchain Limitations in Telecom Systems

In practical telecom applications, blockchain introduces nontrivial constraints, including transaction latency, execution cost, limited throughput, and design trade-offs between public, consortium, and hybrid models [66,67]. Requirements-versus-cost analyses in telecom point out that blockchain's coordination benefits must be weighed against confirmation delay, resource overhead, scalability limits, and execution complexity [8]. Surveys on blockchain in 5G and telecom similarly show that consensus and distributed execution can conflict with stringent latency and throughput requirements [64,66,67]. Resource-management work also identifies smart-contract overhead and transaction delay as important challenges. This thesis addresses these concerns in two ways: Chapter 3 reports measured gas consumption and transaction execution times for each smart contract function on the Ethereum Sepolia test network and develops a queueing-based scalability model to quantify the throughput limits of on-chain spectrum trading, while Chapter 5 mitigates on-chain overhead by separating radio-layer QoS computation and pricing logic off-chain, committing only the minimum verifiable settlement inputs to the blockchain. While studies on mobile smart contracts further emphasise consensus-related bottlenecks and scaling constraints [55, 65]. So the telecom literature treats blockchain as promising but costly: good for trust and automation, but not universally suited to time-critical network functions unless it's carefully scoped. Institutional interest in blockchain for spectrum management is also growing. Ofcom in the United Kingdom has explored the potential of blockchain technology for managing telephone numbers and spectrum resources [68], while the Federal Communications Commission in the United States has shown interest in blockchain as a means of improving real-time spectrum usage and allocation efficiency [69, 70]. These signals reflect a growing recognition that decentralised coordination mechanisms may offer practical advantages over conventional centralised approaches for spectrum governance.

2.5 Blockchain-Based Spectrum Management and Spectrum Trading

2.5.1 Early Blockchain-Based Spectrum Management Concepts

Surveys on blockchain for dynamic spectrum sharing have noted that auction- and market-based models often assume trustworthy handling of transactions and rights transfer without fully solving secure coordination and auditable state management among independent stakeholders [7, 71]. Early studies on blockchain-based spectrum management were therefore motivated by the need to improve transparency, traceability, and distributed trust in wireless resource governance [2, 71]. Architectural work on blockchain-based cellular spectrum management proposed combining blockchain with control frameworks such as SDN to improve trusted coordination, but these studies focused more on demonstrating secure architecture than on implementing realistic short-term trading behaviour [2, 63, 72]. Spectrum-management surveys also show that early blockchain interest was largely driven by the promise of secure record-keeping, tamper-evident coordination, and decentralised trust in wireless resource governance [9, 71, 73]. Broader telecom-blockchain overviews place spectrum management among promising applications, but they similarly indicate that early contributions were often exploratory and design-oriented rather than deeply operational [9, 73].

2.5.2 Blockchain for Spectrum Sharing and Trading

Subsequent studies extended blockchain usage beyond record-keeping into active spectrum sharing and trading. In these frameworks, blockchain became a marketplace infrastructure through which spectrum holders and buyers could coordinate temporary usage transfers [10, 74, 75]. Blockchain-based dynamic spectrum sharing work for 5G and beyond explicitly models operators trading spectrum through consortium-led platforms supported by smart contracts and fair-payment logic [74]. More recent studies on vehicular and coexistence scenarios likewise use consortium blockchain to support dynamic spectrum sharing, secure trading, and cooperative management among multiple participants [76, 77]. Hybrid spectrum-access-system designs also show how permissioned blockchain can be embedded directly into spectrum-access workflows as part of the exchange infrastructure [78]. Paroushev and Ruskov [79] proposed a blockchain-enabled framework for dynamic spectrum usage rights in 5G and 6G networks, in which

usage rights rather than long-term licences become the unit of allocation. Their framework introduces a pricing model that accounts for several factors affecting spectrum value and implements it as a DAML smart contract on the permissioned Canton ledger, reporting gains in efficiency, transparency, and flexibility over static allocation. This work supports the view that dynamic spectrum allocation and pricing logic can be encoded directly in smart contracts. However, it does not draw on measured traffic demand, does not integrate allocation with a radio access network testbed, and does not address performance-aware settlement, which are the aspects addressed in Chapters 4 and 5 of this thesis. Collectively, this literature moves blockchain from passive recording into active transaction support for spectrum coordination.

2.5.3 Smart-Contract-Enabled Market Mechanisms

A more advanced branch of the literature has used smart contracts to encode specific market operations such as spectrum discovery, request submission, price matching, purchase triggering, settlement, and post-transaction logging [2, 80]. In block-chain-based dynamic spectrum sharing frameworks, smart contracts are used to automate fair payment and encode exchange logic between trading parties [74]. Studies on dynamic sharing in vehicular contexts similarly report that smart contracts perform core operational functions within the trading or coordination process [76]. In the context of WiFi and LTE-Unlicensed coexistence, where both technologies must share the same unlicensed band without a centralised coordinator, Yu et al. [77] use smart contracts to manage spectrum access rules and interaction logic among independent operators, ensuring fair channel allocation without requiring either party to trust the other. These developments mirror the broader telecom-blockchain literature, which repeatedly identifies smart contracts as the component that converts blockchain from immutable storage into executable coordination logic [8, 62]. In the spectrum context, this makes them central to any attempt at automated market execution.

2.5.4 Consortium, Hybrid, and Public Blockchain Models

An important discussion in the literature concerns the blockchain deployment model itself, as the choice between public, consortium, and hybrid architectures involves trade-offs that directly affect suitability for telecom coordination.

Public blockchains such as Ethereum provide the strongest decentralisation and transparency guarantees: any entity can join, validate transactions, and verify the full ledger history without requiring permission from a central authority.

This openness works well for spectrum trading because it removes the need for a trusted intermediary, and it means no single operator can control or manipulate the marketplace. That said, public blockchains bring significant practical constraints for telecom applications. Transaction throughput is limited (Ethereum processes approximately 15–30 transactions per second), block confirmation times introduce latency (approximately 12 seconds under proof-of-stake), and all transaction data is publicly visible, which may conflict with commercial confidentiality requirements between operators [66, 67]. Gas costs also introduce variable and sometimes unpredictable execution expenses. Despite these limitations, public blockchains remain relevant for proof-of-concept implementations and for scenarios where maximum transparency and minimal trust assumptions are prioritised over throughput [2]. The spectrum trading framework developed in Chapter 3 of this thesis uses the Ethereum public test network (Sepolia) precisely to evaluate feasibility under these real-world blockchain constraints.

Telecom applications often use consortium or permissioned blockchains rather than public ones. The reason is governance: operators want to know exactly who's on the network. Only pre-approved entities, licensed operators, and regulators can join and validate transactions, and it is the closed membership that makes the difference. Privacy gets better. Latency drops. Throughput ends up higher than public chains typically manage. This preference is especially visible in blockchain-based spectrum studies, where consortium architectures are commonly adopted for operator trading, dynamic sharing, and coexistence management [74, 76, 77]. Performance analysis of permissioned consensus for network sharing reinforces this point by showing that governed participation can achieve sub-second finality and thousands of transactions per second [81]. The trade-off is that consortium blockchains require a governance structure that all participants must agree upon, including who operates the ordering service, who endorses transactions, and who has read access to the ledger.

Hybrid designs have also been proposed to balance the strengths of both models. Veloso et al. [78] proposed a hybrid architecture for spectrum access systems that combines a permissioned blockchain for operational coordination with selective public verifiability for auditability. This approach attempts to retain the privacy and performance benefits of consortium chains while preserving the transparency and independent verification properties of public chains. The wider telecom literature consistently argues that permissioned or hybrid blockchains are better suited than fully public chains for operator-centric coordination because they offer stronger privacy, clearer governance, and more manageable performance characteristics [10, 62].

2.5.5 Critical Comparison of Existing Blockchain-Based Spectrum Studies

Table 2.1 compares representative blockchain-based spectrum studies across six dimensions that are central to the contributions of this thesis: whether the study addresses short-term demand-driven trading, whether it incorporates real or realistic demand data, whether it implements interference-aware feasibility constraints, whether settlement is performance-aware, whether the framework is validated on an operational testbed, and whether the blockchain implementation is deployed and evaluated on a real or test network.

Table 2.1: Comparison of representative blockchain-based spectrum studies against key dimensions addressed in this thesis.

Study	**Short-term trading**	**Real demand data**	**Interference aware**	**Performance-aware settlement**	**5G testbed**	**On-chain evaluation**
Okon et al. [63]	No	No	Partial	No	No	No
Li et al. [74]	Partial	No	No	No	No	Simulated
Shen et al. [76]	Partial	Synthetic	No	No	No	Simulated
Yu et al. [77]	No	No	Yes	No	No	Simulated
Veloso et al. [78]	No	No	Partial	No	No	Partial
Kotobi and Bilén [75]	No	No	Partial	No	No	Simulated
Weiss et al. [2]	No	No	No	No	No	Conceptual
(This-thesis) (Ch. 3–5)	**Yes**	**Yes (Milan)**	**Yes**	**Yes**	**Yes (OAI)**	**Yes (Sepolia)**

Several observations emerge from this comparison. First, architectural proposals such as Okon et al. [63] and Weiss et al. [2] provide foundational design concepts for blockchain-based spectrum coordination, but they remain at the conceptual or design level without implementing executable smart contracts or evaluating on-chain behaviour. Second, studies that do implement smart-contract-enabled trading, such as Li et al. [74] and Shen et al. [76], demonstrate automated exchange logic but do not incorporate real demand data, interference-aware feasibility constraints, or performance-linked settlement. Third, coexistence-oriented

work such as Yu et al. [77] addresses interference management but focuses on long-term coordination instead of short-term utilisation-driven trading. Fourth, none of the reviewed studies couples radio-layer performance measurement with economic settlement logic, which is the specific contribution of Chapter 5 of this thesis.

Across all studies, three recurring gaps stand out. Existing blockchain-based spectrum frameworks aren't driven by empirical evidence of demand, and none of them use real traffic data to identify when and where trading opportunities actually arise. Interference-aware feasibility is either missing or only partially addressed, since no study embeds spatial, temporal, and frequency-domain coexistence checks within the smart contract matching logic itself. And settlement remains static or absent altogether, with no existing framework adjusting economic outcomes based on the radio performance actually delivered to users. This thesis addresses all three gaps through its contributions in Chapters 3, 4, and 5.

2.6 Learning-Assisted Demand Analysis and Spectrum Prediction

2.6.1 Learning for Traffic Prediction and Spectrum Awareness

Machine learning has been widely studied for wireless traffic prediction, spectrum demand estimation, and dynamic resource management [12–14, 82]. By learning recurring temporal patterns or statistical structure in network activity, learning-based methods can help anticipate future congestion, identify periods of lower utilisation, and support more proactive resource allocation. Surveys on machine learning for spectrum sharing and DSA show that learning methods are increasingly used to infer traffic behaviour, predict spectrum conditions, and improve spectrum awareness in complex wireless environments [12, 14, 83, 84]. Traffic-prediction reviews for 5G and beyond likewise show that forecasting methods are now widely used to model heterogeneous temporal and spatial load variation as an input to proactive network management [85–87]. In this literature, learning matters not because it replaces spectrum policy or market design, but because it improves visibility into likely future demand and availability patterns.

2.6.2 Spatial Heterogeneity and Clustering-Based Analysis

In the context of this thesis, one of the most important learning-related ideas is not forecasting alone, but the characterisation of spatial heterogeneity. Wireless demand is uneven across space, and that unevenness is central to identifying mismatches between underutilisation and local scarcity. Studies on proactive resource allocation with real traffic data explicitly show that network demand should be analysed spatially as well as temporally [88]. Furno et al. [86] demonstrated the value of clustering-based analysis by characterising distinct mobile traffic signatures across ten cities, showing that localised high-demand regions can be identified through spatial clustering of real cellular activity data. Similarly, Ozturk et al. [89] used deep learning on real network data to predict spatial demand patterns, confirming that anticipating localised resource pressure requires capturing both temporal dynamics and geographic structure. More general surveys on ML for wireless resource allocation also report that clustering techniques are commonly used to capture structure in decentralised and heterogeneous network environments [82]. This is especially relevant here because short-term spectrum trading opportunities arise not only from temporal fluctuation, but from localised divergence in demand across different parts of the network.

2.6.3 Prediction for Proactive Spectrum Trading

Prediction can reduce reaction delay by making spectrum availability visible before demand arises. Instead of waiting until underutilisation is already observable, a predictive system can prepare contracts, identify likely trading windows, and coordinate access more efficiently. Reviews of dynamic spectrum allocation prediction emphasise that anticipating future allocation-relevant states is preferable to purely reactive adjustment [83]. Sone [88] demonstrated this in a shared-spectrum setting with radar systems, using ML-based time-series forecasting to proactively allocate resources before congestion materialised, showing that earlier allocation decisions reduce interference and improve coexistence outcomes. Ozturk et al. [90] similarly applied machine learning to predict spectrum occupancy and trigger proactive access decisions in cognitive radio networks, confirming that forecasting-driven allocation outperforms reactive approaches in both QoS and spectrum utilisation. For a trading framework, the implication is direct: prediction can expose likely windows of tradable availability before those windows disappear, making pre-authorisation, offer preparation, and market coordination more efficient.

2.6.4 Limits of ML-Only Approaches

Even highly accurate forecasting cannot by itself enforce access rights, authenticate participants, transfer ownership, commit payments, or create an immutable audit trail. The literature treats ML primarily as a predictive and decision-support capability rather than as a mechanism for authorisation or settlement [12, 14, 83, 84]. It can reveal likely underutilisation, estimate future demand, and support proactive resource decisions, but it does not solve the execution problem of trusted exchange among independent actors. A model may predict that a spectrum block will likely be underused in a particular place and time, yet it cannot on its own determine who may sell access, who may buy it, whether payment has been committed, or how the transaction will be logged and audited. Accordingly, ML-only approaches are necessary for anticipation but insufficient for trusted operational spectrum trading. Their main role is to supply predictive intelligence that must be coupled with a separate execution and governance layer.

2.7 Interference, Coexistence, and Regulatory Constraints

2.7.1 Interference-Aware Sharing and Reuse

Spectrum is not a generic digital commodity that can be reassigned freely without regard to location, power, frequency adjacency, and coexistence constraints. Spectrum sharing decisions must account for radio interference, reuse distance, channel conditions, and protection of incumbent or neighboring users. Jeon et al. [29] demonstrated that coordinated dynamic spectrum sharing in 5G requires explicit interference management between co-channel operators to avoid mutual degradation. Sharma et al. [18] showed that full-duplex dynamic sharing introduces additional self-interference and cross-link interference challenges that must be managed through resource-aware coordination. Salani et al. [30] similarly confirmed that dynamic allocation in the C-band must account for coexistence constraints to protect incumbent and adjacent users. Together, these studies show that interference-aware frequency coordination and coexistence management remain central to efficient reuse in dense and heterogeneous wireless systems. In the spectrum-sharing literature, this means that any allocation or trading logic must be constrained by technical feasibility rather than by market demand alone. Even where unused spectrum appears available in an abstract sense, actual reuse may be blocked by interference risk, coexistence obligations, or geographic over-

lap with other active transmissions [1, 18, 29, 91]. Thus, interference awareness is a structural requirement for any operational spectrum market.

2.7.2 Policy and Regulatory Feasibility

In addition to radio coexistence, spectrum transactions must comply with regulatory constraints such as geographic boundaries, license conditions, participant authorisation, and rules regarding temporary or secondary use. Spectrum management studies consistently show that economic flexibility does not eliminate the need for policy control, because temporary reassignment must still respect regulator-defined usage rights and incumbent protections [3, 21, 22, 51]. In practice, operator coordination is always embedded in governance structures that determine who may share, on what terms, and in what locations. The CBRS framework enforces a three-tier access hierarchy governed by the SAS [34], Ofcom's Shared Access licences define per-site power limits and geographic constraints [36], the European LSA model requires incumbent protection conditions set by the national regulator [39], and Massaro showed that LSA governance structures are critical to balancing flexibility with incumbent protection [3]. This means that a usable trading framework cannot treat spectrum as a purely private commodity. It must incorporate participant eligibility, location-aware authorisation, and policy-aware transaction logic so that market outcomes remain legally and operationally valid [91, 92].

2.7.3 Feasibility Constraints

A useful market framework has to combine economic logic with feasibility-aware authorisation. If a system flags a spectrum block as tradable just because it looks underutilised, without checking interference or policy constraints, the resulting transaction could end up technically infeasible or regulatorily invalid. For this reason, feasibility constraints are not secondary implementation details but part of the market logic itself. The strongest recent resource-allocation surveys show that future wireless systems require coordination mechanisms that jointly consider allocation efficiency, interference management, and service constraints [92]. For this thesis, that implies that spectrum trading should be treated as a constrained coordination problem in which radio feasibility and policy authorisation must be verified before access can be transferred.

2.8 Indoor Offloading and Neutral-Host Settlement

2.8.1 Indoor Traffic Growth and Outdoor-to-Indoor Limitations

A substantial fraction of mobile traffic originates indoors, yet conventional macrocellular architectures often serve indoor users from outdoor base stations, leading to penetration loss, shadowing, and variable signal conditions. This mismatch has long motivated indoor-focused architectures, offloading strategies, and dense local deployments that can better match traffic demand to service conditions [93]. Studies of indoor 5G deployment and edge-supported indoor networking similarly show that indoor environments create distinctive performance challenges and cannot be efficiently served by macro-only architectures when quality of service is important [94]. In practical terms, the relevance to this thesis is that indoor traffic concentration creates localised coordination and capacity problems analogous to those seen in dynamic spectrum management, but in a service-delivery and settlement context.

2.8.2 WiFi-Based Traffic Offloading

Early work on traffic offloading focused primarily on WiFi as a lower-cost complement to cellular networks. A 2011 Radisys Corporation white paper examined Internet data offloading for mobile operators, arguing that shifting traffic to WiFi access networks can help accommodate demand growth while reducing service cost [95]. The report described three forms of WiFi offload: hard offload, which is user-driven and relies on manual connection to a private WiFi network, optimised offload, which uses SIM-based authentication to enable secure attachment to public WiFi networks, and integrated offload, which is more tightly embedded within the mobile system and supports deeper core-network integration. Across these approaches, robust mechanisms for authentication, policy execution, and pricing were identified as essential requirements.

Subsequent studies expanded on WiFi offloading from both architectural and user-behaviour perspectives. The authors in [96] reviewed ongoing developments in Internet Protocol (IP) flow mobility from the IETF and 3GPP, and highlighted network-based flow mobility as a promising direction for extending operator capability. Related research in [97] surveyed multiple offloading mechanisms, including Local IP Access (LIPA), Selected IP Traffic Offload (SIPTO), and IP Flow Mo-

bility, identifying open challenges in network management, traffic handling, and energy efficiency, while emphasising the need to guarantee acceptable QoS for offloaded traffic. Coupling between WiFi and cellular networks at the RAN level has been categorised into loose, tight, and very tight integration models [98], with session continuity during vertical handover identified as a key architectural concern. In [99], smartphone user behaviour was studied using data from 298 users who offloaded traffic to WiFi, showing that offloading can help users reduce costs while improving service quality. A WiFi-aware offloading method with deadline constraints was proposed in [100], aiming to complete file transfers before a deadline while reducing the computational burden of scheduling. A dynamic game-theoretic model for crowd-sourced WiFi offloading was presented in [101] to address operator pricing challenges and support expansion of crowd-sourced WiFi capacity under rising demand.

2.8.3 Mobile and Cellular Data Offloading

Beyond WiFi, several studies have addressed offloading within cellular and mobile network contexts. The authors of [102] introduced a Satisfaction-based Dynamic Bandwidth Reallocation (SDBR) scheme for multipath mobile data offloading, intended to improve both network resource utilisation efficiency and user satisfaction, reporting gains in overall revenue and average user experience. In the context of mobile data offloading over opportunistic device-to-device contacts, [103] formulated an NP-hard Target Set Selection (TSS) problem and used heuristic optimisation to design an offloading scheme with polynomial-time complexity, reporting roughly 20% traffic reduction compared to a heuristic-only baseline. For mobile-to-WiFi offloading of commuting users, Lima et al. [104] applied unsupervised learning techniques to model individual mobility patterns and predict when users would encounter WiFi access points, enabling personalised offloading schedules that improved effectiveness while reducing reliance on cellular capacity. Exponential learning with minority game theory was used in [105] to develop a distributed cellular offloading scheme evaluated under parameters including pricing and offered cellular throughput.

2.8.4 HetNet and Neutral-Host Offloading Models

In the context of heterogeneous networks (HetNets), macro-to-small-cell offloading has been studied as a mechanism for improving indoor capacity and reducing macro congestion. Abd El-atty and Gharsseldien [106] proposed a hierarchical HetNet architecture in which femtocells serve indoor traffic while macrocells

handle overflow, demonstrating that small cell integration can significantly reduce blocking probability and improve indoor throughput. Shah [107] addressed the multi-operator dimension of small cell offloading, examining collaborative resource management among competing operators sharing small cell infrastructure and highlighting the importance of service-level agreements in governing inter-operator interactions. More recently, a real-world measurement study of indoor neutral-host networks operating over shared CBRS spectrum [108] confirmed that neutral-host indoor deployments deliver measurable throughput improvements over macro outdoor-to-indoor service, while also exposing practical challenges in subscriber handover, authentication, and inter-operator accounting that remain insufficiently addressed in the current literature.

Neutral-host deployment is attractive because it reduces duplication of infrastructure and allows multiple operators to benefit from shared indoor or local access systems [93, 109]. This shifts the problem from pure radio access provision to inter-operator coordination, accountability, and compensation.

2.8.5 Blockchain-Enabled Offloading

More recently, blockchain technology has been applied to offloading coordination to address trust, automation, and multi-party agreement challenges. Weerasinghe et al. [110] proposed a blockchain-based roaming and offloading mechanism for local 5G operators that leverages blockchain to support agreements between MNOs and facilitate handover execution across subscribers of different operators. A WiFi offloading platform was proposed in [111] that selects WiFi networks based on provider ratings, Internet speed, and signal strength, with offloading decisions executed at the mobile network's software-defined networking layer and triggered automatically without user intervention. The authors noted that unintended or overly aggressive automatic offloading could increase MNO costs, and contrasted this with approaches that require explicit consent and verification among all parties. Beyond WiFi-centric designs, [112] proposed a blockchain-based offloading scheme among base stations using a directed acyclic graph (DAG) and game-theoretic bargaining to reduce cost and delay. Blockchain-based coordination has also been explored in the adjacent domain of computational task offloading: SDBlockEdge [113] introduced a token-based blockchain resource management scheme for collaborative edge computing, while [114] proposed a mobile edge computing framework for optimising task offloading decisions. Although these latter works address computational rather than radio-layer offloading, they demonstrate that blockchain can serve as a coordination and incentive mechanism in multi-stakeholder offloading environments, a principle that the framework in Chapter 5

adapts to the neutral-host settlement context.

2.9 Research Gaps and Thesis Positioning

The reviewed literature shows that conventional spectrum management remains too rigid to support the dynamic, localised, and short-term resource demands emerging in 5G and beyond. Dynamic spectrum access and secondary market research address this rigidity by introducing more flexible and economically responsive allocation mechanisms, but non-blockchain approaches often leave trusted execution, settlement, and multi-party accountability only partially resolved [5, 31–33, 42–46, 48–53, 71]. Blockchain and smart-contract literature in telecommunications addresses some of these issues by providing immutable records, programmable coordination logic, and auditable execution, yet telecom-wide studies also make clear that blockchain introduces nontrivial performance and governance trade-offs [8–10, 54–56, 62–67, 115–117].

A further weakness, particularly visible in the offloading and neutral-host literature, is that settlement is often static, simplified, or insufficiently linked to the service quality actually delivered. Okon et al. [72] explicitly linked offloading, neutral-host operation, billing settlement, and agreed service levels in beyond-5G multi-operator small cell systems, while related work [63] noted that neutral-host and multi-operator settings require better support for coordination, service accounting, and trusted settlement than conventional telecom processes provide. Although offloading and neutral-host architectures are well studied technically, the literature is less mature in linking compensation and settlement directly to delivered QoS in a transparent and auditable way.

Within the more specific blockchain-based spectrum literature, important progress has been made in architecture design, trusted sharing, and smart-contract-enabled coordination, but many representative studies remain focused on conceptual frameworks, coexistence management, or long-horizon sharing logic rather than short-term, utilisation-driven trading under realistic spatio-temporal demand variation [63, 73, 74, 76–78, 81]. At the same time, learning-assisted traffic analysis literature shows strong progress in wireless demand prediction, proactive optimisation, and clustering-based characterisation of heterogeneous network behaviour, but ML-only approaches do not solve the execution problem of access-right transfer, settlement, and auditable coordination among independent actors [12–14, 82–84, 88, 118–120]. Finally, indoor offloading and neutral-host studies show that shared infrastructure is increasingly relevant in practice, especially with Open RAN and indoor densification, yet settlement is still often

treated in a simplified or static way rather than being tightly linked to delivered performance [72, 94, 109, 121–123].

A further gap is visible in the offloading and neutral-host literature. While prior work has made significant progress in offloading techniques spanning WiFi, cellular, and edge computing domains, the existing literature focuses predominantly on the radio-layer or computational aspects of offloading, including when and where to offload, but pays limited attention to how the resulting service should be settled economically between operators. Where settlement is considered, it is typically based on flat fees or traffic-volume accounting that does not reflect the radio performance actually delivered to offloaded users [80]. Although blockchain has been applied to offloading coordination, existing schemes focus on roaming agreements, WiFi selection, or task offloading rather than on coupling radio-layer QoS measurement with automated, performance-aware inter-operator settlement [124].

Against this background, the thesis is positioned around three connected contributions. Chapter 3 presents an initial blockchain-based decentralised spectrum trading framework and proof-of-concept implementation. Chapter 4 extends this through a more comprehensive blockchain-enabled spectrum coordination framework grounded in spatio-temporal traffic analysis, learning-assisted demand characterisation, interference-aware feasibility, and execution-level evaluation. Chapter 5 further extends the thesis into indoor neutral-host Open RAN settings, where blockchain is used to enforce performance-aware settlement based on delivered QoS rather than static charging assumptions. In this way, the thesis addresses not only flexibility in spectrum access, but also trusted execution, feasibility-aware coordination, and accountable service settlement across related telecom coordination problems. Figure 2.2 summarises how the three contributions of this thesis map to the gaps identified in the reviewed literature. Each contribution addresses a specific limitation: Chapter 3 provides automated and auditable trading execution, Chapter 4 adds real-demand evidence and network-level validation, and Chapter 5 introduces performance-aware operator settlement. Together, the three contributions support trusted, feasible, and accountable telecom coordination.

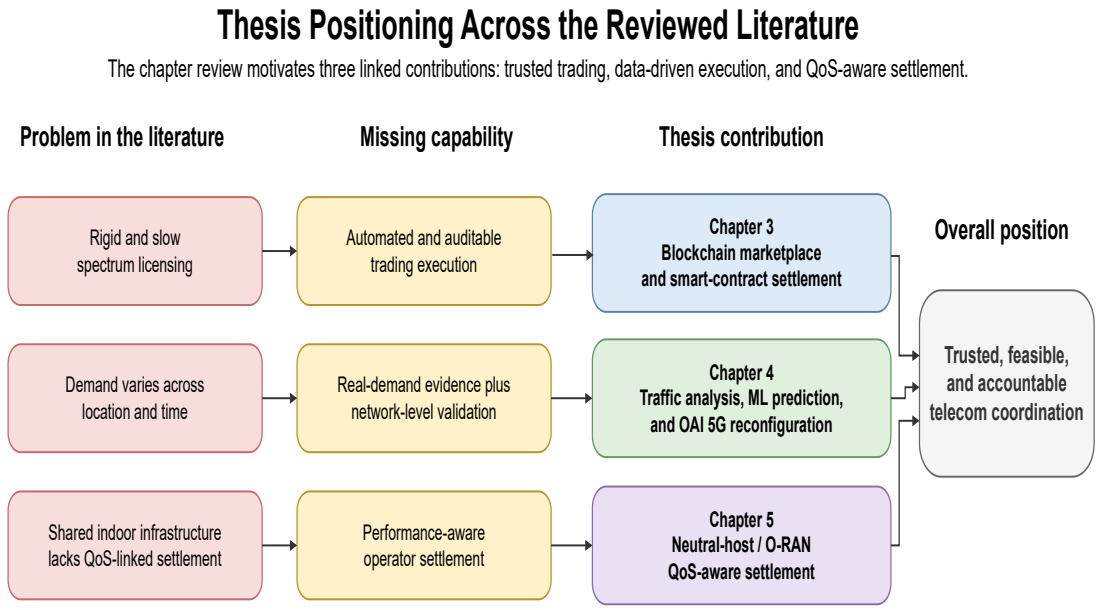


Figure 2.2: Thesis positioning across the reviewed literature. Three gaps identified in the literature are addressed by the contributions in Chapters 3, 4, and 5, respectively.

Chapter 3

Blockchain-Based Spectrum Trading Marketplace

3.1 Introduction

The wireless spectrum is a finite and increasingly valuable resource that underpins the operation of modern wireless communication systems. As demand for mobile broadband services continues to increase, driven by applications such as high-definition video streaming, augmented reality, and the proliferation of connected devices [125], efficient management of spectrum resources has become more critical than ever. However, the traditional approach to spectrum management, whereby national regulatory bodies assign spectrum to mobile network operators (MNOs) through static, desk-based allocation processes, has proven increasingly inadequate in meeting these demands [38, 126].

A key symptom of this inadequacy is underutilisation. Despite the growing scarcity of spectrum, studies have shown that licensed spectrum utilisation rates range between 15% and 85% [24], meaning that significant portions of allocated spectrum remain idle at any given time and location. In the United Kingdom, for example, the shared access spectrum allocation process can take up to six weeks to complete [38,126], creating substantial delays for new entrants seeking access to spectrum resources. Furthermore, MNOs frequently underutilise their allocated spectrum holdings due to inconsistent and spatially uneven traffic patterns [25], worsening the inefficiency of current management approaches. These limitations indicate a fundamental mismatch between the static nature of existing allocation regimes and the dynamic, fluctuating demands of real-world wireless networks.

To address these challenges, this chapter proposes a decentralised marketplace for wireless spectrum trading, built on blockchain technology and realised through Ethereum-based smart contracts. Blockchain's inherent properties, including de-

centralisation, transparency, traceability, and tamper-resistance [127, 128], make it a well-suited foundation for designing a secure and efficient spectrum trading environment. Smart contracts let spectrum transactions run automatically, without a central intermediary, which reduces the time and administrative work usually involved in traditional allocation processes. Spectrum holders with underutilised resources can list and sell or lease their available spectrum to spectrum buyers who require additional capacity, creating a flexible secondary marketplace that adapts to real-time supply and demand.

The marketplace developed in this chapter completes spectrum transactions within seconds, compared to the weeks conventional regulatory processes typically take. All transactions are recorded immutably on the blockchain, giving a transparent, auditable trail that supports regulatory compliance and builds trust among participants. The Telecommunications Regulatory Authority keeps oversight of the marketplace as blockchain administrator, making sure only licensed and authorised operators can take part, without acting as a direct intermediary in individual transactions.

The main contributions of this chapter are as follows:

- A blockchain-based spectrum trading architecture is designed, defining the roles and interactions of spectrum holders, spectrum buyers, and the regulatory authority within a decentralised marketplace.
- A market mechanism is developed, comprising a dynamic pricing function, a rule-based batch matching algorithm, and formal feasibility conditions governing spectrum allocation.
- Interference-aware constraints are incorporated at the market level to ensure spatial, temporal, and frequency-domain coexistence among neighbouring deployments.
- The marketplace is implemented on the Ethereum Sepolia test network using Solidity smart contracts and Python-based automation scripts, and evaluated with respect to execution latency, transaction cost, throughput, and security robustness.

The remainder of this chapter is organised as follows. Section 3.2 positions blockchain technology as an enabling solution. Section 3.3 describes the system architecture and the roles of participating entities. Section 3.4 details the market mechanism design, including the request and offer model, pricing function, and matching algorithm. Section 3.5 addresses interference and coexistence handling.

Section 3.6 justifies the choice of Ethereum and Proof-of-Stake as the underlying blockchain platform. Section 3.7 describes the implementation of the smart contracts and Python scripts. Section 3.8 presents the experimental results and evaluation. Section 3.9 summarises the chapter and motivates the extensions addressed in Chapter 4.

3.2 Positioning Relative to Existing Work

Prior research has explored various applications of blockchain technology to spectrum management and trading. Early works proposed blockchain-based architectures for secure spectrum record-keeping, access control, and stakeholder coordination [2, 129]. These studies mainly focused on architectural design and feasibility, without evaluating execution-level performance. Subsequent studies examined the use of smart contracts to automate spectrum sharing transactions among MNOs [130], and explored economic protocols for securing spectrum trading [131]. More recent consortium and hybrid blockchain frameworks have targeted multi-operator coordination, network slicing, and service level agreement enforcement [132–134]. Other contributions have proposed novel blockchain-based spectrum trading solutions with custom consensus mechanisms [135], and combined blockchain with learning-based dynamic spectrum access strategies [136]. Earlier foundational work also includes blockchain-enabled spectrum access in cognitive radio networks [75, 137] and prototype blockchain radio access network designs [138].

These contributions have advanced our understanding of how blockchain can support spectrum management, but they share a common gap: none offer a practically implemented, experimentally evaluated marketplace for short-term, utilisation-driven spectrum trading under realistic conditions. Much of the existing work focuses on long-term allocation or synthetic demand scenarios, and few studies evaluate execution-level metrics like coordination latency, transaction cost, and robustness against adversarial behaviour. This chapter addresses these gaps with a fully implemented blockchain-based spectrum trading marketplace, evaluated on the Ethereum Sepolia test network under realistic operating conditions.

3.3 System Architecture

This section presents the architecture of the proposed blockchain-based spectrum trading marketplace. The architecture is designed to support secure, transparent,

and efficient spectrum transactions among multiple operators, while preserving regulatory oversight and enabling automated execution through smart contracts. The marketplace operates as a decentralised platform in which authorised participants can list, discover, and purchase spectrum resources without reliance on centralised intermediaries for individual transactions.

3.3.1 Architectural Overview

The proposed marketplace runs on the Ethereum blockchain and involves three principal categories of entities: the Telecommunications Regulatory Authority, spectrum holders, and spectrum buyers. These entities interact through smart contracts deployed on the blockchain, which encode the rules for spectrum listings, ownership transfers, and payment settlements. All transactions are broadcast to the network, validated by Ethereum proof-of-stake validators, and recorded immutably on the distributed ledger, giving a transparent, auditable trail of marketplace activity.

The overall architecture is illustrated in Fig. 3.1. The figure shows the relationships between participating entities, the role of the regulatory authority as blockchain administrator, and the flow of spectrum requests, smart contract execution, payment transfers, and blockchain updates. The marketplace itself operates within the Ethereum environment, while authorised mobile network operators (MNOs) participate as either spectrum holders or spectrum buyers depending on their operational role at any given time. A single operator may act as both a holder and a buyer across different transactions, showing the dynamic nature of real-world spectrum demand.

3.3.2 Participating Entities

- **Telecommunications Regulatory Authority:** The Telecommunications Regulatory Authority (TRA) serves as the administrative entity overseeing the marketplace. The TRA grants access authorisations to MNOs, making sure only licensed and accountable operators can take part in spectrum trading. This authorisation function is enforced at the smart contract and application layers through access control checks, which limit the ability to list, purchase, and manage spectrum resources to approved operator addresses. The TRA also keeps the ability to update the blockchain with transaction records, monitor compliance with spectrum usage regulations, and step in when disputes or non-compliance arise. Importantly, it acts mainly as an overseer rather than a direct participant in individual trans-

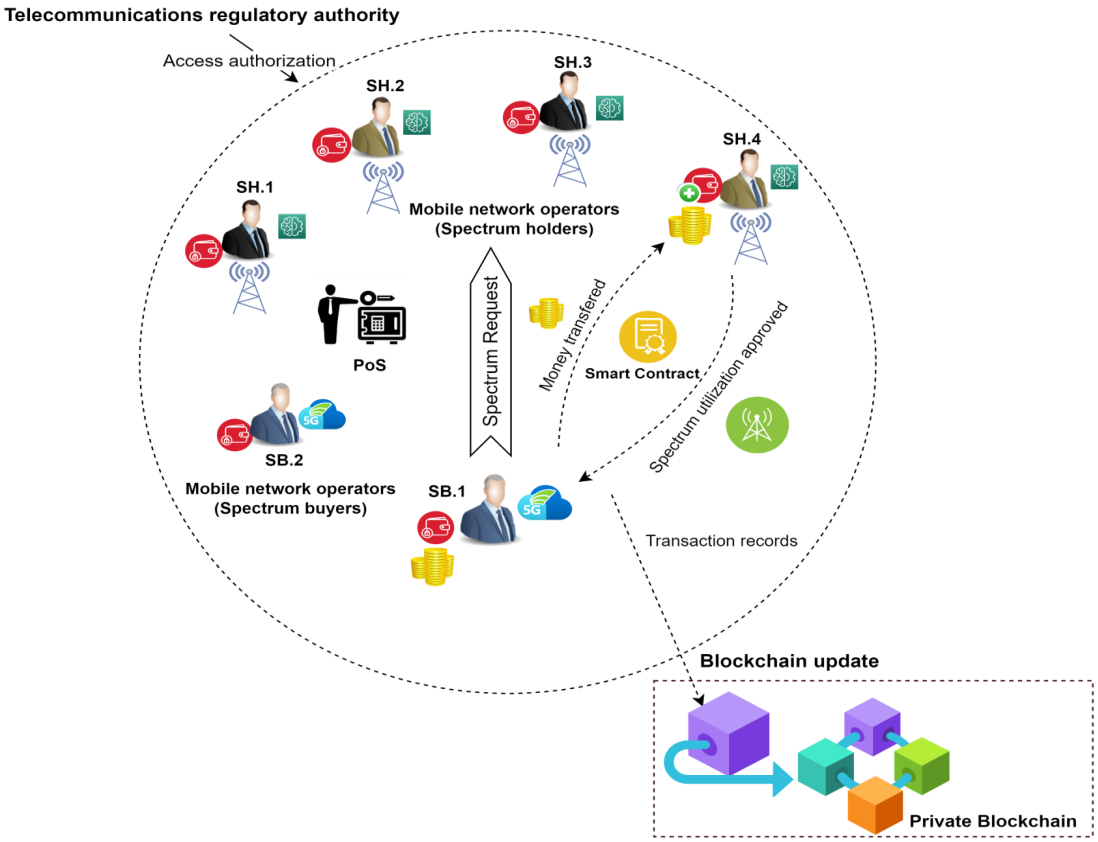


Figure 3.1: Blockchain-based spectrum trading architecture.

actions. Once an operator has been authorised to join the marketplace, day-to-day trading activities proceed without the need for regulatory intervention, allowing the system to retain its decentralised character. This design balances the need for regulatory compliance with the efficiency benefits of automated, peer-to-peer spectrum trading.

- **Spectrum Holders:** Spectrum holders (SHs) are MNOs that possess licensed spectrum and offer portions of their holdings for sale or lease through the marketplace. Each SH manages a set of frequency bands and, recognising that these bands are not fully utilised at all times or in all locations [25], makes its idle spectrum available to other operators in exchange for payment. The motivation for spectrum holders to participate is twofold: they generate additional revenue from otherwise underutilised resources, and they contribute to the more efficient use of spectrum across the wider network ecosystem. Spectrum holders interact with the marketplace by advertising their available spectrum, specifying parameters such as the frequency band, available bandwidth, geographic location, duration of availability, and reserve price. When a suitable match is identified between an advertised listing and a buyer's request, the smart contract is invoked to

formalise the trade.

- **Spectrum Buyers:** Spectrum buyers (SBs) are MNOs or other authorised entities, such as private enterprise networks, that require additional spectrum to meet their operational demands. Spectrum buyers take part in the marketplace by submitting requests that spell out their requirements, including the desired frequency band, bandwidth, geographic location, start time, and lease duration. These requests get broadcast to the network and matched against the available spectrum listings that holders have provided. When a match is found and the terms of the trade are acceptable to both parties, the smart contract executes the transaction, transferring ownership of the spectrum usage rights to the buyer and the corresponding payment to the seller. The spectrum buyer role enables operators to acquire spectrum resources on a flexible, on-demand basis, rather than committing to long-term licences across regions or time periods where the spectrum may not be needed. This is essentially valuable for operators seeking to support localised increases in demand, deploy private networks, or respond to short-term capacity requirements without incurring the costs and delays of conventional licensing processes.

3.3.3 Smart Contracts as the Coordination Mechanism

Smart contracts serve as the central coordination mechanism of the marketplace. A smart contract is a self-executing piece of code deployed on the blockchain that automatically enforces the terms of an agreement when predefined conditions are met. In the proposed system, smart contracts encode the logic for adding spectrum listings, querying available resources, executing purchase transactions, and recording completed trades.

When a spectrum buyer initiates a purchase, the smart contract performs several functions atomically. It verifies that the buyer's payment matches the listed price, transfers ownership of the spectrum from the holder to the buyer, marks the listing as sold to prevent duplicate allocation, and transfers the payment to the holder. All of these actions are executed deterministically and recorded on the blockchain, providing strong guarantees of correctness and transparency. Because smart contracts execute according to their deployed code without intervention, they eliminate the risk of disputes arising from inconsistent or incomplete transaction execution, and they remove the need for trusted intermediaries to mediate between buyers and sellers [130].

The use of smart contracts also enables the marketplace to operate at a speed

that is fundamentally incompatible with manual coordination processes. Spectrum trades that would conventionally require weeks of regulatory and administrative effort can be completed within seconds on the blockchain, subject only to the confirmation latency of the underlying network.

3.3.4 Proof-of-Stake Validators and Blockchain Operation

The proposed marketplace is deployed on the Ethereum public blockchain, which operates under a proof-of-stake (PoS) consensus mechanism. Under PoS, validators are economically incentivised to correctly execute and record transactions according to the protocol's rules, with misbehaviour penalised through the loss of staked Ether. Block proposal and attestation are performed by validators selected by the Ethereum network itself, and the proposed marketplace does not modify or control validator participation.

This architectural choice carries several important implications. First, the marketplace inherits Ethereum's underlying security guarantees, including immutability of recorded transactions, resistance to tampering, and Byzantine fault tolerance. Second, using PoS rather than the earlier proof-of-work mechanism reduces energy consumption and operational overhead in consensus, making the system more sustainable and better suited to high-frequency transaction workloads. Third, separating consensus from application logic means regulatory oversight and access control can be enforced at the smart contract layer, without touching the underlying blockchain protocol. This keeps Ethereum's decentralised security model intact, while still enabling the controlled, accountable participation that spectrum trading requires.

3.3.5 Interaction Flow

The end-to-end interaction within the marketplace works as follows. A spectrum buyer first submits a request specifying its requirements. This request is broadcast to the network, where spectrum holders check it against their available spectrum holdings. When a holder finds a suitable match, it adds the corresponding spectrum listing to the marketplace by invoking the appropriate smart contract function. The buyer can then query the marketplace for available listings, select a match, and initiate a purchase by submitting a payment transaction.

Once the purchase transaction arrives, the smart contract checks its validity, executes the ownership transfer, and settles the payment between the parties. The transaction is then validated by Ethereum PoS validators, included in a block, and recorded on the blockchain. Once confirmed, the buyer holds exclusive usage

rights to the purchased spectrum for the agreed duration, and both parties have access to an immutable record of the transaction for future reference. Section 3.7 presents the detailed workflow of this interaction, including the specific smart contract functions invoked at each step.

This architecture provides a foundation for decentralised, automated, and transparent spectrum trading. The market mechanism that governs how requests are matched to offers, how prices are determined, and how feasibility constraints are enforced is detailed in the following section.

3.4 Market Mechanism Design

This section presents the algorithmic design of the spectrum market that operates on top of the architecture described in Section 3.3. The market mechanism defines how spectrum buyers express their requirements, how spectrum holders advertise available resources, how prices are determined in response to local supply and demand conditions, and how requests are matched to offers under a set of feasibility constraints. The objective is to provide a computationally efficient and demand-responsive market mechanism that is well suited to blockchain-based execution, where deterministic behaviour, bounded computational overhead, and transparency of decision-making are essential.

3.4.1 Request and Offer Model

The marketplace is built around two complementary primitives: requests submitted by spectrum buyers and offers advertised by spectrum holders. Each primitive is structured so that all parameters relevant to allocation, pricing, and feasibility checking are exposed on-chain, enabling smart contracts to enforce trading rules deterministically.

Throughout this chapter, $\mathcal{K}$ denotes the set of spectrum buyers with active (unmatched) requests in the current clearing interval, indexed by k, and $\mathcal{I}$ denotes the set of spectrum holders with active offers, indexed by i. A spectrum buyer $k \in \mathcal{K}$ submits a spectrum request represented as an ordered five-tuple:

$$R_k = (L_k, F_k, B_k, t_k, T_k), \tag{3.1}$$

where $L_k \in \mathcal{L}$ denotes the geographic location or service area, $F_k \in \mathcal{F}$ denotes the requested frequency band, B_k is the requested bandwidth in MHz, t_k is the requested start time, and T_k is the required lease duration. The corresponding requested allocation interval is therefore $[t_k, t_k + T_k)$. This formulation captures

the spatial, spectral, and temporal dimensions of a spectrum request in a way that can be checked against available resources without ambiguity.

Each spectrum holder i advertises its available spectrum capacity $A_{i,t}(L, F) \in \mathbb{R}_+$, measured in MHz, for location $L \in \mathcal{L}$, frequency band $F \in \mathcal{F}$, and discrete clearing interval $t \in \mathbb{T}$. Each advertisement comes with a reserve price $p_i^{\min}(F)$, expressed in tokens per MHz per unit time, setting the minimum price at which the holder is willing to lease the spectrum. Holders can advertise multiple offers covering different bands, locations, and time intervals at once, reflecting how real operators typically manage diverse spectrum portfolios.

When several spectrum offers are available for a given location and time window, the proposed marketplace does not perform implicit radio-level optimisation or physical-layer selection on behalf of buyers. Instead, each offer exposes its contractual and physical attributes directly on the blockchain, including the frequency band, bandwidth, geographic scope, duration, and price. Spectrum buyers choose offers autonomously based on their own objectives and deployment constraints, with pricing acting as the main coordination signal. This design choice avoids hidden optimisation logic, keeps the system transparent, makes allocation decisions fully auditable, and enables automated execution without manual intervention.

3.4.2 Dynamic Pricing Function

Spectrum prices in the marketplace are determined by a posted-price mechanism that adapts to local supply and demand conditions while remaining computationally simple enough for deterministic on-chain settlement. Posted pricing is preferred over auction-based alternatives in this context because it reduces strategic bidding behaviour, eliminates transaction-order dependence, and minimises on-chain computational overhead, all of which are important properties for blockchain-based execution [43, 47].

For each frequency band F, a base rate $p_0(F)$ is defined, expressed in tokens per MHz per unit time. The base price for a request with bandwidth B and duration T is given by

$$P_{\text{base}}(F, B, T) = p_0(F) \cdot B \cdot T. \tag{3.2}$$

To capture the effect of local supply and demand conditions, the aggregated demand and supply at clearing interval t and location L are defined respectively

as

$$D_t(L,F) = \sum_{k \in \mathcal{K}} B_k, \qquad S_t(L,F) = \sum_{i \in \mathcal{I}} A_{i,t}(L,F), \tag{3.3}$$

where $\mathcal{K}$ and $\mathcal{I}$ denote the sets of active buyers and spectrum holders, respectively. From these quantities, a scarcity multiplier is computed as

$$\phi_t(L,F) = \text{clip}\left(\left(\frac{D_t(L,F)+\epsilon}{S_t(L,F)+\epsilon}\right)^{\eta}, \phi_{\min}, \phi_{\max}\right), \tag{3.4}$$

where η controls how sensitive the price is to congestion, and $\phi_{\min}$ and $\phi_{\max}$ bound the multiplier to prevent excessive price volatility. Larger values of η produce more aggressive congestion pricing, and the clipping operation keeps prices within bounds suitable for blockchain execution and predictable for participating operators. The small constant ϵ prevents division by zero when supply or demand is momentarily absent.

To discourage excessive fragmentation of spectrum allocations and to incentivise longer-duration utilisation, a duration adjustment factor is applied:

$$g(T) = \begin{cases} 1 + \kappa_s \left(\dfrac{T_s}{T} - 1\right), & T < T_s, \\ 1 - \kappa_\ell \log\left(1 + \dfrac{T}{T_\ell}\right), & T \geq T_s. \end{cases} \tag{3.5}$$

For short-duration requests below a threshold T_s, the factor increases the price to reflect the administrative and coordination overhead associated with frequent reallocation. For longer-duration requests, the factor introduces a logarithmic discount, encouraging operators to commit to extended utilisation when their needs permit. The parameters κ_s, κ_ℓ, T_s, and T_ℓ are configurable by the marketplace operator to reflect policy preferences.

The final posted price for a request is then

$$P_t(L,F,B,T) = p_0(F) \cdot B \cdot T \cdot \phi_t(L,F) \cdot g(T). \tag{3.6}$$

This pricing function is monotonic in the requested bandwidth B and in the demand-to-supply ratio D/S, bounded by design, and computable from aggregated market statistics. The admissible ranges of the pricing parameters are summarised in Table 3.1. These ranges are design-level policy settings to be chosen by the marketplace operator or the regulator.

The bound on κ_ℓ keeps $g(T)$ positive up to the longest lease $T_{\max}$, and $\phi_{\min} \leq 1 \leq \phi_{\max}$ prices a balanced market ($D = S$) at the base rate.

Table 3.1: Pricing and matching parameters and their admissible ranges.

Parameter	**Meaning**	**Admissible range**
$p_0(F)$	Base rate for band F (tokens per MHz per unit time)	$p_0(F) > 0$
$p_i^{\min}(F)$	Reserve price of holder i	$p_i^{\min}(F) \geq 0$
η	Sensitivity of the price to congestion	$\eta \geq 0$ ($\eta = 0$: no scarcity pricing)
$\phi_{\min}, \phi_{\max}$	Lower and upper bounds of the scarcity multiplier	$0 < \phi_{\min} \leq 1 \leq \phi_{\max}$
ϵ	Small constant preventing division by zero	$\epsilon > 0$
T_s, T_ℓ	Short-lease threshold and long-lease reference duration	$T_s > 0$, $T_\ell > 0$
κ_s	Short-lease surcharge coefficient	$\kappa_s \geq 0$
κ_ℓ	Long-lease discount coefficient	$0 \leq \kappa_\ell < 1/\log(1 + T_{\max}/T_\ell)$
w_1, w_2	Weights of the matching score	$w_1, w_2 \geq 0$, $w_1 + w_2 = 1$

3.4.3 Feasibility Conditions

Not every combination of request and offer constitutes a valid trade. Before a matching is finalised, three feasibility conditions must be satisfied. These conditions ensure that allocations are technically possible, that capacity constraints are respected, and that pricing is acceptable to both parties.

Compatibility. A request R_k can be matched with a spectrum holder i only if the holder supports the requested location and frequency band, that is, the pair (L_k, F_k) is within the holder's advertised capacity.

Capacity Constraint. The holder must have sufficient available bandwidth to accommodate the requested allocation across every clearing interval within the requested duration. Formally,

$$A_{i,t'}(L_k, F_k) \geq B_k, \quad \forall t' \in [t_k, t_k + T_k). \tag{3.7}$$

This constraint ensures that the requested bandwidth can be reserved continuously throughout the lease period without exceeding the holder's available resources at any point.

Price Acceptability. The posted price must meet or exceed the holder's reserve price for the requested resources:

$$P_t(L_k, F_k, B_k, T_k) \geq p_i^{\min}(F_k) \cdot B_k \cdot T_k. \tag{3.8}$$

Only requests that satisfy all three conditions proceed to the matching stage.

3.4.4 Rule-Based Batch Matching Heuristic

The marketplace operates in discrete clearing intervals of duration Δ. Rather than processing each request individually as it arrives, unmatched requests within a clearing interval are accumulated and cleared together through a greedy, rule-based procedure. Batch matching reduces the per-transaction overhead of the blockchain execution and allows the matching algorithm to consider the relative priority of competing requests within each interval.

The matching procedure aims to maximise spectrum utilisation subject to feasibility constraints. Let $\mathcal{W} \subseteq \mathcal{K}$ denote the set of requests accepted (won) in the current clearing interval, where $\mathcal{K}$ is the set of active buyers defined in Section 3.4.1. Formally, the objective is to choose $\mathcal{W}$ so as to maximise

$$\max \sum_{k \in \mathcal{W}} B_k T_k, \tag{3.9}$$

subject to the feasibility conditions described in Section 3.4.3 and the capacity constraints of all participating spectrum holders.

The matching is performed using the procedure outlined in Algorithm 1.

Algorithm 1: Greedy Rule-Based Spectrum Clearing (Heuristic)

```
Collect all unmatched requests submitted within the current clearing interval;
Sort requests in descending order of posted unit price;
foreach request k in sorted order do
    Identify all feasible holders i satisfying compatibility, capacity, and price acceptability conditions;
    Select i* = arg max_i score(k, i);
    Reserve bandwidth B_k over [t_k, t_k + T_k) and update A_{i*,t};
    Instantiate the smart contract C_{k,i*};
Defer any unmatched requests to the next clearing interval;
```

The matching score $\text{score}(k, i)$ combines two considerations: the per-unit value of the trade and the efficiency with which the request packs into the holder's

remaining capacity. Specifically,

$$\text{score}(k,i) = w_1 \cdot \frac{P_t(L_k, F_k, B_k, T_k)}{B_k T_k} + w_2 \cdot \text{fit}(k,i), \tag{3.10}$$

where w_1 and w_2 are configurable weights, and $\text{fit}(k,i)$ measures how efficiently request k utilises the residual capacity of holder i. It is defined as the fraction of the holder's smallest remaining capacity over the requested interval that the request would occupy:

$$\text{fit}(k,i) = \frac{B_k}{\min_{t' \in [t_k,\, t_k+T_k)} A_{i,t'}(L_k, F_k)}, \qquad 0 < \text{fit}(k,i) \leq 1, \tag{3.11}$$

where the upper bound follows from the capacity constraint in Eq. (3.7). A value close to one indicates a tight match that leaves little unused bandwidth, which reduces fragmentation of the holder's spectrum. By prioritising requests with higher posted unit value and rewarding efficient capacity utilisation, the algorithm balances revenue maximisation for holders with efficient use of available spectrum.

3.4.5 Approximation, Clearing, and Stability Properties

The proposed matching algorithm reveals several properties that make it suitable for blockchain-based execution. Sorting the requests incurs $O(n \log n)$ complexity for n requests, while feasibility checking scales as $O(nm \log h)$, where m is the number of spectrum holders and h is the number of discrete time slots considered. These complexities are tractable for the transaction volumes anticipated in a regional spectrum trading marketplace.

The joint time-frequency allocation problem with capacity constraints is NP-hard in general, and exact optimisation is therefore impractical for real-time execution. The greedy procedure proposed here is therefore a heuristic. It seeks a high-utilisation allocation by prioritising requests with higher posted unit value, but it does not guarantee an optimal solution, and no approximation bound is claimed. It trades global optimality for low computational cost and deterministic, auditable behaviour, which are required for blockchain-based execution. Comparing the heuristic with an exact solver on small instances is left for future work.

The design gives two useful guarantees. Capacity gets reserved for the full allocation duration before a contract is finalised, so accepted contracts are always jointly feasible and spectrum can't be oversold. And once a clearing interval ends, whatever requests are still unmatched didn't get left out arbitrarily, they either don't fit the residual capacity or don't clear at the current posted price. In other

words, the market settles into a stable state at the end of each interval, based on whatever supply and demand happen to exist at that point.

Finally, the use of posted pricing combined with greedy batch matching is deliberate. Auction-based mechanisms, while potentially more efficient in some theoretical settings, introduce strategic bidding behaviour, dependence on transaction ordering, and significant on-chain computational overhead. The posted-price, rule-based approach adopted here is more robust to adversarial behaviour, simpler to audit, and better suited to the constraints of blockchain execution [31, 45].

The market mechanism described in this section governs the allocation of spectrum under capacity and pricing constraints. However, spectrum trading must also account for the physical realities of radio interference and coexistence among neighbouring allocations. These considerations are addressed in the following section.

3.5 Interference and Coexistence Handling

The market mechanism described in Section 3.4 governs the allocation of spectrum under capacity and pricing constraints, but it does not on its own guarantee that allocated spectrum can be used without harmful interference. In a real deployment, two spectrum allocations that overlap simultaneously in space, time, and frequency can cause significant degradation to both users, undermining the value of the trade. Dynamic spectrum trading must therefore explicitly account for the risk of interference among neighbouring allocations and incorporate mechanisms to prevent conflicts before they arise.

Physical-layer interference estimation would need detailed channel measurements and a lot of computational overhead, so the system avoids it. Instead, interference-aware constraints are enforced at the market and allocation layer. This lines up with existing regulatory spectrum sharing frameworks and policy-driven access models, most of which already use rule-based geographic, temporal, and frequency separation to keep authorised users from interfering with each other. Because these constraints are built into the matching procedure itself, only allocations that satisfy the policy-defined separation rules are passed on for on-chain instantiation.

It is important to distinguish what these checks guarantee. They are policy-based coexistence checks. They verify that a new allocation respects the separation distance, time overlap, and guard-band rules set by the regulator. They are not physical interference guarantees, because whether harmful interference actually occurs depends on transmit power, antenna patterns, terrain, and prop-

agation conditions, which are not modelled in this chapter. Compliance with the rules therefore reduces the risk of interference, provided the rule parameters are set conservatively, but it does not prove that interference is absent.

3.5.1 Spatial and Temporal Conflict Detection

The first form of interference handling concerns the detection of conflicts that arise when two spectrum allocations overlap simultaneously in space, time, and frequency. Recall that each spectrum request is defined as $R_k = (L_k, F_k, B_k, t_k, T_k)$, where L_k denotes the geographic location, F_k the frequency band, B_k the bandwidth, and T_k the allocation interval. Two allocations k and j are considered to be in conflict if all three of the following conditions are simultaneously true: they are within an interference-relevant geographic distance of one another, their time intervals overlap, and their frequency bands overlap.

This condition is formalised through the following indicator:

$$\mathbb{1}\{\|L_k - L_j\| \leq R_{\text{int}}\} \cdot \mathbb{1}\{T_k \cap T_j \neq \emptyset\} \cdot \mathbb{1}\{F_k \cap F_j \neq \emptyset\} = 1, \tag{3.12}$$

where R_{int} denotes a policy-defined interference radius. This radius is set by the regulatory authority and reflects the maximum geographic separation at which two co-channel allocations are expected to cause harmful interference under typical deployment conditions. If the indicator evaluates to one, the two usage rights are treated as interference-incompatible and can't be granted at the same time. This formulation deliberately sets aside the underlying propagation physics, using a regulatory-defined separation criterion instead of detailed channel modelling, since that's straightforward to evaluate within a smart contract.

3.5.2 Adjacent Channel and Guard-Band Constraints

In addition to co-channel interference, allocations on adjacent frequency bands can interfere with one another through out-of-band emissions and receiver imperfections, even when their nominal frequency assignments do not overlap. To mitigate adjacent channel interference, a minimum guard band is enforced between simultaneously active allocations.

Let f_k^{center} and f_j^{center} denote the centre frequencies of two allocations with bandwidths B_k and B_j. The following constraint must be satisfied:

$$\left| f_k^{\text{center}} - f_j^{\text{center}} \right| \geq \frac{B_k + B_j}{2} + G_{\min}, \tag{3.13}$$

where $G_{\min}$ is a minimum guard band set by the relevant policy or regulation.

The first term on the right-hand side makes sure the two allocations don't overlap in frequency, and the second term adds an extra separation margin to absorb adjacent channel leakage. $G_{\min}$ itself is set by the regulatory authority and can vary depending on the frequency band, expected transmission power, and the spectral characteristics of typical equipment in that band. This constraint prevents harmful adjacent channel interference without needing signal-level measurements, which makes it well suited to evaluation within a smart contract.

3.5.3 Interference-Aware Feasibility Conditions

The capacity and pricing feasibility conditions introduced in Section 3.4.3 establish whether a request can in principle be served by a given spectrum holder. Interference handling adds a further condition: any new allocation must be compatible with all previously accepted allocations currently in force.

Formally, a request R_k is interference-feasible if

$$\mathcal{I}(k, j) = 0, \quad \forall j \in \mathcal{W}, \tag{3.14}$$

where $\mathcal{W}$ denotes the set of accepted allocations that are still active, including those accepted in earlier clearing intervals whose leases have not yet expired, and $\mathcal{I}(k, j)$ is the interference indicator defined in Section 3.5.1. Requests that violate either the spatiotemporal condition or the guard-band constraint defined in Section 3.5.2 are excluded from the matching process during the current clearing interval. This ensures that the marketplace will not instantiate a smart contract that would result in a known coexistence violation.

It is important to note that this feasibility check operates on the set of currently active allocations as recorded on the blockchain. Because the blockchain maintains an authoritative and immutable record of all accepted trades, the set $\mathcal{W}$ is well-defined at any point in time and accessible to the smart contract logic. The interference feasibility check is therefore both deterministic and verifiable.

3.5.4 Interference-Aware Rule-Based Matching

The rule-based batch matching algorithm presented in Section 3.4.4 is extended to incorporate interference feasibility checks. During each clearing interval, candidate buyer-holder pairs that violate spatial, temporal, or frequency-domain interference constraints with respect to the set of previously accepted allocations are filtered out prior to allocation. The remaining candidate pairs are then evaluated using the matching score and the standard greedy selection procedure.

The interference-aware feasibility checks add a linear scan over previously accepted allocations, giving an overhead of $O(|\mathcal{W}|)$ per request. This doesn't change the overall polynomial-time complexity of the matching procedure, and it stays tractable for realistic market sizes. Accepted allocations expire at the end of their lease durations and get removed from the active set $\mathcal{W}$, so the feasibility check overhead scales with the number of currently active allocations, not the total historical volume of trades.

The resulting interference-aware matching procedure ensures that only allocations complying with the policy rules are instantiated on-chain. This guarantees that the marketplace will not grant two usage rights that conflict under those rules, although, as discussed above, it does not by itself guarantee interference-free operation at the physical layer.

3.5.5 Post-Trade Compliance Monitoring

Interference conditions may evolve after an allocation has been finalised, either due to misconfiguration of operator equipment or due to non-compliant operation that deviates from the agreed terms. To address post-trade interference, the proposed system supports compliance monitoring in which operational parameters, such as the activated centre frequency, the actual transmitted bandwidth, and the activation time, are verified against the terms recorded on-chain for each allocation.

Deviations from the agreed terms can trigger a range of policy actions once detected, including warnings to the operator, penalties, or termination of the offending contract. These actions get recorded on the blockchain too, giving an auditable trail of compliance events. The proposed work focuses on market-layer enforcement of interference constraints, but the compliance monitoring mechanism adds a complementary post-trade safeguard that fits with regulatory enforcement practices already used in conventional spectrum management.

By combining pre-trade interference-aware matching with post-trade compliance monitoring, the marketplace provides a layered approach to coexistence handling. Pre-trade constraints prevent the most common forms of interference from being instantiated in the first place, while post-trade monitoring enables the detection and resolution of deviations that emerge during operation. Together, these mechanisms support the integrity of dynamic spectrum trading in a manner that is compatible with both the technical requirements of wireless coexistence and the regulatory requirements of licensed spectrum management.

3.6 Ethereum and the Choice of Proof-of-Stake

The Ethereum blockchain is adopted as the platform for the proposed marketplace, for several reasons. Ethereum supports a mature and widely used smart contract execution environment, and its Solidity programming language provides expressive constructs for implementing the trading logic required in this work. Ethereum's EIP-1559 fee mechanism provides a transparent and predictable model for transaction costs, which is important for operators evaluating the economic overhead of participation in the marketplace.

Crucially, the proposed system is deployed using Ethereum's proof-of-stake consensus mechanism, rather than the earlier proof-of-work approach. Under proof-of-work [127], validators compete to solve computationally intensive puzzles in order to add new blocks to the chain, resulting in high energy consumption and relatively limited transaction throughput. Proof-of-stake replaces this with an economically incentivised selection process in which validators are chosen to propose and attest to blocks based on their staked Ether holdings. This approach significantly reduces the computational and energy overhead of consensus, while preserving the security and integrity guarantees of the underlying blockchain. Furthermore, proof-of-stake systems generally support higher transaction throughput and lower confirmation latency than their proof-of-work counterparts, both of which are important properties for a spectrum trading marketplace that must respond to real-time demand.

Ethereum's consensus layer stays permissionless, but that's not where this marketplace controls who participates. Access control and authorisation happen one layer up, at the smart contract and application level. Only operator addresses the Telecommunications Regulatory Authority has approved can list, purchase, or manage spectrum assets through the deployed contracts. The result is a system that maintains the security and transparency of a public blockchain, yet limits actual participation to licensed, accountable operators.

3.6.1 Layered Blockchain Architecture

The blockchain underlying the marketplace can be described as a set of layers, as shown in Fig. 3.2. The data layer stores blocks, transactions, and the shared ledger state. The network layer is the peer-to-peer network through which nodes exchange transactions and blocks. The consensus layer uses Proof-of-Stake validators to agree on the order of blocks. The execution layer runs the smart contracts in the Ethereum Virtual Machine (EVM). Finally, the application layer contains the participants and tools that interact with the contracts, namely the operators,

the regulator, their wallets, and the Python/Web3 scripts. The proposed marketplace adds the smart contract in the execution layer and the components of the application layer, while the consensus, network, and data layers are provided by the Ethereum network without modification.

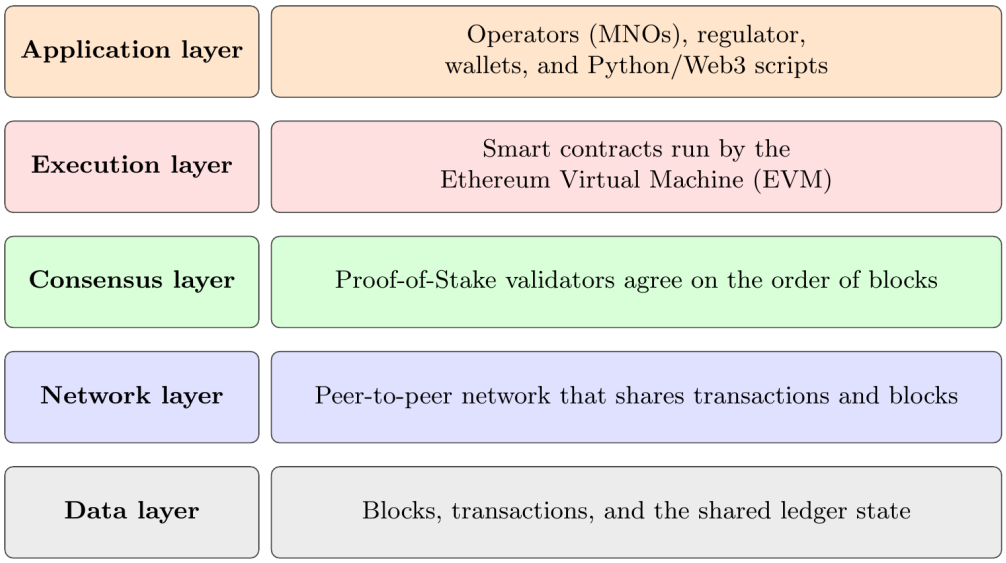


Figure 3.2: Layered architecture of the blockchain used by the spectrum trading marketplace.

3.7 Implementation

This section describes the implementation of the proposed blockchain-based spectrum trading marketplace. The implementation comprises two principal components: a set of Ethereum smart contracts that encode the trading logic on-chain, and a collection of Python automation scripts that interact with the blockchain to manage spectrum listings and extract transaction data. Together, these components realise the architecture described in Section 3.3 and the market mechanism described in Section 3.4, providing a working implementation that can be deployed on the Ethereum network and evaluated under realistic operating conditions. It is important to distinguish what was implemented in this chapter from what is proposed at the design level. The implemented prototype consists of the Solidity smart contract described in Section 3.7.2, which handles listing, discovery, purchase, ownership transfer, and payment, and the Python scripts described in Section 3.7.3, which generate listings, check them against operator criteria, and extract completed trades. The dynamic pricing function (Section 3.4.2), the batch-matching algorithm (Section 3.4.4), and the interference-aware constraints

(Section 3.5) are specified at the design level, while the experimental evaluation in Section 3.8 focuses on the on-chain trading lifecycle. Their full integration and evaluation are left for future work.

In terms of the on-chain and off-chain division, only operations that must be atomic and auditable are executed on-chain: creating listings, transferring ownership, transferring payment, and recording completed trades. Listing generation, matching, price setting, and data extraction are performed off-chain by the Python scripts, and their results are submitted to the contract as transactions.

3.7.1 Marketplace Overview

The objective of the implementation is to develop a decentralised marketplace for trading wireless spectrum rights that addresses the inefficiencies and limitations of current spectrum allocation processes. To achieve this, the implementation must satisfy a number of operational objectives.

Operator Authentication. The marketplace must ensure that participating MNOs, whether acting as spectrum holders or spectrum buyers, are authenticated in a manner that prevents impersonation while preserving their anonymity within the network. Authentication is achieved through the use of blockchain addresses authorised by the Telecommunications Regulatory Authority, which serve as cryptographically verifiable identifiers without requiring the disclosure of personal or commercial information.

Privacy. The marketplace must protect the privacy of participating operators by ensuring that spectrum requests, advertisements, and transactions do not reveal personally identifiable information such as operator names, addresses, or financial details, unless explicitly required by regulatory authorities. While transactions are recorded immutably on the blockchain, they are associated with pseudonymous addresses rather than directly identifiable entities.

Traceability. Even with pseudonymous addresses, the Telecommunications Regulatory Authority still needs a way to trace the identity of participating operators if a dispute, fraud, or non-compliance case comes up. The access authorisation process handles this, linking blockchain addresses to operator identities the regulatory authority already holds. Privacy holds during routine operation, but accountability is there when intervention is needed.

Reliability. The marketplace has to keep working even if network outages or cyberattacks hit individual operators. This comes from the distributed ledger itself, replicated across every participating node, which keeps functioning as long as enough nodes stay available. Each operator can hold a full copy of the blockchain, so transaction records stay intact and accessible.

Data Confidentiality and Integrity. Trading messages must be transmitted securely and protected from unauthorised modification during communication. The cryptographic properties of the Ethereum network ensure that transactions cannot be altered after submission, and that their contents are verifiable by all participants. While privacy concerns the protection of operator identities, confidentiality and integrity concern the protection of transaction contents and their resistance to tampering.

The end-to-end workflow that realises these objectives is illustrated in Fig. 3.3 and proceeds as follows. A spectrum buyer first advertises a spectrum request to the blockchain. Spectrum holders evaluate the request against their available spectrum and add a matched listing to the marketplace by invoking the appropriate smart contract function. The buyer looks through available listings on the marketplace, picks the one that fits best, and kicks off a purchase through MetaMask. From there, the smart contract takes over: it verifies the transaction, transfers spectrum ownership to the buyer, pays the seller, and updates the blockchain record. After that, both sides can refer to the completed transaction, and the buyer holds the spectrum for the duration specified in the contract.

3.7.2 Ethereum Smart Contract

The smart contract is implemented in Solidity, the standard programming language for Ethereum-based smart contracts, and developed using the Remix integrated development environment. Remix is a web-based tool that supports the compilation, testing, and deployment of Solidity smart contracts directly to the Ethereum network. The use of Remix and Solidity reflects standard practice in the Ethereum development ecosystem and ensures compatibility with widely available tools for verification, debugging, and auditing.

The smart contract has two main types of elements: functions and events. Functions are the executable code that carries out the marketplace's operations, and events give a way to notify authenticated nodes when state changes get recorded on the blockchain. Operators interact with the smart contract by calling its functions, and the resulting state changes get propagated to all participants through the blockchain's consensus mechanism.

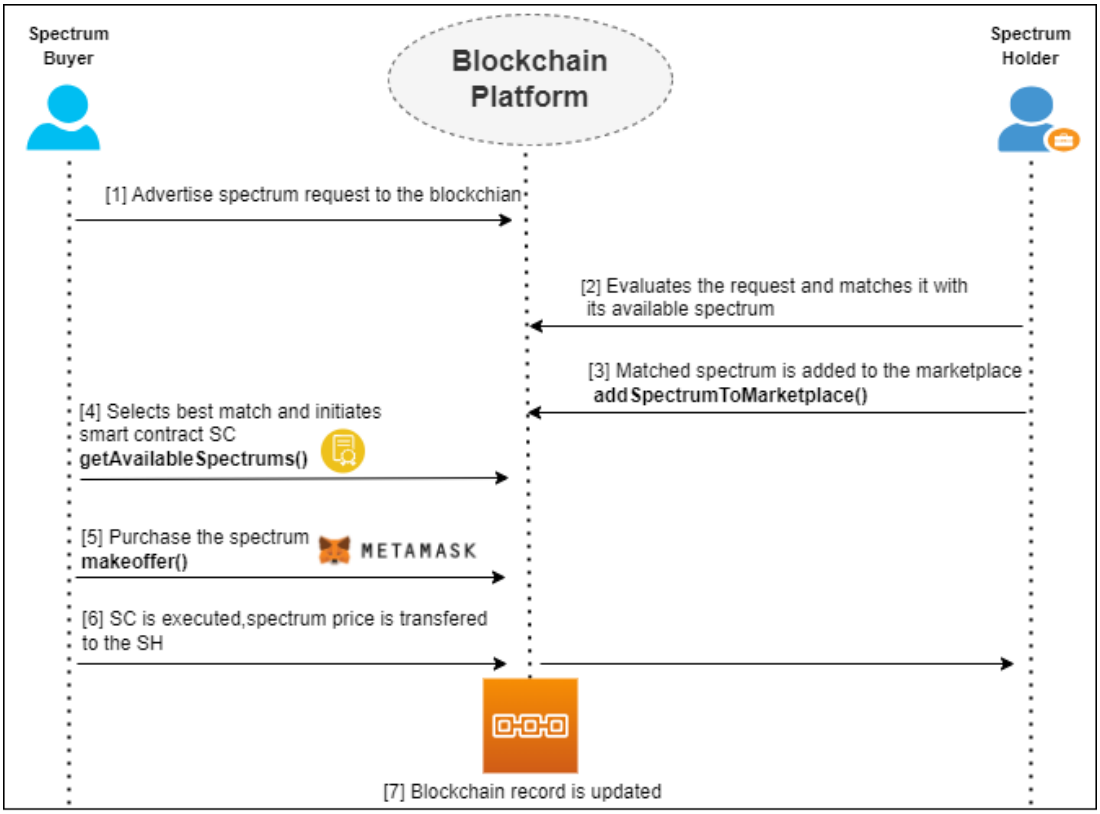


Figure 3.3: Blockchain-based spectrum trading workflow.

The core functions of the smart contract are summarised in Table 3.2, and the smart contract interface as displayed in the Remix environment is illustrated in Fig. 3.4.

These functions implement the complete lifecycle of a spectrum trade, from initial listing through to settlement. The `addSpectrumToMarketplace()` function is invoked by a spectrum holder to publish a new listing, recording all relevant parameters on the blockchain. Once published, the listing becomes visible to everyone through `getAvailableSpectrums()`, which gives a read-only view of the current marketplace state. When a buyer wants to purchase a listing, it calls `makeOffer()` and sends the required payment as part of the transaction. The smart contract checks that the payment matches the listed price, transfers ownership to the buyer, marks the listing as sold to prevent duplicate allocation, and forwards the payment to the seller, all in one atomic operation. The `getSoldSpectrums()` function provides a record of completed trades, which is used by operators to track their spectrum holdings and by external systems to coordinate downstream actions such as network reconfiguration.

The functions are invoked through externally owned accounts using MetaMask as the transaction signing interface. MetaMask is a widely used Ethereum wallet that enables operators to securely sign and submit transactions to the

Table 3.2: Core Smart Contract Functions

Function	Purpose and On-Chain Effect
`addSpectrumToMarketplace()`	Registers a spectrum listing on-chain by storing the frequency band, centre frequency, bandwidth, duration, location, price, and ownership metadata, and emits an event for auditability.
`getAvailableSpectrums()`	Returns the set of unsold spectrum listings, supporting off-chain discovery and operator-side decision-making. This is a read-only function that incurs no transaction cost.
`makeOffer()`	Executes an atomic trade by validating the buyer's payment, transferring ownership of the spectrum to the buyer, marking the listing as sold, and settling the payment to the seller.
`getSoldSpectrums()`	Returns the set of completed trades, supporting monitoring, accounting, and the triggering of downstream network reconfiguration. This is a read-only function.

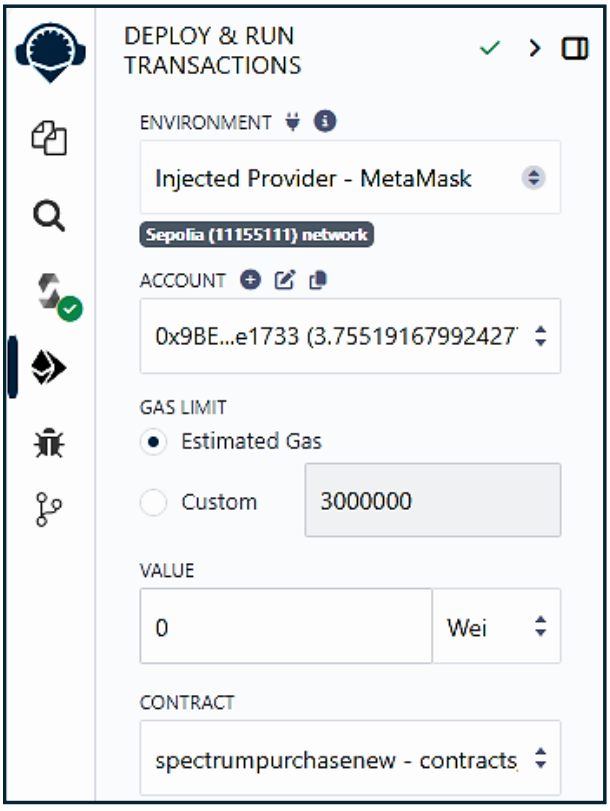


Figure 3.4: Smart contract interface in the Remix IDE with MetaMask injected provider on the Sepolia network.

blockchain. Each transaction is executed on the Ethereum network and incurs a gas fee determined by the computational cost of execution and the prevailing network conditions. These fees are analysed in detail in Section 3.8 to characterise the coordination-layer overhead of the marketplace.

3.7.3 Python Automation Scripts

While the smart contract encodes the trading logic on-chain, day-to-day operation of the marketplace requires off-chain automation to generate spectrum listings, match requests against operator-specified criteria, submit transactions to the blockchain, and extract data from completed trades. To support these functions, two Python scripts have been developed using the Web3 library, which provides a Python interface for interacting with Ethereum-based systems.

The Web3 library is a widely adopted tool for connecting to Ethereum networks and enabling Python programs to perform tasks such as deploying smart contracts, sending transactions, and reading blockchain data. Its use is justified by its ability to simplify complex blockchain operations while remaining fully compatible with Ethereum-based decentralised applications. The two scripts developed for the marketplace serve complementary purposes: the first automates the creation and management of spectrum listings, while the second extracts and stores data on completed trades.

Connecting to Ethereum. Both scripts connect to the Ethereum network through the Web3 library, using an Infura API endpoint to access the Sepolia test network. Sepolia is an Ethereum test network designed for development and experimentation, providing the same operational characteristics as the Ethereum mainnet without the associated financial costs. The endpoint URL, referred to as `sepolia_url`, is the remote procedure call interface through which the scripts submit transactions and query blockchain state. This setup enables secure, cost-effective interaction with deployed smart contracts during development and evaluation.

Generating and Testing Parameters. The first script automates the generation of spectrum listings on behalf of spectrum holders. As shown in Fig. 3.5, the script generates random spectrum parameters including the frequency band, bandwidth, location, duration, and price. These parameters simulate the responses that spectrum holders would provide in a live marketplace, where holders draw on their actual spectrum holdings rather than on randomly generated

values. The script then evaluates the generated parameters against a set of predefined operator criteria to determine whether they satisfy the conditions required to constitute a valid match.

```
Connected to the Ethereum network
Latest Block Number: 4281238
Testing parameters: [3, 60, 'Non-Contiguous', '4 hours', 'G128QW-Learning Hub Building', 650982]
Testing parameters: [2, 60, 'Non-Contiguous', '1 hour', 'G128RS-Adam Smith Building', 524837]
Testing parameters: [2, 20, 'Contiguous', '6 hours', 'G128RS-Adam Smith Building', 206138]
Testing parameters: [2, 20, 'Contiguous', '2 hours', 'G128QW-Learning Hub Building', 863564]
Testing parameters: [2, 60, 'Contiguous', '2 hours', 'G128QW-Learning Hub Building', 391760]
Testing parameters: [3, 20, 'Non-Contiguous', '2 hours', 'G128RS-Adam Smith Building', 493651]
```

Figure 3.5: Random spectrum parameters generation output from the Python script.

Matching Parameters. Following parameter generation, the script compares the generated values against operator-specified criteria. These criteria typically include the desired frequency band, an acceptable bandwidth range, a preference for contiguous or non-contiguous spectrum, a minimum acceptable duration, and a preferred geographic location. If the generated parameters satisfy all of the specified criteria, they are flagged as a match and proceed to the listing stage.

Creating a New Spectrum Listing. When a match is identified, the script constructs and submits a transaction to the smart contract. This transaction invokes the listing function `addSpectrumToMarketplace()` with the matched parameters and registers the new spectrum listing on the blockchain, making it available for discovery and purchase by other operators. The script handles transaction signing, gas fee estimation, and submission to the network, abstracting these blockchain-specific details from the higher-level matching logic.

Extracting Sold Spectrums. The second script connects to the smart contract to retrieve data on completed trades and store this data in a text file for downstream use. The extracted data includes details such as the central frequency and bandwidth of each sold spectrum allocation, along with the identity of the new owner. This information is critical for several purposes. First, it enables operators to track which spectrum bands are currently in use and which remain available, supporting efficient spectrum allocation and planning. Second, it provides the data needed to coordinate network operations and avoid interference among neighbouring allocations. Third, it supports compliance with regulatory requirements and licensing agreements by maintaining an auditable record of

spectrum ownership and usage. The script is designed to run continuously or on demand, ensuring that the local record of sold spectrum remains synchronised with the authoritative record on the blockchain.

Together, the smart contract and the Python automation scripts provide a complete implementation of the proposed marketplace. The smart contract enforces the trading rules on-chain with strong guarantees of correctness and transparency, while the Python scripts handle the off-chain automation needed to manage listings, match requests, and extract transaction data. The performance and security characteristics of this implementation are evaluated in the following section.

3.8 Results and Evaluation

This section presents the experimental evaluation of the proposed blockchain-based spectrum trading marketplace. The evaluation is conducted on the Ethereum Sepolia test network, which provides the same operational characteristics as the Ethereum mainnet without incurring real financial costs. The objective is to characterise the behaviour of the marketplace across several complementary dimensions: the cost and feasibility of deploying and operating the smart contract, the performance of individual trading transactions, the throughput and scalability of the system under load, and the robustness of the marketplace against adversarial behaviour.

The metrics reported in this section relate to coordination-layer performance rather than radio-level metrics. Specifically, gas usage, transaction cost, and throughput characterise the computational and economic overhead of executing spectrum trading operations on the blockchain. These measurements provide empirical parameters for analysing the scalability of the marketplace and for understanding how blockchain-level constraints influence the design of dynamic spectrum trading systems.

3.8.1 Smart Contract Deployment

The deployment of the smart contract is a prerequisite for marketplace operation, marking the transition from theoretical design to practical implementation on the blockchain. Deployment is performed by submitting a transaction that contains the compiled bytecode of the smart contract to the Ethereum network. Once the transaction is included in a block, the smart contract becomes available at a unique address and can be invoked by authorised operators.

Deployment is initiated through MetaMask, which serves as the wallet interface for signing and submitting transactions. The deployment transaction details, including the transaction hash, status, block number, and gas cost, are illustrated in Fig. 3.6. The successful status indicates that the contract was deployed without error and is now operational on the network. The deployment was included in a specific block on the Sepolia test network, and the resulting contract address is used by all subsequent transactions that interact with the marketplace.

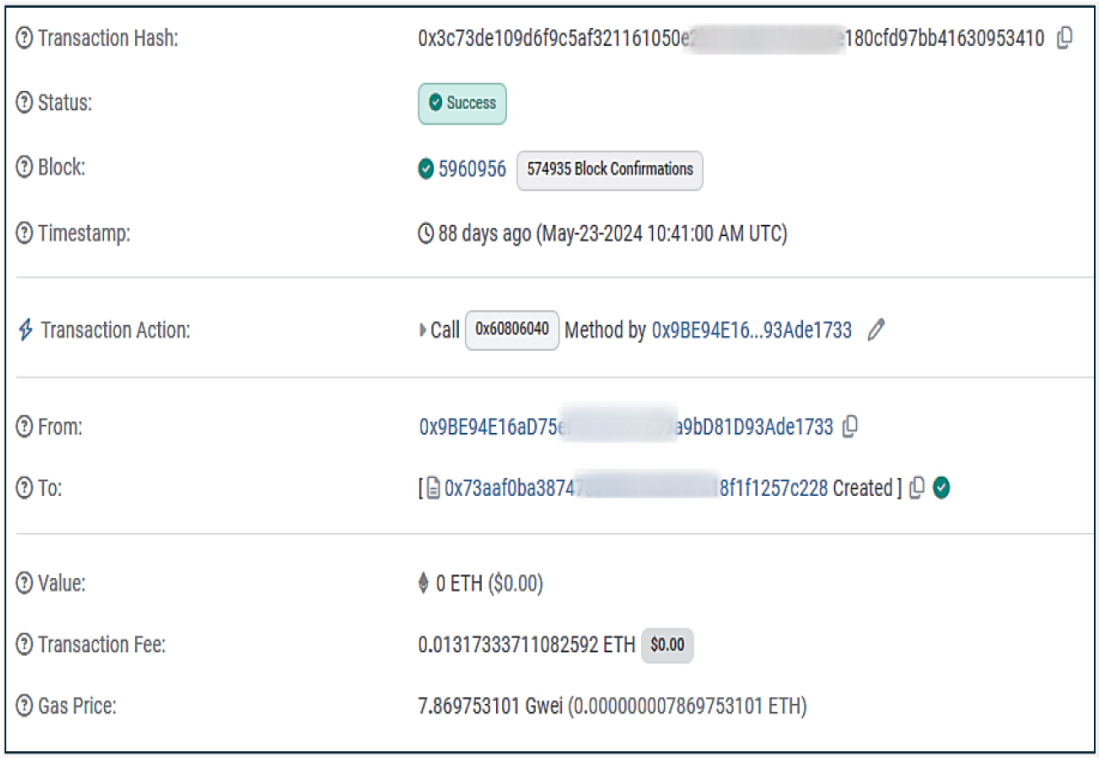


Figure 3.6: Smart contract deployment transaction details in MetaMask.

The cost of deployment is determined by Ethereum's gas fee mechanism. Gas is the unit used to measure the computational resources required to run operations on the blockchain. The total transaction fee comes from two components: the amount of gas the transaction consumes and the gas price, expressed in gwei, where one gwei equals 10^{-9} Ether. Under Ethereum's EIP-1559 fee mechanism, the total transaction fee is computed as

$$F_{\text{total}} = G_u(G_b + G_p), \tag{3.15}$$

where G_u denotes the amount of gas used by the transaction, G_b is the base fee per unit of gas, and G_p is the priority fee per unit of gas. The product $G_u \cdot G_b$ represents the base fee component, which is burned and removed from circulation, while $G_u \cdot G_p$ represents the priority fee that is paid to validators as an incentive

for processing the transaction.

This fee structure does two things at once. Burning the base fee reduces the total Ether supply over time, thereby supporting the network's long-term economic sustainability. The priority fee gives validators an incentive to include transactions in blocks, keeping processing timely even as network load varies. For the proposed marketplace, deployment is a one-time cost, and the recurring transaction costs incurred during normal operation are covered in the following subsections.

3.8.2 Spectrum Trading Transaction

To illustrate the end-to-end operation of the marketplace, this subsection presents a complete spectrum trading transaction, from the initial state of available listings through to the final state in which the spectrum has been transferred to the buyer.

Fig. 3.7 shows the marketplace state before the transaction, displaying a representative spectrum listing available for purchase. The listing includes the frequency band, centre frequency, bandwidth, location, duration, and price, as well as the current owner address and a flag indicating that the spectrum has not yet been sold. These parameters together describe the resource being offered and the terms under which it is available.

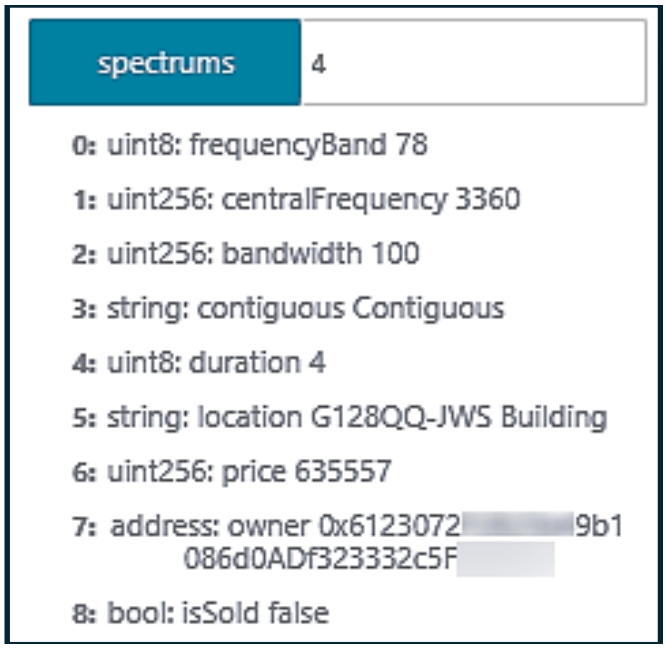


Figure 3.7: Available spectrum listing before the transaction.

When a spectrum buyer wishes to purchase the listing, it initiates a `makeOffer()` transaction through MetaMask. Fig. 3.8 illustrates the MetaMask interface during this transaction, showing the gas limit, the value of the payment being sent,

and the address of the smart contract being invoked. The buyer confirms the transaction, which is then signed and submitted to the network for validation.

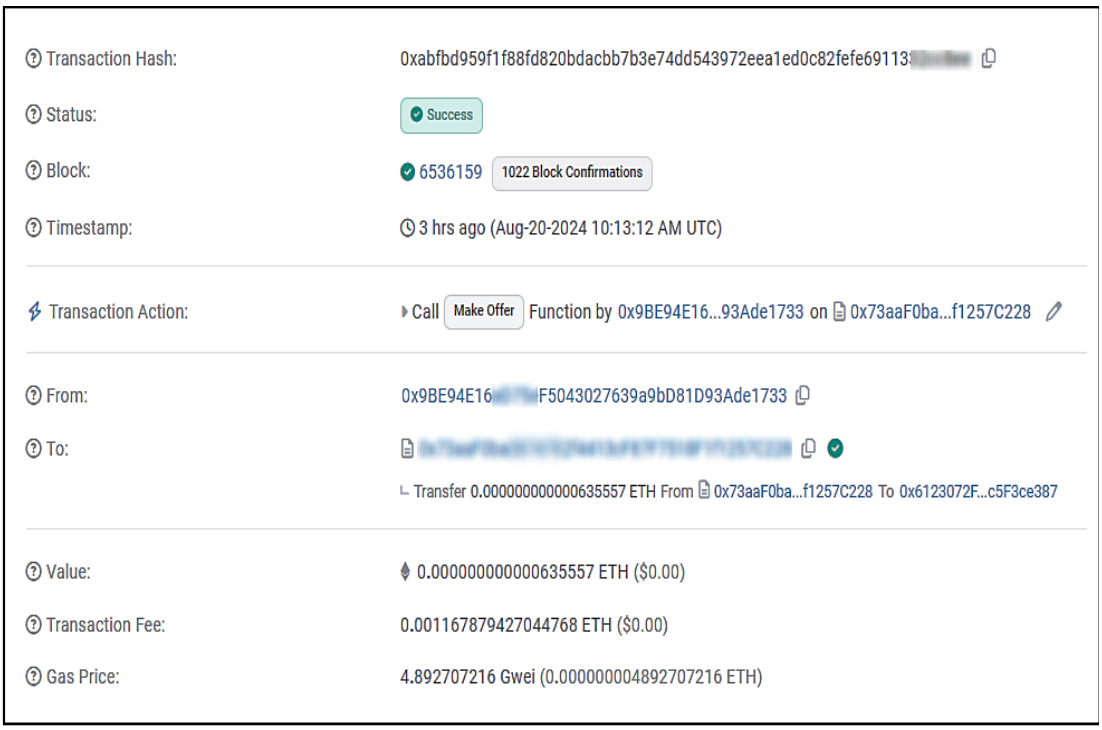


Figure 3.8: Spectrum purchase transaction interface in MetaMask.

The smart contract processes the transaction by verifying that the payment matches the listed price, transferring ownership of the spectrum to the buyer, marking the listing as sold, and forwarding the payment to the seller. All of these actions are executed atomically within a single transaction, providing a strong guarantee that the trade either completes in its entirety or has no effect on the marketplace state. Fig. 3.9 shows the marketplace state after the transaction, with the listing now marked as sold and ownership transferred to the buyer's address. The Ether sent by the buyer has been transferred to the seller, completing the trade.

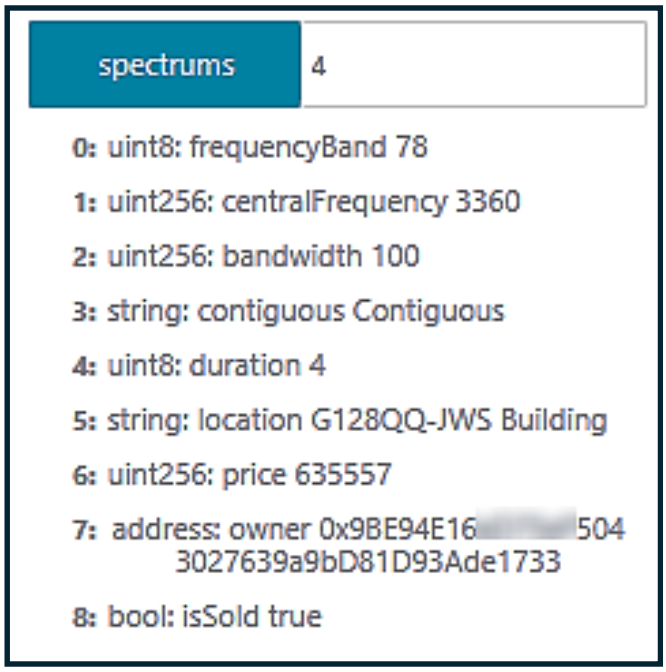


Figure 3.9: Sold spectrum listing after the transaction.

This example illustrates the core operation of the marketplace and provides the basis for the performance and security analysis presented in the following subsections.

3.8.3 Performance Metrics

The performance of the marketplace is characterised by three principal metrics: the gas usage of core transactions, the resulting transaction cost, and the throughput achievable under the constraints of the underlying blockchain. These metrics capture the computational and economic overhead of executing spectrum trading operations and provide the empirical parameters needed for the scalability analysis in Section 3.8.5.

Gas Usage. Fig. 3.10 reports the gas usage of the two spectrum trading transactions on the Ethereum Sepolia network. The `addSpectrumToMarketplace()` transaction, which registers a new listing on the blockchain, exhibits lower gas usage because its primary effect is to write spectrum metadata to contract storage. In contrast, the `makeOffer()` transaction incurs higher gas usage due to the additional execution steps required, including offer validation, ownership transfer, state updates, and payment settlement. Gas is reported in gas units, which measure computational effort and are independent of the price of gas. The monetary cost is obtained by multiplying the gas used by the gas price (Eq. (3.15)).

These measurements characterise the relative computational complexity of the two operations and provide a baseline for understanding the scalability of the marketplace.

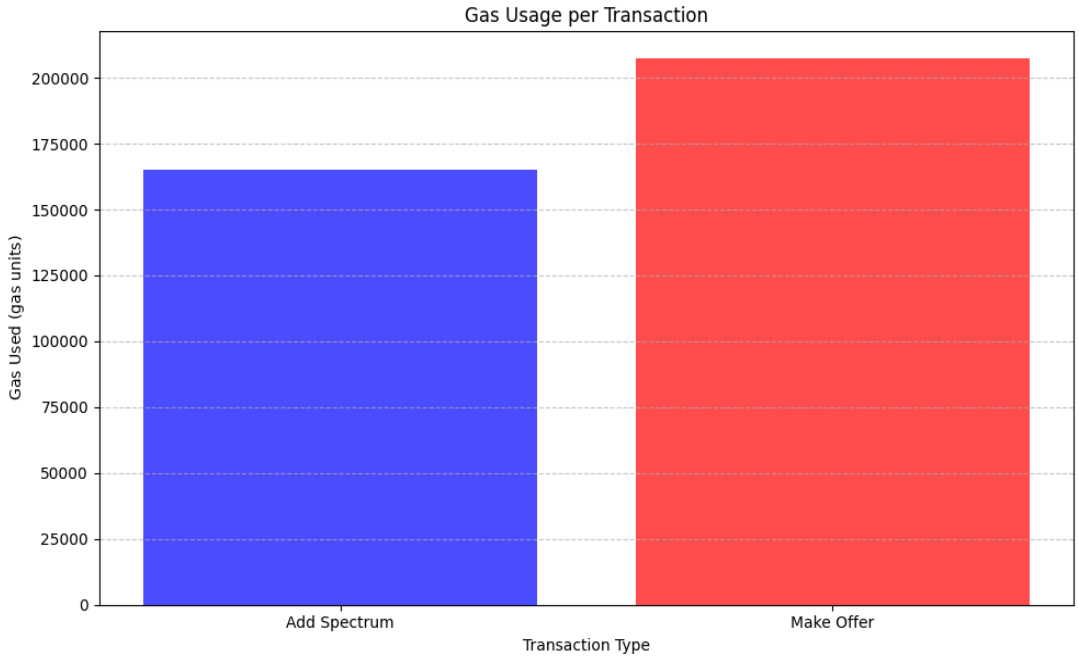


Figure 3.10: Gas usage of core spectrum trading transactions on the Ethereum Sepolia network.

Transaction Cost. The transaction cost results derived from the measured gas usage are shown in Fig. 3.11. As expected, the `makeOffer()` transaction incurs a higher cost than the `addSpectrumToMarketplace()` transaction, reflecting its greater execution complexity and the larger number of on-chain state changes involved. These costs represent the coordination-layer overhead associated with operating the marketplace on the Ethereum network. They are not intended to serve as a direct comparison with traditional spectrum auction pricing or long-term licensing fees, but rather to quantify the relative economic cost of executing spectrum trading operations within the proposed framework.

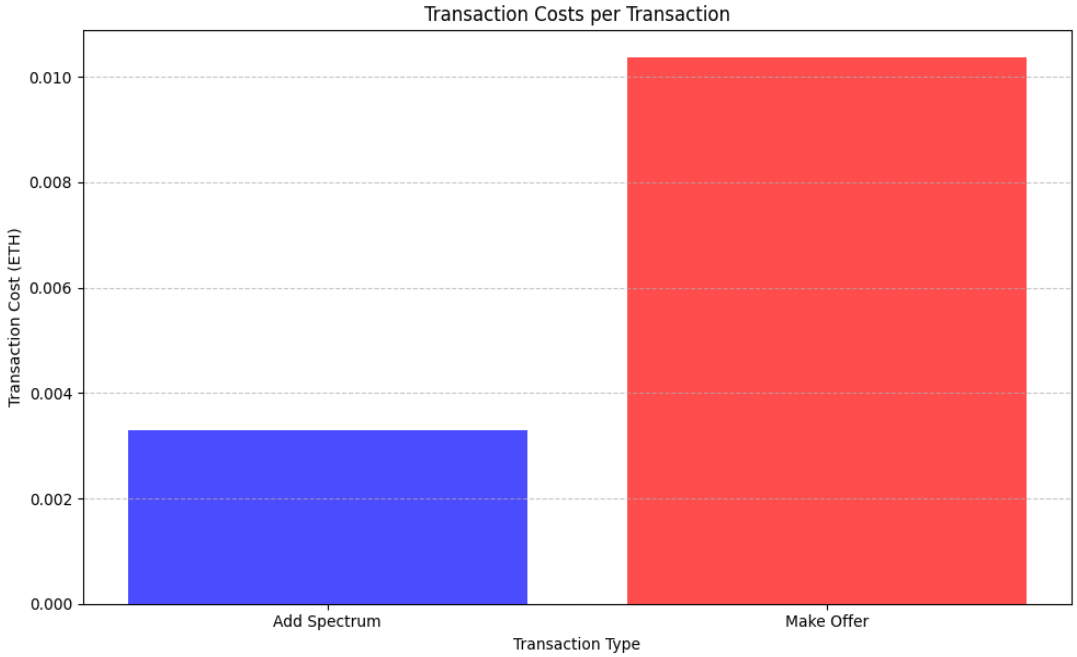


Figure 3.11: Transaction cost of spectrum trading operations derived from measured gas usage.

Throughput. Fig. 3.12 illustrates the throughput characteristics of the marketplace under the measured blockchain configuration. Because its execution logic is simpler, the listing transaction `addSpectrumToMarketplace()` reaches higher throughput, whereas the purchase transaction `makeOffer()` runs at lower throughput because it carries extra validation and state-transition operations. These results reflect coordination-layer throughput characteristics that are constrained by blockchain parameters such as block gas limits and block intervals, rather than by the end-to-end capacity of the application. The throughput measurements are used in Section 3.8.5 to inform the scalability analysis by exposing how transaction complexity influences service rates under increasing load.

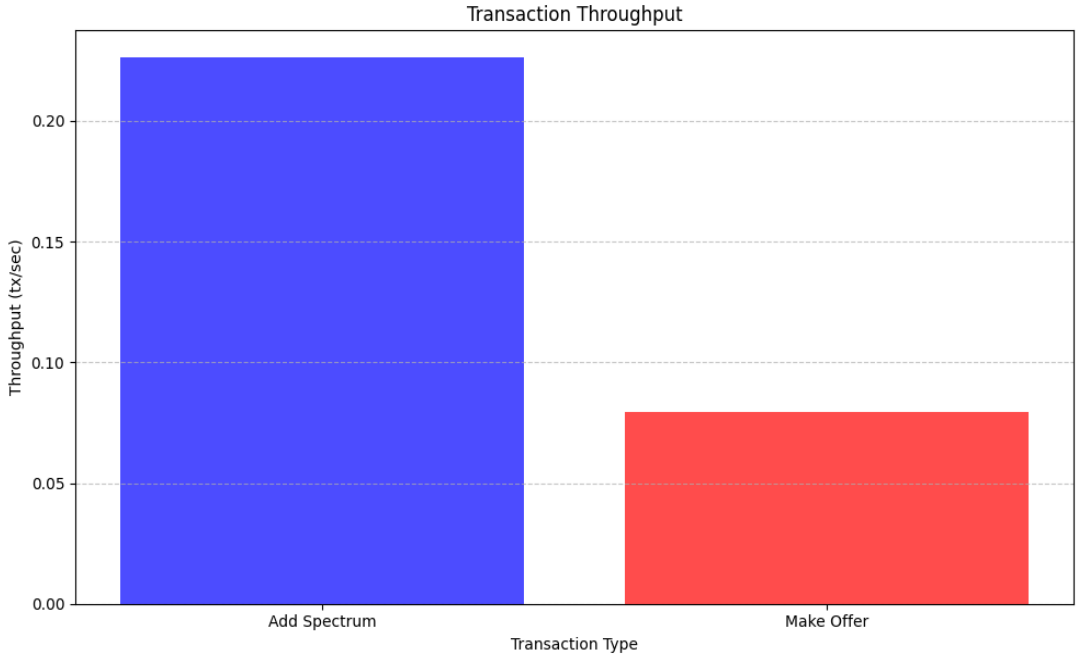


Figure 3.12: Throughput characteristics of the proposed spectrum trading marketplace based on measured blockchain parameters.

Taken together, these measurements provide empirical insight into the execution overhead and coordination behaviour of the proposed marketplace. Rather than demonstrating production-scale operational feasibility in isolation, they serve to parameterise the scalability analysis and to clarify how transaction complexity and blockchain constraints influence coordination latency, cost, and throughput under increasing demand.

3.8.4 Security and Robustness Evaluation

A spectrum trading marketplace must be robust against adversarial behaviour that could undermine the correctness of trades, compromise the integrity of payments, or enable double allocation of spectrum resources. This subsection evaluates the robustness of the proposed marketplace against several classes of adversarial behaviour, including replay manipulation, payment tampering, front-running, and smart contract vulnerabilities. The evaluation focuses on correctness and safety guarantees at the contract-execution and market-coordination layers.

Threat Model. The evaluation considers adversaries acting as malicious spectrum buyers or holders with the ability to observe the public blockchain mempool, submit competing transactions, and attempt replay or manipulation attacks. The adversary is assumed to have no control over Ethereum consensus, cryptographic

primitives, or the execution integrity of deployed smart contracts. Denial-of-service attacks targeting the underlying Ethereum network itself are considered out of scope for this evaluation.

Replay and Double-Purchase Attacks. Replay and double-purchase attempts were empirically evaluated by invoking the `makeOffer()` function on spectrum listings that had already been sold. As demonstrated in Fig. 3.13, all such transactions were deterministically reverted with the error message "Spectrum is already sold". This confirms that the smart contract enforces single-allocation semantics by construction, preventing the same spectrum from being allocated to multiple buyers. The protection is automatic and does not depend on external monitoring or intervention.

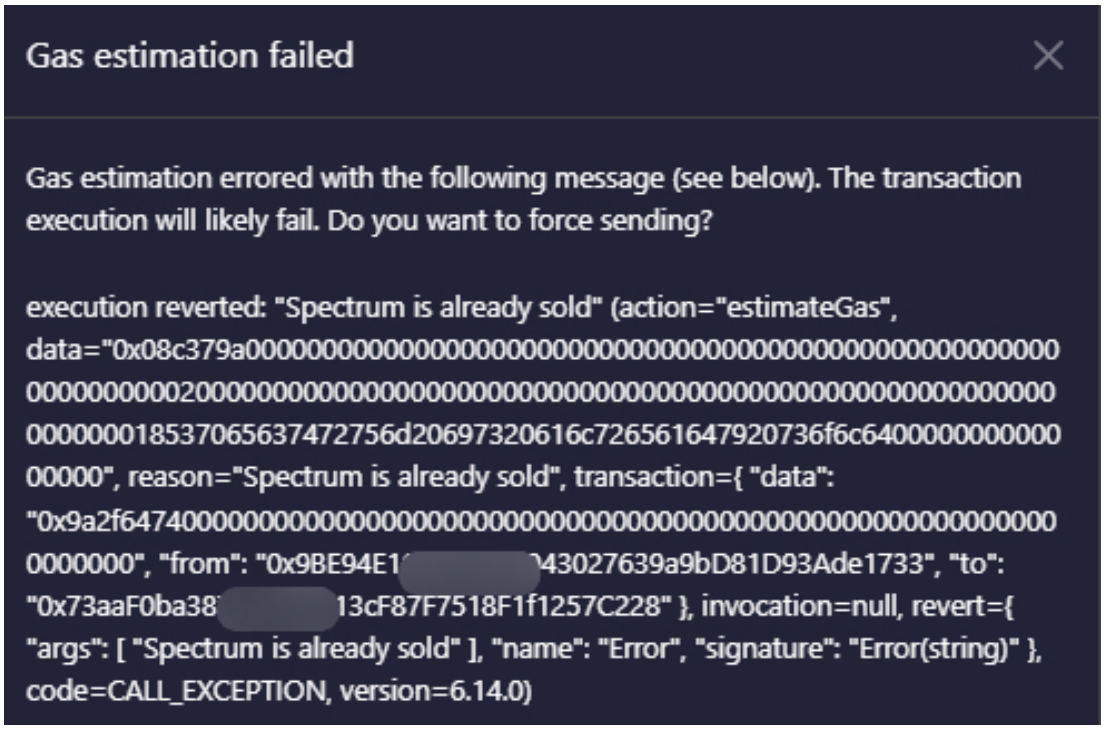


Figure 3.13: Transaction revert when attempting to repurchase an already sold spectrum listing.

Payment Manipulation. To assess robustness against payment manipulation, purchase attempts were submitted with incorrect payment values, where the `msg.value` field of the transaction did not match the listed price of the spectrum. In all tested cases, the transactions reverted with the error message "Incorrect offer price", as illustrated in Fig. 3.14. This demonstrates that the smart contract enforces strict on-chain validation of payment integrity, preventing underpayment, overpayment, or any inconsistency between the agreed price and the

actual settlement amount.



Figure 3.14: Transaction revert caused by submitting an incorrect payment value.

Front-Running. Front-running behaviour was evaluated by submitting competing purchase transactions for the same spectrum listing with different gas prices. In a public blockchain, transactions with higher gas prices are typically prioritised by validators for inclusion in the next block, which creates the possibility for an attacker to observe a pending transaction and submit a competing transaction with a higher gas price in an attempt to acquire the same resource. The experimental results confirm that transaction ordering does influence which buyer successfully acquires the spectrum, but the smart contract still guarantees correctness by allowing only a single successful allocation. All competing transactions revert once the spectrum is sold, preventing double-selling and loss of funds. In other words, front-running can affect transaction ordering, but it does not compromise the safety of the allocation process. This is a safety property, not a fairness property. The contract does not prevent front-running, and an adversary who pays a higher priority fee can still acquire a listing ahead of an honest buyer. Preventing front-running would require additional mechanisms, such as commit–reveal purchasing or batch clearing at a common price, which are left for future work.

Static Analysis of Smart Contract. The smart contract implementation was further analysed using the Slither static analysis tool with Solidity compiler version 0.8.26. Slither applies a suite of one hundred built-in vulnerability detectors that check for common smart contract issues such as re-entrancy, unchecked external calls, integer overflows, and access-control violations. The analysis reported zero findings, indicating that no issues were detected by any of the one hundred detectors. The clean analysis result is shown in Fig. 3.15.

```
C:\Users\2603932A\Desktop\OAI_Intg>slither SpectrumPurchaseNew.sol
'solc --version' running
'solc SpectrumPurchaseNew.sol --combined-json abi,ast,bin,bin-runtime,
srcmap,srcmap-runtime,userdoc,devdoc,hashes --allow-paths
.,C:\Users\2603932A\Desktop\OAI_Intg' running
INFO:Slither:SpectrumPurchaseNew.sol analyzed (1 contracts with
100 detectors), 0 result(s) found
```

Figure 3.15: Slither static analysis output for the spectrum trading smart contract with Solidity v0.8.26, showing zero findings across 100 vulnerability detectors.

All empirical security tests were executed across multiple transactions and accounts on the Ethereum Sepolia test network. The deterministic revert behaviour observed under adversarial transaction attempts confirms that the proposed marketplace is robust against common blockchain-specific attacks within its intended operational model.

3.8.5 Scalability Analysis

The performance metrics reported in Section 3.8.3 characterise the behaviour of individual transactions but do not, on their own, indicate how the marketplace would behave under increasing market load. This subsection presents a scalability analysis that combines measured on-chain parameters with analytical modelling to characterise the throughput, latency, and ledger growth properties of the marketplace as demand increases.

Load Model and Assumptions. Let N_{SB} and N_{SH} denote the number of active spectrum buyers and spectrum holders, respectively. Their combined activity generates spectrum trading transactions, such as `makeOffer()` invocations, with an aggregate arrival rate λ measured in transactions per second. The effect of many buyers and holders is folded into λ, which captures the overall market load. The blockchain itself is characterised by an average block interval τ_b and a block gas limit $G_{\max}$. Transactions are assumed independent, processed first-come-first-served by the mempool and block proposer. A fraction $\alpha \in (0, 1]$ of submitted transactions succeed, and the rest revert due to violated contract conditions, such as duplicate purchase attempts or incorrect payment values.

Measured On-Chain Parameters. Measurements were collected from the Ethereum Sepolia test network to provide empirical values for the analytical model [139]. Table 3.3 summarises the block-level parameters observed for a representative block, while Table 3.4 reports the measured cost of the spectrum purchase operation.

Table 3.3: Measured Block-Level Parameters (Sepolia, Block #6536159)

Metric	**Value**
Block interval τ_b	≈ 12 s
Gas limit $G_{\max}$	30,000,000
Gas used	6,746,257 (22.49%)
Transactions in block	54
Block size	174,575 bytes
Base fee	3.392707216 Gwei

Table 3.4: Measured Transaction Parameters for `makeOffer()`

Metric	**Value**
Gas used G_{offer}	238,698
Transaction fee	0.001167879427 ETH
Transaction type	EIP-1559 (Type 2)

Throughput Under Load. Given the measured gas usage of the `makeOffer()` operation, the maximum number of such transactions that can be included in a single block is approximated by

$$C_b \approx \left\lfloor \frac{G_{\max}}{G_{\mathrm{offer}}} \right\rfloor = \left\lfloor \frac{30{,}000{,}000}{238{,}698} \right\rfloor = 125. \tag{3.16}$$

The corresponding service rate, representing the upper-bound transaction throughput, is

$$\mu \approx \frac{C_b}{\tau_b} \approx \frac{125}{12} \approx 10.42 \text{ tx/s}. \tag{3.17}$$

Letting λ denote the aggregate arrival rate and α the empirical success ratio, the achieved throughput is given by

$$X(\lambda) = \min(\alpha\lambda, \mu). \tag{3.18}$$

Fig. 3.16 illustrates this behaviour. Throughput scales linearly with λ under low load and saturates at the gas-limited capacity μ, independently of the number of buyers or holders. This indicates that the scalability of the marketplace is primarily constrained by the per-block gas limit and the block interval of the underlying blockchain, rather than by the size of the participant pool.

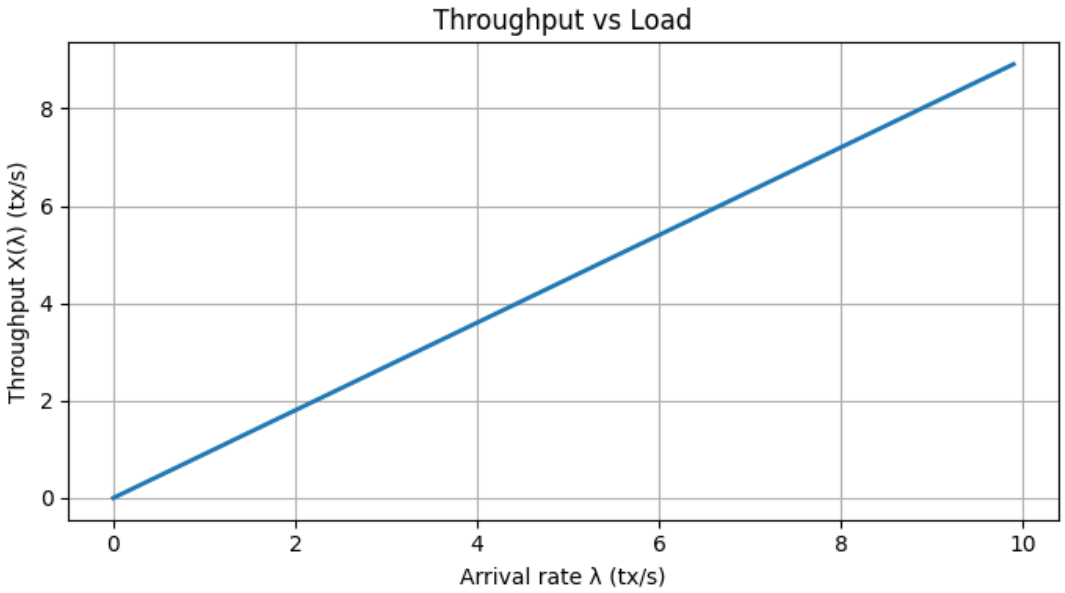


Figure 3.16: Achieved transaction throughput as a function of arrival rate λ. Throughput increases linearly under low load and saturates at the gas-limited service capacity $\mu \approx 10.42$ tx/s.

Latency and Block Confirmation Time. Transaction inclusion in a block is approximated as a queue served at rate μ. For $\lambda < \mu$, the mean waiting time is

$$W(\lambda) \approx \frac{1}{\mu - \lambda}. \tag{3.19}$$

Letting k denote the confirmation depth (the number of blocks required to consider a transaction final), the confirmation delay is

$$T_{\text{conf}} \approx k\tau_b, \tag{3.20}$$

yielding an end-to-end latency

$$L(\lambda) \approx W(\lambda) + T_{\text{conf}}. \tag{3.21}$$

Fig. 3.17 shows the resulting latency for confirmation depths $k = 1$ and $k = 2$. Latency remains close to the confirmation delay under light load and increases sharply as λ approaches the service rate μ, reflecting congestion effects in the mempool. The choice of confirmation depth involves a trade-off between latency and finality: a depth of $k = 1$ provides faster confirmation but slightly weaker guarantees, while $k = 2$ provides stronger finality at the cost of an additional block interval.

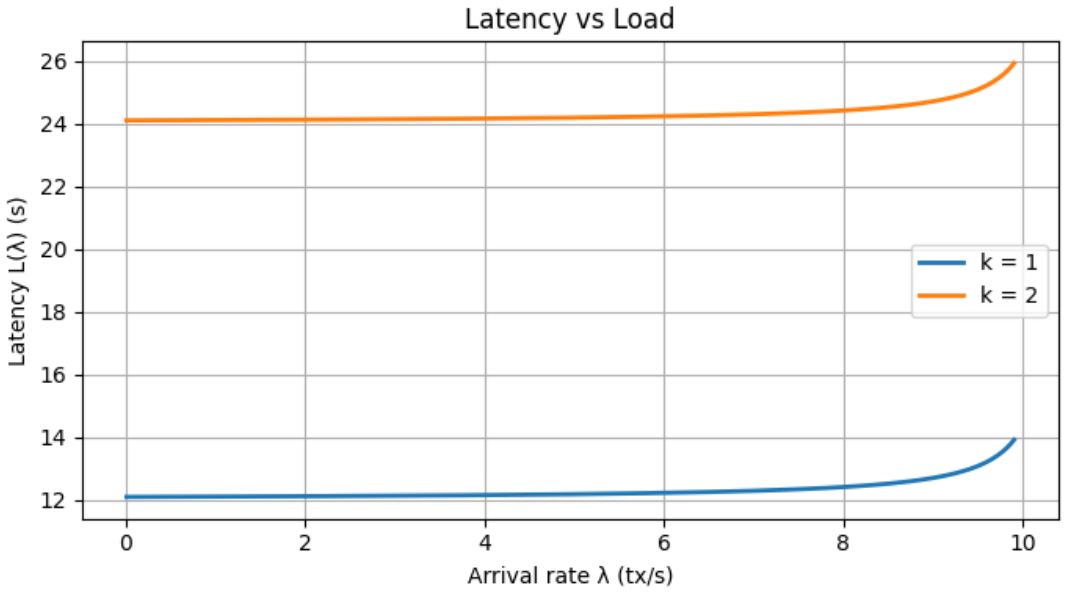


Figure 3.17: End-to-end latency as a function of arrival rate λ for confirmation depths $k = 1$ and $k = 2$.

Regional Partitions and Long-Lived Leases. The proposed architecture can be scaled beyond the per-blockchain capacity through two complementary mechanisms. First, the marketplace can be partitioned across regions through the deployment of independent smart contracts. With R regions and approximately

balanced load $\lambda_r = \lambda/R$, the aggregate service capacity scales as

$$\mu_{\text{tot}} \approx \sum_{r=1}^{R} \mu_r, \tag{3.22}$$

where μ_r is the service rate available to region r. This additive scaling holds only if the regional contracts run on independent execution capacity, such as separate chains or Layer-2 rollups. If all regional contracts are deployed on the same blockchain, they share the same block gas limit and block interval, so that $\sum_{r=1}^{R} \mu_r \leq \mu$, and partitioning does not increase the total transaction capacity. Moreover, on a public blockchain the marketplace shares block space with all other users, so $\mu \approx 10.42$ tx/s, which assumes blocks filled entirely with `makeOffer()` transactions, should be read as an upper bound.

Second, long-lived spectrum leases reduce the frequency of trading transactions by extending the duration of each allocation. Rather than contributing to ongoing transaction throughput, active listings and leases contribute to the persistent on-chain state. Let $n_A(t)$ and $n_L(t)$ denote the number of active listings and active leases, respectively. The contract state footprint scales as

$$S_{\text{state}}(t) \approx S_0 + n_A(t)s_A + n_L(t)s_L, \tag{3.23}$$

where s_A and s_L are the per-listing and per-lease storage requirements, and S_0 is the baseline contract overhead.

Ledger growth is dominated by transaction receipts and logs. Using the measured block size and transaction count, the average ledger footprint per transaction is approximately

$$\ell_{\text{tx}} \approx \frac{174{,}575}{54} \approx 3.23 \text{ KB/tx}. \tag{3.24}$$

At peak capacity, this implies an upper-bound ledger growth rate of approximately 33.6 KB/s, though in practice the marketplace would operate well below saturation, resulting in significantly lower long-term growth.

The scalability analysis indicates that the marketplace can support moderate spectrum trading workloads within the capacity of a single Ethereum-compatible blockchain, which is bounded by the block gas limit and block interval and shared with other users of the chain. Longer-duration leases reduce the required transaction rate, while regional partitioning increases capacity only when partitions are deployed on independent execution layers [140]. These results are analytical and were not validated by a load test at saturation.

3.9 Summary

This chapter presented the design, implementation, and evaluation of a blockchain-based marketplace for wireless spectrum trading. The proposed system uses Ethereum smart contracts to enable authorised spectrum holders and buyers to advertise, discover, and execute spectrum trades transparently and securely under regulatory oversight. A dynamic market mechanism was developed that combines posted pricing, feasibility constraints, batch matching, and interference-aware allocation checks to support spatially, temporally, and spectrally compatible trades. The marketplace was implemented in Solidity and supported by Python/Web3 automation scripts for listing generation, trade execution, and transaction retrieval. Evaluation on the Ethereum Sepolia test network demonstrated correct end-to-end operation, atomic ownership transfer and payment settlement, measurable gas and transaction costs, and robustness against common smart contract risks, with Slither analysis reporting no detected vulnerabilities. These results, taken together with the performance and security evaluations presented earlier in this chapter, indicate that blockchain-based spectrum trading is technically feasible and provides a sound foundation for the extensions developed in Chapter 4. However, the chapter evaluated the marketplace primarily as a coordination and settlement layer. The next chapter therefore extends this work by using real cellular traffic data to demonstrate spatio-temporal spectrum trading opportunities and by integrating the marketplace with an OpenAirInterface-based 5G testbed to validate automated network reconfiguration after completed trades.

Chapter 4

Blockchain-Enabled Dynamic Spectrum Trading for Next-Generation Wireless Networks

4.1 Introduction

Chapter 3 established a decentralised marketplace for wireless spectrum trading, built on Ethereum smart contracts, and demonstrated through experimental evaluation that spectrum transactions can be executed securely, transparently, and within seconds at the coordination layer. That contribution established the technical feasibility of blockchain-based spectrum trading and confirmed its robustness against common forms of adversarial behaviour. However, two important questions were left open. First, the motivation for the marketplace rested on general arguments about spectrum underutilisation rather than on direct, data-driven evidence that short-term and localised trading opportunities actually exist in real cellular networks. Second, the marketplace was evaluated in isolation, without demonstrating that a completed trade could be translated into an operational change within a live radio access network.

This chapter addresses both questions and extends the marketplace of Chapter 3 into a complete, demand-aware, and network-integrated framework suitable for 5G-Advanced, 6G, and beyond networks. The chapter makes three connected advances. The first contribution offers an empirical study of fine-grained cellular activity in an urban setting, using the Milan dataset [141]. From this, a normalised utilisation index is derived to capture how often spectrum demand drops

below its peak, and a two-stage k-means clustering method is applied to reveal spatial variation in demand. A variation that in turn creates opportunities for localised trading. The second contribution presents an ANN-based model for traffic prediction that shifts spectrum trading from reactive to proactive by enabling trading decisions to anticipate periods of low demand before they arise [83, 118]. Third, it integrates the marketplace with an OpenAirInterface (OAI) based 5G testbed [142], demonstrating that a completed on-chain spectrum trade can trigger automated reconfiguration of the radio access network without modifying the underlying physical- or medium-access-layer mechanisms.

It is worth noting that the system proposed here works at the level of spectrum access and coordination, not at the physical or medium-access layer. Its purpose is to make the authorisation, enforcement, and reconfiguration of exclusive spectrum-use rights fast and reliable for all stakeholders involved. As a result, evaluation focuses on coordination latency, execution cost, scalability, and enforcement correctness, rather than on radio-level performance. Positioned this way, the framework complements existing radio access network and PHY/MAC solutions, while filling a key gap in current spectrum-access workflows [7, 53].

The main contributions of this chapter are summarised as follows:

- Using real-world urban cellular activity data from the Milan dataset, this chapter provides demand-side evidence of spatio-temporal variation consistent with spectrum underutilisation, and derives a normalised demand-based utilisation index, a proxy computed from traffic activity rather than a measurement of spectrum occupancy, that reveals short-lived, localised opportunities for spectrum trading.

- Demand patterns are further characterised through learning-assisted analysis: spatial characterisation via a two-stage k-means clustering method, and an illustrative ANN-based traffic prediction example that demonstrates how proactive trading coordination might work.

- The marketplace introduced in Chapter 3 is extended by incorporating feasibility constraints, both interference-aware and policy-driven, at the market level, ensuring exclusive spectrum authorisation while maintaining coexistence with neighbouring deployments.

- Feasibility is validated end-to-end through integration of the trading workflow with an OAI-based 5G testbed, with experimental artefacts and configuration details documented to support reproducibility.

This chapter proceeds as follows. Section 4.2 motivates the case for dynamic spectrum trading through an empirical, spatio-temporal study of urban cellular traffic drawn from the Milan dataset. Section 4.3 presents the ANN-based model used for proactive traffic prediction. Section 4.4 explains how the marketplace from Chapter 3 is positioned and adapted for next-generation networks. Section 4.5 covers the integration of the framework with the OAI-based 5G testbed. Section 4.6 examines the implications, scalability, and limitations of the combined framework. Section 4.7 sets out the reproducibility details and experimental configuration. Section 4.8 closes the chapter with a summary.

4.2 Motivation: Empirical Evidence for Spectrum Trading

Efficient spectrum management is essential to the evolving demands of 6G and beyond, where high-bandwidth, latency-sensitive applications continue to push network requirements higher [21, 125]. Conventional allocation mechanisms depend largely on fixed, static assignments that fail to reflect how demand shifts unevenly across both space and time [1, 24]. The result is that large portions of licensed spectrum sit idle in particular regions or during particular periods, producing inefficient network operation and constraining the flexibility needed to support emerging services such as the Internet of Things, autonomous systems, and immersive applications [18, 25]. The marketplace introduced in Chapter 3 offers a mechanism to exploit this inefficiency; this section instead builds the empirical case that justifies doing so.

To situate the framework within realistic demand conditions, this chapter examines fine-grained urban cellular activity drawn from the Milan city dataset [141]. Although per-operator demand cannot be observed directly, the spatial heterogeneity present in the data establishes the precondition needed for inter-operator spectrum trading. A k-means clustering method is used to capture the dataset's spatial dynamics and show that load profiles vary across different parts of the metropolitan area [86, 89]. This variation confirms that spectrum trading is feasible, since traffic load and spectrum availability differ by location.

The Milan dataset records call, SMS, and internet activity across 10,000 grid squares, each covering an area of $235\,\mathrm{m} \times 235\,\mathrm{m}$ [87]. Recordings are taken at 10-minute intervals, giving 144 data points per day. Because of gaps in the data, the analysis was limited to 5,000 grids over a 21-day window. This allows spectrum utilisation to be assessed across both time and space, and the data shows marked

heterogeneity in traffic demand: activity varies throughout the day and differs from one part of the city to another. Because the grid structure mirrors the real geography of Milan, the resulting demand patterns and clusters can be understood in relation to actual urban locations, rather than treated as abstract groupings alone.

The structure of the dataset is summarised in Table 4.1. Each record of the Telecom Italia *SMS, Call, Internet – Milano* files [141] has $n = 8$ columns. A record exists for every combination of grid square, 10-minute interval, and country code with non-zero activity, so the activity columns were first summed over country codes. The activity values are dimensionless: they are proportional to the number of call detail records in each interval and have been rescaled by the data provider so that the true volumes are not disclosed. Table 4.2 gives the dimensions ($m \times n$) of the data at each stage of the analysis.

Table 4.1: Columns of the Milan telecommunications dataset (one record per grid square, 10-minute interval, and country code).

No.	Column	Type / unit	Description
1	Square ID	Integer, 1–10,000	Identifier of the 235 m × 235 m grid square
2	Time interval	Timestamp (ms)	Start of the 10-minute interval (Unix epoch)
3	Country code	Integer	Country of the interacting party's network
4	SMS-in activity	Dimensionless	Received SMS activity
5	SMS-out activity	Dimensionless	Sent SMS activity
6	Call-in activity	Dimensionless	Received call activity
7	Call-out activity	Dimensionless	Issued call activity
8	Internet traffic activity	Dimensionless	Internet (data session) activity

4.2.1 Spatial Traffic Variations

For clustering purposes, call, SMS, and internet activity were combined into a single traffic-load value per grid, with each grid treated as representing an individual cell. A two-stage clustering approach was used to examine how traffic varies spatially across the city. The results show that some grids consistently carry lower traffic than others, pointing to persistent spatial heterogeneity in demand and, by extension, localised underutilisation of spectrum [82, 88]. These findings underline the need for spectrum trading mechanisms flexible enough to reassign resources dynamically across regions as demand conditions change.

Table 4.2: Dimensions ($m \times n$) of the data used at each stage of the analysis.

Stage	**Size** $m \times n$	**Rows** $\times$ **columns**
Grid–time activity matrix	5,000 × 3,024	grid squares × 10-minute slots (21 days × 144)
First-layer clustering features	5,000 × 144	grid squares × daily profile (averaged over 21 days)
Utilisation index $U_g(t)$	500 × 3,024	top 10% grids × 10-minute slots
ANN data (one grid)	3,024 × 151	samples × input features (144 + 7)
Training window	2,016 samples	14 days × 144
Independent test week	1,008 samples	7 days × 144

First-Layer Clustering

The first clustering stage groups grids according to their average traffic profiles, bringing out clear differences in spectral utilisation between geographic regions. For each grid, traffic values are averaged across the 21 days for every time slot, yielding a 144×1 feature vector at 10-minute resolution. These vectors are fed into the k-means algorithm and applied across all 5,000 grids. The elbow method, with a variance-explained threshold of 0.95, was used to determine the optimal number of clusters. The clusters obtained show that static spectrum allocation cannot capture the pronounced spatial variation in demand seen across the urban area [19]. Fig. 4.1 shows the outcome of this first-layer clustering based on average traffic intensity.

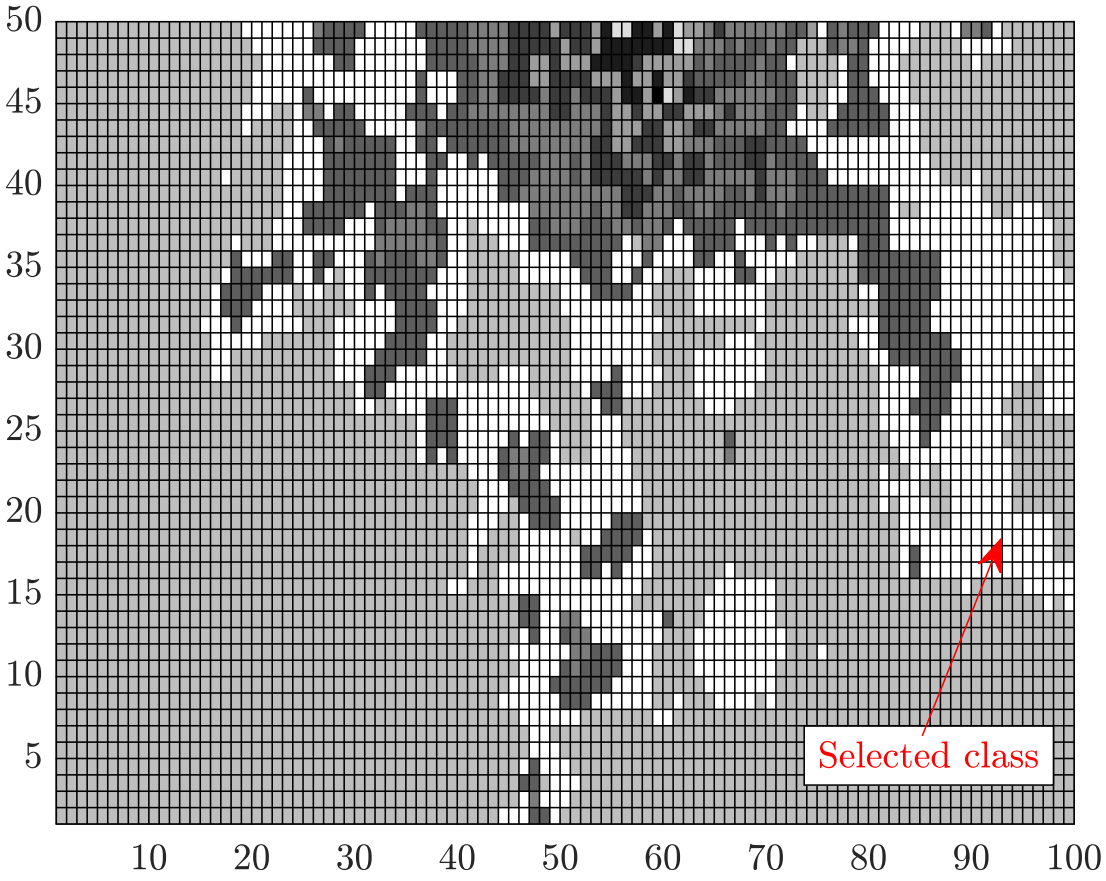


Figure 4.1: First-layer clustering results for Milan city grids, grouped by average traffic intensity, revealing spatial demand heterogeneity across the metropolitan area.

Second-Layer Clustering

Even among grids with similar average demand, further examination shows that meaningful variability in traffic still remains. A second clustering layer is therefore applied to the grids within a chosen first-layer cluster, in order to capture these more subtle differences. The method used mirrors that of the first layer, differing only in that it operates on this smaller subset of grids. This second layer reveals micro-level spatial variation in demand within regions that had appeared largely uniform, reinforcing the case that static spectrum allocation falls short [28, 29]. Fig. 4.2 presents the outcome of this refinement. Taken together, this hierarchical approach allows traffic patterns across the urban environment to be distinguished at a much finer level of granularity. The colour coding used in Figs. 4.1 and 4.2 is explained in Table 4.3.

Table 4.3: Colour coding of the clustering maps in Figs. 4.1 and 4.2.

Colour	**Meaning**
First layer (Fig. 4.1)	
White	Selected first-layer class, refined in the second layer
Light to dark grey	Other first-layer clusters. Darker shades indicate higher average traffic intensity
Second layer (Fig. 4.2)	
Grey shades	First-layer clusters not refined (as in Fig. 4.1)
Dark blue	Sub-cluster with the lowest average traffic
Blue	Low average traffic
Cyan	Low to moderate average traffic
Green	Moderate average traffic
Orange	High average traffic
Yellow	Sub-cluster with the highest average traffic

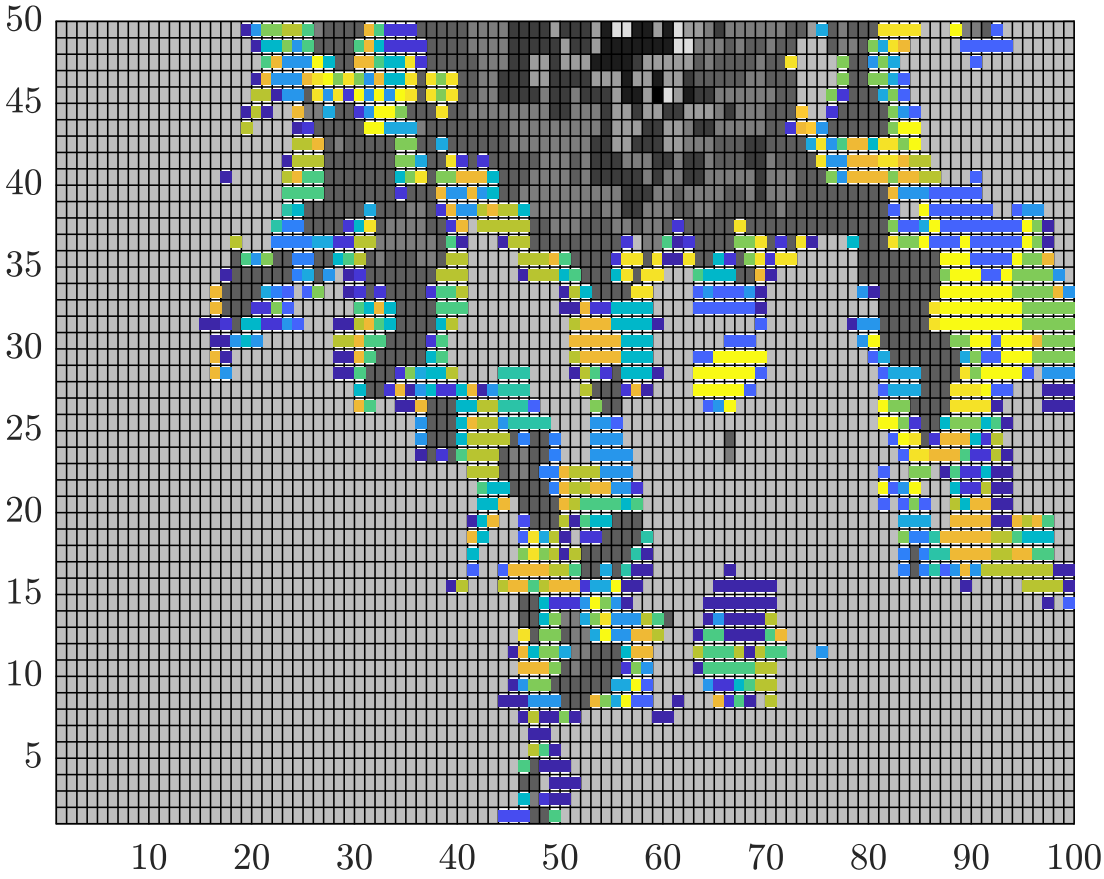


Figure 4.2: Second-layer clustering results for the selected grids, showing micro-level spatial variation in traffic demand within regions of similar average traffic.

4.2.2 Temporal Traffic Variations

Alongside spatial variation, individual grids also show substantial changes in spectrum usage over time. Traffic rises during peak hours and falls during off-peak

periods throughout the day. Because static spectrum allocation does not adapt to these time-based shifts, spectrum sits unused during quieter periods while becoming congested at peak times [23, 26]. Fig. 4.3 illustrates this temporal variation in traffic activity for two different days within a grid selected at random from the dataset.

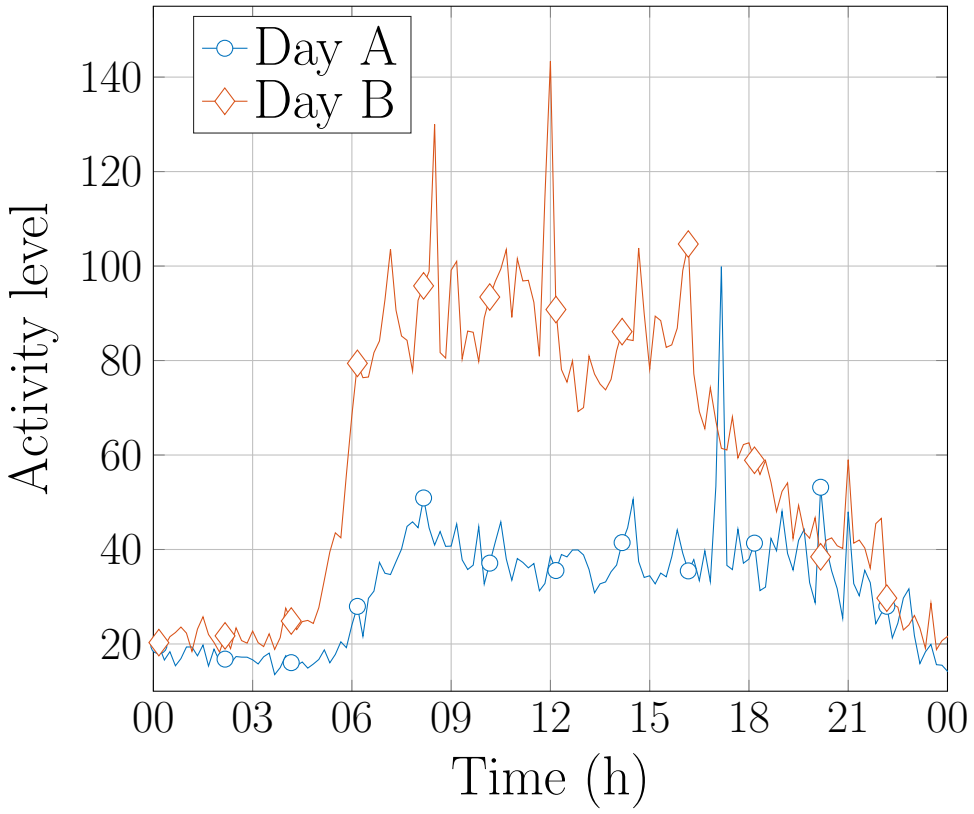


Figure 4.3: Daily fluctuations in traffic demand within a selected grid, showing how temporal variation affects spectrum usage over the course of the day.

This spatio-temporal analysis points to two main conclusions. First, traffic differs considerably from one grid to another, meaning spectrum allocation should be tailored to the specific demands of each region. Second, traffic changes throughout the day, meaning spectrum should be allocated dynamically in real time to prevent underutilisation during off-peak hours and congestion during peak hours [1, 30]. Together, these findings expose clear shortcomings in static spectrum management and build a strong case for a dynamic spectrum trading system. Such a system would allow spectrum holders to sell or lease their unused spectrum to buyers who need it, enabling allocation to adjust efficiently across both time and space. This approach would also support the economic sustainability of MNOs, since holders gain additional revenue while buyers obtain spectrum at a lower cost. Even when the price difference is modest, buyers benefit from being able to acquire spectrum for specific regions rather than paying for a licence

covering all regions at once.

4.2.3 Quantifying Temporal and Spatial Trading Opportunities

The preceding analysis shows qualitative evidence of spatial and temporal heterogeneity in urban traffic demand, but spectrum trading decisions call for a more precise, quantitative measure of how often demand falls below peak levels in a given region [24, 25]. Here, spectrum utilisation is not measured directly at the PHY/MAC layer but inferred from spatio-temporal traffic patterns, which serve as a proxy for identifying spectrum that is going underused. Building on this, an inferred utilisation index is derived from the Milan dataset to estimate opportunities for short-term spectrum trading.

Let $A_g(t)$ denote the aggregated internet traffic activity observed in grid g during time interval t. For each grid, $\min_g$ and $\max_g$ denote the minimum and maximum activity levels recorded over the observation period. From this, a demand-based utilisation index, used throughout this chapter as a proxy for spectrum utilisation, is defined as

$$U_g(t) = \frac{A_g(t) - \min_g}{\max_g - \min_g}, \tag{4.1}$$

where $U_g(t) \in [0, 1]$. When $U_g(t)$ is low, this signals a period where traffic demand sits well below the grid's peak, which in turn points to a larger margin for short-term spectrum trading. It should be stressed that $U_g(t)$ is a normalised proxy derived from traffic data alone and it does not represent physical-layer spectrum occupancy or the utilisation of physical resource blocks. The trading opportunities it identifies reflect demand-side capacity margin only, independent of any specific PHY/MAC configuration, which is instead handled through the policy and interference constraints carried over from Chapter 3.

To avoid drawing trivial conclusions from regions with sparse load, the analysis is restricted to grids falling within the top 10% of average traffic demand citywide. Fig. 4.4 shows the empirical cumulative distribution function (CDF) of $U_g(t)$ for one representative high-demand grid. The results indicate that utilisation stays at moderate or lower levels for a considerable share of the time, suggesting that trading opportunities recur even in the busiest parts of the city.

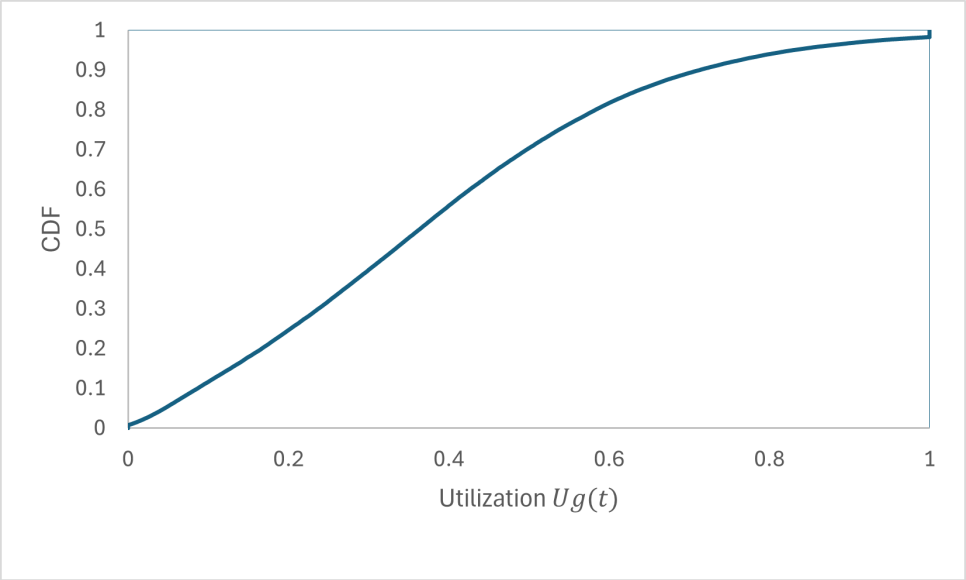


Figure 4.4: Empirical CDF of $U_g(t)$ for a representative high-demand grid, drawn from the top 10% of grids by average traffic demand. The curve traces the proportion of time the grid spends operating below its peak demand.

Fig. 4.5 further compares the temporal distributions of $U_g(t)$ across several representative high-demand grids. Clear differences in medians and spread across these grids point to substantial spatial heterogeneity even among the busiest regions, supporting the case for spectrum trading mechanisms that account for both location and time, rather than relying on uniform static allocation [38, 126].

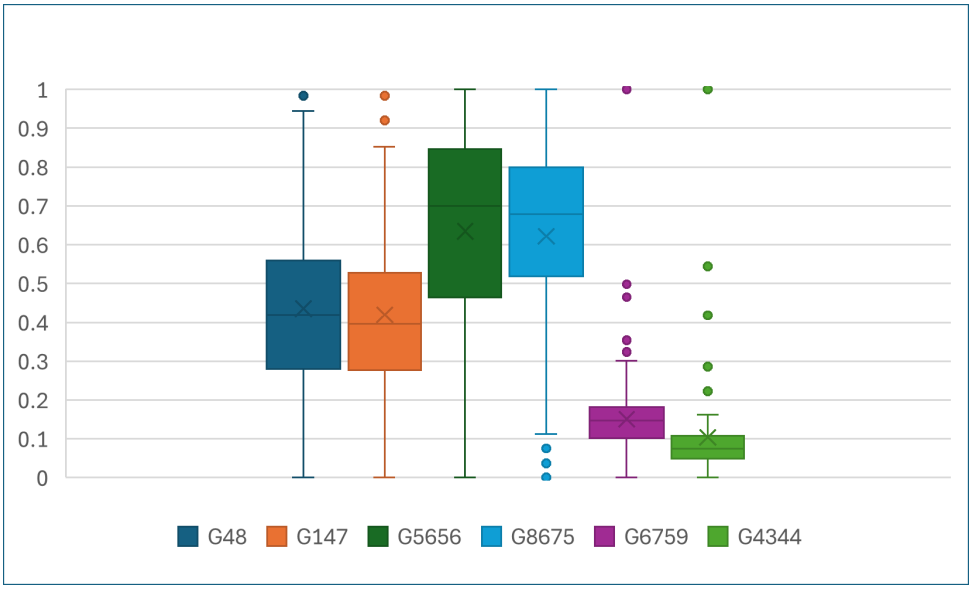


Figure 4.5: Spatial heterogeneity in $U_g(t)$ across representative high-demand grids. The boxplots show that utilisation distributions differ considerably by location, pointing to localised spectrum trading opportunities even within high-demand regions.

4.3 ANN-Based Traffic Prediction for Proactive Trading

The empirical analysis of Section 4.2 establishes that recurring, localised spectrum trading opportunities exist. If cellular traffic loads can be predicted, blockchain-based spectrum trading can be coordinated proactively and more efficiently. Machine learning has been widely applied to wireless traffic and spectrum prediction [85,90,143,144], and recent surveys confirm the growing relevance of learning-assisted approaches for anticipating demand fluctuations and supporting proactive resource decisions in 5G and beyond networks [12–14]. To demonstrate this, an ANN-based algorithm is applied to a selected cell, showing how anticipating periods of low demand can inform proactive trading decisions. It is emphasised that the ANN is used here to characterise and forecast demand rather than to serve as the primary contribution. Its role is to supply predictive intelligence that is coupled with the execution and governance layer provided by the marketplace.

4.3.1 Role and Structure of the ANN

The ANN is used for a single, narrowly defined task: given the position of a 10-minute interval within the day and within the week, it predicts the traffic activity of one grid square for that interval. It is therefore a regression model that produces a continuous output, not a classifier. Its role in the framework is illustrative. It shows that, because urban traffic is strongly periodic, the low-demand windows identified in Section 4.2.3 can be anticipated in advance, so that a spectrum holder could list capacity before an idle period begins. The forecast is not fed automatically into the smart contract or into a closed-loop trading controller in this thesis as such integration is identified as future work.

The workflow of the ANN is shown in Fig. 4.6. The input layer receives two groups of features: a 144-element time-of-day encoding, one element per 10-minute slot of the day, and a 7-element day-of-week encoding. The full model therefore has 151 inputs, while the ablated model of Section 4.3.3 uses only the 144 time-of-day inputs. The inputs are connected to a single hidden layer of 20 neurons with hyperbolic-tangent sigmoid activation (the MATLAB `fitnet` default), and the hidden layer is connected to one output neuron with linear activation, which produces the normalised activity forecast $\hat{y}_{\text{norm}}$. The forecast is mapped back to activity units by multiplying by $y_{\max}$, and the error is computed against the observed activity. The network is trained with Bayesian regularisation (MATLAB's `trainbr`), which minimises the squared prediction error together

with a penalty on the network weights.

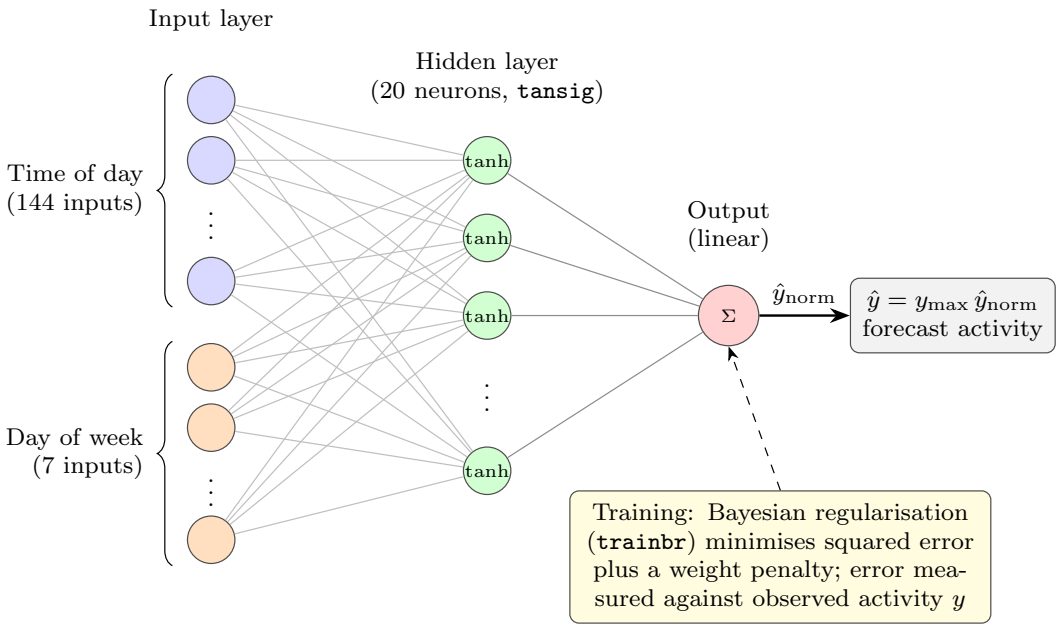


Figure 4.6: Workflow of the ANN forecaster: time-of-day and day-of-week encodings feed a single hidden layer of 20 `tansig` neurons and a linear output neuron, whose normalised output is rescaled to activity units. The ablated model omits the day-of-week inputs.

4.3.2 Model Configuration

The 21-day observation window of the selected grid is divided chronologically into a 14-day training window and a 7-day test week that follows it and is held out completely. Only the training window is used to fit the model. Within it, the 2,016 samples are divided at random into training (50%), validation (25%), and internal test (25%) subsets. This ratio differs from the 80/20 or 70/30 splits common in classification, for three reasons. First, the decisive evaluation in this chapter is the chronologically separate test week, which provides 1,008 unseen samples. The random split inside the training window serves only to monitor generalisation during training, and larger validation and internal test shares make that monitoring more reliable. Second, because traffic is strongly periodic, a 50% share of the training window (about 1,008 samples) still represents each of the 144 time-of-day slots about seven times and each day of the week about 144 times, so the model sees the full daily and weekly cycle during fitting. Third, Bayesian regularisation, described below, controls overfitting through a penalty on the network weights rather than through the size of the training set, which makes a smaller training share acceptable. Performance on the test week is assessed using

the mean squared error (MSE) between predicted activity levels and the observed values, together with the regression metrics reported in Section 4.3.3.

To keep gradient descent stable and stop features with larger scales from dominating training, all telecommunications traffic data was scaled using min-max normalisation. Raw activity levels y were rescaled to the range $[0, 1]$ using the global maximum value $y_{\max}$ identified within the 14-day training window, so that $y_{\text{norm}} = y/y_{\max}$. This same normalisation was applied to both the training data and the independent test set before being passed into the ANN predictive engine.

Training is carried out using Bayesian regularisation (MATLAB's `trainbr` routine), which generalises better on city-scale telecommunications data than standard backpropagation. The architecture used is a shallow feedforward network (MATLAB's `Fitnet` configuration). The baseline model uses 20 hidden units for the primary prediction task and 10 units for the ablation benchmark.

4.3.3 Full and Ablated Models

Two models are implemented to isolate the contribution of weekly seasonality. The full model is built to capture both daily and weekly periodicities in telecommunications traffic, drawing on two groups of input features: a time-of-day index, given as a normalised 1×144 vector spanning 10-minute intervals across 24 hours, and a day-of-week index, given as a normalised 1×7 vector spanning the weekly cycle. The ablated model acts as a control, with the day-of-week feature removed to represent a scenario lacking weekly seasonality.

Prediction performance on the independent test week is shown in Fig. 4.7. Comparing the full and ablated models confirms that including weekly seasonality meaningfully reduces prediction error, supporting the use of the full feature set for proactive trading coordination.

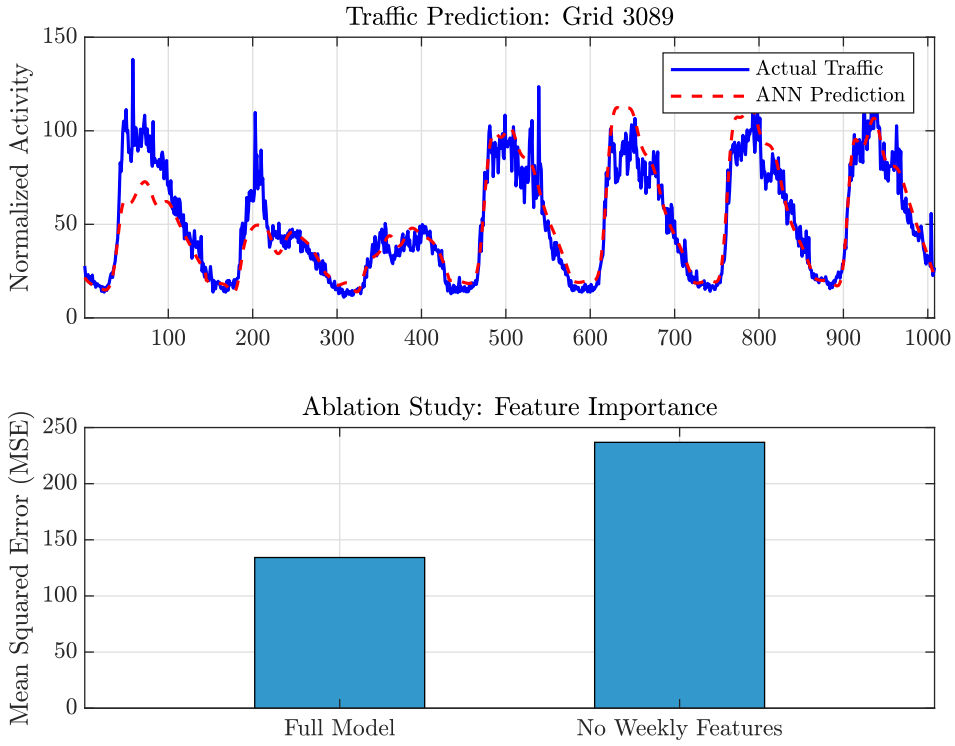


Figure 4.7: ANN-based traffic load prediction for a representative Milan grid (Grid 3089), trained on 14 days of data and evaluated over an independent 7-day test week. Upper panel: observed and forecast activity. The horizontal axis is the index of the 10-minute slot in the test week (1–1,008), and the vertical axis is activity in the dimensionless units of the Milan dataset. Lower panel: mean squared error over the test week, in squared activity units, for the full model (20 neurons) and the ablated model without day-of-week features (10 neurons).

Because the task is regression, classification metrics such as accuracy, precision, recall, F1-score, and the confusion matrix are not defined for it. They require discrete class labels, whereas the ANN predicts a continuous activity value. Prediction accuracy is instead quantified with standard regression metrics, computed over the $N = 1{,}008$ ten-minute slots of the independent test week:

$$\begin{aligned} \mathrm{MSE} &= \frac{1}{N}\sum_{n=1}^{N}(\hat{y}_n - y_n)^2, \qquad \mathrm{RMSE} = \sqrt{\mathrm{MSE}}, \\ \mathrm{MAE} &= \frac{1}{N}\sum_{n=1}^{N}|\hat{y}_n - y_n|, \qquad \mathrm{MAPE} = \frac{100}{N}\sum_{n=1}^{N}\left|\frac{\hat{y}_n - y_n}{y_n}\right|, \\ R^2 &= 1 - \frac{\sum_n(\hat{y}_n - y_n)^2}{\sum_n(y_n - \bar{y})^2}, \end{aligned} \tag{4.2}$$

where y_n and $\hat{y}_n$ are the observed and forecast activity in slot n and $\bar{y}$ is the mean

observed activity. The results are listed in Table 4.4.

Table 4.4: Regression accuracy of the ANN forecaster on the independent 7-day test week (Grid 3089, N = 1,008 ten-minute slots). Activity is in the dimensionless units of the Milan dataset. Only the MSE was recorded for the ablated model; its RMSE follows from the MSE.

Metric	**Full model** (all features, 20 neurons)	**Ablated model** (no day of week, 10 neurons)
MSE (activity^2)	134.2	236.8
RMSE (activity)	11.59	15.39
MAE (activity)	7.72	–
MAPE (%)	15.3	–
R^2	0.84	–

For reference, the observed activity in the test week ranges from 11.0 to 138.2 with a mean of 50.2, so the RMSE of the full model corresponds to about 9% of the observed range, and the model explains 84% of the variance of the test-week traffic. The largest errors occur at the sharp daily peaks, notably on the first test day, while the timing of the daily rise and fall, which determines the low-demand windows relevant to trading, is captured closely. Removing the day-of-week features increases the MSE by about 76%. This comparison should be interpreted with care, because the two models differ in two respects: the ablated model lacks the day-of-week inputs and also has a smaller hidden layer (10 rather than 20 neurons). Part of the difference in error may therefore be due to the reduced capacity of the network rather than to the missing weekly features alone. A controlled ablation with equal hidden-layer sizes is left for future work, and the result should be read as consistent with, rather than a precise measure of, the contribution of weekly seasonality.

4.3.4 From Reactive to Proactive Trading

Integrating prediction with the blockchain-based marketplace shifts spectrum trading from a reactive to a proactive paradigm, offering several advantages [83, 88]. First, forecasting spectrum availability allows trading to begin ahead of actual demand, cutting down the time required to gain spectrum access. Second, secondary MNOs gain visibility into how long the spectrum is expected to remain available, allowing them to plan their operations around the predicted lease windows. Third, if these predictions were linked to the smart contracts from Chapter 3, listings could be created automatically as soon as a forecast idle period approaches [127, 130]. In this thesis, the prediction component is illustra-

tive. It is demonstrated separately and is not connected to the smart contract in a closed loop, which is left for future work.

4.4 Extended Marketplace for Next-Generation Networks

The marketplace architecture, smart contract functions, and Python automation scripts used in this chapter are those developed in Chapter 3, and are not repeated here in full. The system still consists of the Telecommunications Regulatory Authority acting as administrator, along with spectrum holders and spectrum buyers, all interacting through the Ethereum smart contracts described in Section 3.7. This section describes how that marketplace is positioned for next-generation networks and the extensions introduced to support dynamic, demand-aware operation.

4.4.1 Deployment and Consensus Considerations

Rather than relying on proof-of-work, as many traditional public blockchains do, the proposed system is deployed on the Ethereum public blockchain via the Sepolia test network, operating under a proof-of-stake consensus mechanism [127]. Ethereum itself remains permissionless at the consensus layer, but access control and participant authorisation are instead enforced at the smart-contract and application layers, as set out in Chapter 3. Only authorised operator addresses, namely spectrum holders and spectrum buyers, are able to list, purchase, and manage spectrum assets through the deployed contracts, which ensures regulatory compliance without any need to alter the underlying blockchain protocol [2, 128]. Compared with proof-of-work, this proof-of-stake mechanism offers better energy efficiency and lower operational overhead, while still preserving the security and integrity guarantees of the public Ethereum network [8]. Transaction validation and block production remain the responsibility of Ethereum's proof-of-stake validators, carried out according to the network's own consensus rules. The proposed system neither modifies nor controls validator selection or consensus participation.

4.4.2 Marketplace Objectives in the 5G and Beyond Context

When situated within next-generation networks, the marketplace must meet a set of operational objectives that build on the functional requirements already introduced in Chapter 3. Operator authentication ensures that MNOs, whether acting as spectrum holders or buyers, are verified in a way that prevents impersonation while still preserving anonymity within the network. Privacy requires that spectrum requests, announcements, and transactions do not reveal personally identifiable information such as names, addresses, or financial details, except in cases where regulatory authorities explicitly require this. Traceability allows the regulatory authority to identify MNOs when disputes or unexpected events arise, such as fraudulent transactions or non-compliance, which ensures accountability while still balancing privacy [7, 74]. Reliability means that each operator can maintain a full copy of the blockchain, which preserves the integrity and transparency of all transactions and keeps the system operational even during network outages or cyberattacks. Finally, data confidentiality and integrity require that trading messages are transmitted securely, without alteration, and that their contents remain hidden from any unauthorised parties.

4.4.3 Interference-Aware and Policy-Driven Constraints

The interference and coexistence handling introduced in Chapter 3 is retained and applied at the market and allocation layer. Spatial, temporal, and frequency-domain conflict detection, the minimum guard-band constraint, and the interference-aware feasibility conditions all continue to govern which candidate trades may be instantiated on-chain [91, 92]. In the dynamic setting of this chapter, these constraints take on additional importance, since the demand-driven and predictive triggering of trades increases the frequency with which new allocations are evaluated against the set of active allocations. As in Chapter 3, these are policy-based checks rather than physical interference guarantees, so that demand-responsive trading does not grant usage rights that conflict under the regulator's separation rules.

4.5 Integration with an OAI-Based 5G Testbed

The proposed framework is aimed at next-generation wireless networks, where spectrum coordination that is dynamic, localised, and auditable will only grow in importance [27, 53]. To confirm its practical feasibility, the framework is in-

tegrated with an OpenAirInterface (OAI) based 5G testbed [142, 145], offering a mature, reproducible platform on which automated network reconfiguration following a successful spectrum transaction can be demonstrated. Although the validation platform itself is 5G-based, the mechanism proposed here operates at the spectrum access and coordination layer, meaning it is equally applicable to 5G-Advanced, 6G, and future networks.

The blockchain-based marketplace manages dynamic spectrum authorisation and settlement, while the automated configuration process updates the relevant radio parameters on the testbed in response. This illustrates how on-chain spectrum transactions translate into concrete network-side actions, such as updating the operating frequency and bandwidth of the gNB and the user equipment (UE).

4.5.1 OAI Architecture and Integration Point

Fig. 4.8 shows the OAI 5G architecture used in this work and where the proposed framework is integrated. The testbed consists of three main parts: the user equipment (UE), the gNB, and the 5G core network, which includes the Access and Mobility Management Function (AMF), the Session Management Function (SMF), and the User Plane Function (UPF). The gNB connects to the core network through the standard N2 and N3 interfaces. Because OAI is run in RF-simulator mode, the radio link between the UE and the gNB is emulated in software, so no radio hardware is required.

The proposed framework does not modify any OAI component. Instead, it is integrated at the configuration level. When a spectrum trade is completed on the blockchain, the purchased centre frequency and bandwidth are read from the smart contract and used to update the configuration of the gNB and the UE. Both are then restarted with the new spectrum parameters, while the core network continues to run unchanged. In this way, an on-chain trade is translated directly into a network-side action.

4.5.2 Dynamic Configuration of Testbed Components

Once a spectrum transaction is completed, the testbed components are reconfigured dynamically to use the newly acquired spectrum parameters. This involves updating the gNB and UE with the allocated centre frequency and bandwidth taken from the blockchain transaction record. The aim is to show that blockchain-based spectrum trading can drive automated network reconfiguration without altering the underlying PHY/MAC mechanisms.

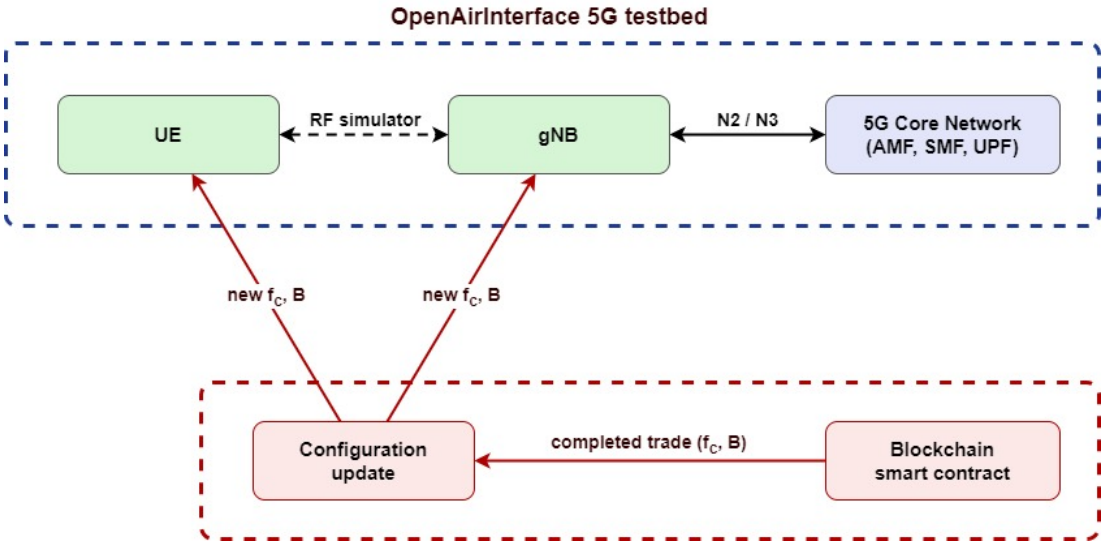


Figure 4.8: OAI 5G testbed architecture and the integration point of the proposed framework (red). A completed on-chain trade updates the carrier frequency and bandwidth of the gNB and UE without modifying OAI.

Network Operation With Newly Allocated Spectrum. Operating the OAI-based 5G testbed with the updated spectrum parameters confirms the end-to-end integration between the blockchain marketplace and the network configuration workflow. Fig. 4.9 shows the testbed running with the updated frequency, confirming that the newly acquired spectrum parameters have been applied successfully. Fig. 4.10, showing the network status, further confirms that the relevant components remain operational following the update.

```
 ⠿ Container rfsim5g-mysql       Running                                    0.0s
 ⠿ Container rfsim5g-oai-nrf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-amf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-smf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-spgwu   Running                                    0.0s
 ⠿ Container rfsim5g-oai-ext-dn  Running                                    0.0s
 ⠿ Container my-oai-gnb          Started                                    2.8s
freq_read = 3360
tx_rx_freq = 3360 MHz
absoluteFrequencySSB = 624000.0
dl_absoluteFrequencyPointA = 622728.0
sha256:9a19a4ad7feea8a2dd439fa6cc5112082de110861b9bdea9df73b48b3cf3f7ef
my-oai-gnb
my-oai-gnb
[+] Running 7/7
 ⠿ Container rfsim5g-oai-nrf     Running                                    0.0s
 ⠿ Container rfsim5g-mysql       Started                                    0.3s
 ⠿ Container rfsim5g-oai-amf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-smf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-spgwu   Running                                    0.0s
 ⠿ Container rfsim5g-oai-ext-dn  Running                                    0.0s
 ⠿ Container my-oai-gnb          Started                                    3.1s
          USE_ADDITIONAL_OPTIONS: -E --sa --rfsim -r 106 --numerology 1 -C 3360000000 --nokrnmod
[+] Running 8/8
 ⠿ Container rfsim5g-oai-nrf     Running                                    0.0s
 ⠿ Container rfsim5g-mysql       Running                                    0.0s
 ⠿ Container rfsim5g-oai-amf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-smf     Running                                    0.0s
 ⠿ Container rfsim5g-oai-spgwu   Running                                    0.0s
 ⠿ Container rfsim5g-oai-ext-dn  Running                                    0.0s
 ⠿ Container my-oai-gnb          Running                                    0.0s
 ⠿ Container rfsim5g-oai-nr-ue   Started                                    3.3s
freq_read = 3360
freq_read = 3360
freq_read = 3360
```

Figure 4.9: Network running with a newly allocated frequency following a completed on-chain spectrum trade.

```
CONTAINER ID   IMAGE                      COMMAND                  CREATED          STATUS
9e804b12ca21   oai-nr-ue:develop          "/opt/oai-nr-ue/bin/…"   28 seconds ago   Up 24 seconds (healthy)
39dbe7e20494   my-gnb:v1                  "/opt/oai-gnb/bin/en…"   34 seconds ago   Up 30 seconds (healthy)
ec16ea24861d   ubuntu:bionic              "/bin/bash -c ' apt …"   41 seconds ago   Up 37 seconds (healthy)
0bb932c2d0eb   oai-spgwu-tiny:latest      "/openair-spgwu-tiny…"   41 seconds ago   Up 37 seconds (healthy)
78f423753df1   oai-smf:latest             "/bin/bash -c '/open…"   41 seconds ago   Up 38 seconds (healthy)
a2108755f3c1   oai-amf:latest             "/bin/bash /openair-…"   41 seconds ago   Up 39 seconds (healthy)
e34a0355eb9c   mysql:5.7                  "docker-entrypoint.s…"   41 seconds ago   Up 33 seconds (health: starting)
6837afccacf0   oai-nrf:latest             "/bin/bash /openair-…"   41 seconds ago   Up 40 seconds (healthy)
```

Figure 4.10: Network status following spectrum integration, showing the updated parameters and confirming healthy operation across all components.

gNB Configuration. The integration process includes real-time updates to the gNB configuration, so that the network runs using the new spectrum. Fig. 4.11 presents the gNB logs, which show that the configuration was updated successfully to reflect the new spectrum allocation. These logs confirm the gNB has been updated correctly, pointing to seamless network operation.

```
DL frequency 3360000000: band 78, UL frequency 3360000000
[PHY]   DL frequency 3360000000 Hz, UL frequency 3360000000 Hz: band 78, uldl offset 0 Hz
[PHY]   Configuring MIB for instance 0, : (Nid_cell 0,DL freq 3360000000, UL freq 3360000000)
[PHY]   Initializing frame parms for mu 1, N_RB 106, Ncp 0
[PHY]   Init: N_RB_DL 106, first_carrier_offset 900, nb_prefix_samples 108,nb_prefix_samples0 132
[PHY]   Doing symbol rotation calculation for gNB TX/RX, f0 3360000000.000000 Hz, Nsymb 28
[PHY]   Symbol rotation 0/28 => (32767,0)
[PHY]   Symbol rotation 1/28 => tl 0.000036 (32767,-1) (0.000000)
[PHY]   Symbol rotation 2/28 => tl 0.000072 (32767,-1) (0.000000)
[NR_PHY]   Number of bad PUCCH received: 0
[NR_MAC]   Frame.Slot 384.0
UE ID 0 RNTI 0982 (1/1) PH 0 dB PCMAX 0 dBm, average RSRP 0 (0 meas)
UE 0: dlsch_rounds 18/0/0/0, dlsch_errors 0, pucch0_DTX 0, BLER 0.02542 MCS 9
UE 0: dlsch_total_bytes 2214
UE 0: ulsch_rounds 144/0/0/0, ulsch_DTX 0, ulsch_errors 0
UE 0: ulsch_total_bytes_scheduled 16704, ulsch_total_bytes_received 16588
UE 0: LCID 1: 74 bytes TX
```

Figure 4.11: gNB configuration logs confirming the application of the newly traded spectrum parameters.

End-User Equipment Configuration. For end-to-end network functionality, the UE must also be updated to operate on the new spectrum. Fig. 4.12 presents the UE logs, confirming that the UE has been configured to use the new frequency and completing the integration process. This brings the entire network, from core components through to end-user devices, into alignment with the new spectrum allocation.

```
== Starting NR UE soft modem
Additional option(s): -E --sa --rfsim -r 106 --numerology 1 -C 3360000000 --nokrnmod
/opt/oai-nr-ue/bin/nr-uesoftmodem.Rel15 -O /opt/oai-nr-ue/etc/nr-ue-sim.conf -E --sa --rfsim -r 106 --numerology 1 -C 3360000000 --nokrnmod
[CONFIG] get parameters from libconfig /opt/oai-nr-ue/etc/nr-ue-sim.conf , debug flags: 0x00000000
[LIBCONFIG] Path for include directive set to: /opt/oai-nr-ue/etc
[CONFIG] function config_libconfig_init returned 0
[CONFIG] config module libconfig loaded
[LIBCONFIG] config: 1/1 parameters successfully set, (1 to default value)
Send signal 35 to display resource usage...
[LIBCONFIG] log_config: 3/3 parameters successfully set, (2 to default value)
[LIBCONFIG] log_config: 53/53 parameters successfully set, (53 to default value)
[LIBCONFIG] log_config: 53/53 parameters successfully set, (53 to default value)
[LIBCONFIG] log_config: 16/16 parameters successfully set, (16 to default value)
[LIBCONFIG] log_config: 16/16 parameters successfully set, (16 to default value)
[PHY]   [UE 0] Measured Carrier Frequency 3360000000 Hz (offset 0 Hz)
[PHY]   [AbsSFN 0.0] Slot0 Symbol 2 type 0: Pilot/Data extraction  0.86
[PHY]   [AbsSFN 0.0] Slot0 Symbol 2: Channel Scale   0.14
[PHY]   [AbsSFN 0.0] Slot0 Symbol 2 first_symbol_flag 0: Channel Level   0.02
[PHY]   [AbsSFN 0.0] Slot0 Symbol 2 log2_maxh 0 channel_level 0: Channel Comp   0.76
[PHY]   [AbsSFN 0.0] Slot0 Symbol 2: Channel Combine   0.14
[PHY]   [AbsSFN 0.0] Slot0 Symbol 2: LLR Computation   0.21
```

Figure 4.12: End-user equipment configuration logs confirming operation on the newly allocated frequency.

4.6 Discussion

Taken together, the results show that blockchain technology can be integrated effectively into spectrum trading workflows to tackle long-standing inefficiencies in spectrum access and coordination [2, 128]. By enabling exclusive spectrum-

use rights to be transferred through executable smart contracts, the proposed framework supports dynamic spectrum trading that responds to demand as it varies across both space and time. This is especially relevant in dense, heterogeneous environments, where static allocation policies often result in underutilised spectrum [24, 25], a point confirmed by the empirical analysis in Section 4.2.

A key outcome of this implementation is the ability to enforce spectrum trading transactions automatically, transparently, and in a way that can be verified. Because smart contract execution ensures that ownership transfer, authorisation, and settlement happen atomically, the risk of double allocation or inconsistent state is eliminated [130]. In contrast to traditional spectrum management processes, which depend on manual coordination and intermediaries [38, 126], this approach reduces the need for human intervention and cuts down on operational errors, improving reliability at the spectrum access layer.

The integration with a 5G testbed underscores the practical feasibility of the framework further still. That network components can be reconfigured automatically following a successful spectrum trade shows that blockchain-based coordination can drive the configuration of a standards-based 5G stack [142]. Because the experiment used the OAI RF simulator on a single machine, it validates the path from an on-chain trade to the configuration of the gNB and UE, but it does not measure over-the-air radio performance or behaviour under live traffic. Crucially, none of this requires modifying PHY- or MAC-layer mechanisms; instead, it makes spectrum authorisation and reconfiguration faster and more reliable, letting existing radio access network implementations operate within newly acquired spectrum without disruption.

The scalability analysis from Chapter 3 carries over directly to this chapter's framework, and its conclusions matter even more in this dynamic setting. System performance is limited mainly by blockchain-level parameters, such as block gas limits and block intervals, rather than by the number of spectrum buyers or holders participating [8]. This carries an important design implication: adding more market participants need not degrade performance, provided transaction arrival rates remain within the blockchain's service capacity. As demand approaches this capacity, congestion can be relieved by partitioning trading contracts by region and by favouring long-lived spectrum trading agreements, both of which reduce transaction frequency without compromising correctness or exclusivity.

These findings also point to the value of data-driven demand analysis in guiding spectrum trading decisions [82, 118]. By drawing on spatio-temporal traffic patterns, the framework can identify when and where exclusive spectrum-use rights are likely to go underused, supporting more informed and efficient trading

decisions. This demand analysis is intended as motivation rather than a closed-loop control mechanism, but it nonetheless demonstrates how empirical traffic data can complement blockchain-based coordination to improve spectrum utilisation. Compared with existing blockchain-based spectrum sharing approaches that focus primarily on architecture or long-term sharing logic [74, 76, 78], the present work uniquely combines real demand evidence, predictive intelligence, and live testbed validation. The principal limitations of the present work are that physical-layer interference is not directly measured but instead handled through policy-level constraints, that the testbed validation is conducted on a single platform, and that the ANN prediction is illustrative rather than embedded in a closed-loop control system. These limitations define natural directions for future work.

4.7 Reproducibility and Experimental Setup

This section documents the artefacts, parameters, and experimental environment required to reproduce the framework and its integration with the 5G testbed.

4.7.1 Released Artefacts and Configurable Parameters

The blockchain components consist of Ethereum smart contracts, their associated application binary interfaces (ABIs), and client-side interaction scripts. Released as artefacts are the smart contracts that implement the spectrum listing, purchase, and settlement logic, the compiled ABIs that off-chain applications use to interact with these contracts, and the Web3-based Python scripts used to submit transactions, monitor events, and collect on-chain metrics. The parameters of the proposed market mechanism, namely the pricing parameters ($p_0(F)$, η, and the duration factors) and the market parameters (Δ, $G_{\min}$, and R_{int}), are documented in Chapter 3, together with blockchain parameters such as the confirmation depth k and the network endpoint.

4.7.2 Blockchain Environment

The layered architecture of the blockchain used in this work is shown in Fig. 3.2 in Chapter 3. All experiments were run on the Ethereum Sepolia test network. Smart contracts were compiled using Solidity version 0.8.26 and deployed through standard Ethereum tooling. Transaction receipts and logs were collected to measure gas consumption, confirmation delay, and success ratios, while static analysis

was carried out using the Slither tool to assess smart contract robustness, following the same approach reported in Chapter 3.

4.7.3 5G Testbed Hardware and Software Stack

The 5G environment was deployed on a single commodity machine: a Dell XPS 13 7390 laptop fitted with an Intel Core i7-10710U CPU at 1.10 GHz and 16 GB of RAM, running Ubuntu Kylin 18.04. The 5G core network, gNB, and UE were all implemented using OpenAirInterface in RF simulator mode [145]. Specifically, the OAI 5G RF simulator used was built from the develop branch, commit `3ea34db9a7` (December 2021). The resulting network topology is shown in Fig. 4.13. All components of the testbed run on a single host machine. The UE, the gNB, and the 5G core network are deployed as Docker containers, and the radio link between the UE and the gNB is emulated by the OAI RF simulator. The gNB connects to the core network through the N2 and N3 interfaces. The same host runs the Python/Web3 scripts, which connect to the Ethereum Sepolia test network over the Internet through an Infura RPC endpoint and apply the traded carrier frequency and bandwidth to the gNB and UE.

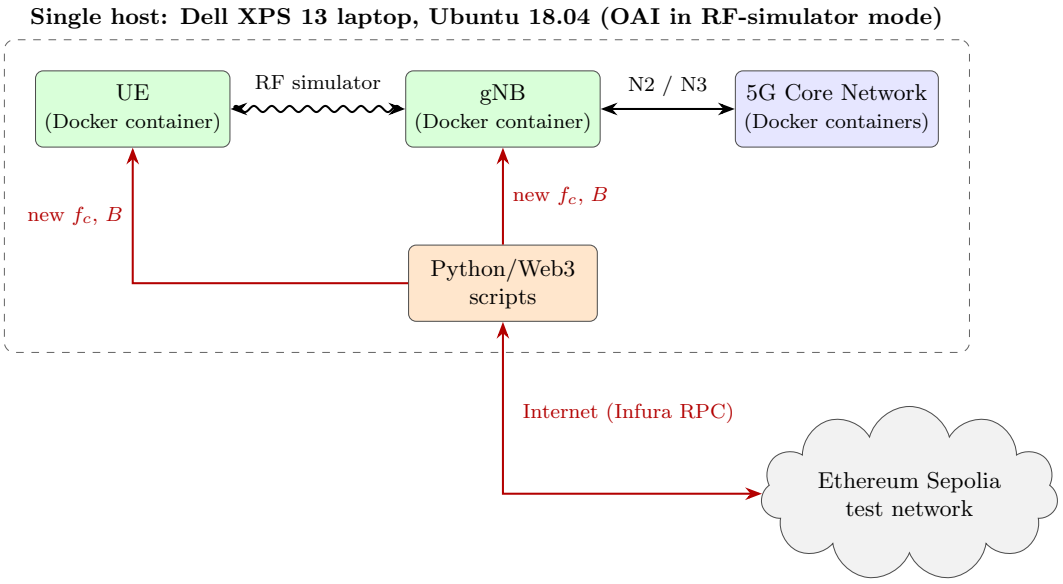


Figure 4.13: Network topology of the testbed: the UE, gNB, and 5G core network run as containers on a single host, and the Python/Web3 scripts link the host to the Ethereum Sepolia network.

4.8 Summary

This chapter extended the blockchain-based spectrum trading marketplace of Chapter 3 into a complete, demand-aware, and network-integrated framework for next-generation wireless networks. Three connected advances were presented. First, an empirical study of the Milan city dataset [141] provided data-driven, demand-side evidence of spatio-temporal variation consistent with spectrum underutilisation. A normalised demand-based utilisation index was derived to quantify how frequently demand falls below peak levels, and a two-stage k-means clustering procedure exposed spatial heterogeneity, both across and within regions, creating localised and short-term trading opportunities. These results supplied the empirical motivation that was absent from Chapter 3.

Second, an ANN-based traffic prediction model was introduced to enable the shift from reactive to proactive spectrum trading [83, 90]. By using time-of-day and day-of-week features to forecast periods of low demand, the chapter also compared a full model against an ablated variant. This showed how predictive intelligence can be coupled with the executable marketplace, so trading decisions can anticipate availability instead of only reacting to it. Third, the framework was integrated with an OpenAirInterface-based 5G testbed [142], demonstrating that a completed on-chain spectrum trade can trigger automated reconfiguration of the gNB and UE without modifying the underlying PHY/MAC mechanisms. This provided the end-to-end validation that was identified as a key limitation of the standalone marketplace.Run in RF-simulator mode, this validated the trade-to-configuration workflow that was missing from the standalone marketplace, rather than over-the-air radio performance.

Together with Chapter 3, this chapter shows that blockchain-enabled spectrum trading is technically feasible, driven by real demand heterogeneity, and compatible with operational 5G deployments. The next chapter takes this coordination framework further, into shared indoor and neutral-host environments, where blockchain enforces performance-aware settlement based on delivered QoS.

Chapter 5

Blockchain-Based Settlement for Indoor Neutral-Host Offloading

5.1 Introduction

Most cellular networks today are deployed by public Mobile Network Operators (MNOs) over high-power, exclusively licensed spectrum and primarily serve mobile consumers through smartphone connections. However, recent reports suggest that the growth rate of mobile data usage is slowing, while also indicating that a dominant fraction of traffic (often reported at around 80%) originates indoors [146]. This imbalance is a good reason to rethink the traditional macro-centric, outdoor-first design philosophy, especially at mid-band frequencies, where outdoor-to-indoor (O2I) propagation can take a serious hit due to building penetration loss and shadowing.

The International Telecommunications Society (ITS) has highlighted mobile offloading as an important mechanism for alleviating macro-network congestion and generating additional revenue through shared access infrastructure [147]. Their analysis points to several dimensions that shape how viable offloading is in practice: the physical setting of the deployment, which determines propagation conditions and user density, the ownership model, which decides whether the offloading infrastructure is operator-owned, venue-owned, or run by a neutral third party, the service provisioning approach, which sets how subscribers are authenticated and authorised, the revenue model, which determines how costs and revenues are distributed among stakeholders, and the offloading process management, which covers session handover, traffic steering, and quality assurance. Studies have shown that indoor user experience can be significantly improved by offloading cellular subscribers from outdoor macro O2I service to a local, higher-throughput access network operated by a neutral-host operator (NHO) [97, 98].

In many venues, macro-to-indoor coverage is impaired by penetration loss and log-normal shadowing, which leads to a pronounced low-throughput tail for indoor users even when outdoor macro coverage is strong [148,149]. A neutral-host network can improve downlink SINR and deliver more predictable throughput in challenging indoor conditions [97,98].

A first motivation for offloading is macro load relief and capacity improvement. When an NHO serves indoor traffic locally, demand is partially removed from the macro layer, which reduces congestion pressure during peak periods while improving indoor data rates [95,97]. A second motivation is service continuity in coverage-challenged indoor locations. Even if an MNO provides broad outdoor coverage, indoor environments can experience severe O2I attenuation (e.g., due to low-E glass and energy-efficient materials) and shadowing variability, leading to unstable throughput. Offloading such users to an indoor NHO cell mitigates these impairments by avoiding the O2I path and improving signal conditions [98,148]. A third motivation is short-duration venue usage, where users enter a location for a limited time (e.g., shopping centres, airports, campuses) and expect immediate high-speed connectivity without requiring manual actions or duplicative indoor deployments by each MNO [99,146].

In a neutral-host deployment, an MNO subscriber entering an indoor venue should be able to offload seamlessly onto the neutral-host infrastructure and return to macro service upon leaving, without manual intervention. However, this requires automated procedures for subscriber authentication across operator boundaries, session-level authorisation that governs which subscribers may use the neutral-host radio resources, and a settlement mechanism that accounts for the volume and quality of service consumed during the offloading period. In practice, such inter-operator procedures remain largely ad hoc or absent, which limits the adoption of neutral-host models despite their radio-layer benefits.

While neutral-host deployments can improve indoor radio performance, the economic relationship between NHOs and MNOs is typically governed by fixed monthly fees, flat-rate infrastructure access charges, or coarse traffic-volume accounting that does not differentiate between high-throughput and degraded service [80]. While such arrangements may account for the volume of data consumed, they do not reflect the radio-layer quality actually experienced by offloaded users. In practice, downlink throughput can vary significantly with macro load, cross-tier interference leakage, and NHO capacity, meaning that the same neutral-host infrastructure may deliver strong performance under favourable conditions and degraded performance under congestion or high leakage. Without settlement mechanisms that account for these variations, accountability is weakened and

incentives are misaligned: an NHO that delivers consistently high throughput is compensated identically to one whose indoor coverage is impaired, reducing the incentive to maintain service quality. Moreover, the absence of standardised methods to immutably record and audit service-level outcomes makes it difficult to resolve disputes when stakeholders disagree over the performance actually delivered [124].

Motivated by the above, this chapter introduces an end-to-end Quality of Service (QoS)-aware automated settlement framework for indoor neutral-host operation. Blockchain technology enables participating entities to establish trust by executing smart contracts. When supported by an appropriate consensus mechanism, the blockchain digitally facilitates, verifies, and enforces agreements between two or more parties [61]. Building on this capability, this work focuses on achievable downlink throughput as a radio-centric QoS indicator and quantifies the benefit (or degradation) of offloading indoor users from macro O2I service to an NHO small cell. To this end, a system-level evaluation is employed that produces throughput distributions for (i) macro-served indoor users (O2I baseline) and (ii) NHO-served indoor users, which incorporates penetration loss and log-normal shadowing consistent with 3GPP TR 38.901 large-scale effects [148]. Throughput traces are consumed by an off-chain middleware layer that computes throughput-centric QoS, derives reward–penalty pricing adjustments, and meters usage. Settlement is then enforced on-chain via a blockchain smart contract that records usage and automatically forwards payment from the MNO to the NHO, enabling transparent and auditable settlement. The end-to-end workflow is implemented and evaluated on the Ethereum Sepolia test network.

Together with the spectrum trading contributions of Chapters 3 and 4, this chapter demonstrates that blockchain-based coordination can be extended beyond inter-operator spectrum exchange to encompass performance-aware settlement in shared indoor environments. The three contributions of this thesis, taken together, address trust, accountability, and operational flexibility across complementary telecom coordination problems: dynamic spectrum trading (Chapters 3 and 4) and performance-aware neutral-host settlement (this chapter). The concluding synthesis of these contributions, along with limitations and future directions, is presented in Chapter 6.

5.2 System Model

We consider an indoor neutral-host deployment where an NHO operates a small cell within a building, serving subscribers from multiple MNOs. Indoor users may

be served either by an outdoor macro base station through an outdoor-to-indoor (O2I) propagation path, or by the NHO indoor cell via a short-range indoor link. The deployment scenario is illustrated in Fig. 5.1.

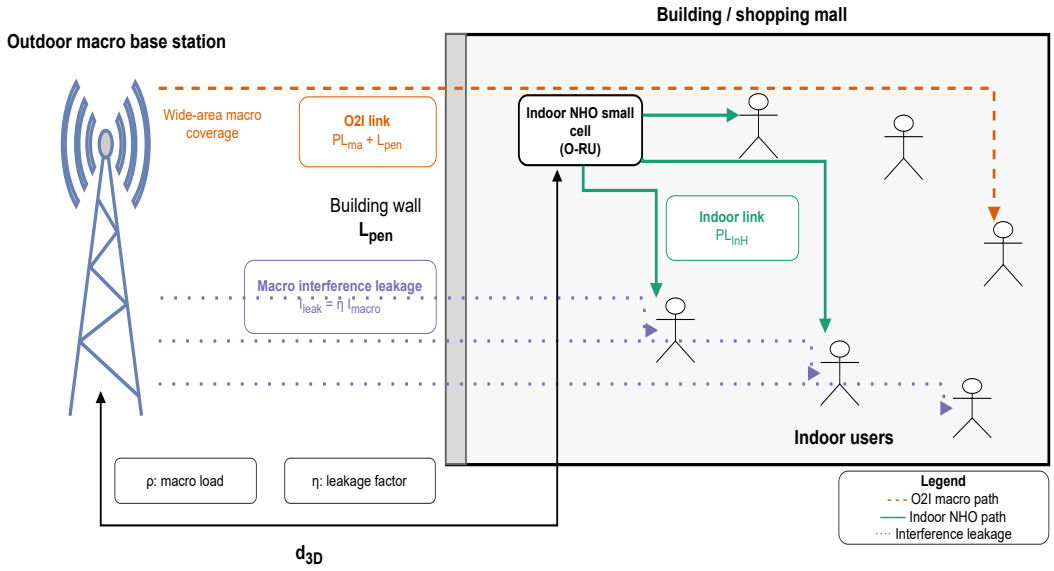


Figure 5.1: Deployment scenario showing indoor users served by either an outdoor macro base station (O2I path with penetration loss L_{pen}) or an indoor neutral-host small cell (InH path). The macro load factor ρ controls inter-cell interference, while the leakage factor η captures cross-tier coupling into the indoor layer.

The outdoor macro cell is included as a radio-performance baseline representing indoor service in the absence of an indoor neutral-host deployment. To quantify the radio performance benefit, a Monte Carlo system-level simulation is employed that models indoor users under both serving conditions. The simulation parameters and the source of each radio model are listed in Table 5.1.

Table 5.1: System-level simulation parameters and model sources.

Parameter	**Value / source**
Simulation area	1200 m × 1200 m
Building footprint	100 m × 100 m, centred in the area
Macro base stations	3, at 400 m from the building centre (0°, 120°, 240°)
Antenna heights	Macro 25 m; NHO small cell 3 m; UE 1.5 m
Carrier frequency f_c	3.5 GHz
Bandwidth B (macro and NHO)	50 MHz each
Transmit power	Macro 46 dBm; NHO 27 dBm
Antenna gain	Macro 18 dBi; NHO 5 dBi
NHO pathloss	3GPP InH LOS/NLOS, Eq. (5.3) [148]
Macro pathloss	3GPP TR 36.814 macro model with frequency correction, Eq. (5.4) [150]
Building entry loss L_{pen}	Log-normal, mean 25 dB, std. dev. 8 dB, after [149]
Shadowing std. dev.	Indoor $\sigma_{\text{In}} = 3$ dB; outdoor (macro) 6 dB
Noise PSD, noise figure	-174 dBm/Hz; $NF = 7$ dB
Implementation loss	2 dB
Users per realisation	400
Monte Carlo realisations	1000 per (ρ, η) configuration
Macro load ρ, leakage η	$\{0.3, 0.5, 0.7\}$ each
NHO contention parameter β	2
Risk sensitivity λ; penalty amplifier κ	0.8; 1.5
Base price P_0	0.001 ETH/GB

A system-level large-scale propagation abstraction is adopted where received power is determined by distance-dependent pathloss, building penetration loss (for O2I), and log-normal shadowing. Small-scale fading is not explicitly modelled, as the objective is to compare long-term throughput trends, QoS differentiation, and settlement outcomes rather than instantaneous link adaptation. This abstraction is aligned with system-level evaluation practices in 3GPP channel modelling studies [148].

For user u served by transmitter b, the received power (dBm) is

$$P_{r,u}^{(b)} = P_t^{(b)} + G^{(b)} - PL_u^{(b)} - X_\sigma, \tag{5.1}$$

where $P_t^{(b)}$ is transmit power (dBm), $G^{(b)}$ is antenna gain (dBi), $PL_u^{(b)}$ is large-scale pathloss (dB), and X_σ is log-normal shadowing (dB). Because powers add

only in the linear domain, all received powers, interference terms, and noise are converted to milliwatts before computing the SINR:

$$P_{\mathrm{mW}} = 10^{P_{\mathrm{dBm}}/10}. \tag{5.2}$$

The SINR is converted back to dB, $\mathrm{SINR}_{\mathrm{dB}} = 10\log_{10}(\mathrm{SINR})$, only for reporting.

Indoor users served by the NHO are modelled using the 3GPP Indoor Hotspot (InH) large-scale pathloss model [148]:

$$PL_{\mathrm{InH}}(d_{3D}) = \begin{cases} 32.4 + 17.3\log_{10}(d_{3D}) + 20\log_{10}(f_c), & d_{3D} < 15 \text{ m}, \\ 38.3\log_{10}(d_{3D}) + 17.3 + 24.9\log_{10}(f_c), & d_{3D} \geq 15 \text{ m}, \end{cases} \tag{5.3}$$

where d_{3D} is the 3D distance (m) and f_c is the carrier frequency (GHz). The two branches are the line-of-sight (LOS) and non-line-of-sight (NLOS) InH expressions of 3GPP TR 38.901 [148]. As a simplification, the LOS expression is applied within 15 m of the indoor cell and the NLOS expression beyond it. Log-normal shadowing with standard deviation σ_{In} captures indoor variability.

Indoor users served by outdoor macro base stations are modelled using a 3GPP UMa pathloss formulation augmented with building penetration loss to represent O2I propagation:

$$PL_{\mathrm{UMa}}(d_{3D,\mathrm{km}}) = 128.1 + 37.6\log_{10}(d_{3D,\mathrm{km}}) + 20\log_{10}\left(\frac{f_c}{2}\right), \tag{5.4}$$

where $d_{3D,\mathrm{km}}$ is the 3D distance in kilometres and f_c is in GHz. The first two terms are the macro-cell model of 3GPP TR 36.814, defined at 2 GHz [150], and the last term scales it to the simulated carrier frequency. For indoor users, an additional penetration loss L_{pen} is applied, drawn from a log-normal distribution consistent with ITU-R P.2109 [149]. The resulting O2I large-scale loss captures the degradation experienced by indoor users relying on outdoor macro coverage.

Thermal noise over bandwidth B is computed as

$$N_{\mathrm{dBm}} = -174 + 10\log_{10}(B) + NF, \tag{5.5}$$

where B is in Hz and NF is the receiver noise figure (dB). The noise power is converted to milliwatts using Eq. (5.2). The downlink SINR for user u is

$$\mathrm{SINR}_u = \frac{P_{s,u}}{I_u + N_0}, \tag{5.6}$$

with $P_{s,u}$ the serving-signal power, I_u the aggregate interference, and N_0 the noise

power in linear scale.

The aggregate interference I_u depends on the serving scenario and is constructed from two components. Inter-cell macro interference, arising from non-serving macro base stations, is defined as

$$I_{\text{macro}} = \rho \sum_{b \neq b^\star} P_{r,u}^{(b)}, \tag{5.7}$$

where $b^\star$ denotes the serving macro and $\rho \in (0, 1]$ is the macro load factor that captures the fraction of time–frequency resources occupied by neighbouring macro transmitters. When the NHO operates in co-channel mode, leaked macro interference entering the indoor layer is captured via a leakage factor η:

$$I_{\text{leak}} = \eta\, I_{\text{macro}}, \tag{5.8}$$

where η controls the degree of cross-tier coupling. The aggregate interference used in the SINR computation is then

$$I_u = \begin{cases} I_{\text{macro}}, & \text{user served by macro,} \\ I_{\text{leak}}, & \text{user served by NHO.} \end{cases} \tag{5.9}$$

Thus, ρ controls macro congestion and occupancy, while η controls the strength of cross-tier interference coupling into the neutral-host indoor layer.

5.2.1 Achievable Rate Proxy from SINR

To obtain tractable radio-centric throughput estimates without full PHY simulation, an achievable-rate proxy is computed from the downlink SINR using the Shannon capacity expression:

$$T_u^{\text{peak}} = B \log_2(1 + \text{SINR}_u)\,, \tag{5.10}$$

where B is the system bandwidth and SINR_u is in linear scale. This provides an idealised upper-bound rate estimate used for comparative evaluation between macro O2I service and neutral-host indoor service, rather than an exact NR PHY throughput model [151]. Because a practical link supports only a finite set of modulation and coding schemes, the throughput values reported in Section 5.5 are obtained by mapping each user's SINR to a discrete spectral efficiency. The

SINR is compared with the Shannon requirement $2^{SE} - 1$ of each level in

$$SE \in \{0,\ 0.2,\ 0.8,\ 1.8,\ 3.0,\ 4.5,\ 6.0\} \text{ bit/s/Hz}, \tag{5.11}$$

the highest level whose requirement is met is selected, and the rate is $T_u = B \cdot SE$. With $B = 50$ MHz, this gives seven throughput levels, $\{0, 10, 40, 90, 150, 225, 300\}$ Mbps, which explains the step-like shape of the throughput CDF and its saturation at 300 Mbps.

In addition to cross-tier interference, the NHO experiences intra-system contention when multiple users are served concurrently. This is modelled by defining an NHO load factor

$$\rho_{\text{nho}} = \min\left(0.95, \frac{N_{\text{off}}}{K_{\text{nho}}}\right), \tag{5.12}$$

where N_{off} is the number of offloaded users admitted to the neutral-host during the accounting interval and K_{nho} is the neutral-host service capacity (users per interval). The achievable throughput for a user u served by the neutral-host network is reduced by contention and is scaled as

$$T_{u,\text{nho}}^{\text{del}} = \frac{T_{u,\text{nho}}^{\text{peak}}}{1 + \beta\rho_{\text{nho}}}, \tag{5.13}$$

where $\beta \geq 0$ controls the severity of contention. When $\beta = 0$, offloaded users receive the peak proxy rate regardless of NHO load. As β increases, contention effects become more pronounced. For example, at $\beta = 2$ and $\rho_{\text{nho}} = 0.95$, the delivered proxy rate drops to approximately 34% of the peak rate. In this work, $\beta = 2$ is used to represent a moderately loaded indoor small-cell environment. This contention-adjusted rate proxy $T_{u,\text{nho}}^{\text{del}}$ is used as T_u in the QoS computation (Section 5.3.1) for all users served by the neutral-host network. For macro-served indoor users, T_u is the peak proxy rate obtained directly from (5.10) without contention scaling, since the effect of macro congestion is already captured through the load-scaled interference I_{macro} in the SINR computation (Equation 5.7), which reduces the macro throughput as ρ increases.

5.3 QoS Metric and Settlement Mechanism

5.3.1 Throughput-Centric QoS

QoS is defined using the achievable-rate proxy as the sole performance indicator. Latency and queueing are intentionally excluded to avoid assumptions about traffic models and scheduler implementations. Let $\mathcal{U}$ denote the set of all indoor

users evaluated within an accounting interval. For user $u \in \mathcal{U}$, normalised QoS is defined as

$$QoS_u = \frac{T_u}{\max_{v \in \mathcal{U}} T_v}, \tag{5.14}$$

where v ranges over all users in $\mathcal{U}$ and the denominator is the maximum achievable throughput observed across the entire user set, yielding $QoS_u \in [0, 1]$. The set $\mathcal{U}$ pools the macro-served and NHO-served users of the same Monte Carlo realisation, so both serving conditions are normalised by one common reference. This is necessary because normalising each condition by its own maximum would map both to the same $[0, 1]$ scale and hide the difference between them.

5.3.2 Serving-Condition Aggregation and QoS Gain

To compare radio performance across the two serving conditions, users are grouped according to whether they are served by the outdoor macro base station ($s =$ macro) or by the indoor neutral-host cell ($s =$ nho). For each serving condition $s \in \{\text{macro, nho}\}$, the mean QoS is

$$\overline{QoS}_s = \frac{1}{|\mathcal{U}_s|} \sum_{u \in \mathcal{U}_s} QoS_u, \tag{5.15}$$

and the QoS gain attributable to neutral-host offloading is

$$\Delta QoS = \overline{QoS}_{\text{nho}} - \overline{QoS}_{\text{macro}}. \tag{5.16}$$

A positive ΔQoS indicates that NHO service improves the delivered rate relative to the macro O2I baseline.

5.3.3 Usage Accounting

Traffic served by the NHO on behalf of an MNO is accounted over an interval τ (in seconds) by converting the aggregate neutral-host throughput from a rate (Mbps) into a data volume (GB). The throughput T_{nho} (Mbps) is first converted to bits per second by multiplying by 10^6, then multiplied by the accounting interval τ to obtain the total number of bits transferred. This is divided by 8 to convert bits to bytes and by 10^9 to convert bytes to gigabytes. For on-chain settlement, the result is scaled to milli-GB to avoid floating-point arithmetic in the smart contract:

$$U_{\text{mGB}} = \left\lfloor 1000 \cdot \frac{T_{\text{nho}} \cdot 10^6 \cdot \tau}{8 \cdot 10^9} \right\rfloor, \tag{5.17}$$

where T_{nho} is the aggregate neutral-host throughput allocated to the offloaded subset (Mbps), averaged per accounting interval, and $\lfloor \cdot \rfloor$ denotes the floor operation. The result U_{mGB} represents the metered usage in milli-GB that is committed on-chain as part of the settlement record.

5.3.4 Reward–Penalty Pricing (Risk-Aware)

The purpose of the reward–penalty pricing mechanism is to adjust the wholesale price that the MNO pays the NHO based on the radio performance actually delivered to offloaded users. Unlike static pricing, where the NHO receives the same payment regardless of service quality, this mechanism increases the price when offloading consistently improves user throughput (rewarding the NHO for good performance) and decreases the price when a significant fraction of offloaded users experience degraded throughput (penalising the NHO for poor performance). The mechanism operates in three steps: first, it quantifies the probability that individual users are worse off after offloading. Second, it combines this degradation risk with the mean QoS gain into a single risk-adjusted metric. Third, it maps this metric to a pricing multiplier that scales the base price up or down.

Since degradation can occur for a subset of offloaded users even when the mean gain is positive, the penalty probability is defined as

$$p_{\text{pen}} = \Pr(\Delta QoS_u < 0), \tag{5.18}$$

where ΔQoS_u is the per-user QoS gain for the offloaded subset. A high p_{pen} indicates that a large fraction of users experienced worse throughput on the NHO than on the macro baseline, even if the mean gain is positive. A risk-aware effective gain is then defined as

$$Q_{\text{eff}} = \Delta QoS - \lambda\, p_{\text{pen}}, \tag{5.19}$$

where $\lambda \geq 0$ controls risk sensitivity. The term $\lambda\, p_{\text{pen}}$ acts as a penalty deduction: even when the mean gain ΔQoS is positive, a high degradation probability reduces Q_{eff}, reflecting the fact that a service that benefits most users but harms others should not be fully rewarded. The reward–penalty multiplier is

$$\alpha = \begin{cases} 1 + Q_{\text{eff}}, & Q_{\text{eff}} \geq 0, \\ \max(0, 1 - \kappa |Q_{\text{eff}}|), & Q_{\text{eff}} < 0, \end{cases} \tag{5.20}$$

where $\kappa > 1$ amplifies penalties so that negative performance outcomes reduce

the price more aggressively than positive outcomes increase it. For the on-chain implementation, $\lambda = 0.8$ and $\kappa = 1.5$.

Three pricing quantities are distinguished in the settlement mechanism. P_0 is the **base price per GB**, a fixed reference rate agreed between the MNO and the NHO before any performance adjustment is applied. It represents the price the MNO would pay under neutral conditions where the NHO neither improves nor degrades user throughput relative to the macro baseline. P is the **adjusted price per GB** after the reward–penalty multiplier α is applied:

$$P = P_0\,\alpha. \tag{5.21}$$

When the NHO delivers strong performance ($\alpha > 1$), the adjusted price exceeds the base price and the NHO earns a premium. When performance is poor ($\alpha < 1$), the adjusted price falls below the base price and the NHO receives a discount. C is the **total settlement cost** for the accounting interval, obtained by multiplying the adjusted price per GB by the total data volume consumed:

$$C = U_{\mathrm{GB}} \cdot P. \tag{5.22}$$

This is the actual amount of ETH transferred from the MNO wallet to the NHO wallet via the smart contract. In summary, P_0 is what the MNO would pay per GB under baseline conditions, P is what the MNO actually pays per GB after performance adjustment, and C is the total payment for all data consumed during the interval.

The benefit of this mechanism is threefold. First, it introduces accountability into neutral-host operation by linking revenue directly to delivered performance rather than to static contractual terms. Second, it creates a financial incentive for the NHO to maintain high service quality, since better radio performance translates to higher per-GB revenue. Third, it protects the MNO from overpaying for degraded service, since the penalty term automatically discounts the price when a significant fraction of offloaded subscribers are worse off than they would have been on the macro network.

5.3.5 On-Chain Settlement via Smart Contract

Settlement is enforced via a smart contract that records metered usage and executes payment transfers according to the off-chain pricing output. Radio KPI processing and pricing computation remain off-chain for scalability, while the final settlement is executed on-chain for auditability and non-repudiation. This sep-

aration allows physics-aware KPI processing without embedding radio analytics inside the EVM.

5.4 System Architecture and Settlement Workflow

5.4.1 Framework Overview

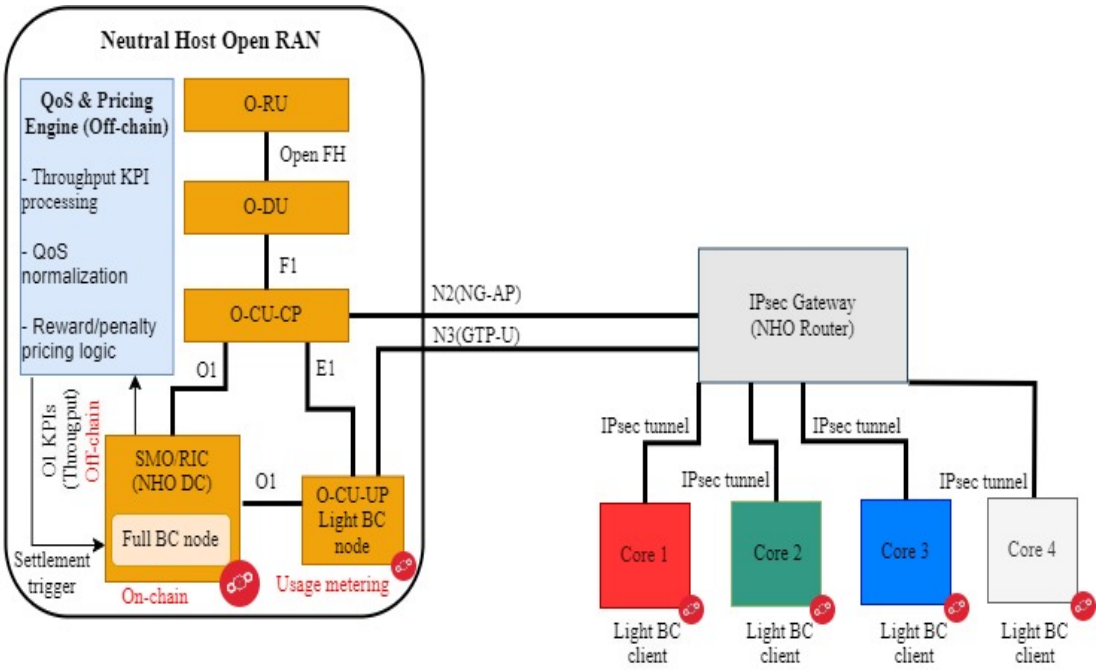


Figure 5.2: Throughput-aware settlement architecture for an indoor neutral-host Open RAN deployment. Throughput KPIs are collected via O1 and processed off-chain by the QoS and pricing engine to compute reward/penalty adjustments. Settlement is executed on-chain via a smart contract maintained by the NHO and verified by lightweight audit clients at participating MNO cores.

An indoor NHO deployment in a shopping mall environment is considered, where the NHO operates a shared Open RAN infrastructure serving subscribers from multiple MNOs. The proposed architecture, illustrated in Fig. 5.2, separates (i) the radio and KPI collection plane (Open RAN functions and management interfaces), (ii) the off-chain QoS and pricing plane (middleware processing), and (iii) the on-chain settlement plane (smart-contract enforcement and auditability).

Within the NHO domain, the Open RAN stack comprises O-RU, O-DU, O-CU-CP, and O-CU-UP. Network KPIs are collected via the O1 interface toward the SMO/RAN intelligent controller (RIC) in the NHO data centre, which hosts the off-chain QoS and pricing engine. Multi-operator connectivity is supported

via an IPsec gateway that terminates per-MNO IPsec tunnels and forwards N2 (NG-AP) control-plane signalling to the O-CU-CP and N3 (GTP-U) user-plane traffic to the O-CU-UP. Settlement is executed using a blockchain smart contract maintained by the NHO (full node), while lightweight audit clients at MNO cores independently verify settlement outcomes.

5.4.2 Radio Simulation Model

To quantify the radio performance benefit of indoor neutral-host service, a system-level simulation is employed that models indoor users under two serving conditions: (i) indoor users served by an outdoor macro base station (baseline), and (ii) indoor users served by the NHO indoor cell. The outdoor macro cell is included only as a radio-performance baseline, and represents indoor service in the absence of an indoor neutral-host deployment.

For each Monte Carlo iteration, user locations are sampled uniformly in a target indoor region, and large-scale channel effects (pathloss and shadowing) are generated to compute received power, interference, SINR, and achievable throughput. The simulation outputs empirical distributions of throughput samples for both macro-served indoor users and NHO-served indoor users. These throughput traces are then consumed by the middleware layer for QoS processing, pricing, and settlement.

5.4.3 Middleware QoS Processing and Usage Accounting

The middleware performs two main functions: (i) throughput-centric QoS computation, and (ii) usage accounting. Throughput KPIs are ingested from the simulation (or from O1 KPI streams in a deployment setting) and normalised to produce a dimensionless QoS score. Usage is accounted for over an operator-defined interval by aggregating transmitted bytes (or, equivalently, throughput over the session duration). Importantly, QoS and pricing computations are executed off-chain to preserve scalability and avoid placing radio-analytics logic on-chain.

5.4.4 Blockchain-Based Settlement Layer and Implementation

Settlement between the MNO and the neutral-host operator (NHO) is enforced via a blockchain smart contract that records metered usage and executes payment transfers based on pricing outputs generated by the off-chain middleware.

The on-chain layer provides immutability and auditability, which enables MNOs to verify settlement outcomes through transaction traces without requiring full participation in consensus. To preserve scalability, the blockchain stores only the minimum information required for settlement verification (e.g., usage units and payment transfer), while radio KPI processing, QoS computation and pricing inputs are computed off-chain.

The settlement layer is implemented on an Ethereum-compatible blockchain (EVM) and evaluated on the Ethereum Sepolia test network. The smart contract is written in Solidity and deployed using Remix IDE, with transaction signing and key management handled via MetaMask. Off-chain middleware (Web3.py) submits settlement transactions on behalf of the MNO, while the NHO maintains the contract address and receives forwarded payments.

Three entities are considered: (i) the MNO payer account (operator wallet) that triggers settlement and provides ETH, (ii) the NHO payee account that receives forwarded payments, and (iii) the smart contract that stores settlement parameters (e.g., base price and risk factors) and enforces payment forwarding. QoS and usage computation are intentionally performed off-chain to avoid embedding radio analytics inside the EVM, and only the minimum verifiable settlement inputs are committed on-chain.

For each accounting interval, the middleware computes the mean gain ΔQ, the penalty probability p_{pen}, and the usage U_{mGB}, and the MNO wallet submits them to `settleUsage` together with a payment. The contract recomputes the effective gain $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$ and the unit price on-chain using its stored parameters, and derives the required payment from the usage. The transaction is rejected if the attached `msg.value` is lower than the required payment. Otherwise, the contract forwards exactly the required amount to the NHO address, refunds any excess to the MNO, and emits a `Settled` event that records the usage, effective gain, unit price, and payment, producing an immutable record that can be audited by all parties.

Table 5.2: Smart-contract interface for risk-aware settlement and payment forwarding.

Function	Caller	Purpose / Inputs
`setNeutralHost(_nh)`	Owner (NHO admin)	Sets the NHO payee address.
`setParams(_basePriceWeiPerGB, _penaltyFactorPPM, _lambdaPPM)`	Owner (NHO admin)	Sets P_0 (Wei/GB), κ, and λ (in ppm).
`quotePriceWeiPerGB(meanDQ_ppm, pPen_ppm)`	Middleware (read)	Computes $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$ and α, and returns the unit price P in Wei/GB. Integer ppm scaling avoids floating point.
`settleUsage(usage_mGB, meanDQ_ppm, pPen_ppm)`	MNO wallet (tx, `payable`)	Recomputes Q_{eff} and the unit price on-chain, computes the required payment, rejects underpayment, forwards the required amount to the NHO, refunds any excess, and emits a `Settled` event recording the usage and payment.

Table 5.2 summarises the callable interface. Fig. 5.3 provides the corresponding Solidity implementation excerpt.

```
function setNeutralHost(address payable _nh) external onlyOwner {    undefined gas
    neutralHost = _nh;
    emit NeutralHostUpdated(_nh);
}

function setParams(uint256 _basePriceWeiPerGB, uint256 _penaltyFactorPPM, uint256 _lambdaPPM)
    external
    onlyOwner
{
    basePriceWeiPerGB = _basePriceWeiPerGB;
    penaltyFactorPPM = _penaltyFactorPPM;
    lambdaPPM = _lambdaPPM;
    emit ParamsUpdated(_basePriceWeiPerGB, _penaltyFactorPPM, _lambdaPPM);
}


function quotePriceWeiPerGB(int256 meanDQ_ppm, uint256 pPen_ppm)    infinite gas
    public
    view
```

Figure 5.3: Solidity excerpt illustrating configuration and pricing functions (e.g., neutral-host payee update, parameter update, and the read-only pricing function `quotePriceWeiPerGB` used by the middleware).

To make the offloading and settlement lifecycle explicit at the message level, Table 5.3 summarises the end-to-end phases and the responsible entity at each step.

Table 5.3: End-to-end offloading and settlement phases (message-level view).

Phase	**Entity**	**Action**
1	NHO O-RAN	Serve indoor users. Collect KPIs (e.g., SINR/throughput proxy) via O1.
2	Middleware	Compute QoS, ΔQ, p_{pen}, and usage U_{mGB}.
3	Middleware → Contract	(Optional) call `quotePriceWeiPerGB` to preview the unit price.
4	MNO wallet → Contract	Submit `settleUsage`$(U_{\text{mGB}}, \Delta Q, p_{\text{pen}})$ with `msg.value`; the contract recomputes the price and rejects underpayment.
5	Contract → NHO / MNO	Forward the required ETH to the NHO, refund any excess to the MNO, and emit a `Settled` event.

Settlement correctness is verified by the presence of an internal ETH transfer from the contract to the NHO address in the Sepolia transaction traces, which demonstrates automated payment forwarding under both penalty and reward regimes.

The responsibilities for settlement validation are divided as follows. The off-chain middleware is responsible for collecting the throughput KPIs and computing QoS, the mean gain ΔQ, the penalty probability p_{pen}, and the usage U_{mGB}. The smart contract is responsible for everything that follows from these inputs: it recomputes the effective gain and the unit price on-chain, derives the required payment, rejects underpayment, forwards the exact amount to the NHO, refunds any excess, and records the settlement in an event log. Changes to the pricing parameters and to the NHO address are restricted to the owner. As a result, neither party can alter the pricing or payment rules at settlement time. The contract cannot, however, verify that the KPIs and usage supplied by the middleware are correct; these off-chain inputs remain a trust assumption. This risk is mitigated because each MNO can cross-check the reported values against its own network records, and because every settlement is permanently recorded, so that misreporting can be detected and disputed afterwards.

5.5 Performance Evaluation and Results

Monte Carlo simulations are conducted with 1000 independent realisations per parameter configuration. A total of 400 users are randomly distributed, with indoor users concentrated within a building footprint. Key parameters include macro load $\rho \in \{0.3, 0.5, 0.7\}$ and macro-to-NHO leakage $\eta \in \{0.3, 0.5, 0.7\}$. For each run, throughput distributions, QoS values, and settlement outcomes are recorded.

To provide a statistically robust comparison between macro-served indoor users and neutral-host (NHO) served users, the empirical cumulative distribution functions (CDFs) of received power (RSRP proxy), SINR, and downlink throughput are evaluated. CDF-based analysis enables visualisation of both median performance and tail behaviour, which is critical for evaluating user experience and settlement fairness as shown in Figures 5.4, 5.5, and 5.6.

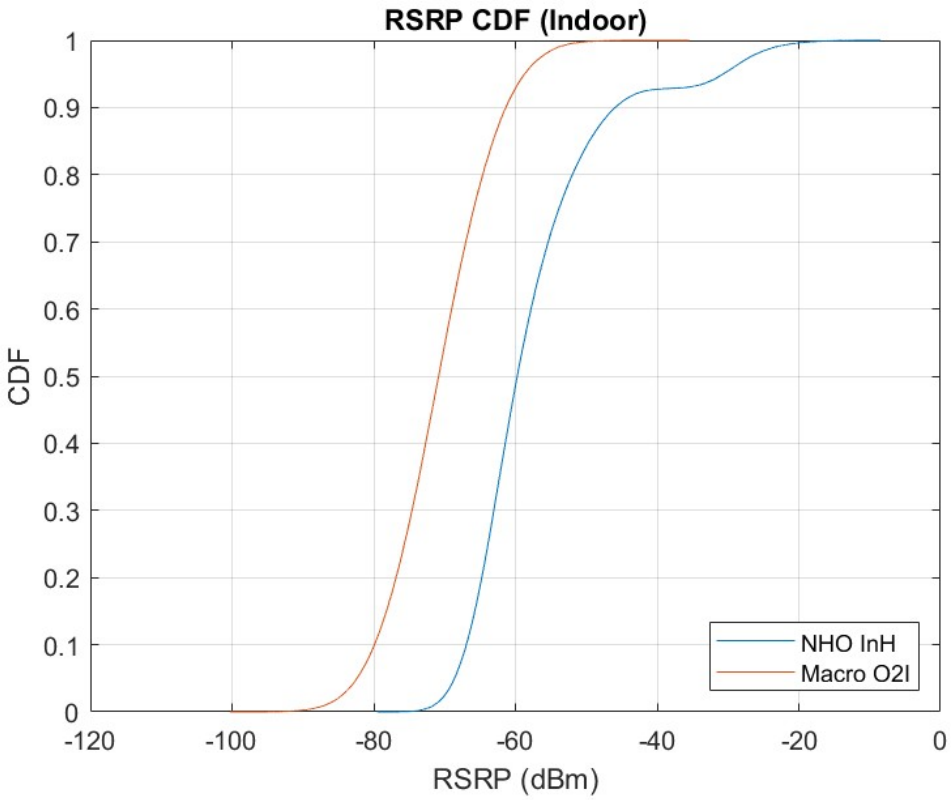


Figure 5.4: Empirical CDF of received signal strength (RSRP proxy) for indoor users under macro O2I and neutral-host indoor service.

Figure 5.4 compares the large-scale received power distributions for indoor users under macro O2I service and neutral-host (NHO) indoor service. The NHO curve is consistently right-shifted, confirming a stronger indoor link budget due to reduced penetration loss and more favourable indoor geometry. Quantita-

tively, the median RSRP improves from −70.85 dBm (macro O2I) to −59.71 dBm (NHO), representing a +11.13 dB gain. At the 10th percentile the improvement is +13.02 dB (from −79.92 dBm to −66.90 dBm), and at the 90th percentile it is +15.29 dB (from −61.39 dBm to −46.10 dBm). These distributional shifts provide the physical justification for indoor neutral-host offloading.

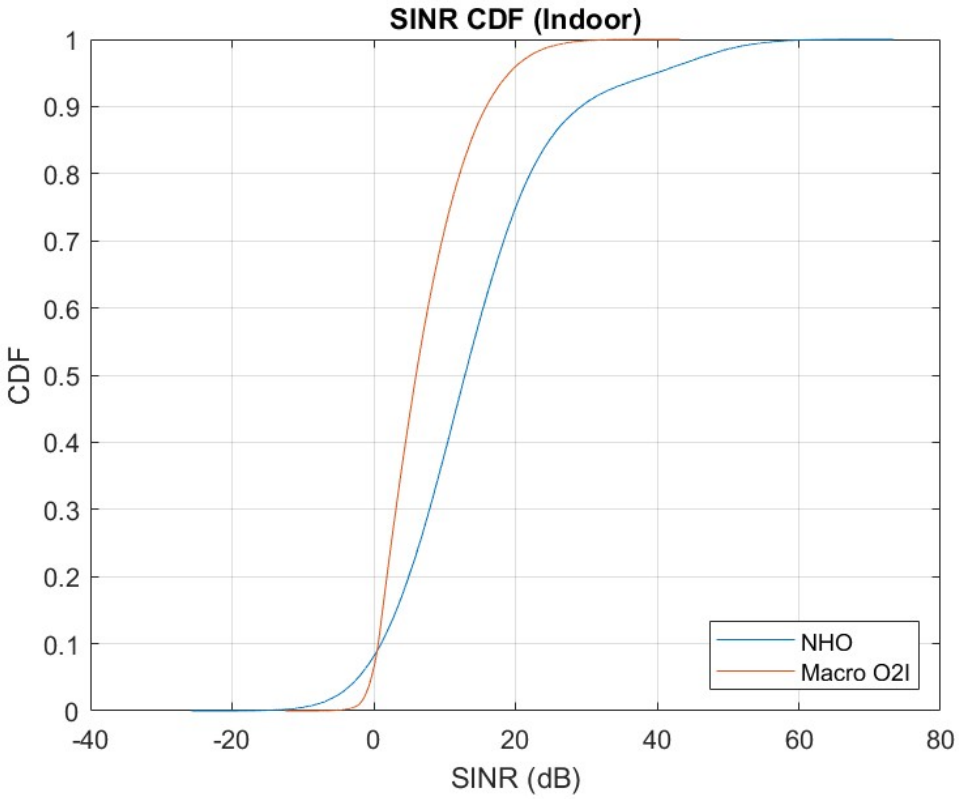


Figure 5.5: Empirical CDF of downlink SINR for indoor users under macro and neutral-host operation.

Figure 5.5 presents the SINR distributions after accounting for load-scaled macro interference and cross-tier leakage. The neutral-host (NHO) scenario exhibits a clear right-shift relative to macro O2I, indicating improved interference-limited indoor performance due to stronger desired signal and reduced effective macro coupling for most users. Quantitatively, the 90th-percentile SINR increases from 15.55 dB (macro O2I) to 27.66 dB (NHO), representing a +12.11 dB improvement. At the median, SINR improves by +5.83 dB (from 5.53 dB to 11.36 dB). These shifts confirm that indoor neutral-host service improves SINR distributionally. However, gains can saturate under high macro load as the macro layer becomes interference-limited and leakage coupling constrains the achievable improvement.

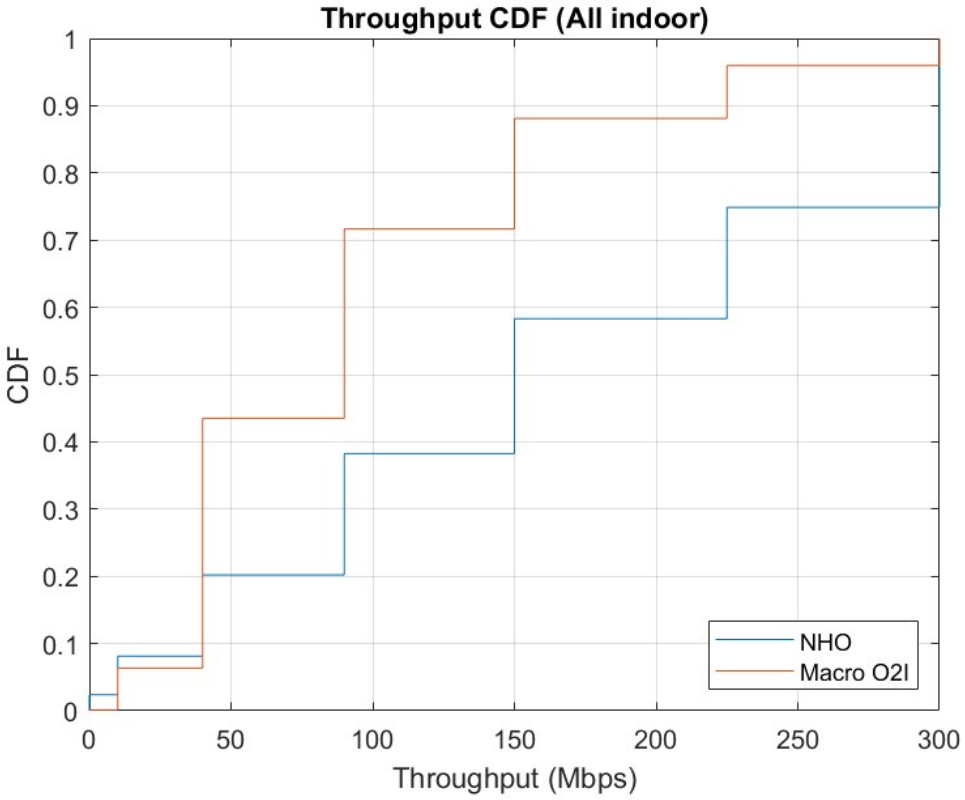


Figure 5.6: Empirical CDF of achievable downlink throughput derived from SINR-based spectral efficiency mapping.

Figure 5.6 reports the achievable downlink throughput distributions for indoor users under macro O2I and neutral-host (NHO) service. The NHO curve is right-shifted, indicating higher throughput across much of the user population. Quantitatively, the 90th-percentile throughput increases from 225.0 Mbps (macro O2I) to 300.0 Mbps (NHO), representing a +75.0 Mbps gain. At the median, throughput increases by +60.0 Mbps (from 90.0 Mbps to 150.0 Mbps), and the mean increases by +69.084 Mbps (from 95.282 Mbps to 164.366 Mbps). The upper tail saturation observed in the NHO curve is expected under the adopted abstraction: with fixed bandwidth and a bounded spectral-efficiency/rate proxy, additional SINR yields diminishing throughput gains. Overall, these distributional shifts provide quantitative support for reward regimes in performance-coupled settlement.

5.5.1 On-Chain Settlement Validation

Settlement is validated on Sepolia by confirming that the smart contract forwards the computed payment to the neutral-host operator (NHO) address via an internal ETH transfer. Figures 5.7 and 5.8 show two executed settlement transactions

corresponding to penalty and reward regimes, respectively. In the penalty regime, a negative effective gain yields a discounted unit price and a smaller forwarded payment. In the reward regime, a positive effective gain yields a higher unit price and a larger forward payment. These immutable ledger records provide non-repudiable evidence that settlement is automatically enforced according to radio-derived performance outcomes.

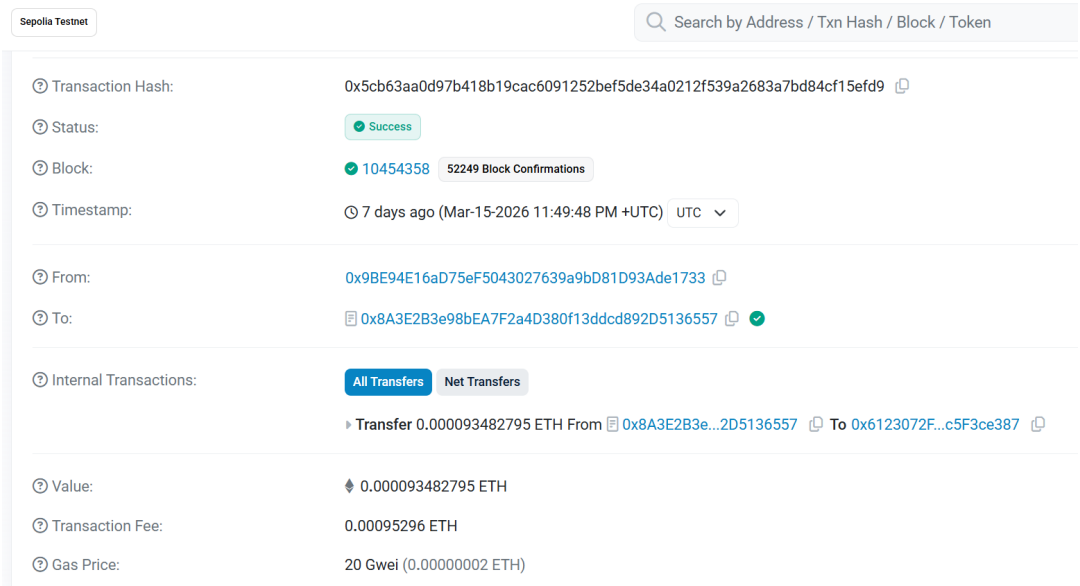


Figure 5.7: Penalty-regime settlement on Sepolia: internal transfer of 9.3482795×10^{-5} ETH from the contract to the NHO address (Tx: `0x5cb6...efd9`).

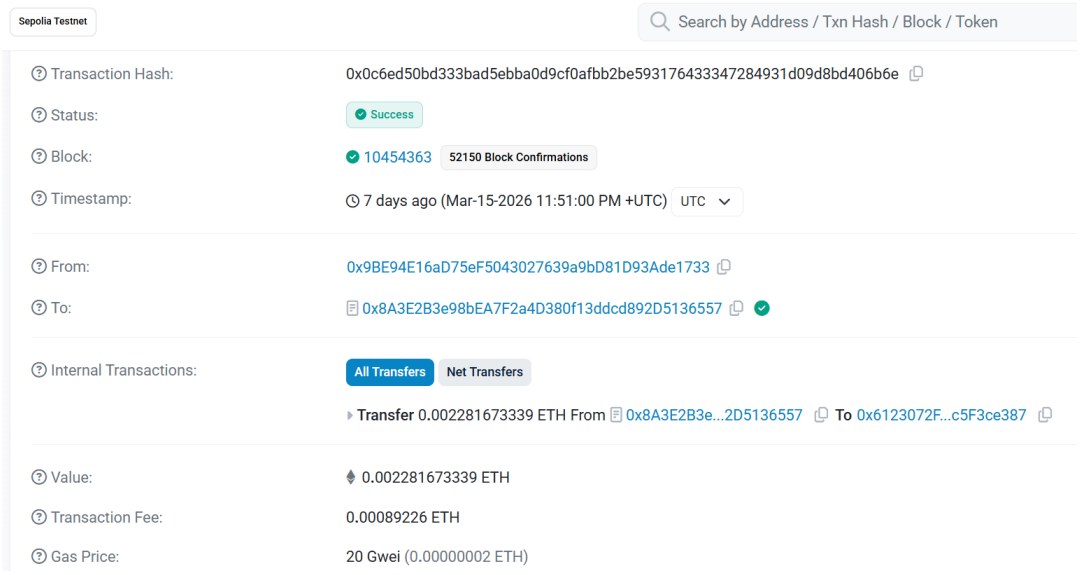


Figure 5.8: Reward-regime settlement on Sepolia: internal transfer of $2.281673339 \times 10^{-3}$ ETH from the contract to the NHO address (Tx: `0x0c6e...b6e`).

5.5.2 Sensitivity Results under Worst vs Random Admission Policies

Sensitivity to macro load ρ and cross-tier leakage η is evaluated under two admission policies: (i) *Worst* admission (reward-biased steering), which admits the weakest delivered macro users among the candidate set, and (ii) *Random* admission (stress-test steering), which admits a random subset of candidates. For each (ρ, η) point, the mean performance gain ΔQ for the offloaded subset and the penalty probability $p_{\text{pen}} = \Pr(\Delta Q_u < 0)$ are reported, where ΔQ_u denotes per-user gain within the offloaded subset. Settlement uses the risk-aware effective gain $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$ with $\lambda = 0.8$ and $\kappa = 1.5$ (base price $P_0 = 0.001$ ETH/GB), matching the deployed smart contract implementation.

Worst Admission (Reward-Biased Steering)

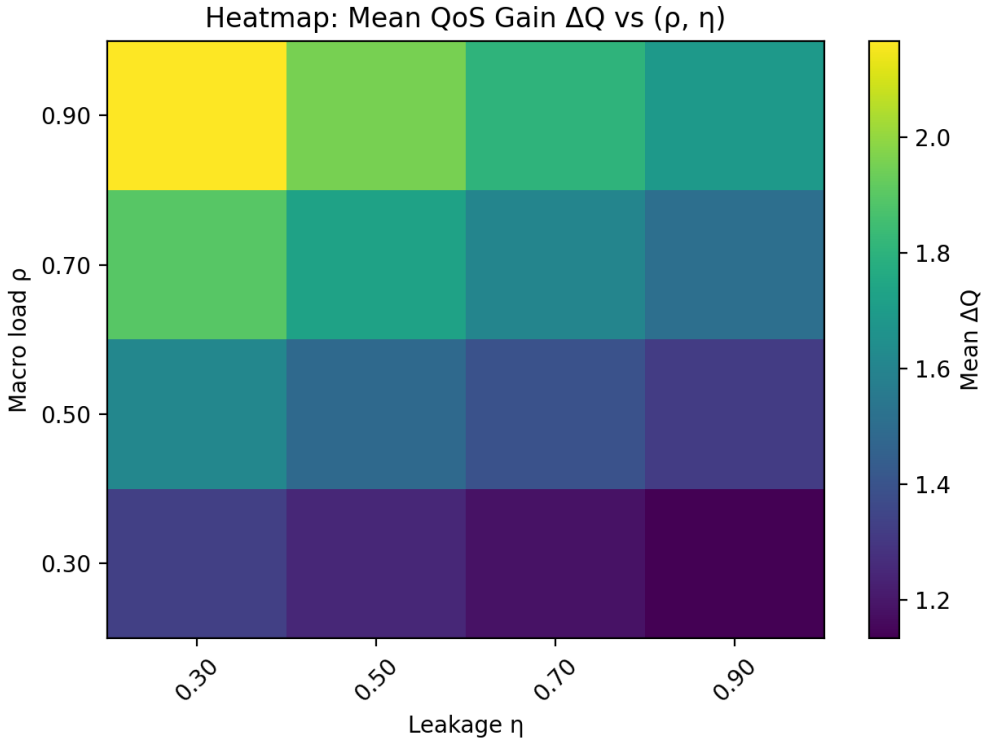


Figure 5.9: Worst admission: heatmap of mean gain ΔQ versus macro load ρ and leakage η.

Figure 5.9 shows that Worst admission yields positive mean gain across the (ρ, η) grid because it preferentially offloads the weakest macro users. The mean gain increases with macro load ρ (the macro O2I baseline becomes more constrained under higher utilisation) and decreases with leakage η (stronger cross-tier coupling reduces the neutral-host advantage).

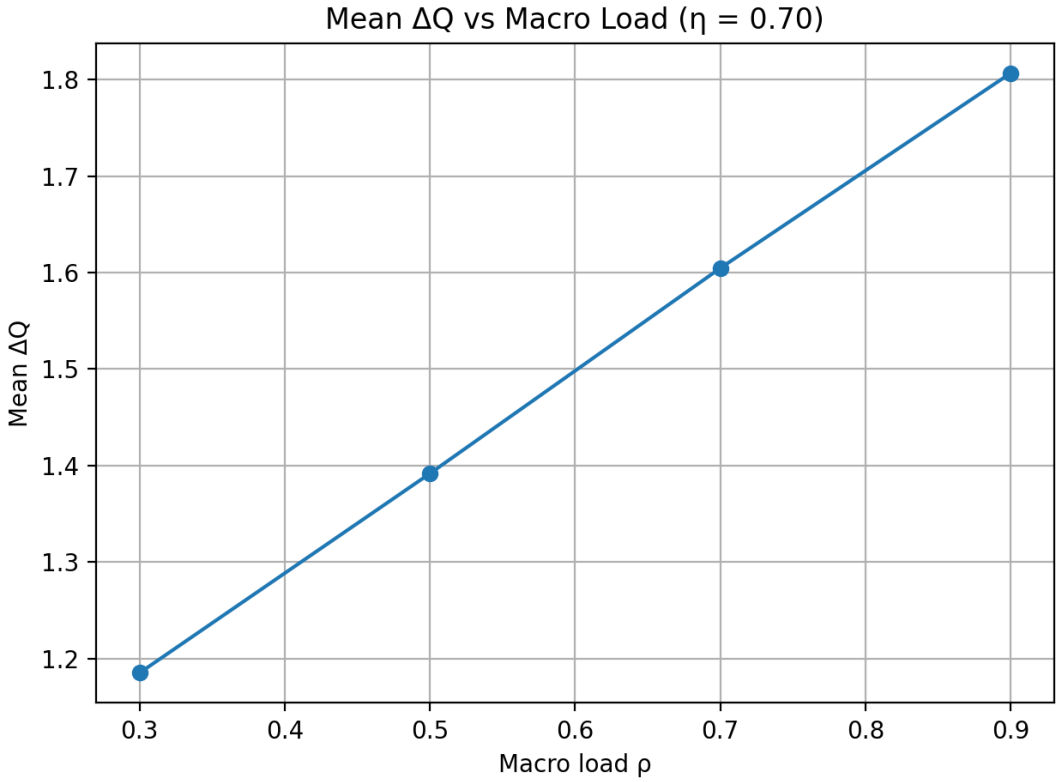


Figure 5.10: Worst admission: mean gain ΔQ versus macro load ρ at fixed leakage $\eta = 0.70$.

Figure 5.10 shows the ρ trend explicitly: as macro load increases, offloading becomes more valuable because macro-served indoor users become increasingly interference- and resource-limited, whereas the neutral-host indoor layer benefits from proximity and improved link conditions. The gain growth eventually saturates, reflecting practical limits imposed by neutral-host contention and the bounded rate/throughput proxy.

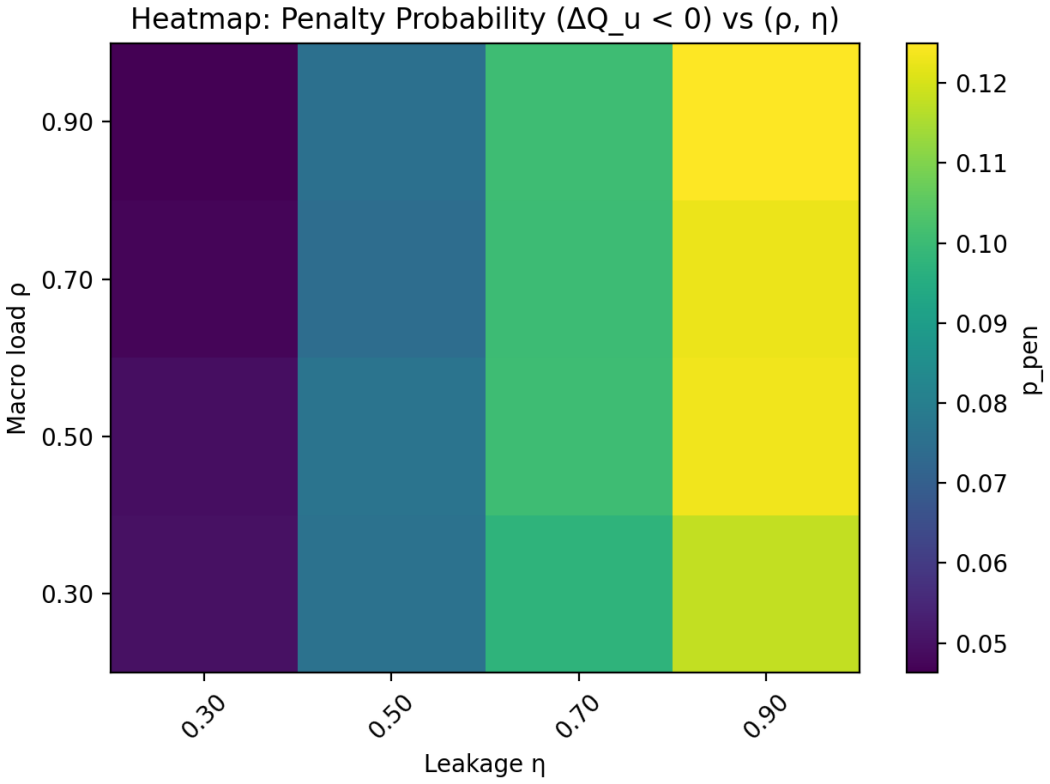


Figure 5.11: Worst admission: penalty probability $p_{\text{pen}} = \Pr(\Delta Q_u < 0)$ versus (ρ, η).

Figure 5.11 shows that, even under informed steering, a non-zero fraction of offloaded users can experience degradation due to randomness in user placement, shadowing, and leakage coupling. The penalty probability remains relatively low overall but increases with η, indicating that cross-tier coupling is the dominant driver of tail risk.

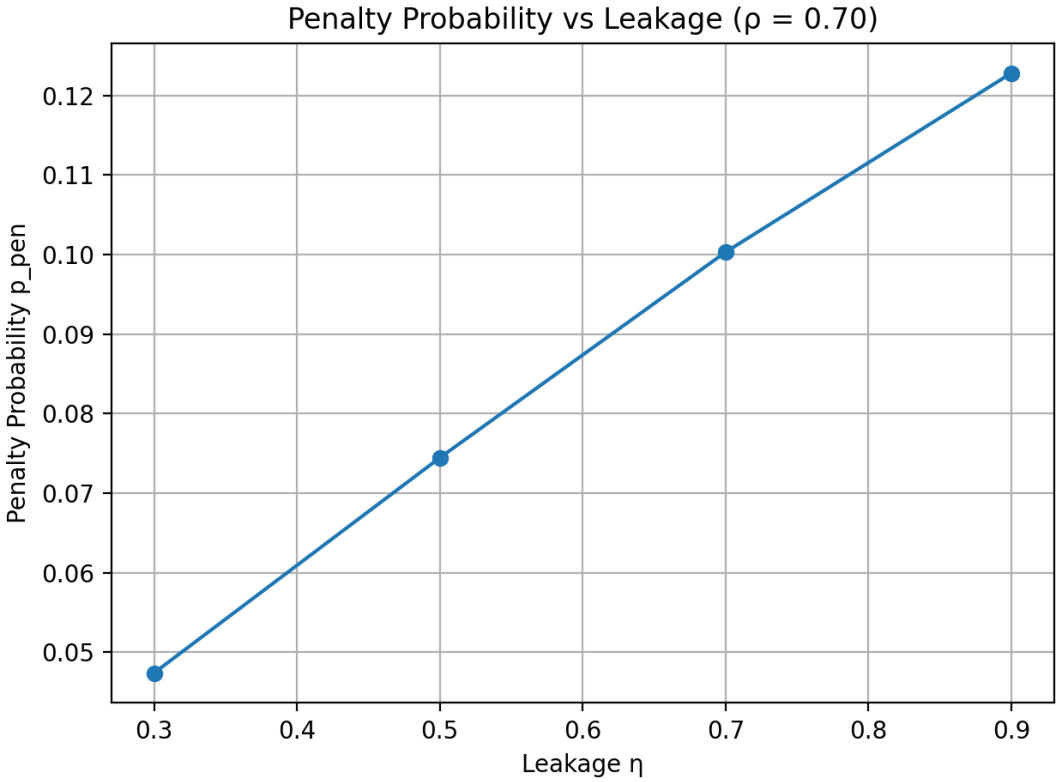


Figure 5.12: Worst admission: penalty probability versus leakage η at fixed macro load $\rho = 0.70$.

Figure 5.12 shows that leakage η is the dominant driver of tail risk under Worst admission: as η increases, more macro interference couples into the indoor neutral-host layer, increasing the probability that a subset of offloaded users experiences degradation even when the mean gain remains positive.

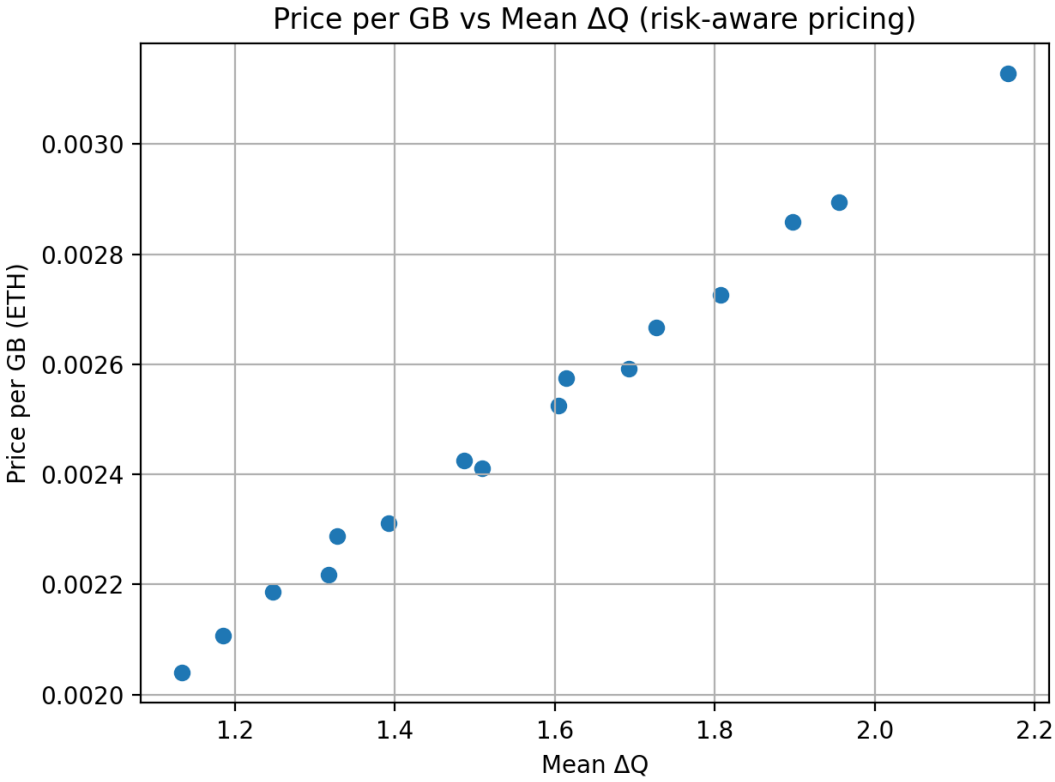


Figure 5.13: Worst admission: price per GB versus mean gain ΔQ under risk-aware settlement using $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$ with $\lambda = 0.8$ and penalty factor $\kappa = 1.5$ (base price $P_0 = 0.001$ ETH/GB).

Figure 5.13 shows that the settlement unit price generally increases with the mean gain ΔQ, but is ultimately determined by the risk-adjusted effective gain $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$. As a result, operating points with similar ΔQ can still yield different prices when the penalty probability differs, reflecting the intended discounting of high-risk regimes.

Random Admission (Uncertainty/Stress-Test Steering)

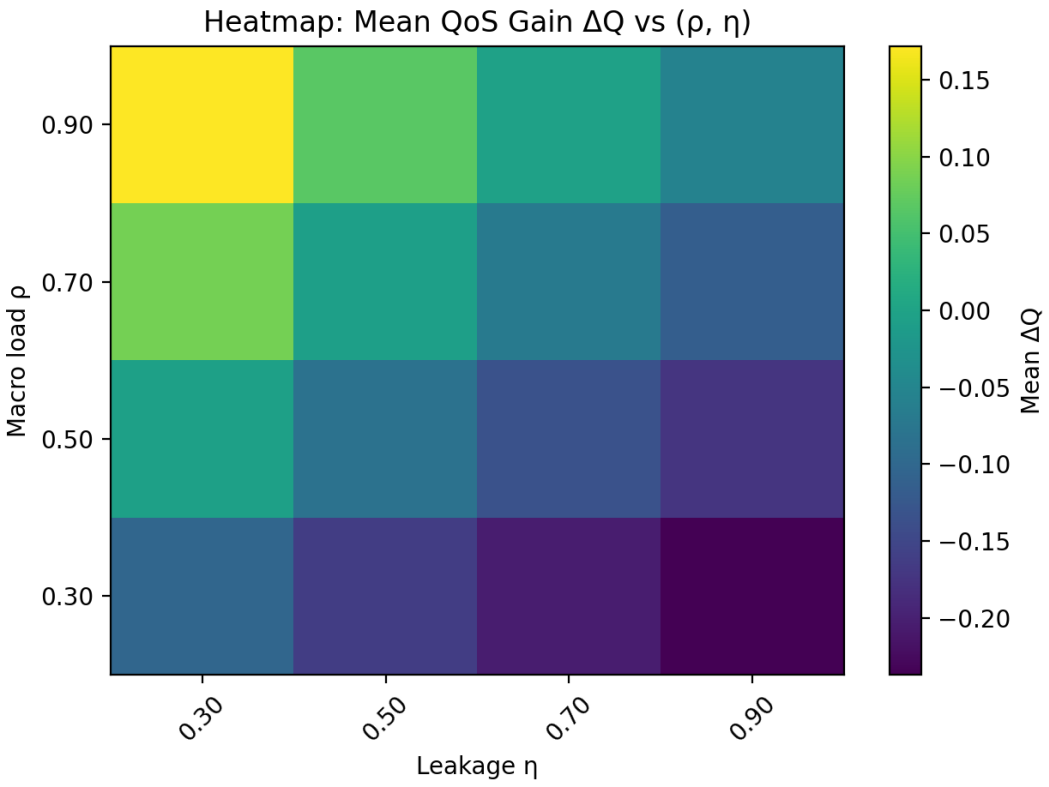


Figure 5.14: Random admission: heatmap of mean gain ΔQ versus macro load ρ and leakage η.

Figure 5.14 shows that Random admission can yield negative mean gains because users are admitted without ensuring that neutral-host service dominates the macro baseline under contention and leakage. This outcome reflects steering uncertainty rather than an inherent failure of the neutral-host network.

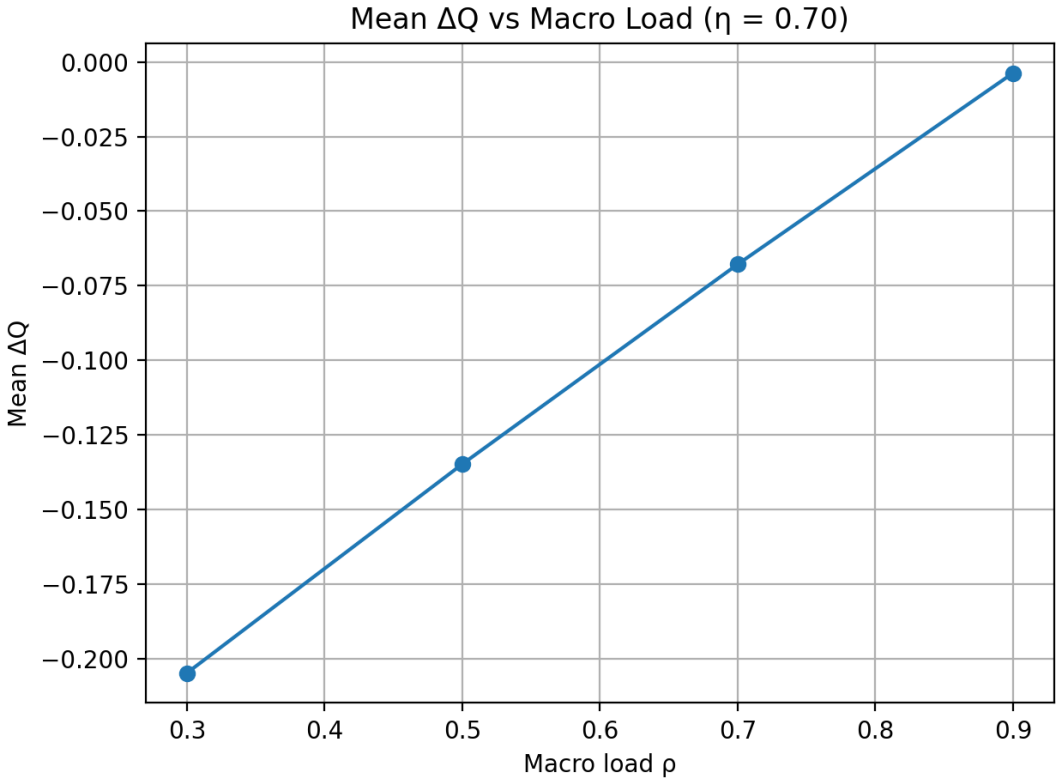


Figure 5.15: Random admission: mean gain ΔQ versus macro load ρ at fixed leakage $\eta = 0.70$.

Figure 5.15 shows that, as macro load ρ increases, the mean gain under Random admission tends to move upward because the macro O2I baseline degrades under heavier utilisation. However, the gains remain smaller and more variable than in Worst admission, since Random admission admits users without prioritising those most likely to benefit.

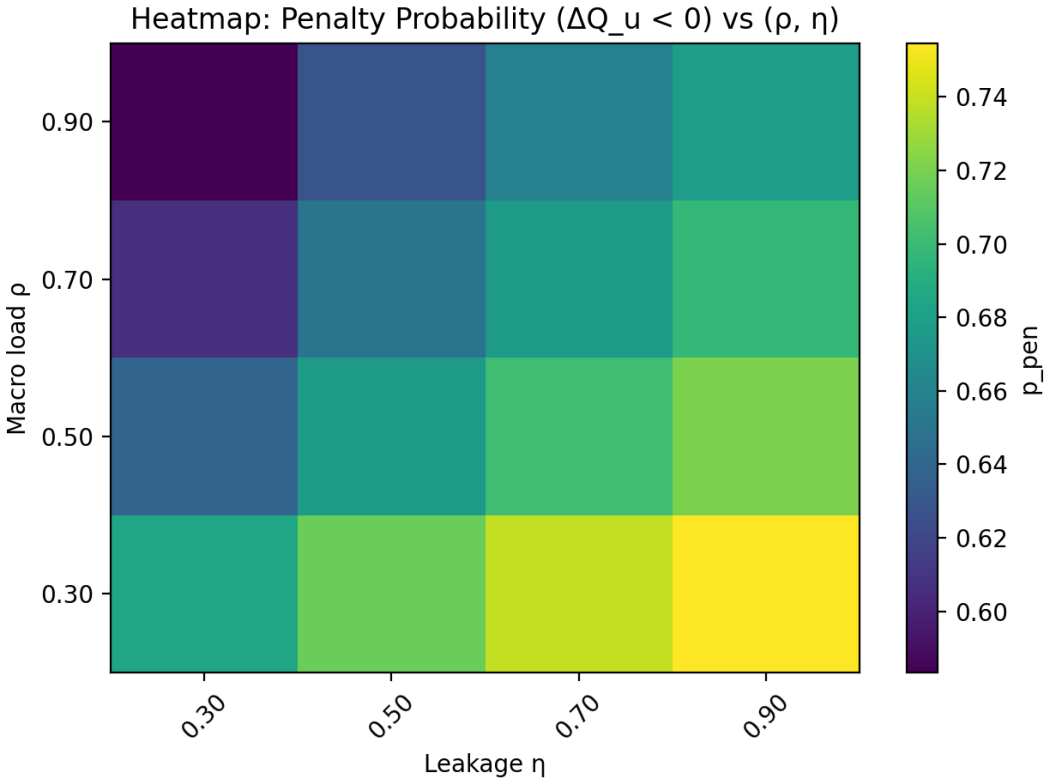


Figure 5.16: Random admission: penalty probability $p_{\text{pen}} = \Pr(\Delta Q_u < 0)$ versus (ρ, η).

Figure 5.16 shows that the penalty probability under Random admission is substantially higher than under Worst admission, indicating that steering uncertainty can expose many offloaded users to degradation. The highest p_{pen} values occur under larger leakage η, where stronger macro-to-NHO coupling increases the likelihood that offloaded users do not benefit despite being admitted.

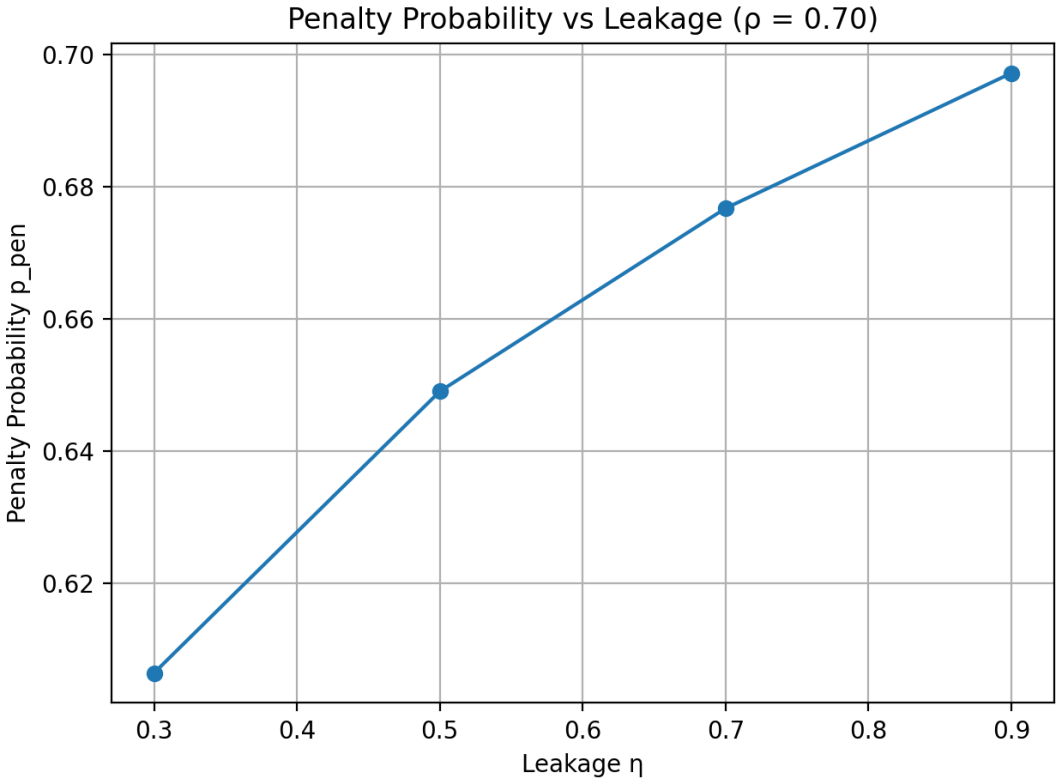


Figure 5.17: Random admission: penalty probability versus leakage η at fixed macro load $\rho = 0.70$.

Figure 5.17 shows that the penalty probability increases steeply with leakage η under Random admission, confirming that cross-tier coupling strongly amplifies degradation risk when steering is uncertain.

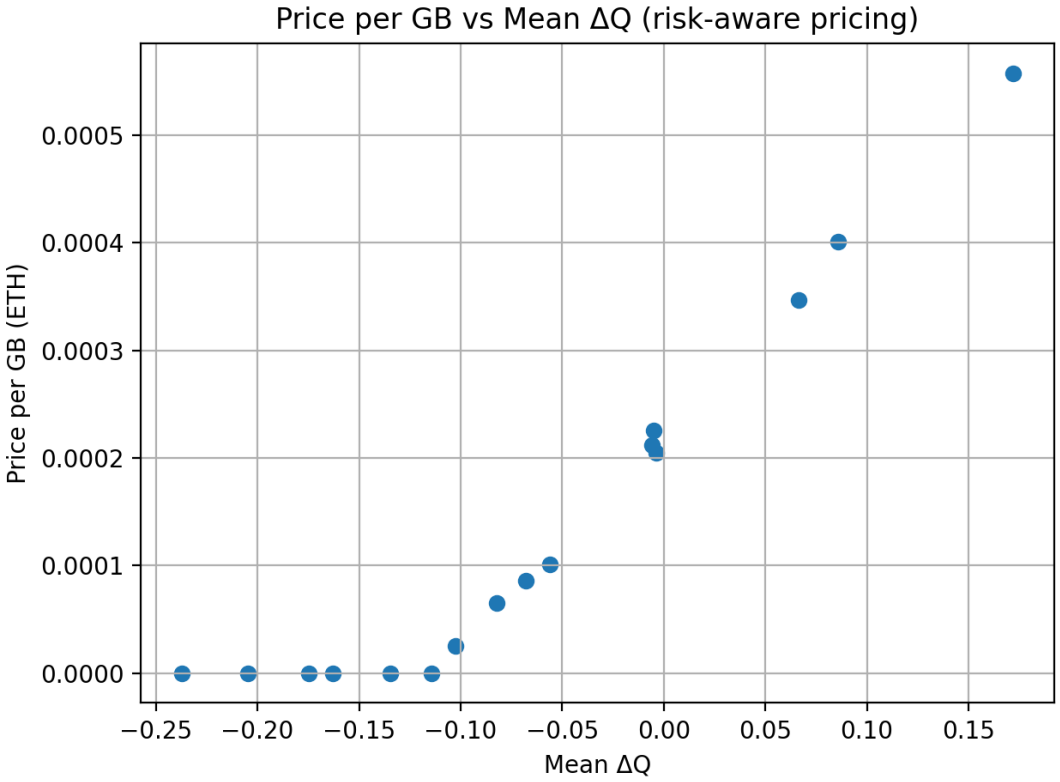


Figure 5.18: Random admission: price per GB versus mean gain ΔQ under risk-aware settlement.

Figure 5.18 shows that the unit price is sharply discounted in Random admission when the penalty probability is high. Since the contract pricing is driven by the risk-adjusted effective gain $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$, large p_{pen} can push Q_{eff} negative even when the mean gain ΔQ is near zero, thereby reducing the settlement price. This confirms the mechanism is accountability-driven: unstable offloading is permitted but economically discouraged.

Worst admission highlights the best-case benefit of steering the macro tail to the neutral host, while Random admission captures uncertainty and exposes penalty regimes. The settlement mechanism responds accordingly by coupling pricing to the risk-adjusted gain $Q_{\text{eff}} = \Delta Q - \lambda p_{\text{pen}}$, thereby aligning neutral-host revenue with delivered performance and degradation risk.

5.6 Summary

This chapter presented a QoS-aware, usage-based settlement framework for indoor neutral-host offloading in multi-operator deployments. Using a system-level propagation abstraction, distributional improvements in received power, SINR, and achievable downlink throughput were demonstrated for neutral-host indoor service compared to a macro O2I baseline. Radio performance was then coupled to economic accountability by computing gain and degradation risk for the offloaded subset and enforcing settlement via an Ethereum-compatible smart contract that forwards payments to the neutral-host wallet with immutable audit trails. Sensitivity analysis across macro load ρ and leakage η under Worst and Random admission policies showed that informed steering can deliver robust gains, while steering uncertainty can expose penalty regimes. Risk-aware pricing addresses this by discounting unstable regimes through the effective gain term.

Chapter 6

Conclusions and Future Directions

6.1 Summary of the Thesis

This thesis investigated how blockchain technology, smart contracts, and learning-assisted demand analysis can be combined to support more flexible, trusted, and accountable coordination and settlement mechanisms for wireless spectrum management. The work was motivated by three interconnected limitations of current approaches. First, conventional licensed spectrum allocation relies on long-term, static assignment structures that do not reflect the short-term and location-dependent variations in actual network demand, leading to significant underutilisation of valuable spectrum resources. Second, while dynamic spectrum access and spectrum trading models have been proposed to improve flexibility, many existing approaches do not fully address trusted authorisation, multi-party coordination, auditable settlement, and implementation-level accountability. Third, in shared indoor and neutral-host environments, settlement between operators remains largely static and decoupled from the radio performance actually delivered to users, weakening accountability and misaligning financial incentives.

To address these limitations, the thesis made three connected contributions: dynamic spectrum trading, demand-aware coordination, and performance-aware settlement. These contributions were developed progressively across Chapters 3, 4, and 5, and are summarised below.

6.1.1 Blockchain-Based Spectrum Trading Marketplace (Chapter 3)

Chapter 3 presented the design, implementation, and evaluation of a decentralised marketplace for wireless spectrum trading based on blockchain technology. The marketplace was built on Ethereum smart contracts and deployed on the Sepolia test network under a proof-of-stake consensus mechanism. The architecture defined the roles of three principal categories of entities: the Telecommunications Regulatory Authority as administrator, spectrum holders as providers of licensed spectrum, and spectrum buyers as consumers seeking to acquire spectrum on a flexible basis.

A market mechanism was developed comprising a dynamic posted-price function that adapts to local supply and demand conditions through a scarcity multiplier and a duration adjustment factor, a greedy rule-based batch-matching heuristic with deterministic behaviour, and three feasibility conditions covering compatibility, capacity, and price acceptability. Policy-based coexistence rules were specified at the market level to express spatial, temporal, and frequency-domain separation among neighbouring allocations. These elements were specified at the design level, while the deployed contract implemented the listing, purchase, and settlement lifecycle. The smart contract implementation comprised four core functions covering the complete lifecycle of a spectrum trade, supported by Python automation scripts for off-chain listing management and data extraction.

The experimental evaluation demonstrated that spectrum transactions can be executed within seconds at the coordination layer, with ownership transfer, payment settlement, and state updates performed atomically within a single transaction. The security evaluation confirmed that the smart contract enforces single allocation and rejects incorrect payments, and that front-running can change which buyer succeeds but cannot cause double allocation or loss of funds; static analysis raised no findings across one hundred detectors. The analytical scalability model showed that throughput is bounded by the block gas limit and block interval, which are shared with other users of the chain, and that longer leases, and regional partitioning on independent execution layers, can relieve congestion.

6.1.2 Blockchain-Enabled Dynamic Spectrum Trading (Chapter 4)

Chapter 4 extended the marketplace of Chapter 3 into a complete, demand-aware, and network-integrated framework for next-generation wireless networks.

The chapter addressed two limitations of the standalone marketplace: the absence of data-driven evidence for spectrum trading opportunities, and the lack of integration with an operational radio access network.

First, an empirical analysis of fine-grained urban cellular activity from the Milan city dataset provided direct evidence of spatio-temporal spectrum underutilisation. provided demand-side evidence of spatio-temporal variation consistent with spectrum underutilisation. A normalised demand-based utilisation index, used as a proxy rather than a measurement of spectrum occupancy, was derived to quantify how frequently demand falls below peak levels, and a two-layer k-means clustering procedure exposed the spatial heterogeneity that creates localised trading opportunities, both across and within regions. These results established that spectrum demand exhibits sufficient spatial and temporal variation to justify short-term, location-aware spectrum trading.

Second, an ANN-based traffic prediction model was introduced to enable the shift from reactive to proactive spectrum trading. By forecasting low-demand intervals using time-of-day and day-of-week features, and by comparing a full model against an ablated variant, the chapter illustrated how predictive intelligence can be coupled with the executable marketplace so that trading decisions anticipate availability rather than merely react to it.

Third, the framework was integrated with an OpenAirInterface-based 5G testbed, demonstrating that a completed on-chain spectrum trade can trigger automated reconfiguration of the gNB and UE without modifying the underlying PHY/MAC mechanisms. Because the testbed ran in RF-simulator mode on a single machine, this validated the path from an on-chain trade to the configuration of a 5G stack, rather than over-the-air radio performance.

6.1.3 Blockchain-Based Settlement for Indoor Neutral-Host Offloading (Chapter 5)

Chapter 5 extended the blockchain-based coordination logic beyond spectrum trading into indoor neutral-host and Open RAN environments, addressing the problem of performance-aware settlement between MNOs and neutral-host operators. The chapter introduced an end-to-end QoS-aware automated settlement framework that explicitly links radio-layer quality of service to economic charging.

A Monte Carlo system-level evaluation demonstrated distributional improvements in received power, SINR, and achievable downlink throughput for neutral-host indoor service compared to a macro outdoor-to-indoor baseline. A throughput-centric QoS metric was defined, and a reward–penalty pricing mechanism was

developed that adjusts the price per GB based on the measured QoS gain and the probability of per-user degradation. Settlement was enforced on-chain via an Ethereum smart contract that recomputes the risk-aware price, verifies the payment, and forwards it to the neutral-host wallet with an immutable audit trail. The radio results are simulation-based, the KPI inputs to the contract are trusted off-chain data, and the O-RAN integration is presented as a target architecture.

Sensitivity analysis under varying macro load and cross-tier leakage conditions, evaluated under both Worst (reward-biased) and Random (stress-test) admission policies, showed that informed steering can deliver robust gains, while steering uncertainty exposes penalty regimes. The risk-aware pricing mechanism responds by discounting unstable regimes through the effective gain term, thereby aligning neutral-host revenue with delivered performance and degradation risk.

6.2 Revisiting the Research Questions

The thesis was guided by five research questions set out in Chapter 1. This section discusses how each has been addressed through the theoretical development, implementation, and evaluation carried out in the thesis.

RQ1: How can licensed spectrum coordination be made more flexible and responsive to short-term, location-dependent demand variation? This question was addressed primarily in Chapters 3 and 4. Chapter 3 designed a marketplace that enables spectrum holders to list and sell underutilised spectrum to buyers on a flexible, on-demand basis, showing that the on-chain part of a trade completes within seconds, compared with the weeks-long administrative processes of conventional licensing Chapter 4 provided empirical evidence from the Milan city dataset that short-term, location-dependent trading opportunities genuinely exist, and introduced a two-layer clustering procedure and a normalised utilisation index that expose the spatial and temporal heterogeneity motivating dynamic spectrum coordination.

RQ2: How can blockchain and smart contracts be used to support trusted authorisation, transaction execution, and auditable settlement in dynamic spectrum trading? This question was addressed across all three contribution chapters. Chapter 3 demonstrated that Ethereum smart contracts can enforce single-allocation semantics, validate payments, transfer ownership, and settle transactions atomically, while the Telecommunications Regulatory Authority retains oversight through application-layer access control. The security

evaluation confirmed correct behaviour under replay and payment-manipulation attempts, and showed that front-running can affect transaction ordering but not allocation safety. Chapter 4 retained and applied these mechanisms in the context of demand-driven trading. Chapter 5 extended the same blockchain-based coordination to a different domain, using smart contracts to enforce performance-aware settlement between MNOs and neutral-host operators with immutable audit trails.

RQ3: How can spatio-temporal traffic analysis and learning-assisted prediction improve the identification of proactive spectrum trading opportunities? This question was addressed in Chapter 4. The two-layer k-means clustering analysis of the Milan dataset revealed persistent spatial heterogeneity in traffic demand at both the city and sub-regional levels. Temporal analysis showed significant daily fluctuations in utilisation. Together, these analyses identified the recurring patterns that create spectrum trading opportunities. The ANN-based prediction model demonstrated that low-demand intervals can be forecast in advance, which would enable proactive spectrum trading. The forecast was illustrative and was not connected to the smart contract in a closed loop.

RQ4: How can interference and regulatory feasibility constraints be integrated into dynamic spectrum coordination mechanisms? This question was addressed in Chapters 3 and 4. Chapter 3 formalised spatial, temporal, and frequency-domain conflict detection through a deterministic interference indicator, and introduced a minimum guard-band constraint for adjacent channel protection. These policy-based feasibility conditions were specified as part of the matching heuristic, so that only allocations complying with the separation rules would be instantiated on-chain. They are rule-based checks rather than physical interference guarantees, and were specified at the design level. Post-trade compliance monitoring was also described as a complementary safeguard. Chapter 4 retained these constraints in the demand-driven setting, where the increased frequency of trade evaluation under predictive triggering makes interference-aware feasibility particularly important.

RQ5: How can blockchain-based coordination be extended to indoor neutral-host and O-RAN environments to support performance-aware settlement? This question was addressed in Chapter 5. The chapter demonstrated that the blockchain-based coordination logic developed for spectrum trad-

ing can be adapted to a fundamentally different coordination problem: settlement between MNOs and neutral-host operators in shared indoor deployments. By defining a throughput-centric QoS metric, deriving a risk-aware effective gain that accounts for both mean performance improvement and per-user degradation probability, and enforcing settlement on-chain with automated payment forwarding, the chapter showed that blockchain can support performance-aware accountability in neutral-host environments. The sensitivity analysis under Worst and Random admission policies confirmed that the pricing mechanism appropriately rewards beneficial offloading and discounts uncertain or degraded regimes.

6.3 Limitations

While the thesis has demonstrated the feasibility and value of blockchain-based coordination for spectrum trading and settlement, several limitations should be acknowledged.

First, the interference and coexistence handling in Chapters 3 and 4 operates at the market and policy layer rather than at the physical layer. Interference is managed through regulatory-defined separation criteria and guard-band constraints rather than through direct channel measurements or signal-level estimation. While this approach is consistent with policy-driven spectrum sharing frameworks and is well suited to blockchain-based execution, it does not capture the full complexity of real-world propagation conditions, and compliance with these rules does not by itself guarantee interference-free operation.

Second, the blockchain-based marketplace was evaluated on the Ethereum Sepolia test network rather than on the Ethereum mainnet or a production-grade consortium blockchain. While Sepolia provides the same operational characteristics as the mainnet without financial costs, the gas prices, network congestion, and validator behaviour on a test network may differ from production conditions. The scalability analysis presented in Chapter 3 addresses this gap analytically, but its throughput figures are upper bounds, since block capacity is shared with other users of the chain, and empirical validation on a production network remains a direction for future work.

Third, the ANN-based traffic prediction model introduced in Chapter 4 is illustrative rather than embedded in a closed-loop control system. The prediction model was used to demonstrate the concept of proactive trading rather than to serve as a fully optimised forecasting engine. It was evaluated on a single grid, and its ablation also changed the hidden-layer size. In addition, the utilisation index is a demand-based proxy derived from traffic activity rather than a measurement

of spectrum occupancy. More advanced prediction architectures, such as deep recurrent networks or transformer-based models, and their integration into closed-loop trading workflows, represent natural extensions.

Fourth, the 5G testbed integration in Chapter 4 was conducted using OpenAirInterface in RF simulator mode on a single computing platform. While this demonstrates the end-to-end path from an on-chain trade to network configuration, validation on hardware-based testbeds with over-the-air transmission and multiple distributed nodes would strengthen the practical claims.

Fifth, the neutral-host settlement framework in Chapter 5 uses a system-level large-scale propagation abstraction without explicit small-scale fading or full PHY simulation. The Monte Carlo evaluation captures long-term throughput trends and QoS differentiation but does not model instantaneous link adaptation, scheduling, or queueing effects. Moreover, the KPI and usage inputs to the settlement contract are trusted off-chain data that the blockchain cannot verify. Additionally, the settlement framework was evaluated for a single NHO serving multiple MNOs. Multi-NHO and multi-venue scenarios introduce additional challenges for attribution and fairness.

6.4 Future Directions

The contributions of this thesis open several promising avenues for future investigation.

Advanced Interference-Aware Coordination. Future work could incorporate more sophisticated interference modelling into the marketplace, including propagation-aware feasibility checks based on estimated channel conditions or measurement reports from participating operators. This would enable the marketplace to make finer-grained coexistence decisions while preserving the transparency and auditability properties of blockchain-based execution.

Scalable Consortium Blockchain Architectures. The use of the public Ethereum blockchain provides strong security guarantees but introduces constraints on throughput, latency, and transaction cost. Future work could explore consortium or hybrid blockchain architectures tailored to telecom operator environments, where controlled participation, predictable performance, and lower execution overhead can be achieved without sacrificing the auditability and coordination benefits demonstrated in this thesis.

Tighter Integration with Programmable Network Control. The 5G-testbed integration demonstrated in this thesis operates through configuration-file updates following a completed trade. Future work could explore tighter integration with programmable network control platforms, such as O-RAN RAN Intelligent Controllers (RICs) or SMO frameworks, to enable real-time, API-driven spectrum reconfiguration triggered directly by on-chain settlement events. This would further reduce the latency between trade completion and network-side activation.

Richer Multi-Metric Settlement Models. The neutral-host framework in Chapter 5 uses downlink throughput as the sole QoS indicator. Future work could incorporate additional verifiable KPIs, such as outage probability, percentile-based latency constraints, uplink power reduction, or energy efficiency metrics, provided they remain tied to measurable quantities and avoid unverifiable traffic assumptions. Multi-metric settlement models would enable more comprehensive performance accountability in shared infrastructure deployments.

Multi-Operator Fairness and Attribution. When multiple MNOs share the same neutral-host infrastructure, fairness and attribution become important considerations. Future work could develop mechanisms for accounting separately for each MNO's usage and performance, ensuring that settlement outcomes are equitable and that cross-operator interference effects are appropriately attributed.

Privacy-Preserving Settlement. The on-chain settlement layer can be further strengthened through privacy-preserving commitments and event-only logging to improve scalability while maintaining auditability. Techniques such as zero-knowledge proofs or homomorphic encryption could enable operators to verify settlement correctness without revealing commercially sensitive information about their traffic volumes or pricing arrangements to other participants.

Bibliography

[1] A. Gupta, R. Kausar, S. Tanwar, A. Alabdulatif, V. Vimal, and S. Aluvala, "Efficient spectrum sharing in 5g and beyond network: a survey," *Telecommunication Systems*, vol. 88, no. 1, p. 29, 2025.

[2] M. B. Weiss, K. Werbach, D. C. Sicker, and C. E. C. Bastidas, "On the application of blockchains to spectrum management," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 2, pp. 193–205, 2019.

[3] M. Massaro, "Next generation of radio spectrum management: Licensed shared access for 5g," *Telecommunications Policy*, vol. 41, no. 5-6, pp. 422–433, 2017.

[4] Y.-C. Liang, *Dynamic spectrum management: From cognitive radio to blockchain and artificial intelligence*. Singapore: Springer Nature, 2020.

[5] S. Mondal, M. P. Dutta, and S. K. Chakraborty, "Dynamic spectrum management and spectrum sharing techniques in 5g networks: A survey," *Journal of Mobile Multimedia*, vol. 20, no. 5, pp. 997–1038, 2024.

[6] M. Sheharyar, A. Naushad, I. Khan, M. A. Hussain, and F. Khan Kakar, "Study on fixed and dynamic spectrum access models for cellular networks," *Journal of Telecommunications and the Digital Economy*, vol. 10, no. 2, pp. 267–287, 2022.

[7] L. Perera, P. Ranaweera, S. Kusaladharma, S. Wang, and M. Liyanage, "A survey on blockchain for dynamic spectrum sharing," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 1753–1802, 2024.

[8] N. Afraz, F. Wilhelmi, H. Ahmadi, and M. Ruffini, "Blockchain and smart contracts for telecommunications: Requirements vs. cost analysis," *IEEE Access*, vol. 11, pp. 95 653–95 666, 2023.

[9] M. Moussaoui, E. Bertin, and N. Crespi, "Blockchain and smart contracts for telecommunications: The whys and wherefores," in *Blockchain and Smart-Contract Technologies for Innovative Applications*. Springer, 2024, pp. 259–279.

[10] R.-V. Tkachuk, D. Ilie, K. Tutschku, and R. Robert, "A survey on blockchain-based telecommunication services marketplaces," *IEEE Transactions on Network and Service Management*, vol. 19, no. 1, pp. 228–255, 2022.

[11] A. Arnaz, J. Lipman, M. Abolhasan, and M. Hiltunen, "Toward integrating intelligence and programmability in open radio access networks: A comprehensive survey," *Ieee Access*, vol. 10, pp. 67 747–67 770, 2022.

[12] A. Kaur and K. Kumar, "A comprehensive survey on machine learning approaches for dynamic spectrum access in cognitive radio networks," *Journal of Experimental & Theoretical Artificial Intelligence*, vol. 34, no. 1, pp. 1–40, 2022.

[13] E. Lykakis, I. O. Vardiambasis, and E. Kokkinos, "Data traffic prediction for 5g and beyond: Emerging trends, challenges, and future directions: A scoping review," *Electronics*, vol. 14, no. 23, p. 4611, 2025.

[14] F. R. Guimarães, J. M. B. da Silva Jr, C. Casimiro Cavalcante, G. Fodor, M. Bengtsson, and C. Fischione, "Machine learning for spectrum sharing: A survey," *Foundations and Trends in Networking*, vol. 14, no. 1-2, pp. 1–159, 2024.

[15] S. Wijethilaka and M. Liyanage, "Survey on network slicing for internet of things realization in 5g networks," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 2, pp. 957–994, 2021.

[16] N. Slamnik-Kriještorac, H. Kremo, M. Ruffini, and J. M. Marquez-Barja, "Sharing distributed and heterogeneous resources toward end-to-end 5g networks: A comprehensive survey and a taxonomy," *IEEE Communications Surveys & Tutorials*, vol. 22, no. 3, pp. 1592–1628, 2020.

[17] M. Saunders, P. Lewis, and A. Thornhill, *Research Methods for Business Students*, 8th ed. Harlow, UK: Pearson, 2019.

[18] S. K. Sharma, T. E. Bogale, L. B. Le, S. Chatzinotas, X. Wang, and B. Ottersten, "Dynamic spectrum sharing in 5g wireless networks with

full-duplex technology: Recent advances and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 1, pp. 674–707, 2017.

[19] M. M. Buddhikot, "Understanding dynamic spectrum access: Models, taxonomy and challenges," in *2007 2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks*. IEEE, 2007, pp. 649–663.

[20] M. Mustonen, "Analysis of recent spectrum sharing concepts in policy making," Ph.D. dissertation, University of Oulu, 2017.

[21] M. Parvini, A. H. Zarif, A. Nouruzi, N. Mokari, M. R. Javan, B. Abbasi, A. Ghasemi, and H. Yanikomeroglu, "Spectrum sharing schemes from 4g to 5g and beyond: Protocol flow, regulation, ecosystem, economic," *IEEE Open Journal of the Communications Society*, vol. 4, pp. 464–517, 2023.

[22] W. Webb, A. Medeisis, and L. F. Minervini, "Evolved spectrum usage rights: A catalyst for liberal spectrum management reform," *Telecommunications Policy*, vol. 48, no. 3, p. 102715, 2024.

[23] P. Ren, Y. Wang, Q. Du, and J. Xu, "A survey on dynamic spectrum access protocols for distributed cognitive wireless networks," *EURASIP Journal on Wireless Communications and Networking*, vol. 2012, no. 1, p. 60, 2012.

[24] P. Kolodzy and I. Avoidance, "Spectrum policy task force," *Federal Commun. Comm., Washington, DC, Rep. ET Docket*, vol. 40, no. 4, pp. 147–158, 2002.

[25] D. Schoolar, "mmwave still remains underutilized as operators thirst for spectrum," 2023.

[26] F. Hu, B. Chen, and K. Zhu, "Full spectrum sharing in cognitive radio networks toward 5g: A survey," *IEEE access*, vol. 6, pp. 15 754–15 776, 2018.

[27] W. S. H. M. W. Ahmad, N. A. M. Radzi, F. S. Samidi, A. Ismail, F. Abdullah, M. Z. Jamaludin, and M. Zakaria, "5g technology: Towards dynamic spectrum sharing using cognitive radio networks," *IEEE access*, vol. 8, pp. 14 460–14 488, 2020.

[28] M. A. Kamal, H. W. Raza, M. M. Alam, M. M. Su'ud, and A. b. A. B. Sajak, "Resource allocation schemes for 5g network: A systematic review," *Sensors*, vol. 21, no. 19, p. 6588, 2021.

[29] J. Jeon, R. D. Ford, V. V. Ratnam, J. Cho, and J. Zhang, "Coordinated dynamic spectrum sharing for 5g and beyond cellular networks," *IEEE Access*, vol. 7, pp. 111 592–111 604, 2019.

[30] W. Salani, L. Mfupe, and O. O. Oyerinde, "Dynamic spectrum allocation in the c-band: An overview," *Applied Sciences*, vol. 15, no. 17, p. 9762, 2025.

[31] M. R. Hassan, G. C. Karmakar, J. Kamruzzaman, and B. Srinivasan, "Exclusive use spectrum access trading models in cognitive radio networks: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2192–2231, 2017.

[32] S. Sengupta and M. Chatterjee, "An economic framework for dynamic spectrum access and service pricing," *IEEE/ACM Transactions on Networking*, vol. 17, no. 4, pp. 1200–1213, 2009.

[33] R. H. Tehrani, S. Vahid, D. Triantafyllopoulou, H. Lee, and K. Moessner, "Licensed spectrum sharing schemes for mobile operators: A survey and outlook," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 4, pp. 2591–2623, 2016.

[34] Federal Communications Commission, "3.5 GHz band / citizens broadband radio service (CBRS)," FCC Part 96 Rules, 2015, updated 2020.

[35] OnGo Alliance, "CBRS deployment status and market update," Industry Report, 2025.

[36] Ofcom, "Enabling wireless innovation through local licensing: Shared access to spectrum supporting mobile technology," Statement, Office of Communications, UK, 2019.

[37] ——, "Evolution of the shared access licence framework," Consultation and Statement, Office of Communications, UK, 2024.

[38] ——, "Shared access licence guidance," Ofcom, Tech. Rep., Jul. 2024, accessed: 2026-06-02. [Online]. Available: https://www.ofcom.org.uk/siteassets/resources/documents/consultations/category-1-10-weeks/129951--enabling-opportunities-for-innovation/associated-documents/shared-access-licence-guidance-july-2024.pdf?v=381319

[39] ETSI, "Reconfigurable radio systems (RRS); system architecture and high level procedures for operation of licensed shared access (LSA) in the 2300–2400 MHz band," ETSI TS 103 235, 2015.

[40] European Commission, "Promoting the shared use of radio spectrum resources in the EU," Commission Communication, 2022.

[41] Bundesnetzagentur, "Administrative regulation for frequency allocations for local frequency utilisation in the frequency range 3,700–3,800 MHz," BNetzA, Germany, 2019.

[42] J. Pastirčák, J. Gazda, and D. Kocur, "A survey on the spectrum trading in dynamic spectrum access networks," in *Proceedings ELMAR-2014*. IEEE, 2014, pp. 1–4.

[43] D. Niyato, E. Hossain, and Z. Han, "Dynamics of multiple-seller and multiple-buyer spectrum trading in cognitive radio networks: A game-theoretic modeling approach," *IEEE Transactions on Mobile Computing*, vol. 8, no. 8, pp. 1009–1022, 2008.

[44] F. Li, Z. Sheng, J. Hua, and L. Wang, "Preference-based spectrum pricing in dynamic spectrum access networks," *IEEE Transactions on Services Computing*, vol. 11, no. 6, pp. 922–935, 2016.

[45] F. Benedetto, L. Mastroeni, and G. Quaresima, "Auction-based theory for dynamic spectrum access: A review," in *2021 44th international conference on telecommunications and signal processing (TSP)*. IEEE, 2021, pp. 146–151.

[46] J. Jia, Q. Zhang, Q. Zhang, and M. Liu, "Revenue generation for truthful spectrum auction in dynamic spectrum access," in *Proceedings of the tenth ACM international symposium on Mobile ad hoc networking and computing*, 2009, pp. 3–12.

[47] X. Zhou and H. Zheng, "Trust: A general framework for truthful double spectrum auctions," in *IEEE INFOCOM 2009*. IEEE, 2009, pp. 999–1007.

[48] H. Wang, E. Dutkiewicz, G. Fang, and M. D. Mueck, "Spectrum sharing based on truthful auction in licensed shared access systems," in *2015 IEEE 82nd Vehicular Technology Conference (VTC2015-Fall)*. IEEE, 2015, pp. 1–5.

[49] R. Abozariba, M. Asaduzzaman, M. Patwary, M. K. Naeem, S. J. Nawaz, and S. K. Sharma, "Secondary spectrum allocation framework via concurrent auctions for 5g and beyond networks," *Wireless Networks*, vol. 28, no. 4, pp. 1489–1504, 2022.

[50] M. Khadem, M. Ansarifard, N. Mokari, M. R. Javan, H. Saeedi, and E. A. Jorswieck, "Dynamic fairness-aware spectrum auction for enhanced licensed shared access in uav-based networks," *IEEE Transactions on Communications*, vol. 73, no. 5, pp. 3076–3092, 2024.

[51] A. Basaure, "Regulation for dynamic spectrum management," Ph.D. dissertation, Pontificia Universidad Catolica de Chile (Chile), 2016.

[52] W. Lehr, R. Berry, C. Caicedo, I. Kadota, K. Mu, Z. Xie, and I. Tamim, "Automating spectrum sharing from the ground up," 2024.

[53] W. Lehr, R. Berry, I. Kadota, C. E. C. Bastidas, K. Mu, Z. Xie, and I. Tamim, "Automating spectrum sharing: A bottom-up approach and research agenda," *Telecommunications Policy*, vol. 49, no. 5, p. 102937, 2025.

[54] K. Yue, Y. Zhang, Y. Chen, Y. Li, L. Zhao, C. Rong, and L. Chen, "A survey of decentralizing applications via blockchain: The 5g and beyond perspective," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 4, pp. 2191–2217, 2021.

[55] S. A. Boamah, J. D. Louis, J. McNair, and J. E. Gilbert, "A survey for blockchain-enabled resource management in edge networks," *IEEE Access*, 2025.

[56] M. Wu, X. Ling, J. Wang, Y. Le, K. Huang, B. Cao, Y. Huang, D. Niyato, Z. Ding, and X. You, "Blockchain-driven resource management in wireless communications and networks: Models, approaches, and applications," *IEEE Communications Surveys & Tutorials*, 2025.

[57] S. Jiang, J. Cao, H. Wu, Y. Yang, M. Ma, and J. He, "Blochie: a blockchain-based platform for healthcare information exchange," in *2018 ieee international conference on smart computing (smartcomp)*. IEEE, 2018, pp. 49–56.

[58] X. Zhang and X. Chen, "Data security sharing and storage based on a consortium blockchain in a vehicular ad-hoc network," *Ieee Access*, vol. 7, pp. 58 241–58 254, 2019.

[59] Y. Wang, Z. Su, Q. Xu, and N. Zhang, "Contract based energy blockchain for secure electric vehicles charging in smart community," in *2018 IEEE 16th Intl Conf on Dependable, Autonomic and Secure Computing, 16th Intl Conf on Pervasive Intelligence and Computing, 4th Intl Conf on Big Data*

Intelligence and Computing and Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech). Ieee, 2018, pp. 323–327.

[60] Y. Zhang and J. Wen, "The iot electric business model: Using blockchain technology for the internet of things," *Peer-to-Peer Networking and Applications*, vol. 10, no. 4, pp. 983–994, 2017.

[61] S. Wang, L. Ouyang, Y. Yuan, X. Ni, X. Han, and F.-Y. Wang, "Blockchain-enabled smart contracts: architecture, applications, and future trends," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, vol. 49, no. 11, pp. 2266–2277, 2019.

[62] Q. Mamun, Z. Pan, and J. Wu, "Blockchain in communication networks: A comprehensive review," *IET Blockchain*, vol. 6, no. 1, p. e70031, 2026.

[63] A. A. Okon, I. Elgendi, O. S. Sholiyi, J. M. Elmirghani, A. Jamalipour, and K. Munasinghe, "Blockchain and sdn architecture for spectrum management in cellular networks," *Ieee Access*, vol. 8, pp. 94 415–94 428, 2020.

[64] G. Praveen, V. Chamola, V. Hassija, and N. Kumar, "Blockchain for 5g: A prelude to future telecommunication," *Ieee Network*, vol. 34, no. 6, pp. 106–113, 2020.

[65] V. Deval, V. K. Dwivedi, A. Dixit, A. Norta, S. A. Shah, R. Sharma, and D. Draheim, "Mobile smart contracts: exploring scalability challenges and consensus mechanisms," *IEEE Access*, vol. 12, pp. 34 265–34 288, 2024.

[66] M. Tahir, M. H. Habaebi, M. Dabbagh, A. Mughees, A. Ahad, and K. I. Ahmed, "A review on application of blockchain in 5g and beyond networks: Taxonomy, field-trials, challenges and opportunities," *IEEE access*, vol. 8, pp. 115 876–115 904, 2020.

[67] A. Chaer, K. Salah, C. Lima, P. P. Ray, and T. Sheltami, "Blockchain for 5g: Opportunities and challenges," in *2019 IEEE Globecom Workshops (GC Wkshps)*. IEEE, 2019, pp. 1–6.

[68] Ofcom, "How blockchain technology could help to manage UK telephone numbers," https://uk.pcmag.com/encryption/117859/uk-regulator-wants-to-manage-phone-numbers-with-blockchain, 2018, accessed: 2026-06-02.

[69] FCC, "FCC's Rosenworcel talks up 6G," https://www.nexttv.com/news/fccs-rosenworcel-talks-up-6g, 2018, accessed: Nov. 22, 2024.

[70] S. Han and X. Zhu, "Blockchain based spectrum sharing algorithm," in *2019 IEEE 19th International Conference on Communication Technology (ICCT)*. IEEE, 2019, pp. 936–940.

[71] M. K. Luka, O. U. Okereke, E. E. Omizegba, and E. C. Anene, "Blockchains for spectrum management in wireless networks: A survey," *arXiv preprint arXiv:2107.01005*, 2021.

[72] A. Okon, N. Jagannath, I. Elgendi, J. M. Elmirghani, A. Jamalipour, and K. Munasinghe, "Blockchain-enabled multi-operator small cell network for beyond 5g systems," *Ieee Network*, vol. 34, no. 5, pp. 171–177, 2020.

[73] Z. Ullah, F. Al-Turjman, N. Saeed, and A. Coronato, "A comprehensive survey of blockchain applications in communication networks," Research Square preprint rs-8622215/v1, 2026.

[74] Z. Li, W. Wang, and Q. Wu, "Blockchain-based dynamic spectrum sharing for 5g and beyond wireless communications," in *International Conference on Blockchain and Trustworthy Systems*. Springer, 2020, pp. 575–587.

[75] K. Kotobi and S. G. Bilén, "Blockchain-enabled spectrum access in cognitive radio networks," in *2017 wireless telecommunications symposium (WTS)*. IEEE, 2017, pp. 1–6.

[76] X. Shen, M. Li, J. Yang, and J. Yi, "Research on dynamic spectrum sharing in the internet of vehicles based on blockchain and game theory," *Sensors*, vol. 26, no. 4, p. 1190, 2026.

[77] Z. Yu, Z. Chang, T. Mikkonen, V. Frascolla, and S. Mumtaz, "Blockchain cooperative spectrum management for wi-fi and lte-unlicensed coexistence networks," *IEEE Transactions on Communications*, 2026.

[78] A. Veloso, J. Sousa, D. Abreu, F. Saraiva, and A. Abelém, "Hybrid architecture for spectrum access systems with permissioned blockchain," *IEEE Access*, 2026.

[79] D. Paroushev and P. Ruskov, "Blockchain-enabled dynamic spectrum usage rights: A framework for 5g/6g spectrum allocation and pricing," in *2024 32nd National Conference with International Participation (TELECOM)*, 2024, pp. 1–4.

[80] B. R. Kunthu, R. N. Taware, and S. K. Anumula, "Blockchain-anchored audit trail model for transparent inter-operator settlement," 2025.

[81] E. Zeydan, J. Mangues-Bafalluy, S. S. Arslan, Y. Turk, and K. Antevski, "Analysis and performance evaluation of blockchain consensus mechanisms for network sharing," *Distributed Ledger Technologies: Research and Practice*, vol. 5, no. 2, pp. 1–25, 2026.

[82] D. G. S. Pivoto, F. A. de Figueiredo, C. Cavdar, G. R. de Lima Tejerina, and L. L. Mendes, "A comprehensive survey of machine learning applied to resource allocation in wireless communications," *IEEE Communications Surveys & Tutorials*, 2025.

[83] A. C. Cullen, B. I. Rubinstein, S. Kandeepan, B. Flower, and P. H. Leong, "Predicting dynamic spectrum allocation: a review covering simulation, modelling, and prediction." *Artificial Intelligence Review*, vol. 56, no. 10, p. 10921, 2023.

[84] A. Gbenga-Ilori, A. L. Imoize, K. Noor, and P. O. Adebolu-Ololade, "Artificial intelligence empowering dynamic spectrum access in advanced wireless communications: A comprehensive overview," *AI*, vol. 6, no. 6, p. 126, 2025.

[85] A. Azzouni and G. Pujolle, "A long short-term memory recurrent neural network framework for network traffic matrix prediction," *arXiv preprint arXiv:1705.05690*, 2017.

[86] A. Furno, M. Fiore, R. Stanica, C. Ziemlicki, and Z. Smoreda, "A tale of ten cities: Characterizing signatures of mobile traffic in urban areas," *IEEE Transactions on Mobile Computing*, vol. 16, no. 10, pp. 2682–2696, 2016.

[87] F. Xu, Y. Lin, J. Huang, D. Wu, H. Shi, J. Song, and Y. Li, "Big data driven mobile traffic understanding and forecasting: A time series approach," *IEEE transactions on services computing*, vol. 9, no. 5, pp. 796–805, 2016.

[88] S. P. Sone, "Proactive resource allocation in spectrum sharing with radar systems via ml-based wireless network time series forecasting," Ph.D. dissertation, University of Oulu, 2022.

[89] M. Ozturk, M. Gogate, O. Onireti, A. Adeel, A. Hussain, and M. A. Imran, "A novel deep learning driven, low-cost mobility prediction approach for 5g cellular networks: The case of the control/data separation architecture (cdsa)," *Neurocomputing*, vol. 358, pp. 479–489, 2019.

[90] M. Ozturk, M. Akram, S. Hussain, and M. A. Imran, "Novel qos-aware proactive spectrum access techniques for cognitive radio using machine learning," *IEEE Access*, vol. 7, pp. 70 811–70 827, 2019.

[91] S. Iyer, A. Kalla, O. A. López, and C. De Alwis, *Intelligent Spectrum Management: Towards 6G*. John Wiley & Sons, 2024.

[92] A. M. Jaradat, M. Alayedi, and H. Arslan, "A survey of radio resource scheduling for 6g and future wireless networks," *IEEE Open Journal of the Communications Society*, vol. 6, pp. 10 191–10 218, 2025.

[93] G. J. Sutton, J. Zeng, R. P. Liu, W. Ni, D. N. Nguyen, B. A. Jayawickrama, X. Huang, M. Abolhasan, Z. Zhang, E. Dutkiewicz *et al.*, "Enabling technologies for ultra-reliable and low latency communications: From phy and mac layer perspectives," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2488–2524, 2019.

[94] N. Jawad, "Mobile edge cloud: Intelligent deployment and services for 5g indoor network," Ph.D. dissertation, Brunel University London, 2021.

[95] G. Brown, "Internet offload for mobile operators," 2011.

[96] A. De la Oliva, C. J. Bernardos, M. Calderon, T. Melia, and J. C. Zuniga, "Ip flow mobility: smart traffic offload for future wireless networks," *IEEE Communications Magazine*, vol. 49, no. 10, pp. 124–132, 2011.

[97] K. Samdanis, T. Taleb, and S. Schmid, "Traffic offload enhancements for eutran," *IEEE Communications Surveys & Tutorials*, vol. 14, no. 3, pp. 884–896, 2011.

[98] Y. Khadraoui, X. Lagrange, and A. Gravey, "A survey of available features for mobile traffic offload," in *European Wireless 2014; 20th European Wireless Conference*. VDE, 2014, pp. 1–4.

[99] S. Husnjak, D. Peraković, and I. Forenbacher, "Data traffic offload from mobile to wi-fi networks: Behavioural patterns of smartphone users," *Wireless communications and mobile computing*, vol. 2018, no. 1, p. 2608419, 2018.

[100] W. Tang, C. Wu, L. Qi, X. Zhang, X. Xu, and W. Dou, "A WiFi-aware method for mobile data offloading with deadline constraints," *Concurrency and Computation: Practice and Experience*, vol. 33, no. 7, p. e5318, 2021.

[101] S. Hao and L. Duan, "To help or disturb: introduction of crowdsourced wifi to 5g networks," *IEEE Transactions on Mobile Computing*, vol. 22, no. 9, pp. 5583–5596, 2022.

[102] N. Bhooanusas, S.-I. Sou, and K.-C. Cheng, "Satisfaction-based dynamic bandwidth reallocation for multipath mobile data offloading," *Computer Networks*, vol. 185, p. 107594, 2021.

[103] P. Sharma, S. Shukla, and A. Vasudeva, "Data offloading via optimal target set selection in opportunistic networks," *Mobile Networks and Applications*, vol. 26, no. 3, pp. 1270–1280, 2021.

[104] E. Lima, A. Aguiar, P. Carvalho, and A. C. Viana, "Human mobility support for personalised data offloading," *IEEE Transactions on Network and Service Management*, vol. 19, no. 2, pp. 1505–1520, 2022.

[105] B. Majumder and T. Venkatesh, "Mobile data offloading based on minority game theoretic framework," *Wireless Networks*, vol. 28, no. 7, pp. 2967–2982, 2022.

[106] S. M. Abd El-atty and Z. Gharsseldien, "Mobile traffic offloading in heterogeneous networks-based small cell technology," *Arabian Journal for Science and Engineering*, vol. 41, no. 2, pp. 555–567, 2016.

[107] S. Shah, S. Kittipiyakul, Y. Lim, and Y. Tan, "Collaborative resource management for negotiable multi-operator small cell networks," *Sensors*, vol. 19, no. 16, p. 3550, 2019.

[108] J. R. Palathinkal, M. I. Rochman, V. Sathya, M. Yavuz, and M. Ghosh, "Indoor neutral-host networks over shared spectrum and shared infrastructure: A comparison study of real-world deployments," 2026.

[109] E. J. Oughton and W. Lehr, "Surveying 5g techno-economic research to inform the evaluation of 6g wireless technologies," *IEEE Access*, vol. 10, pp. 25 237–25 257, 2022.

[110] N. Weerasinghe, T. Hewa, M. Dissanayake, M. Ylianttila, and M. Liyanage, "Blockchain-based roaming and offload service platform for local 5g operators," in *2021 IEEE 18th Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 2021, pp. 1–6.

[111] P. Fernando, L. Gunawardhana, W. Rajapakshe, M. Dananjaya, T. Gamage, and M. Liyanage, "Blockchain-based wi-fi offloading platform for 5g," in *2020 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2020, pp. 1–6.

[112] V. Hassija, V. Chamola, V. Gupta, and G. S. Chalapathi, "A blockchain based framework for secure data offloading in tactile internet environment," in *2020 International Wireless Communications and Mobile Computing (IWCMC)*. IEEE, 2020, pp. 1836–1841.

[113] Z. Latif, C. Lee, K. Sharif, and S. Helal, "Sdblockedge: Sdn-blockchain enabled multihop task offloading in collaborative edge computing," *IEEE Sensors Journal*, vol. 22, no. 15, pp. 15 537–15 548, 2022.

[114] A. Naouri, H. Wu, N. A. Nouri, S. Dhelim, and H. Ning, "A novel framework for mobile-edge computing by optimizing task offloading," *IEEE Internet of Things Journal*, vol. 8, no. 16, pp. 13 065–13 076, 2021.

[115] S. Bayraktar, S. Gören, and T. Serif, "Blockchain interoperability for future telecoms," *Telecom*, vol. 6, no. 1, p. 20, 2025.

[116] T. Rathod, N. K. Jadav, M. D. Alshehri, S. Tanwar, R. Sharma, R.-A. Felseghi, and M. S. Raboaca, "Blockchain for future wireless networks: A decade survey," *Sensors*, vol. 22, no. 11, p. 4182, 2022.

[117] D. Cuellar, M. Sallal, and C. Williams, "Bsm-6g: Blockchain-based dynamic spectrum management for 6g networks: Addressing interoperability and scalability," *Ieee Access*, vol. 12, pp. 59 643–59 664, 2024.

[118] B. Ma, W. Guo, and J. Zhang, "A survey of online data-driven proactive 5g network optimisation using machine learning," *IEEE access*, vol. 8, pp. 35 606–35 637, 2020.

[119] D. Janu, K. Singh, and S. Kumar, "Machine learning for cooperative spectrum sensing and sharing: A survey," *Transactions on Emerging Telecommunications Technologies*, vol. 33, no. 1, p. e4352, 2022.

[120] E. Dritsas and M. Trigka, "Machine learning in information and communications technology: A survey," *Information*, vol. 16, no. 1, p. 8, 2024.

[121] H. H. Mahmoud, D. Mi, A. Aneiba, M. Idrissi, G. Sharma, C. H. Foh, and P. Xiao, "A review of open-ran intelligence: Opportunities, challenges, real-life applications and impacts," *ACM Computing Surveys*, vol. 58, no. 8, pp. 1–36, 2026.

[122] L. Zhao, M. M. Bala, W. Gang, P. Chengkang, Y. Yannan, T. Zhigang, Y.-C. Tseng, C. Xiang, B. Shen *et al.*, "A survey on open-source-defined

wireless networks: Framework, key technology, and implementation," *arXiv preprint arXiv:2209.01891*, 2022.

[123] A. Kumar and A. Pal, "A survey on computation offloading and current trends," *IEEE Access*, 2025.

[124] P. Wang, J. Meng, J. Chen, T. Liu, Y. Zhan, W.-T. Tsai, and Z. Jin, "Smart contract-based negotiation for adaptive qos-aware service composition," *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 6, pp. 1403–1420, 2018.

[125] T. Barnett, S. Jain, U. Andra, and T. Khurana, "Cisco visual networking index (vni) complete forecast update, 2017–2022," *Americas/EMEAR Cisco Knowledge Network (CKN) Presentation*, vol. 1, no. 1, p. 2018, 2018.

[126] Ofcom UK, "Shared access licence application form," Ofcom, Tech. Rep., 2024, accessed: 2026-06-02. [Online]. Available: https://www.ofcom.org.uk/siteassets/resources/documents/spectrum/spectrum-information/ofw589-shared-access-licence-application-form.pdf?v=396350

[127] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," https://bitcoin.org/bitcoin.pdf, 2008.

[128] M. K. Luka, O. U. Okereke, E. E. Omizegba, and E. C. Anene, "Blockchains for spectrum management in wireless networks: A survey," *arXiv preprint arXiv:2107.01005*, 2021.

[129] W. Wang and Y. Zhao, "Blockchain-based spectrum management architecture and trading mechanism design for space-air-ground integrated network," *IEEE Communications Letters*, vol. 27, no. 10, pp. 2692–2696, 2023.

[130] S. Zheng, T. Han, Y. Jiang, and X. Ge, "Smart contract-based spectrum sharing transactions for multi-operators wireless communication networks," *IEEE access*, vol. 8, pp. 88 547–88 557, 2020.

[131] H. Zhang and Y. Zhao, "Enhancing economic security in blockchain-based spectrum trading: A stablecoin approach," in *2023 9th International Conference on Computer and Communications (ICCC)*. IEEE, 2023, pp. 2503–2507.

[132] H. Li, P. Gao, Y. Zhan, and M. Tan, "Blockchain technology empowers telecom network operation," *China Communications*, vol. 19, no. 1, pp. 274–283, 2022.

[133] S. T. Muntaha, Q. Z. Ahmed, F. A. Khan, Z. D. Zaharis, and P. I. Lazaridis, "Hybrid blockchain-based multi-operator resource sharing and sla management," *IEEE Open Journal of the Communications Society*, vol. 6, pp. 362–377, 2024.

[134] S. T. Muntaha, P. I. Lazaridis, M. Hafeez, Q. Z. Ahmed, F. A. Khan, and Z. D. Zaharis, "Blockchain for dynamic spectrum access and network slicing: A review," *Ieee Access*, vol. 11, pp. 17 922–17 944, 2023.

[135] L. Xue, W. Yang, W. Chen, and L. Huang, "Stbc: A novel blockchain-based spectrum trading solution," *IEEE Transactions on Cognitive Communications and Networking*, vol. 8, no. 1, pp. 13–30, 2021.

[136] S. Wang, X.-L. Huang, F. Hu, and S. Yu, "A fast-convergence, induced dynamic spectrum access based on accelerated q-learning for cognitive radio networks," *IEEE Transactions on Vehicular Technology*, 2025.

[137] K. Kotobi and S. G. Bilen, "Secure blockchains for dynamic spectrum access: A decentralized database in moving cognitive radio networks enhances security and user access," *ieee vehicular technology magazine*, vol. 13, no. 1, pp. 32–39, 2018.

[138] Y. Le, X. Ling, J. Wang, and Z. Ding, "Prototype design and test of blockchain radio access network," in *2019 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2019, pp. 1–6.

[139] L. Zhang, B. Lee, Y. Ye, and Y. Qiao, "Ethereum transaction performance evaluation using test-nets," in *Euro-Par 2019: Parallel Processing Workshops*, ser. Lecture Notes in Computer Science, vol. 11997. Cham: Springer, 2020, pp. 179–190.

[140] T. T. A. Dinh, J. Wang, G. Chen, R. Liu, B. C. Ooi, and K.-L. Tan, "Blockbench: A framework for analyzing private blockchains," in *Proceedings of the 2017 ACM international conference on management of data*, 2017, pp. 1085–1100.

[141] G. Barlacchi, M. De Nadai, R. Larcher, A. Casella, C. Chitic, G. Torrisi, F. Antonelli, A. Vespignani, A. Pentland, and B. Lepri, "A multi-source dataset of urban life in the city of milan and the province of trentino," *Scientific data*, vol. 2, no. 1, p. 150055, 2015.

[142] N. Nikaein, M. K. Marina, S. Manickam, A. Dawson, R. Knopp, and C. Bonnet, "OpenAirInterface: A flexible platform for 5G research," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 5, pp. 33–38, 2014.

[143] H. Rutagemwa, A. Ghasemi, and S. Liu, "Dynamic spectrum assignment for land mobile radio with deep recurrent neural networks," in *2018 IEEE International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2018, pp. 1–6.

[144] L. Yu, J. Chen, G. Ding, Y. Tu, J. Yang, and J. Sun, "Spectrum prediction based on taguchi method in deep learning with long short-term memory," *IEEE Access*, vol. 6, pp. 45 923–45 933, 2018.

[145] OpenAirInterface Software Alliance, "OpenAirInterface: 5G software alliance for democratising wireless innovation," https://openairinterface.org, 2024, accessed: Nov. 22, 2024.

[146] Ericsson, "Ericsson mobility report," https://www.ericsson.com/4ae12c/assets/local/reports-papers/mobility-report/documents/2023/ericsson-mobility-report-november-2023.pdf, 2023, november 2023. Accessed: May 2025.

[147] T. W. Kaiser, S. Verbrugge, M. Van der Wee, and D. Colle, "An overview of different business models for public wi-fi and their implications on indirect revenue," 2017.

[148] 3GPP, "Study on channel model for frequencies from 0.5 to 100 GHz (3GPP TR 38.901, release 18)," ETSI, Sophia Antipolis, France, Tech. Rep., 2024. [Online]. Available: https://www.etsi.org/deliver/etsi_tr/138900_138999/138901/18.00.00_60/tr_138901v180000p.pdf

[149] ITU-R, "Prediction of building entry loss (recommendation ITU-R P.2109-2)," International Telecommunication Union, Geneva, Switzerland, Tech. Rep., Aug. 2023. [Online]. Available: https://www.itu.int/dms_pubrec/itu-r/rec/p/R-REC-P.2109-2-202308-I!!PDF-E.pdf

[150] 3GPP, "Further advancements for E-UTRA physical layer aspects," 3rd Generation Partnership Project, Technical Report TR 36.814 V9.0.0, 2010.

[151] C. E. Shannon, "A mathematical theory of communication," *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948.