



# **An Exploration of the Spectrum Management Resource Allocation Problem under Realistic Conditions**

by

Bailey Heading

In fulfilment of the requirements for the degree of

**Doctor of Philosophy**

November 2025

Department of Engineering

School of Electrical and Mechanical Engineering

**The University of Adelaide**







# Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

*B. Harding*





# Acknowledgements

I would firstly like to begin by acknowledging the support I have received through my time as a PhD candidate from my primary and co-supervisors Dr. Matthew Sorell and Dr. Madhulika Tripathi, as well as my external industry supervisors Dr. Don Gossink, and Dr. Madhulika Tripathi (again). I also wish to extend my acknowledgements to University of Adelaide and to my industry partner Consunet, as well as the team at Consunet that I have collaborated with during my time on placement as an intern applied scientist at the company.





# Abstract

Radio frequency (RF) *spectrum* is rapidly becoming one of the hottest commodities powering our increasingly digitised societies. Highly contested access to this resource has motivated a need to improve spectral efficiency to meet future demand. Industry challenges in dividing and allocating spectrum among competing users have revealed unclear rules and criteria as a major barrier to progress. A core insight is that applying the limitations imposed by reality may be necessary to refine and solve the problem. The challenge can be characterised as the large-scale coordination of time-sensitive geospatial access to a contested resource under complicated realistic conditions. It is proposed that a *digital twin* model may provide a suitable solution.

Three applications of advanced RF propagation modelling in a 3D environment are showcased to implement a concept referred to as *digital twin* spectrum management. A *digital twin* can be defined as an ultra realistic digital world model, suitable for real-time applications due to providing a sufficiently accurate and detailed understanding of the environment comparable to reality. Application 1 considers the impact of RF user mobility for allocated spectrum, using time-weighted reachability in conjunction with RF propagation modelling to produce a result known as a RF coverage shortest path tree (SPT). Application 2 introduces multi-user coordination via mutual signal-to-interference (SIR) analysis, enabling interference-aware allocation in shared and dynamic environments. Application 3 is *vertical* spectrum management, focusing on the scenarios of drone-to-drone and drone-ground co-interference between flight paths.

Spectrum management is tightly regulated by slow-changing national and international agreements. In general, incumbency is prioritised while new entrants must prove through RF analyses that their access will not interfere. This system provides baseline long term access but fails to exploit short term access opportunities to fill *gaps* and support the advancement of diverse, novel, and innovative applications of spectrum. However, manual RF analyses are too slow for these proposed complex, multi-user scenarios, so advanced and realistic simulations may be required to achieve this. To realistically simulate the resource allocation problem, An understanding of the most impactful real-world factors is required, including the relationship between life patterns and spectrum demand, allocation mechanisms and metrics, the benefits of decentralised management, and RF propagation modelling. With access to this kind of simulation, it becomes feasible to explore the possibilities for both planning and prediction.



# Executive Summary

Radio frequency (RF) *spectrum* is increasingly treated as one of the most valuable resources in modern digital societies. The purpose of this thesis is to define the spectrum management resource allocation problem under realistic conditions. The research originated from an industry challenge of how to divide and allocate spectrum among multiple competing users. While investigating this, a significant barrier emerged. Resource allocation is a technical problem, but the criteria that govern spectrum allocation, and what makes an allocation effective or ineffective, are poorly defined. The difficulty is not in producing a single definition of the problem, but in the abundance of possible definitions and evaluation criteria, many of which lack realism. This insight highlights that real world applicability is the key filter for reducing the problem space to a manageable set of problem definitions. Once refined, these definitions can be meaningfully addressed and tested. The strategy pursued in this thesis is: (1) to identify the aspects of spectrum management that are shaped by real world constraints, and (2) to develop methods to filter relevant problem definitions through modelling or simulation as needed.

Chapter 1 introduces spectrum management as the foundation for understanding how spectrum is allocated and regulated, showing that while oversight has shifted from commercial control to national regulators and the International Telecommunications Union (ITU), the process remains static, centralised and slow, with changes taking years due to complex interference analyses and international coordination. This rigidity produces an “artificial spectrum scarcity” by failing to reflect the spatial and temporal variation in demand. The chapter reviews key research areas including dynamic spectrum management (DSM), cognitive radio (CR) and spatial temporal spectrum prediction that seek to overcome these limits through adaptive, intelligent and predictive approaches. However, it is found that much of this work is tested only under idealised conditions, limiting its real world applicability. It is concluded that scalable, realistic simulations rather than manual measurements and surveying are essential to evaluate and implement dynamic spectrum management (DSM) solutions in real world scenarios.

In Chapter 2, the focus is on modelling dynamic spectrum demand to improve allocation in dynamic spectrum management (DSM). The chapter identifies the shortcomings of static allocation, which fails to reflect the temporal and spatial variability of real-world communication patterns. To address this, it explores how synthetic data can be used to model demand

across large geospatial areas with many users. Several strategies are reviewed, including Markov processes, multi-agent simulations, abstraction methods, modelling life patterns and timetable-driven agent modelling with large language model (LLM)s. These approaches provide ways to generate realistic datasets that link user behaviour to spectrum demand, enabling the testing of allocation strategies. Demonstrations show how region-level Markov chains and agent-level simulations can replicate daily demand cycles, offering scalable and adaptable tools for evaluating practical spectrum management solutions.

In Chapter 3, spectrum allocation mechanisms are examined within the context of dynamic spectrum management (DSM), focusing on how strategies must balance efficiency, fairness, and computational feasibility in managing spectrum access as a scarce and interference-prone resource. Fairness metrics such as Jain's index are highlighted as essential tools for evaluating allocation outcomes. Spatial clustering is explored in the literature, with a focus on Voronoi tessellations to provide a mathematically grounded method for partitioning users spatially. Demonstrations illustrate how these approaches adapt to user mobility and shifting demand, showing both their benefits and limitations. The chapter concludes that Voronoi-based methods are dynamic and scalable strategies for spatially clustering users in dynamic spectrum management (DSM), supporting the allocation process. This finding may be combined with fairness metrics to assess the quality of allocation decisions, though challenges remain in balancing adaptability, computational cost, and real-world constraints.

In Chapter 4, the practical challenges of deploying real-time dynamic spectrum management (DSM) across large and diverse geospatial environments are explored. The chapter highlights constraints such as latency, interference, and infrastructure limitations that complicate large-scale spectrum coordination under realistic conditions. To address these, decentralisation, hierarchical spectrum management, and mobile ad hoc network (MANET)s are explored, with a focus on decentralised clustering protocols that can group spectrum users together, providing a way to coordinate their spectrum access with reduced communication overheads and latency. Techniques such as isochrone and shortest path tree (SPT) modelling are introduced as tools for predicting mobility and assessing user reachability within digital twin environments. Demonstrations show how decentralised clustering and mobility-aware modelling could provide real world systems with genuinely useful information to manage spectrum allocations in real-time. While initial results are promising, it is noted that real-world applications require further evaluation in detailed three-dimensional (3D) simulations with realistic RF propagation and user behaviour modelling.

Chapter 5 addresses the central challenge of spectrum allocation of determining whether simultaneous RF spectrum access by multiple users will cause harmful interference under realistic conditions. The chapter highlights the limitations of distance-only models and emphasises the importance of incorporating environmental context. Advanced radio frequency (RF)



propagation modelling techniques that integrate terrain and urban data into 3D analyses are explored. Key approaches include 3D city modelling for dense urban areas, digital elevation model (DEM) integration for terrain-dominated rural and hilly environments, and ray tracing methods for capturing reflection, diffraction, and shadowing effects. By combining these methods with openly available datasets and emerging tools such as Sionna RT, the chapter demonstrates how realistic, scalable simulations can be achieved to support dynamic spectrum management. Four demonstrations show practical applications, each balancing accuracy, data availability, and computational cost, while also indicating areas for future improvement such as hybrid models and enhanced material property data.

In Chapter 6, three applications of advanced radio frequency (RF) propagation modelling in a 3D digital twin environment for spectrum management and resource allocation are presented. Building on the concept of an advanced digital world model, the chapter relates this to the idea of a *digital twin*, defined as an ultra realistic digital world model, suitable for real-time applications due to providing a sufficiently accurate and detailed understanding of the environment comparable to reality. Application 1 combines reachability modelling (via a SPT) with RF propagation modelling to enable interference-aware temporal spectrum planning, constructing a time-weighted radio frequency (RF) coverage SPT that could integrate possible mobility into the spectrum allocation process. Application 2 develops multi-user coordination using mutual signal-to-interference (SIR) analysis to anticipate and mitigate interference risks in complex shared environments, demonstrating how digital twin simulations can shift spectrum management from static two-dimensional allocation to dynamic, data-driven, and mobility-aware strategies. Application 3 extends this approach into *vertical* spectrum management, where altitude is modelled explicitly, focusing on drone-to-drone and drone-to-ground interactions through flight path and 3D RF propagation analysis.

The *Additional Applications* extend the thesis beyond terrestrial spectrum management to highlight the adaptability of digital twin and RF propagation methods in applied domains. The section on digital forensics introduces how telecommunications data, combined with historical digital twins, can reconstruct past network conditions for evidentiary purposes. Two case studies show how forensic analysis requires integrating legacy technologies, environmental reconstruction, and RF modelling to test alibis and device movements under conditions that may no longer exist. The space spectrum section is exploratory and outlines how RF propagation and ray tracing could support analysis of satellite links and potential interference, with a simple Sionna RT demonstration on a low resolution Earth model. Maritime spectrum models vessel clustering, reachability, and hybrid satellite and coastal systems to manage interference and plan coverage along shipping routes. Together, these examples show how digital twins can be adapted for historical reconstruction, extraterrestrial communication, and maritime coordination, broadening the scope of spectrum management research.



# Contents

<b>Abstract</b>	<b>vii</b>
<b>Executive Summary</b>	<b>ix</b>
<b>Acronyms</b>	<b>xx</b>
<b>0 Overview</b>	<b>1</b>
0.1 Research Opening . . . . .	1
0.2 Problem Background . . . . .	1
0.3 Approach . . . . .	3
0.4 Literature . . . . .	4
0.5 Contributions . . . . .	6
<b>1 Introduction to Spectrum Management</b>	<b>9</b>
1.1 Introduction . . . . .	10
1.2 Approach . . . . .	11
1.2.1 Proposition . . . . .	11
1.2.2 Questions . . . . .	11
1.2.3 Topics . . . . .	11
1.3 Literature . . . . .	12
1.3.1 Spectrum Management . . . . .	12
1.3.2 Dynamic Spectrum Management (DSM) . . . . .	14
1.3.3 Cognitive Radio . . . . .	18
1.3.4 Spatio Temporal Spectrum Prediction . . . . .	21
1.4 Background . . . . .	25
1.4.1 Spectrum . . . . .	25
1.4.2 Value . . . . .	26
1.4.3 Multiple Access . . . . .	27
1.4.4 Conventional Management . . . . .	28
1.4.5 Commercial Usage . . . . .	32
1.4.6 Satellite Communications . . . . .	33
1.5 Discussion . . . . .	34



1.5.1	Summary . . . . .	34
1.5.2	Key Findings . . . . .	36
1.5.3	Limitations . . . . .	37
1.5.4	Future Work . . . . .	38
1.6	Next Steps . . . . .	39
<b>2</b>	<b>Modelling Dynamic Spectrum Demand</b>	<b>41</b>
2.1	Introduction . . . . .	42
2.2	Approach . . . . .	43
2.2.1	Proposition . . . . .	43
2.2.2	Question . . . . .	43
2.2.3	Topics . . . . .	43
2.3	Literature . . . . .	44
2.3.1	Markov Chains & Processes . . . . .	44
2.3.2	Multiagent Modelling . . . . .	46
2.3.3	Abstracted Agent Modelling . . . . .	48
2.3.4	Life Patterns & Timetables . . . . .	51
2.4	Demonstrations . . . . .	56
2.4.1	Context . . . . .	56
2.4.2	Summary . . . . .	56
2.4.3	Structure . . . . .	57
2.4.4	Demonstration 1: Region-level Spectrum Demand Model . . . . .	58
2.4.5	Demonstration 2: Agent-level Spectrum Demand Model . . . . .	62
2.5	Discussion . . . . .	70
2.5.1	Summary . . . . .	70
2.5.2	Key Findings . . . . .	71
2.5.3	Limitations . . . . .	72
2.5.4	Future Work . . . . .	73
2.6	Next Steps . . . . .	74
<b>3</b>	<b>Spectrum Allocation Mechanisms</b>	<b>75</b>
3.1	Introduction . . . . .	76
3.2	Approach . . . . .	77
3.2.1	Proposition . . . . .	77
3.2.2	Question . . . . .	77
3.2.3	Topics . . . . .	77
3.3	Literature . . . . .	78
3.3.1	Resource Distribution & Allocation . . . . .	78
3.3.2	Spatial Clustering Techniques . . . . .	83
3.3.3	Fairness Metrics . . . . .	84
3.3.4	Voronoi Tessellations . . . . .	89



3.4	Demonstrations . . . . .	92
3.4.1	Context . . . . .	92
3.4.2	Summary . . . . .	92
3.4.3	Structure . . . . .	93
3.4.4	Demonstration 1: Spectrum Allocation Metrics . . . . .	94
3.4.5	Demonstration 2: Voronoi-Based Clustering . . . . .	97
3.5	Discussion . . . . .	108
3.5.1	Summary . . . . .	108
3.5.2	Key Findings . . . . .	109
3.5.3	Limitations . . . . .	110
3.5.4	Future Work . . . . .	111
3.6	Next Steps . . . . .	112
<b>4</b>	<b>Dynamic Spectrum Allocation in the Real World</b>	<b>113</b>
4.1	Introduction . . . . .	114
4.2	Approach . . . . .	115
4.2.1	Propositions . . . . .	115
4.2.2	Question . . . . .	115
4.2.3	Topics . . . . .	115
4.3	Background . . . . .	116
4.3.1	MANET Protocols . . . . .	116
4.3.2	Isochrones & SPT routing Engines . . . . .	117
4.4	Literature . . . . .	120
4.4.1	Decentralisation and Hierarchical Spectrum Management . . . . .	120
4.4.2	Mobile Ad-hoc Network (MANET) . . . . .	122
4.4.3	MANET Clustering . . . . .	126
4.4.4	Isochrones & Shortest Path Trees (SPTs) . . . . .	129
4.5	Demonstrations . . . . .	132
4.5.1	Context . . . . .	132
4.5.2	Summary . . . . .	133
4.5.3	Structure . . . . .	133
4.5.4	Demonstration 1: Decentralised Clustering Protocols . . . . .	134
4.5.5	Demonstration 2: Visualisation of Decentralised Clustering . . . . .	136
4.5.6	Demonstration 3: Reachability Modelling with Shortest Path Trees . . . . .	141
4.6	Discussion . . . . .	148
4.6.1	Summary . . . . .	148
4.6.2	Key Findings . . . . .	150
4.6.3	Limitations . . . . .	151
4.6.4	Future Work . . . . .	152
4.7	Next Steps . . . . .	153



<b>5</b>	<b>Modelling Radio Spectrum in 3-Dimensional Space</b>	<b>155</b>
5.1	Introduction . . . . .	156
5.2	Approach . . . . .	157
5.2.1	Proposition . . . . .	157
5.2.2	Questions . . . . .	157
5.2.3	Topics . . . . .	157
5.3	Literature . . . . .	158
5.3.1	3D City Modelling . . . . .	158
5.3.2	Digital Elevation Model (DEM) . . . . .	160
5.3.3	Radio Frequency (RF) Ray Tracing . . . . .	162
5.4	Background . . . . .	165
5.4.1	Sionna Differentiable RT . . . . .	165
5.4.2	Radio Materials in Mitsuba 3 . . . . .	166
5.5	Demonstrations . . . . .	167
5.5.1	Context . . . . .	167
5.5.2	Summary . . . . .	167
5.5.3	Tools . . . . .	169
5.5.4	Structure . . . . .	169
5.5.5	Demonstration 1: LOS Mapping from Elevation Data . . . . .	170
5.5.6	Demonstration 2: Ray Tracing Buildings from OSM Data . . . . .	174
5.5.7	Demonstration 3: Ray Tracing Buildings and Terrain from Google 3D Tiles	177
5.5.8	Demonstration 4: Ray Tracing 3D Meshes from Elevation Data . . . . .	181
5.6	Discussion . . . . .	186
5.6.1	Summary . . . . .	186
5.6.2	Key Findings . . . . .	187
5.6.3	Limitations . . . . .	188
5.6.4	Future Work . . . . .	189
5.7	Next Steps . . . . .	190
<b>6</b>	<b>Coordinated Spectrum Access via Digital Twin Simulation</b>	<b>191</b>
6.1	Introduction . . . . .	192
6.2	Approach . . . . .	194
6.2.1	Proposition . . . . .	194
6.2.2	Questions . . . . .	194
6.2.3	Topics . . . . .	194
6.3	Literature . . . . .	195
6.3.1	Digital Spaces, Partial and Complete Digital Twins . . . . .	195
6.3.2	Drones' Spectrum Access . . . . .	198
6.4	Demonstrations . . . . .	201
6.4.1	Context . . . . .	201
6.4.2	Summary . . . . .	202



6.4.3	Tools . . . . .	203
6.4.4	Structure . . . . .	203
6.4.5	Demonstrations 1-3: RF Coverage Reachability . . . . .	204
6.4.6	Demonstration 4: Multi-user Coordinated Spectrum Access . . . . .	218
6.4.7	Demonstration 5: Coordinating Spectrum Access with Drones . . . . .	236
6.5	Discussion . . . . .	251
6.5.1	Summary . . . . .	251
6.5.2	Key Findings . . . . .	252
6.5.3	Limitations . . . . .	253
6.5.4	Future Work . . . . .	254
6.6	Conclusion . . . . .	255
<b>Additional Applications</b>		<b>257</b>
	Digital Forensics . . . . .	257
	Space Spectrum . . . . .	260
	Maritime Spectrum . . . . .	261
<b>Future Work</b>		<b>263</b>
<b>Conclusion</b>		<b>265</b>
<b>Bibliography</b>		<b>267</b>



# Acronyms

<b>2D</b>	two-dimensional.	<b>CBRS</b>	citizens broadband radio service.
<b>3D</b>	three-dimensional.	<b>CDMA</b>	code division multiple access.
<b>3G</b>	third generation.	<b>CDS</b>	connected dominating set.
<b>4G</b>	fourth generation.	<b>CGI</b>	computer generated imagery.
<b>5G</b>	fifth generation.	<b>CNR</b>	carrier-to-noise ratio.
<b>5G-ALLSTAR</b>	fifth generation - agile and flexible integration of satellite and cellular.	<b>CPU</b>	central processing unit.
<b>6G</b>	sixth generation.	<b>CR</b>	cognitive radio.
<b>ACMA</b>	Australian Communications and Media Authority.	<b>CRAHN</b>	cognitive radio ad hoc network.
<b>AI</b>	artificial intelligence.	<b>CRD</b>	connected dominating set routing.
<b>AODV</b>	ad hoc on-demand distance vector.	<b>CRN</b>	cognitive radio network.
<b>API</b>	application programming interface.	<b>CSI</b>	channel state information.
<b>AV</b>	autonomous vehicle.	<b>CTMC</b>	continuous-time Markov chain.
<b>AWGN</b>	additive white gaussian noise.	<b>D2D</b>	device-to-device.
<b>BWP</b>	bandwidth part.	<b>DBSCAN</b>	density-based spatial clustering of applications with noise.
<b>CA</b>	carrier aggregation.	<b>DEM</b>	digital elevation model.
<b>CAD</b>	computer-aided design.	<b>DQN</b>	deep Q-networks.
<b>CBD</b>	city business district.	<b>DRA</b>	distributed resource allocation.
		<b>DRF</b>	dominant resource fairness.
		<b>DRL</b>	deep reinforcement learning.



<b>DRR</b> deficit round robin.	<b>H2BMM</b> high-order hidden bivariate Markov model.
<b>DSA</b> dynamic spectrum access.	<b>HCIS</b> hierarchical cell identification scheme.
<b>DSM</b> dynamic spectrum management.	<b>HF</b> high frequency.
<b>EDGR</b> enhanced directional greedy routing.	<b>HMM</b> hidden Markov model.
<b>EDT</b> energy detection thresholds.	<b>HPC</b> high-performance computing.
<b>EMGR</b> enhanced message-aware greedy routing.	<b>IMT</b> international mobile telecommunications.
<b>EV</b> electric vehicle.	<b>ISAC</b> integrated sensing and communication.
<b>FCC</b> Federal Communications Commission.	<b>ISM</b> industrial, scientific and medical.
<b>FCD</b> floating car data.	<b>ITS</b> intelligent transportation systems.
<b>FDMA</b> frequency division multiple access.	<b>ITU</b> International Telecommunications Union.
<b>FDS</b> fuzzy decision system.	<b>JSE</b> joint SOP exploration.
<b>FFR</b> fractional frequency reuse.	<b>LBT</b> listen-before-talk.
<b>FHSS</b> frequency hopping spread spectrum.	<b>LEO</b> low Earth orbit.
<b>FPGA</b> field programmable gate array.	<b>LiDAR</b> light detection and ranging.
<b>FPV</b> first-person view.	<b>LLM</b> large language model.
<b>GA</b> genetic algorithm.	<b>LOS</b> line of sight.
<b>GEO</b> geostationary orbit.	<b>LTE</b> long-term evolution.
<b>GeoTIFF</b> geographic tagged image file format.	<b>MANET</b> mobile ad hoc network.
<b>GIS</b> geographic information system.	<b>MC</b> multi-connectivity.
<b>GNSS</b> global navigation satellite system.	<b>MC-TRACE</b> multicasting through time reservation using adaptive control for energy efficiency.
<b>GPS</b> global positioning system.	<b>MEO</b> medium Earth orbit.
<b>GPSR</b> greedy perimeter stateless routing.	
<b>GPU</b> graphical processing unit.	





**MH-TRACE** multi-hop time reservation using adaptive control for energy efficiency.

**MILP** mixed-integer linear programming.

**MIMO** multiple input multiple output.

**ML** machine learning.

**mmWave** millimeter-wave.

**MST** minimum spanning tree.

**NAV** networked autonomous vehicle.

**NB-TRACE** network-wide broadcasting through time reservation using adaptive control for energy efficiency.

**NLP** natural language processing.

**NP** non-deterministic polynomial-time.

**NR** new radio.

**NTIA** National Telecommunications and Information Administration.

**NTN** non-terrestrial network.

**OFDM** orthogonal frequency-division multiplexing.

**OFDMA** orthogonal frequency-division multiple access.

**ORS** opportunistic relay selection.

**OSM** open street map.

**OSRM** open source routing machine.

**P2P** peer-to-peer.

**PL** path loss.

**POMDP** partially observable Markov decision process.

**PRB** physical resource block.

**PSO** particle swarm optimisation.

**PU** primary user.

**QoS** quality of service.

**RAL** radio access layer.

**RF** radio frequency.

**RL** reinforcement learning.

**RRM** radio resource management.

**SAS** spectrum access system.

**SATCOM** satellite communications.

**SDN** software defined network.

**SDR** software defined radio.

**SH-TRACE** single-hop time reservation using adaptive control for energy efficiency.

**SIR** signal-to-interference.

**SMAP** spectrum management architecture and protocol.

**SNR** signal-to-noise ratio.

**SOP** spectrum opportunity exploration.

**SOTA** state-of-the-art.

**SPT** shortest path tree.

**SRTM** shuttle radar topography mission.

**SSA** shared spectrum access.

**SU** secondary user.



**SVM** support vector machine.

**TCP** transmission control protocol.

**TDMA** time division multiple access.

**TRACE** time reservation using adaptive control for energy efficiency.

**TV** television.

**UAS** unmanned aerial systems.

**UAV** unmanned aerial vehicle.

**UHF** ultra high frequency.

**UN** United Nations.

**UPT** user-perceived throughput.

**V2V** vehicle-to-vehicle.

**WCDMA** wideband code division multiple access.

**WRC** World Radiocommunication Conference.

**XR** extended reality.



# 0 Overview

## 0.1 Research Opening

This research began with an industry problem of “*how to allocate radio spectrum between users with competing needs at the primary allocation level [1]*”. In the process of exploring this question, numerous gaps were identified in the literature about how to address spectrum management. Each chapter of this work intends to address one of these gaps, each framed through a specific proposition (a statement) and a corresponding question. Structurally, each chapter begins with introductory and approach sections that outline the problem, followed by a literature review to discover the state-of-the-art (SOTA). This is followed by demonstrative work, which is then analysed and discussed, followed by key findings.

## 0.2 Problem Background

Radio spectrum is one of the hottest commodities in an increasingly digitised and wireless society. Access is highly limited due to co-interference or *pollution* [2] of the resource that occurs as the number of users increases. Due to its critical importance and scarcity, spectrum access is a very valuable commodity. Rising demand and higher data rate requirements have motivated the shift toward higher frequencies where more bandwidth is available, sustaining progress across successive generations of networks.

This trend cannot continue indefinitely, as propagation properties at very high frequencies become unsuitable for reliable telecommunications. With demand accelerating, more efficient use of existing spectrum bands is essential. Research directions such as *dynamic* and *cognitive* spectrum management [3, 4] highlight the importance of adaptability and intelligent decision making in allocation. Yet, the challenge of deploying such approaches lies in testing and validation under realistic conditions.

This work therefore shifts focus to digital twin spectrum management, where advanced RF propagation modelling is applied in a 3D environment to create ultra realistic digital replicas of the physical world. These digital twins can capture mobility, interference, and multi-user coor-



dination effects at the fidelity required for real-time applications. They provide a platform to test and refine allocation strategies in conditions that mirror reality, bridging the gap between theoretical spectrum management proposals and practical deployment.

### 0.3 Approach

The approach taken for this thesis was to establish what specifically is required to make a dynamic spectrum management (DSM) system as distinct from a static spectrum management system. In static systems, tasks might only be carried out occasionally in years or even decades, while in dynamic spectrum management they may be repeated as often as needed, even in intervals shorter than a second. This shift turns previously rare and complex processes into continuous, rapid tasks. Each chapter of this thesis addresses one of the identified sub-problems to provide the benefits of a DSM system previously unavailable in static spectrum management systems. These sub-problems are as follows as per Table 1.

Table 1: Propositions and questions related to DSM

#	Proposition	Question
1	Spectrum management is conventionally static with changes tending to occur in the range of years.	How does conventional spectrum management work and which parts could become dynamic?
2	DSM systems could adapt for variable demand for radio spectrum.	How can a DSM be aware of the live demand across large urban and rural regions?
3	There are many possible allocations for radio spectrum.	Which allocation mechanisms are suitable for DSM solutions often with constraints such as low latency?
4	Spectrum management is a large geospatial problem with tight temporal constraints.	In a realistic setting, which techniques are available to implement real DSM solutions known for requiring significant levels of distributed control and coordination?
5	Effective ways of modelling spectrum usage are needed to assess the impact of allocations.	What options are available to model usage of radio spectrum in DSM for large urban and rural environments?
6.a	Having access to advanced 3D radio spectrum modelling has applications to <i>deconflict</i> allocations ahead of time.	How can the possibilities of what an allocated user may do with their access to radio spectrum <i>deconflict</i> allocations to multiple users for an allocation time interval?
6.b	Conventional spectrum management allocates in 2D space (latitude and longitude).	How can 3D modelling of radio spectrum be used to coordinate access of radio spectrum between the ground and drones?

## 0.4 Literature

Literature review has been central to the progress of this PhD. The DSM systems that form a major focus of the thesis are primarily theoretical, and their scale makes real-world experimentation impractical. To overcome this, experimental work has been carried out in simulations, which provide a more effective way to represent the large-scale geospatial challenges of the DSM problem. These simulations draw on theoretical foundations and prior literature, which were then adapted and extended to address the DSM-specific scenarios.

In chapter 1, the literature on spectrum management establishes what is required to move from *static* to *dynamic* spectrum management. Spectrum management research highlights the core challenges for DSM. Cognitive radio (CR) is consistently identified as a central concept, though its interpretation varies. Research also highlights the role of spatial and temporal prediction as essential to anticipating demand and availability.

In chapter 2, the literature introduces modelling techniques for radio spectrum demand. Markov chains & processes are widely used to capture stochastic behaviour across regions, while multi-agent modelling is applied to explore lower-level interactions. Abstracted agent modelling appears in the literature as a way to balance complexity with tractability. Studies on life patterns & timetables link spectrum demand to realistic daily activities.

In chapter 3, the focus in the literature is on spectrum allocation mechanisms. Work on resource distribution provides high-level frameworks, while spatial clustering and Voronoi tessellations offer practical methods for geographic partitioning. Fairness metrics are repeatedly used in allocation studies to evaluate outcomes.

In chapter 4, implementation-related research covers decentralisation, n-tier spectrum management, and MANETs. The literature highlights decentralisation as critical for scalability, rapid deployment, and resilience. Hierarchical approaches are shown to support both large-scale coordination and local, low-latency decisions. MANET clustering is presented as a direct method for structuring decentralised DSM. Literature here also introduces isochrones and shortest-path trees (SPTs) as techniques to incorporate user movement and reachability into allocation processes.

In chapter 5, the literature expands into radio spectrum modelling in 3D environments. 3D city modelling, digital elevation model (DEM), and radio frequency (RF) ray tracing are shown to be essential tools for capturing coverage and interference in both urban and rural settings. Beyond RF coverage, 3D city models also provide a structured way to represent user movements, vertical positioning such as activity in multi-storey buildings, and localised demand for spectrum in dense environments. Digital elevation model (DEM)s play a similar role in rural



or hilly urban areas where terrain strongly affects propagation, enabling accurate reflection of geographical constraints. Radio frequency (RF) ray tracing extends these models by generating data that can inform allocation decisions with propagation-aware detail.

In chapter 6, further applications of 3D spectrum modelling are explored. The literature on digital spaces and digital twins shows how high-fidelity simulations can be used for interference-aware allocation and multi-user coordination. Studies on drones' and satellites' spectrum access demonstrate how *vertical* spectrum management adds another layer of complexity, extending allocation into three dimensions.

## 0.5 Contributions

This thesis contributes research to the spectrum management problem, providing a scaffolded framework for how the problem may be approached. It starts from a simpler understanding and representation, progressing toward a more complex representation. The more complex models shown towards the end of this thesis are arguably on the cutting-edge of this research space, offering a glimpse into the future of spectrum management. This study started from an industry problem. The question was “*How can radio spectrum be allocated between users at the primarily allocation level?*”. Pursuing this question revealed gaps in the dynamic spectrum management (DSM) problem definitions. These gaps are expressed in pairs of a proposition with a corresponding question as detailed in the Approach subsections. Each chapter of this thesis addresses one of these gaps.

Chapter 1 provides an analysis for how spectrum management looked in the past, how it looks now, and which aspects could benefit from becoming more *dynamic*. The problem space is defined before the chapters that are more focussed on technical problems. This chapter’s contribution is to establish what the existing state of the spectrum management problem is to confirm the problem being solved really exists.

Chapter 2 contributes a framework for modelling dynamic spectrum demand by integrating statistical, behavioural, and hybrid approaches. It reviews and synthesises four modelling strategies, then demonstrates their application through two scalable models: a region-level Markov chain method and an agent-level hybrid simulation using statistical data and timetable-driven behaviour. The chapter’s contribution lies in showing how these methods can generate synthetic yet realistic demand datasets, highlighting trade-offs between efficiency and realism, and establishing hybrid approaches as a robust foundation for testing and developing dynamic spectrum management systems.

Chapter 3 addresses the mechanics of the initial allocation problem in DSM and contributes a framework combining fairness metrics, clustering methods, and Voronoi tessellations. It links theoretical models to practical demonstrations, showing where each method succeeds or fails. The chapter demonstrates that fairness metrics can formally guide allocation, clustering provides scalable spectrum reuse, and Voronoi tessellations offer a robust basis for dynamic spatial partitioning. It also highlights the limits of these approaches under mobility, uncertainty, and boundary effects, pointing to the need for adaptive mechanisms in real-world DSM systems.

Chapter 4 presents an analysis of spectrum management techniques for implementing DSM systems in real-world settings. The main contribution is a demonstration of how a mobile ad hoc network (MANET)-based protocol can group spectrum users into fully decentralised clusters to manage communication overheads, alongside the use of shortest path tree (SPT)s to





model realistic user reachability. Together, clustering and SPTs provide complementary methods for structuring users and capturing mobility constraints, offering a practical foundation for decentralised spectrum management, under realistic conditions.

Chapter 5 analyses techniques for 3D city modelling, digital elevation model (DEM)s, and RF ray tracing to support large-scale DSM in urban and rural environments. It reviews the wider problem, literature, and demonstrations using open-source tools such as Sionna RT. While basic urban ray tracing is not novel, the chapter's main contribution is advancing methods for generating high-resolution urban models and terrain-aware rural models, where landscape data is critical for RF propagation. The broader aim is to establish the background and tools needed to enable advanced spectrum and digital twin modelling within a DSM system

Chapter 6 applies advanced spectrum and digital twin modelling to coordinated spectrum access in DSM systems. It builds on earlier methods by synthesising results together, producing new outcomes. Application 1 combines time-weighted reachability with RF coverage to form an RF coverage SPT, capturing interference risk across mobile users. Application 2 demonstrates multi-user coordination through mutual signal-to-interference (SIR) analysis, quantifying and mitigating interference in urban and rural scenarios. Application 3 extends spectrum management vertically, using 3D modelling and ray tracing to evaluate drone-to-drone, drone-to-ground, and elevated-user interactions.





# Chapter 1

## Introduction to Spectrum Management

*The goal of this chapter is to provide necessary background information on spectrum management which is relevant to the following chapters that are more focused on solutions to the spectrum management problem. The chapter begins with an in-depth look at literature in this research space. The motivation for interest in specific areas of research is examined along with the roles these areas may play in providing solutions for spectrum management. Relevant papers of interest are summarised, discussed, and compared. The intended outcome of this chapter is to establish the current conventional or contemporary approach to spectrum management, enabling a clearer definition of the problem that is to be addressed. From these findings, the research gaps that will be explored in the remainder of this thesis are identified.*

## 1.1 Introduction

Spectrum management began as a commercially controlled process, treating the transmission of RF communications as a licensed technology rather than a shared commodity managed by the government [5, 2]. Commercial interests often prioritised profit over societal benefit, blocking technological breakthroughs that would render existing systems obsolete. This approach can be significantly at odds with the public interest when providing radio spectrum access in defence, health, and safety. To this day, balancing commercial and public interests remains a very hot topic of debate.

The modern era of spectrum management arguably began once the importance and value of radio spectrum was better understood and appreciated. To get better value out of the available radio spectrum and to protect the public interest, radio spectrum was carved up into frequency ranges to be allocated for applications suited to their frequency. Early allocations were for government, defence, radio usage, and eventually television. As newer technologies have emerged over the years such as commercial mobile phone networks, the squeeze for radio spectrum allocations has only gotten more contested.

Today, radio spectrum allocation can be described as a slow, manual, and inflexible process due to fixed, long-term allocations. For each allocation change, the potential impact of interference must be carefully analysed to ensure compliance with regulations such as those provided by the International Telecommunications Union (ITU). It is well known that user spectrum demand varies both spatially and temporally, unlike the fixed spectrum allocations intended to address the demand. While allocating spectrum based on expected peak traffic demands helps optimise the use of fixed allocations, it is an imperfect solution with significant compromises. Tripathi [2] describes the effect of this phenomenon as an “artificial spectrum scarcity” due to inefficient allocations preventing spectrum being used when it is available.

There is a significant amount of related work in the area of research covered by this thesis. The main topics covered in this section are about conventional spectrum management, dynamic management techniques, smart radio technologies such as cognitive radio (CR), spectrum access techniques, and simple to complex modelling techniques relevant to finding solutions for the spectrum management problem. The purpose of this section is to establish what other research already exists in this space. The section covers the research that is most relevant to the problem that this thesis intends to solve.



## 1.2 Approach

### 1.2.1 Proposition

Spectrum management is conventionally static with changes generally in the range of years.

### 1.2.2 Questions

Q1) How does conventional spectrum management work and what could become dynamic?

### 1.2.3 Topics

**Spectrum Management** - A thorough understanding is essential before engaging with dynamic spectrum management (DSM), as it provides the regulatory, technical, and policy foundations behind all spectrum use. This includes knowledge of institutional frameworks, legacy command-and-control models, signal characteristics, and spectrum allocation mechanisms. This foundational insight enables a critical appreciation of the limitations in traditional approaches and the motivations for transitioning to more flexible, market-driven, and decentralised models. Without this background, it's difficult to fully understand the operational challenges or policy implications that DSM aims to address.

**Dynamic Spectrum Management (DSM)** - Research focused on DSM provides a strong foundation for understanding its key challenges and how it connects to related areas in the field. By understanding what would make a DSM system better than its equivalent conventional spectrum management system, a justification to change it can be established.

**Cognitive Radio (CR)** - CR is widely recognised as a foundational concept for translating theoretical DSM frameworks into practical implementations. Even when a CR is not directly part of the solution, many of its key characteristics, such as autonomy, flexible access to multiple frequency bands, and the ability to make intelligent, context-aware decisions are likely to feature prominently in most proposed DSM approaches.

**Spatio Temporal Spectrum Prediction** - A substantial body of research within spectrum management explores low-latency spectrum sensing, with promising applications in multiple access techniques such as channel hopping and interference avoidance. Spectrum allocation introduces additional complexity, as decisions must account not only for immediate conditions but also for future intervals of use. This necessitates some form of spectrum *prediction*, positioning this field as a potentially critical technology for enabling DSM-based solutions.

## 1.3 Literature

### 1.3.1 Spectrum Management

#### Summary

*Spectrum management is a foundational yet increasingly complex aspect of modern communication systems, and it forms a core motivation for this thesis. While the field is well-researched, structured and supported by regulatory bodies like the ITU at the international level, and regional specific bodies like the FCC, and NTIA for the USA or the ACMA for Australia, it faces significant challenges in adapting to emerging technologies and growing spectrum demands. Recent literature points to a necessary evolution from traditional command-and-control models toward more flexible, technologically informed, and market-driven approaches, such as dynamic spectrum access, decentralised and hierarchical control, and hybrid regulatory frameworks. These evolving strategies aim to enhance spectrum efficiency, accommodate innovation, and better balance public, commercial, and defence interests.*

#### Review

Spectrum management forms the central motivation of this thesis. It is argued that the currently available and in-use systems for spectrum management are not going to be sufficient for future spectrum-using technologies and systems. However, this does not mean that spectrum management is a poorly understood or under-researched field. Stine and Portugal [5] provide an excellent introduction to the area of spectrum management, covering essential aspects of electromagnetic spectrum utilisation and its management. They explain foundational concepts such as signal characteristics, modulation methods, and propagation phenomena affecting communication systems. They outline the historical evolution of spectrum regulation from initial chaos to structured international and national regulatory frameworks managed by organisations like the International Telecommunications Union (ITU), Federal Communications Commission (FCC), and National Telecommunications and Information Administration (NTIA) (in the US context). Furthermore, they discuss advancements in technology including dynamic spectrum management and spatial reuse which enhance spectrum efficiency and facilitate more flexible, adaptive spectrum use. Finally, it acknowledges practical challenges posed by legacy spectrum allocations and equipment constraints, highlighting the complexity of modern spectrum management and the necessity of balancing technological innovation with operational realities.

Much of the work in spectrum management is focused on technology-based solutions to solve problems such as how to make fast and effective spectrum allocations. Spectrum management arguably requires additional technological solutions such as new antennas and spectrum sensing systems. Even if such systems become available, the rules of spectrum management itself

may require modification to allow these systems to be used to their full potential. Peha [6] examines different approaches to spectrum management policy, comparing traditional central planning methods to alternative market-based solutions. They outline various methods for allocating and accessing spectrum, such as exclusive licensed access, unlicensed open access, and hybrid models. They discuss the trade-offs between regulated approaches, ensuring interoperability, supporting public interests, avoiding interference and flexible market-oriented models that maximise economic efficiency and innovation. Lastly, they explore temporary versus permanent rights, highlighting the role of government in balancing control, flexibility, and efficiency, ultimately suggesting a multifaceted approach to effectively manage spectrum resources.

Historically and conventionally, spectrum management at the primary allocation level [1] has been a centralised process following a command-and-control approach [2] to manage spectrum access as a resource. However, as has been seen in many other fields which involve the management of a resource or access to a resource across a large area, decentralisation is being considered as a potentially valid option for the future. Anker [7] argues for a shift from centralised spectrum management to a decentralised spectrum governance model, where governments transition from strict controllers to facilitators of market-driven coordination. They detail how traditional command-and-control mechanisms limit efficient spectrum use and responsiveness to technology advancements and changing market conditions. They propose a hybrid approach integrating property rights, commons approaches, and government oversight to protect public interests. By providing a case study from the Netherlands involving cognitive radio (CR) technology, they demonstrate how a government can successfully transition to a facilitator role, managing spectrum through market design, monitoring, and supporting decentralised collaboration among stakeholders.

Understanding contemporary approaches to spectrum management requires examining both the economic and regulatory tools that have shaped current global practices. Cave et al. [8] provide an overview of contemporary spectrum management strategies and evaluates tools such as auctions, administrative pricing, spectrum trading, property rights, and commons models. They highlight shortcomings in traditional spectrum management practices such as inflexibility and inefficiency, discussing the economic, technical, and regulatory considerations necessary for adopting market-oriented reforms. By including global case studies, they emphasise the strengths and weaknesses of various regulatory approaches, advocating for reforms that balance market incentives with governmental responsibilities to ensure efficient spectrum utilisation while addressing competition and public service requirements.

### 1.3.2 Dynamic Spectrum Management (DSM)

#### Summary

*Dynamic spectrum management (DSM) addresses the limitations of traditional static spectrum allocation by enabling more flexible and adaptive use of radio frequencies. It has become a central concept in modern spectrum research, driven by advancements in technologies like cognitive radio (CR)s, the emergence of fifth generation (5G) and the future of sixth generation (6G) networks. These innovations have made dynamic approaches more feasible and practical, inspiring a wide range of studies that explore optimisation, interference mitigation, and intelligent allocation methods using AI and blockchain. DSM not only enhances network efficiency and adaptability, but also aims to bridge high-level policy with real-time access needs, making it a key enabler of future communication systems.*

#### Review

Spectrum management has often been described as an inflexible and slow process where spectrum is allocated *statically* over time intervals in the range of decades. Likely due to the pre-existing use of this terminology, research into the idea of making spectrum management more flexible and adaptive in a shorter interval of time has often been referred to as *dynamic* spectrum management. Dynamic spectrum management (DSM) is a key concept of many proposed future spectrum management plans. The purpose of DSM is to provide a way to dynamically optimise the allocation and utilisation of radio frequencies to enhance telecommunications network efficiency and performance. Most goals of DSM arguably require spectrum modelling in some form to be able to make good allocations and optimisation decisions. DSM has been a very active research area for many years [9], with new waves of research and publications emerging with technological advancements.

There are many aspects of spectrum management that could be made more dynamic. One such example which takes place at the device level is the concept of a *spectrum handoff*. Conventional approaches to *spectrum handoff* rely heavily on protocol-based methods to achieve consistent and reliable performance with minimal undesirable outcomes such as *dropped calls* or reduced data rates for users. Babalola et al. [10] survey dynamic spectrum networks with a focus on CR technology. They emphasise *spectrum handoff*, which is crucial for efficiently managing ongoing transmissions by shifting frequency bands to avoid interference with primary users. Various proactive and reactive sensing techniques for channel selection and switching are discussed, alongside analytical approaches using Markov chains and queuing models. They provide comparative insights into the effects of *handoff* on data delivery time, channel utilisation, and latency, suggesting avenues for future research.

The primary focus of this thesis is the primary allocation level [1] of spectrum management



where spectrum is allocated between the biggest users such as government, commercial users, and for industrial, scientific and medical (ISM) applications. Many proposed plans across the literature generally provide solutions that bring the spectrum allocation processes much closer to the lower level spectrum access management processes. In other words, the spectrum allocation process would become more acutely aware of what the real-time and future spectrum access requirements are. Bohge [11] addresses dynamic resource allocation challenges for orthogonal frequency-division multiple access (OFDMA)-based packet-oriented cellular systems. They develop mathematical optimisation models to explore the potential of dynamically allocating sub-channels and transmission power among users, significantly improving throughput, fairness, and reliability. They introduce heuristic algorithms suitable for real-world equipment limitations, including packet scheduling, dynamic power allocation, and inter-cell power coordination, demonstrating substantial performance improvements compared to traditional methods.

In response to the growing challenges of spectrum congestion and interference in wireless communication, especially in defense applications, several innovative approaches have been explored. Suchanski et al. [12] propose DSM to mitigate interference issues, particularly for military communication systems facing spectrum scarcity. A central frequency broker dynamically assigns frequencies to radio networks based on real-time channel conditions, reducing interference and enhancing communication robustness. They further suggest applying game theory for improving frequency assignment strategies. Their approach is validated in both simulated and real environments.

DSM systems would require interference analyses and similar processes to be repeated frequently, making the problem highly computationally demanding. However, Finding spectrum management solutions quickly with manual analyses is considered to be a very challenging problem. Even with technology available, many variants of how the problem to be solved is defined when written mathematically have been shown to be non-deterministic polynomial-time (NP)-hard. One of the common problem definitions used to describe the spectrum management problem is focused on how *pieces* of spectrum can be divided between a population of potential spectrum users. Luo and Zhang [9] explore the computational complexity of dynamically allocating spectrum among multiple users sharing a common frequency band. They characterise the problem as NP-hard under practical conditions, exploring both discrete and continuous formulations.

Dynamic solutions for spectrum management require information about the live or recent state of spectrum. Information specifically about spectrum channels is known as channel state information (CSI). Kaur and Kumar [13] provide a study about integrating machine learning (ML) with CR technology to intelligently manage spectrum allocation under imperfect CSI conditions. A cooperative reinforcement learning (RL) framework is introduced within a decentralised multi-agent system, further enhanced by cloud computing resources. This robust design



improves network capacity, reduces outage probabilities, and increases convergence speeds, effectively managing the inherent imperfections of CSI in practical scenarios.

Spectrum management is not a straightforward technological challenge solvable by introducing a new generation of technology. The current systems in place must continue operating for some period of time with a smooth and well-planned out transition to the newer technologies with agreement from all parties involved. As a result of this, there is a significant amount of *operational technical debt* or *availability-induced technical debt*. Tripathi [2] propose a comprehensive *bring your own spectrum* (BYOS) framework, addressing the challenges of existing spectrum allocation policies through an operator-agnostic, multi-level approach. They incorporate historical policy analysis, propose novel trading mechanisms based on interference metrics (such as spectrum *pollution*) and introduce token-based management for fair competition. The framework includes blockchain-based implementation suggestions and analyses spectrum allocation in terms of policy diffusion, providing guidelines for future spectrum trading environments.

The research project 5G-ALLSTAR gained a significant amount of attention in the spectrum management space, as well as contributing a substantial amount of published research relevant to DSM generally. Kim et al. [14] summarise the 5G-ALLSTAR project's work on integrating cellular and satellite networks to provide seamless, reliable, and ubiquitous broadband services. Key technologies include mmWave 5G NR cellular access, satellite-based NR access, spectrum sharing, and multi-connectivity solutions. The project demonstrates a fully integrated satellite-cellular prototype, addressing technical, economic, and standardisation challenges while supporting diverse use-cases like vehicular communications and high-density events.

Building on the foundational work described in the initial stages of the 5G-ALLSTAR project, further investigations sought to address the practical implementation and integration challenges of multi-connectivity systems. Cassiau et al. [15] detail the outcomes of the 5G-ALLSTAR project, focusing on multi-connectivity (MC) solutions that integrate satellite and terrestrial communications directly at the user-equipment level, enhancing 5G availability everywhere. They outline customisations in the 5G physical layer, spectrum sharing methods, beam-forming technologies, and a multipath transmission control protocol (TCP) protocol optimised for hybrid satellite-terrestrial connectivity. The developed tools demonstrated some successes, showing significant performance improvements in terms of reliability, latency reduction, and robustness for next-generation applications like automotive or vehicular communications.

A common theme across DSM literature is the idea that all that is needed to make spectrum management dynamic is a very good cognitive radio (CR). Liang [16] systematically discusses dynamic spectrum management (DSM), highlighting CR as a fundamental enabling technology allowing secondary user (SU)s to opportunistically access licensed spectrum. They integrate blockchain for policy enforcement and transparency, alongside artificial intelligence (AI) methods to enhance adaptability and efficiency. They review opportunistic and concurrent spectrum access techniques, providing a comprehensive guide for policy-making.

### 1.3.3 Cognitive Radio

#### Summary

*Cognitive radio (CR) emerged from early academic work as a concept for intelligent, adaptive radios that respond dynamically to their environment using techniques like machine learning (ML) and decision-making cycles. This foundational idea has since inspired a wide array of research into using software defined radio (SDR)s for more effective and flexible spectrum management. Applications range from spectrum sensing and predictive modelling to agile transmission control, with particular interest in enabling real-time interference avoidance and resilient communication in complex environments. Current studies continue to explore CR in both theoretical and applied contexts, including anti-jamming strategies, spectrum sensing under noisy conditions, and advanced optimisation for resource allocation in cognitive radio network (CRN)s.*

#### Review

Many sources across the literature point to the origins of cognitive radio (CR) (as a term at least) being from Mitola [17]. They present the cognitive radio (CR) concept as an intelligent software-defined radio capable of adapting to user context and network demands through real-time decision-making. They introduce the radio knowledge representation language (RKRL), a cognitive cycle including observation, planning, and action, integrating machine learning (ML) and natural language processing (NLP) to make radios context-aware and autonomously adaptive. They demonstrate these principles through a prototype (CR1), emphasising CR's potential for improving radio resource management, network efficiency, and user experience. While there is to our knowledge no complete system quite like this in operation currently, the ideas behind it have remained a pillar of research mostly under the wider umbrella of DSM research.

Many of the proposed applications of CR in spectrum management fall into one of two major variants. The first is the use of spectrum sensing to collect information about the live state of spectrum usage in the transmission environment. This may involve some form of predictive loop including other information such as simulation of the known environment and historical data. The second is where the transmitters themselves have access to a CR which can change its frequency, power, and direction to avoid or minimise interference. Babalola et al. [10] discuss a range of methods used for channel selection and spectrum handoff in cognitive radio network (CRN)s. They detail key processes involved in spectrum management. These include spectrum sensing, spectrum decision, spectrum sharing, and spectrum mobility. Particular attention is given to both proactive and reactive spectrum handoff schemes. They demonstrate how Markov processes and queuing models can be effective in reducing latency and cumulative delay. These methods also help improve channel utilisation.

A major application of CR is to address signal processing problems that human operators could not feasibly solve due to the combined challenges of high complexity and the extreme speeds at which signals vary. These problems often arise in contexts such as spectrum management, where advanced processing is critical for interference analyses. Having access to a smart radio that can break patterns from attacks attempting to disrupt communications is of huge interest to both governments and the commercial sector, which can also face such attacks. Tian et al. [18] investigate a game-theoretic approach to jamming and anti-jamming strategies in space communications, considering a powerful *cognitive jammer* with simultaneous multi-channel spectrum sensing capabilities. Using sliding window energy detection (SWED), the jammer employs a *detect and jam* strategy, whereas the transmitter-receiver side uses frequency hopping spread spectrum (FHSS) for mitigation. They formulate this scenario as a zero-sum game and explore the impacts of various parameters such as propagation delay and signal-to-noise ratio (SNR) on the jamming effectiveness, providing guidance on configuration strategies for resilient space communication networks.

The concept of a CR has arguably always been intended to work as part of a greater network of similarly equipped CRs. In recent years the term cognitive radio network (CRN) has become quite prevalent. Thabit and Ziboon [19] present a new CR detection system designed and implemented using statistical feature selection methods for identifying the presence of primary user signals under noisy conditions. Evaluated in additive white gaussian noise (AWGN) and fading environments, this FPGA-based system demonstrates their detection performance, claiming to achieve a detection probability of 100% at low signal-to-noise ratios (-18dB). The results suggest that their approach significantly outperforms traditional detection methods, providing practical improvements for real-world spectrum sensing applications, with potential applications in spectrum management.

Continuing on the topic of CRNs, Teekaraman et al. [20] introduce a hybrid optimisation algorithm aimed at resource allocation in heterogeneous CRNs. They address the challenges of efficient spectrum allocation in the presence and absence of primary user (PU)s, focusing on maximising network capacity, data rate, and reducing interference. The proposed method combines clustering techniques and relay selection mechanisms to enhance spectrum sharing among secondary user (SU)s. Their simulations demonstrate promising results for the proposed hybrid optimisation model, which efficiently balances capacity, interference, and spectrum utilisation. The results show significant QoS improvements compared to conventional approaches, especially under higher interference scenarios.

Lastly, applying artificial intelligence (AI) to CRNs, Robert and Vidya [21] propose a genetic algorithm (GA) optimised fuzzy decision system (FDS) to address the issues of spectrum allocation, routing, and the deafness problem in multi-hop CRNs. The developed decision support system (DSS) applies fuzzy logic to manage uncertain data about channel availabil-



ity, interference, and transmission power. The GA optimises membership functions within the fuzzy logic system, enabling efficient spectrum allocation and channel selection. The proposed approach significantly improves performance metrics like throughput, channel utilisation time, and reduces channel-switching delays and deafness problems compared to traditional fuzzy-logic based approaches, as validated through simulation experiments.

### 1.3.4 Spatio Temporal Spectrum Prediction

#### Summary

*Spatial-temporal spectrum prediction is arguably a critical enabler for dynamic spectrum management (DSM), allowing cognitive systems to anticipate future spectrum usage and improve decision-making. Various studies demonstrate the benefits of accurate prediction models, such as hidden Markov model (HMM)s and machine learning (ML) techniques, in enhancing throughput for secondary user (SU)s and minimising disruption to primary user (PU)s. Recent research also addresses practical challenges, including missing or corrupted data, by proposing robust, real-time forecasting frameworks that integrate tensor-based recovery and joint optimisation methods. These advancements bring attention to the growing importance of intelligent, adaptive spectrum management strategies, especially in the context of 5G and future wireless networks (such as 6G) through the integration of prediction, sensing, and database-driven approaches.*

#### Review

Spatio temporal spectrum prediction is a significant aspect of the problem this thesis explores. The challenges of providing satisfactory spectrum management solutions arguably need systems that are cognitive and predictive. In general, it could be said that a primary benefit of cognition is the ability to make decisions regarding the future rather than just the present. This is arguably a key challenge in making DSM practical for real-world spectrum management, where users typically need time to plan their usage. Barnes and Maharaj [22] investigate the impact of prediction-based channel allocation on cognitive radio network (CRN)s, focusing on secondary user (SU) performance and primary user (PU) disruption. Using real-world spectrum measurements from Pretoria, South Africa, and a channel switching simulator, a hidden Markov model (HMM) is applied to predict PU behaviour, enabling proactive channel selection. Their results indicate that accurate PU behaviour modelling significantly enhances CR performance, doubling SU throughput and halving PU disruption under heavy traffic conditions. They explore trade-offs between SU throughput and PU disruption rate, with improved performance observed as prediction accuracy increases. Practical measurements show high similarity with theoretical simulations, confirming the validity of the approach.

A broader view of statistical methods used for spectrum occupancy prediction has also been explored in the context of improving dynamic spectrum access. Eltom et al. [23] provide a detailed review and classification of statistical spectrum opportunity exploration (SOP) methods used for dynamic spectrum access (DSA) in wireless networks. They highlight the critical role of spectrum occupancy prediction in optimising CR operations, emphasising methods such as memoryless stochastic sources (such as Bernoulli and Poisson), finite-order Markov chains, Bayesian models, and linear regression models, as well as ML approaches like neural networks

and support vector machine (SVM)s. They categorise and assess these techniques based on a sequential prediction framework, illustrating their suitability for various scenarios, extending the discussion to cooperative prediction scenarios, and addressing both theoretical and practical challenges. They express the importance of empirical validation through measurement campaigns and stress challenges such as model validity, complexity, and the balance between prediction accuracy and implementation feasibility.

There is another application of spectrum prediction which is arguably more immediately useful for contemporary spectrum management systems due to the ease of which it could be integrated. In general, forecasting spectrum usage is difficult, as only predictable patterns can be modelled. This is arguably not the case for much ad-hoc spectrum usage, which is unplanned by definition. This form of prediction could be called an extrapolation in time forwards. Another variant of spectrum prediction is to interpolate spatially in between the known spectrum sensing points. Ding et al. [24] propose a robust online spectrum prediction (ROSP) framework to accurately predict spectrum states while handling missing and anomalous data. Existing spectrum prediction methods often fail to account for anomalies, scalability, and missing data. To address these issues, ROSP integrates time-series forecasting with matrix completion and recovery to form a joint optimisation problem, solved using an alternating direction optimisation method. The approach is validated using real-world spectrum datasets, demonstrating superior robustness and accuracy over existing methods, with around 10 dB improvement in prediction performance. The framework's online nature ensures real-time adaptability, making it suitable for applications like adaptive spectrum sensing and dynamic spectrum access (DSA).

Some research has focused on the broader concept of spectrum inference as a complement or alternative to real-time sensing. Ding et al. [25] provide a comprehensive survey of spectrum inference (or spectrum prediction) in cognitive radio network (CRN)s. Spectrum inference aims to predict occupied and available frequency bands by leveraging past spectrum occupancy data, reducing reliance on real-time spectrum sensing. They categorise spectrum inference methods based on time, frequency, and spatial dimensions, detailing various algorithms including linear prediction, Markov models, neural networks, Bayesian inference, and pattern mining. They discuss how these techniques improve spectrum utilisation, reduce interference, and enable more efficient dynamic spectrum access (DSA). Furthermore, they highlight applications in fifth generation (5G) and high frequency (HF) communications, while addressing challenges like model accuracy, computational complexity, and real-world deployment constraints. They also compare existing inference techniques and bring attention to the increasing importance of multi-dimensional spectrum inference for future wireless networks.

A significant challenge for many prediction systems is the reliability of historical data, especially when anomalies or missing entries are common. Li et al. [26] address the challenge of spectrum prediction when historical observations are incomplete or corrupted by anomalies.



They introduce a novel tensor-based recovery approach that uses a hankelised time-structured spectrum tensor, which preserves spectral and temporal dependencies. They formulate the data recovery problem as a tensor completion task, employing a robust online optimisation method using the alternating direction method. They introduce the concept of maximum predictability to assess the impact of missing data and anomalies, demonstrating that their approach significantly improves both recovery accuracy and computational efficiency while enhancing the predictability of frequency bands. Similar to the previous work on prediction robustness, Li et al. [27] propose a method for predicting spectrum availability in wireless networks despite missing or corrupted historical data. By modelling the spectrum as a 3rd-order tensor and using tensor completion techniques, the proposed robust online spectrum map prediction (ROSMP) algorithm can reconstruct missing values efficiently. The approach exploits spatial-temporal-spectral structures in spectrum data and applies alternating direction minimisation for real-time updates. Experimental evaluations using real-world spectrum data confirm that ROSMP significantly outperforms baseline methods in prediction accuracy and robustness.

The primary user (PU) and secondary user (SU) concept for DSM is one of the main models for dynamic spectrum access (DSA) proposed across the literature. This is arguably because it is considered to be the most compatible with the current spectrum management processes in which the PUs (or incumbents) are guaranteed their regular service, and the DSA for SUs is only ever used in the spectrum gaps. Yang and Zhao [28] propose a novel frame structure to improve the throughput of SUs in cognitive radio network (CRN)s. Instead of randomly selecting channels for sensing, the approach utilises spectrum prediction based on historical data to choose only channels predicted to be idle. They evaluate the impact of imperfect predictions, traffic intensity, and the number of channels on throughput. Simulation results show significant throughput improvements, even with prediction errors, as the method reduces the probability of selecting busy channels. The findings highlight that while prediction errors can affect performance, the proposed method performs well in high-traffic scenarios, suggesting it may be a viable enhancement for CRNs.

Conversely in some ways to the PU and SU strategy which intends to maintain the status quo while *squeezing* DSM in the gaps, some proposed systems across the literature look at overhauling the conventional approach to replace spectrum management entirely with DSM data-driven predictive technologies. Zhang et al. [29] propose a smart spectrum management (SSM) framework that integrates spectrum knowledge (SK) and real-time observation (RO) to optimise spectrum utilisation in 5G and future wireless networks. They introduce a three-layered model to categorise spectrum-related data into spectrum data (SD), spectrum information (SI), and spectrum knowledge (SK), with SK acting as an intelligent enabler for spectrum opportunity exploration (SOP). The proposed SKRO-enabled SSM utilises both historical spectrum data and real-time sensing to improve spectrum efficiency while ensuring minimal interference with PUs. Extensive simulations demonstrate that this approach significantly enhances SOP utilisation.

In highly contested RF transmission environments, they achieve up to a 16.55% improvement in spectrum usage efficiency. The findings highlight the potential of AI-driven spectrum management in addressing the increasing demands of wireless communication.

Expanding on the idea of hybrid systems, additional research has explored how joint sensing and prediction-enabled databases can provide more adaptive spectrum access solutions. Zhang et al. [30] introduces a novel approach to dynamic spectrum management (DSM) by integrating spectrum sensing, a spectrum database, and spectrum prediction to enhance spectrum opportunity exploration (SOP). The proposed joint SOP exploration (JSE) scheme uses the predictive capabilities of a spectrum database alongside real-time sensing to optimise spectrum access for SUs while minimising interference with PUs. Unlike traditional methods that rely solely on geo-location databases or direct sensing, this hybrid approach balances accuracy and efficiency by dynamically selecting the best method based on real-time conditions. Simulation results demonstrate that JSE adapts well to varying predictive and sensing capabilities, improving spectrum utilisation while maintaining low interference levels.

The idea of applying spectrum prediction to a DSM problem has been explored at many different scales. Tonnemacher et al. [31] explore techniques to enhance spectrum sharing in the citizens broadband radio service (CBRS) band. They address the challenge of maximising spatial spectrum reuse while minimising interference between different priority users in the three-tiered CBRS system. They propose listen-before-talk (LBT) schemes that enable secondary (general authorised access (GAA)) users to opportunistically access the spectrum while protecting higher-tier users (priority access and incumbents). Through simulations and real-world testbed experiments using SDRs, they demonstrate that these LBT schemes significantly improve user-perceived throughput (UPT) for SUs. Additionally, they introduce a Q-learning-based reinforcement learning (RL) approach to dynamically adjust energy detection thresholds (EDT)s for opportunistic transmitters, mitigating hidden and exposed terminal problems. The results show that adaptive EDT can improve SU throughput by up to 350% while causing only a 4% reduction in PU throughput, highlighting the potential of ML for efficient spectrum sharing in CBRS networks.

## 1.4 Background

### 1.4.1 Spectrum

Effective spectrum management depends on a clear understanding of spectrum itself. For this reason, its key characteristics are defined to establish a foundation for analysis.

Property	Description
<b>Fragmentation</b>	How <i>complete</i> is the spectrum?
<b>Linearity</b>	Which properties change linearly with frequency?
<b>Non-linearity</b>	Which properties change non-linearly with frequency?
<b>Persistence</b>	How long does spectrum last when transmitted?

Table 1.1: Spectrum characteristic properties

**Fragmentation** tells us how *complete* the spectrum is. Radio spectrum exists along a continuous range, but interference and margins mean it is very fragmented. There are limitations with how spectrum can be used and channelised for different technologies.

**Linearity** for spectrum generally refers to how it changes with increasing frequency. *Spectrum* in the context of using it for some telecommunications purpose is defined as a range of frequencies bounded by a lower and upper frequency. Many applications will attempt to treat the entire range as being the same, but every part of the spectrum is different. A crucial detail to know is which characteristics change predictably with increasing frequency.

**Non-linearity** for spectrum generally refers to the characteristics of a lower and upper bounded piece of spectrum which change with frequency, but not specifically with increasing frequency. These characteristics tend to be caused by some physics phenomena such as how it may interact with particles in air at extremely specific wavelengths. These characteristics can be far more complex to understand and account for. This is especially true if they relate to complex environmental factors which may occur and change almost unpredictably.

**Persistence** is the time duration and area of effect that a transmitted frequency will continue to effect once transmitted. The factors that affect this are related to the frequency of the spectrum, the way it was transmitted, and the environment in which it is transmitted. For example, higher frequency spectrum in a dense environment will tend to *fade* very quickly while the lower frequency spectrum will pass directly through this environment affecting a wider area. However, the higher frequency will experience more *secondary* effects such as *multipath* causing it to *linger* or persist within the smaller area for longer.

### 1.4.2 Value

Quantifying the value of radio spectrum allocations is a hard-to-solve problem in the spectrum allocation space. Globally, spectrum has generally been valued using two approaches. Approach one is by merit, as judged by a central authority such as a government. Approach two is by market. These approaches certainly do quantify the value of spectrum in some way, but more scientific methods are also needed, especially for those who intend to make huge offers based on a perceived value of the spectrum. Many metrics can be used to measure particular aspects of value for spectrum, but they generally cannot all be optimised at once. The value of spectrum is mostly about its usefulness for a task, which is generally telecommunications. The exact same spectrum could have a different value based on where and how it would be used, which varies between potential users.

Property	Description
<b>Usability</b>	Can the spectrum accomplish the desired use case?
<b>Fungibility</b>	Can the spectrum replace other spectrum?
<b>Uniqueness</b>	Can the spectrum be replaced?
<b>Aggregability</b>	Can the spectrum be combined with other spectrum?

Table 1.2: Spectrum value properties

**Usability** includes many technical characteristics of both the radio frequencies and the technology that would use it. It can vary a lot for different use cases, but for competitors it will often be very similar. Due to this, the usability of spectrum that has commercial applications can contribute to the market value very sharply.

**Fungibility** is a common term when talking about spectrum and describes the degree to which spectrum in different frequency bands can be interchanged. Some spectrum has high utility for substituting other bands because it shares characteristics of both low and high frequencies, falling roughly between the two.

**Uniqueness** in this context is essentially the inverse of the concept of fungibility. If there is no suitable replacement for a spectrum band, it will naturally have a higher value. Similarly to usability, if there is competition for the spectrum, its market value could explode because whoever gets exclusive access could have a significant market advantage.

**Aggregability** is a characteristic of spectrum that is specific to the technological solutions available and potentially limited by the use case, or even budget. There are numerous examples of telecommunications' technology that can transmit data using multiple spectrum bands in some way. Some examples are dual-band, frequency hopping, and multispectral waveforms.

### 1.4.3 Multiple Access

Multiple access schemes for radio spectrum are one of the most critical concepts in spectrum management.

Property	Description
FDMA	Frequency Division Multiple Access
TDMA	Time Division Multiple Access
CDMA	Code Division Multiple Access
SDMA	Space Division Multiple Access

Table 1.3: Spectrum value properties

**FDMA** represents a high-level access method used in conventional spectrum management. Spectrum allocation within this approach relies on the two key strategies of frequency separation and frequency reuse. Frequency separation refers to assigning different frequency bands to geographically distinct services or users to prevent interference. This ensures that incompatible signals are kept apart by distance and frequency. Frequency reuse involves allocating the same frequency bands to multiple users in different locations, provided those locations are sufficiently spaced to avoid significant interference. This is commonly implemented in cellular networks through planned reuse patterns, such as tessellating hexagonal cells, allowing efficient use of limited spectrum.

**TDMA** tends to work with extremely small time increments. The purpose of this choice is so that although multiple users are never accessing the spectrum at the same time, they are able to access it frequently enough that there is relatively little latency or delay in their access. If the time increments were much bigger, phone calls for example would have a perceivable amount of delay

**CDMA** on the other hand provides further benefits beyond multiple access to spectrum. There is also a security and privacy aspect to it. This is well demonstrated in Bluetooth where despite such a high number of users being able to access it, it is very difficult to listen in on another user's communications.

**SDMA** introduces spatial separation as a method for enabling multiple users to share the same frequency band at the same time. Instead of dividing the spectrum by frequency, time, or code, it uses directional antennas and beamforming to isolate users based on their physical location. Smart antennas, particularly those used in newer mobile phone networks and massive multiple input multiple output (MIMO) systems, can steer beams toward individual users while reducing interference toward others. This allows the same frequency resources to be reused within very small geographic areas, which improves spectral efficiency.

#### 1.4.4 Conventional Management

Conventional spectrum management can be described as the traditional, centralised approach to regulating radio frequency (RF) spectrum, typically administered by national or international authorities. In this model, the spectrum is divided into specific frequency bands that are pre-allocated and licensed to particular users or services such as television broadcasters or mobile network operators within defined geographic areas. The International Telecommunications Union (ITU) was established in 1865 as the International Telegraph Union to standardise and facilitate cross-border telegraph communication. It became a specialised agency of the United Nations (UN) in 1947, expanding its role to encompass all aspects of global telecommunications, including radio, satellite, and broadband technologies. One of the ITU's critical functions is spectrum allocation, ensuring the efficient and interference-free use of radio frequencies worldwide. Through its radio regulations and the World Radiocommunication Conference (WRC), the ITU coordinates spectrum use between countries, balancing the needs of governments, industries, and emerging technologies to promote global connectivity and innovation [32, 33].

Spectrum allocation typically occurs in the spatial domain and is sometimes framed as an *envy-free fairness* or *fair cake-cutting* problem [34]. In the case of spectrum management, access to the resource is being divided or allocated, rather than the resource itself, given that radio spectrum is essentially just the collection of possible RF transmissions that could be made. When countries do this process, they generally represent the allocations as a combination of geographical spatial divisions and radio frequency (RF) spatial divisions. Many countries have a similar static spectrum allocations system but tend to only provide a document with hundreds of pages of tables. Australia and the USA are two examples where the government departments or entities responsible for spectrum allocations provide better visualisations of their allocation system. For this reason they are used as two case studies for what conventional spectrum management at the licensing level looks like.



## Hierarchical Cell Identification Scheme (HCIS)

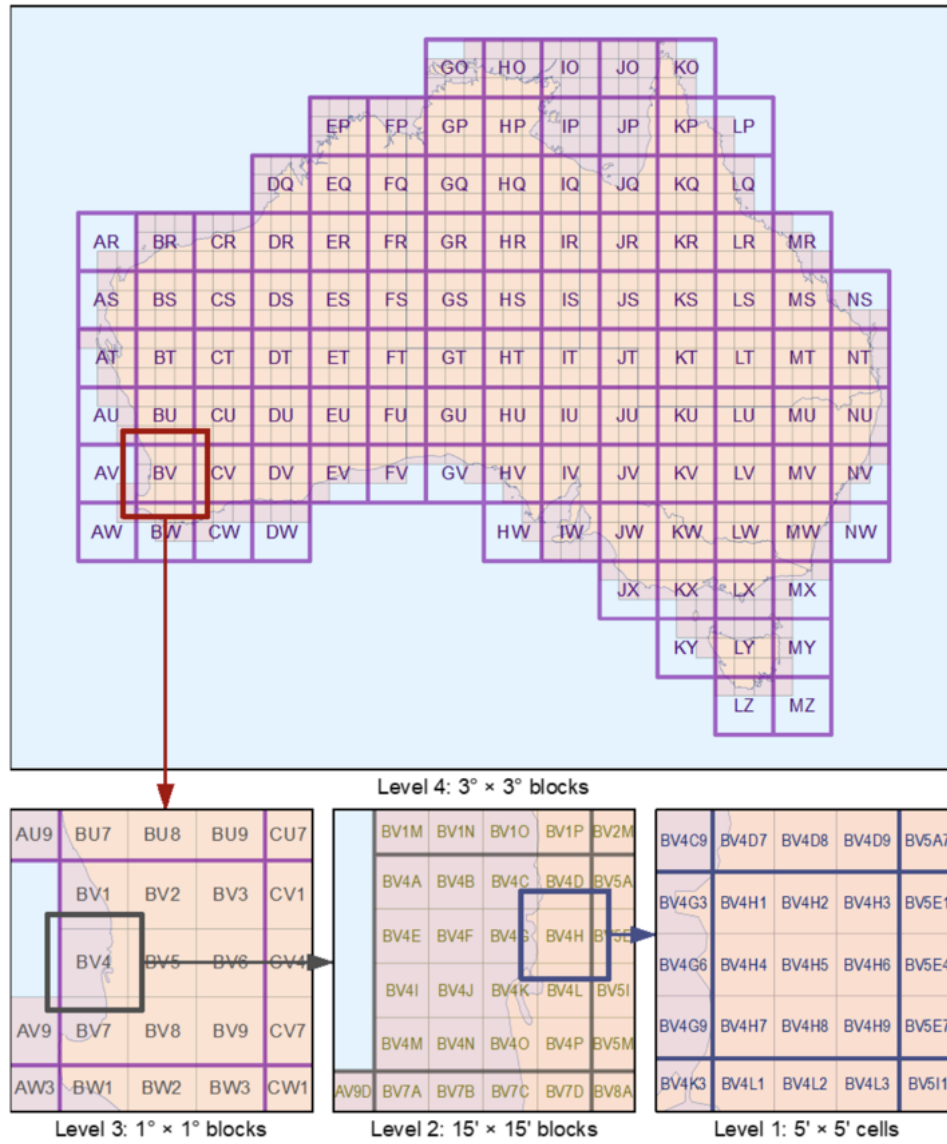


Figure 1.4.4.2: Illustration by the ACMA of HCIS from [35]





### 1.4.5 Commercial Usage

There are two stages to the commercial spectrum usage life cycle. There is the spectrum acquisition stage and then the spectrum usage stage. Given the limited chances to gain further spectrum access due to the infrequency of spectrum auctions, the acquisition stage is very critical for commercial spectrum users. It follows that whatever spectrum the commercial user has must be used as effectively as possible. When reviewing the standards for the generational commercial phone networks such as fourth generation (4G) long-term evolution (LTE) and fifth generation (5G) new radio (NR), it becomes clear that many protocols exist or are being developed that enable more dynamic use of the spectrum allocated to mobile network operators within their assigned bands.

In modern commercial mobile phone networks, spectrum is managed by breaking down the allocated spectrum using standardised channel widths and spacings. This is largely done to maintain orthogonality as is required to implement the spectrum sharing technique orthogonal frequency-division multiplexing (OFDM). The main benefit being that the noisy side lobes of the transmissions have a cancelling effect with their neighbours. It also simplifies the allocation process by working with a relatively small number of possible bandwidths. The 4G LTE and 5G NR standards both define a process called carrier aggregation (CA) where the resource blocks can be combined. There are multiple ways in which the aggregation can occur, but each of them maintains the orthogonality requirement because of the use of standard channel widths.

5G NR defines the smallest size spectrum assignment as a physical resource block (PRB) with its width depending on the bandwidth part (BWP) value chosen. In 5G, the OFDM symbol rates allow BWP configurations of 180kHz, 360kHz, 720kHz, 1440kHz, and 2880kHz to provide some adaptability for different multipath conditions [40, 41]. The smallest time step in 5G is called a sub-frame and is defined as a 1ms interval [40, 41]. Carriers requiring a duplex transmission scheme can be assigned for use in frequency division multiple access (FDMA) as an uplink or downlink component and may also be assigned for use in TDMA where it alternates between downlink and uplink. For each user that is assigned aggregated channels, one is the primary channel called a PCell (primary cell) and the others are called SCells (secondary cells). In the case where an uplink and downlink channel are required for a user, any aggregated channels are called supplementary uplink and downlink (SUL and SDL) [41].

### 1.4.6 Satellite Communications

Although satellite communications (SATCOM) is not the primary focus of this thesis, it is included as background due to its distinct characteristics within the conventional spectrum management system. Unlike terrestrial spectrum allocation, typically allocated at the national level, SATCOM spectrum involves cross-border considerations, adding significant complexity. As a result, the management processes (overseen by the International Telecommunications Union (ITU)) are more adaptive and modernised to address the global nature of satellite operations. These allocations must account for both terrestrial networks (ground-based systems) and orbital communication systems, referred to as non-terrestrial network (NTN)s in the context of SATCOM. Coordination between these systems requires careful planning to avoid RF interference.

One key method for resolving spectrum conflicts in SATCOM is the calculation of the carrier-to-interference (C/I) ratio, a core part of interference analysis. Much of the international collaboration on such issues occurs at the World Radiocommunication Conference (WRC), where countries work together to agree on mutually beneficial solutions. A notable example from the 2012 conference provided insights into the technical procedures employed in SATCOM spectrum management [42]. While the current processes remain slow and manually intensive, the availability of high-quality spectrum data and advanced analytical systems could enable near real-time solutions. Such developments may offer valuable insights for strategies to improve spectrum usage and *deconfliction* in terrestrial networks.




Figure 1.4.6.1: Interference analysis diagram between an incumbent SATCOM system (Intersputnik) and a proposed SATCOM system (Thaicom) by the ITU from the presentation slides [42]

## 1.5 Discussion

### 1.5.1 Summary

Chapter 1 has four sections: *Introduction*, *Approach*, *Literature*, and *Background*.

The *Introduction* sets the scene by tracing spectrum management from early commercially driven control to the present government-managed allocation systems. Allocations are described as typically fixed for long periods, making the process slow and inflexible. As real spectrum demand changes across time and location, static allocations are shown to result in underuse in some places and congestion in others, which is described as an “artificial spectrum scarcity” by Tripathi [2].

The *Approach* section states the central proposition that conventional spectrum management is mostly static and changes at the pace of years. A guiding question is posed about how the conventional system works and which parts could be made dynamic. Four significant areas in the literature are named as containing the majority of recent research with this goal. These are spectrum management, dynamic spectrum management (DSM), cognitive radio (CR) and spatio temporal spectrum prediction.

The *Literature* section reviews research across these four areas. It is found that traditional command-and-control models may give way to more flexible and market-aware approaches supported by technology and regulation. DSM research looks at fourth generation (4G) long-term evolution (LTE) and fifth generation (5G) new radio (NR) capabilities, artificial intelligence (AI), and blockchain to support faster allocation and interference mitigation. Practical limits and non-deterministic polynomial-time (NP)-hard problems are found to be notable challenges across DSM research. CR is described as a promising concept for sensing, decision making, handover and anti-jamming in cognitive radio network (CRN)s. Spatial and temporal prediction methods such as hidden Markov model (HMM)s, machine learning (ML) and tensor-based recovery are surveyed, with findings showing that prediction can increase secondary user (SU) throughput while protecting incumbents, even with missing or noisy data.

The *Background* section provides foundational spectrum management theory and examples. Key spectrum properties such as fragmentation, linearity, non-linearity and persistence are defined, followed by value concepts including usability, fungibility, uniqueness and aggregability. Multiple access schemes frequency division multiple access (FDMA), time division multiple access (TDMA) and code division multiple access (CDMA) are summarised. Conventional management is described through the International Telecommunications Union (ITU) and World Radiocommunication Conference (WRC), with national case studies involving the Australian Communications and Media Authority (ACMA) in Australia and the National Telecommunica-



tions and Information Administration (NTIA) in the United States. Commercial systems such as 4G LTE and 5G NR are explained as organising and aggregating channels using orthogonal frequency-division multiplexing (OFDM), physical resource block (PRB)s and carrier aggregation to enable efficient use of licensed bands. Lastly, Satellite communications (SATCOM) is discussed, focussing on cross-border coordination, ITU processes and carrier-to-interference ratio analysis as core tools for *deconfliction*.

## 1.5.2 Key Findings

### **Finding 1: Conventional spectrum management creates artificial scarcity**

The review shows that conventional spectrum management is slow, rigid, and based on long-term centralised allocations. These static processes cause spectrum to remain underutilised in some places and congested in others, resulting in an artificial spectrum scarcity. International and national authorities, such as the ITU, NTIA and ACMA etc, enforce stability through fixed rules and licences, but this prevents spectrum use from adapting to local or time-varying demand. A dynamic model has been widely proposed in the literature as an alternative, where rights and settings could change rapidly in response to conditions, improving efficiency while still protecting priority users through sensing, prediction, and short-duration licences.

### **Finding 2: Solutions need to be tested under realistic conditions**

A major finding of this chapter is that much DSM research provides brilliant technical solutions (such as devices and algorithms) but the scenarios and conditions they are tested under tend to be mostly idealised or undefined. In early research and development work, this makes sense as these technologies may play enabling roles in many future problem definitions and in the context of specific DSM scenarios. This broad scope makes progress difficult to measure, and many advancements will be overlooked unless demonstrated through real-world or simulated applications.

### **Finding 3: Simulating realism is difficult**

Testing simulated realism is also difficult. As alluded to in the previous paragraph, a common issue identified across much the reviewed literature was the difficulty to demonstrate or quantify how the research results would function under a real world scenario. Arguably, the biggest reason for this is because it is extremely challenging to set up such a test. In response to this possible gap in DSM research, a major focus for the remainder of this thesis is developing and demonstrating scalable and computationally feasible methods and approaches for modelling large scale telecommunications scenarios at full scale for detailed dense urban areas and large sparse rural areas. To be most useful for researchers, the simulation approach needs to have adaptive levels of detail so that it can provide fast results during rapid and experimental development, but more detailed and complex results for the final results or to better be able to differentiate the performance of the top candidate solutions.

### 1.5.3 Limitations

The scope of the research reviewed in this section is constrained by the volume of literature that could be feasibly covered. Spectrum management intersects with a broad range of disciplines including telecommunications, computer science, economics, defence, and regulatory policy. A substantial body of relevant work exists within these adjacent domains, however differences in terminology, conceptual frameworks, and problem definitions make it difficult to locate and incorporate all relevant studies. As a result, there is potential for influential approaches and findings from outside the immediate body of literature on spectrum management to remain unexamined, limiting opportunities for cross-disciplinary application.

The review has also focused on specific geographic regions, with detailed case studies centred on Australia and the USA for conventional spectrum management. While these examples provide valuable insight into current allocation mechanisms, they are not fully representative of global practices. Other jurisdictions operate under different regulatory, technical, and market conditions, which may influence the applicability of the findings presented here. Consequently, conclusions drawn from the regional case studies may not generalise to regions where governance structures, spectrum usage patterns, and infrastructure development differ significantly.

A further limitation is that the section is based solely on review and synthesis of published research without experimental validation, simulations, or mathematical modelling to support or test the concepts discussed. Without quantitative analysis or empirical demonstrations, the mechanisms and relationships described cannot be confirmed under real-world conditions. This absence of supporting models or test results constrains the ability to predict the performance, scalability, or operational viability of the approaches examined, and leaves open questions that would require further investigation to resolve. This is a key differentiator between this chapter, and the ones that follow, which have a significant focus on modelling, simulation, and demonstrative work. This chapter reviewed literature and real-world practice to identify existing problems, define those requiring solutions, and implement dynamic spectrum management (DSM) approaches.

#### 1.5.4 Future Work

Future work could address the gaps caused by the limited scope of research coverage. Given the multi-disciplinary nature of spectrum management, further investigations could explore methods and findings from adjacent fields such as distributed systems, economics, and artificial intelligence that have not yet been widely applied to this problem space. Cross-disciplinary studies could provide novel approaches for allocation optimisation, interference management, and governance models that have not been considered in the core literature. Establishing consistent terminology and shared frameworks between disciplines would help to make such integration more effective.

Expanding the geographic scope of case studies would also be of value. Comparative research across a wider range of national and regional spectrum management regimes could identify alternative practices, regulatory innovations, and technical solutions not present in the Australian and US contexts. Future work could analyse how different governance structures, spectrum market mechanisms, and infrastructure maturity levels influence the success or failure of conventional and dynamic management strategies. This would also support the development of adaptable models that can be customised for local conditions while retaining general principles.

Finally, there is a clear need to test and confirm the concepts discussed in the literature. Simulation studies, mathematical models, and prototype systems allow quantitative evaluation of proposed approaches to spectrum allocation, prediction, and management. For example, models of dynamic allocation integrated with cognitive radio (CR) capabilities could be tested under varying traffic conditions and interference environments to assess scalability and robustness. This would help move proposed solutions from conceptual designs to practical systems, while providing measurable performance data to guide both technical and policy decisions. This thesis adopts this approach to establish the practical relevance of these concepts.



## 1.6 Next Steps

The next steps are to take the gaps identified between the current static, plan driven spectrum management model and the requirements for achieving dynamic spectrum management (DSM), and address them through targeted modelling, simulation, and more focused research. Building on the key findings, this will involve developing realistic, scenario driven evaluations that demonstrate how proposed DSM technologies and processes such as short duration licences, sensing, and prediction perform under conditions that mirror present day operational environments. The focus will be on creating scalable, computationally efficient simulation methods capable of representing both dense urban and sparse rural areas, with adaptive levels of detail to suit different stages of research. This will enable faster iterative development of concepts while preserving the capacity for high fidelity evaluation of top candidate solutions, closing the gap between idealised technical proposals and deployable real world systems.



## Chapter 2

# Modelling Dynamic Spectrum Demand

*This chapter focuses on the techniques and motivations for modelling radio spectrum demand across an urban environment, considering different levels of complexity for various applications. The discussion begins by reviewing techniques historically used in telecommunication technologies and the reasons why modern networks necessitate more advanced modelling methods. Building on these insights, the chapter explores techniques for mapping human urban life patterns to expected (or predicted) RF device usage, resulting in a predictive spectrum demand model. Key areas of investigation include statistical methods such as Markov processes, stochastic agent simulations, and physics- or reality-based techniques for accurately modelling urban environments and traffic movement, and how these dynamics affect radio spectrum demand over time and space.*

## 2.1 Introduction

In the field of electromagnetics, the term spectrum is used to refer to a continuous range of electromagnetic radiation waves. In telecommunications, the range to radio spectrum is reduced to the frequencies able to travel freely through free-space or the air, making them suitable for telecommunications. As the demand for spectrum access increases, there are few options available to fill the demand gap. It is widely agreed across industry and academia that spectrum management is one of the most promising yet largely unsolved problems [43, 44]. If solved, it could provide additional capacity for telecommunications technologies to keep up with the sharply increasing demand for access.

It is very common in the field of spectrum management to refer to the entire radio frequency range simply as spectrum to simplify the concept. However, the consequence is that it can be easy to overlook how much the utility or applications of radio spectrum vary across the usable range for telecommunications. These differences are due to how frequencies interact with their transmission environment. This can result in spectrum management solutions being researched or even implemented in simulation that would be impractical or essentially useless if used with real radio spectrum. For this reason, it is important to define exactly what is meant by spectrum in the context of dynamic spectrum management (DSM) solutions. Additionally, a modeling approach should have the right level of complexity to provide meaningful insight into realistic solutions for practical spectrum management challenges.

The problem this chapter intends to address is how to define spectrum demand in the context of the spectrum allocation problem. The overall goal of dynamic spectrum management (DSM) research is to make spectrum usage more efficient by making the allocation process more dynamic as distinct from the more static command-and-control approach conventionally used globally. A key insight from chapter 1 was that it is unclear what should guide changes in allocations [45], making it hard to know which part of the process to make more dynamic. On one hand, 100% allocation efficiency is achievable by simply ensuring all spectrum is always allocated. However, spectrum usage efficiency has not necessarily been achieved in this process. If the times and locations of spectrum transmission were always known, achieving such efficiency might be possible. However, this situation is only likely to occur in a scenario where all communication is constant and planned. The problem lies here as this is not generally how user generated demand tends to be. It is hard to predict changes in spectrum demand, yet these changes should determine how the allocation is adjusted, so understanding them is essential.

## 2.2 Approach

Static allocation of radio spectrum is considered to be the conventional approach to spectrum management globally. There are arguably many limitations to this approach, but it does provide solutions to the core problems that spectrum management was initially designed to solve. This section examines the core problems, presenting them below as *propositions* for this chapter with questions also provided which guide the purpose of this chapter. This chapter aims to show how a dynamic spectrum management (DSM) system can address problems largely overlooked by conventional spectrum management, which typically allocates radio spectrum long-term or *statically*.

### 2.2.1 Proposition

Radio spectrum is allocated to satisfy demand while accounting for interference constraints associated with increasing usage. In an ideal spectrum management system, the radio spectrum allocations would perfectly match the demand for radio spectrum across all radio spectrum access dimensions such as frequency, space, time, and code. If the demand for radio spectrum changes with time, the allocations would change with it.

### 2.2.2 Question

How can demand for radio spectrum be represented and modeled to simulate a DSM system?

### 2.2.3 Topics

**Markov Chains & Processes** - Modelling a problem as a series of interconnected states with variable transition weights has applications for spatial and temporal problems.

**Multiagent Modelling** - Modelling a problem at the lowest levels can be an effective way to accurately model higher-level problems, such as regional level spectrum demand in a DSM problem.

**Abstracted Agent Modelling** - The complexity and scale of DSM problems can be very immense leading to computational problems during simulation and latency problems for real-world applications. Abstraction is a common technique to address this problem.

**Life Patterns & Timetables** - By considering life patterns of individual agents, a complex but generalised dataset can be created to help develop and test more catered solutions.

## 2.3 Literature

### 2.3.1 Markov Chains & Processes

#### Summary

*Markov processes offer a powerful framework for modelling dynamic systems, particularly in scenarios with limited or coarse data, such as wireless network environments. These models are often applied to mobile communication systems to manage spectrum handoff as users move between cells, helping optimise mobility predictions and network configurations. In more complex cognitive radio network (CRN)s, Markov models support strategies like hybrid spectrum access and proactive channel allocation, ensuring efficient spectrum use and maintaining quality of service (QoS) for different spectrum user types. By predicting user behaviour and spectrum availability, these methods improve performance metrics such as throughput, delay, and reliability, showcasing the adaptability and effectiveness of Markovian modelling in modern network and spectrum management.*

#### Review

Markov chains provide a way to model change over time or in response to events or actions. For modelling in particular, it is especially useful when the available data is coarse as it can be represented quite well as higher-level state changes. One of the most common problems faced by mobile phone networks is how to manage mobility of spectrum users in the network. Most mobile phone networks consist of many cells which are responsible for transmitting and receiving communications for the users of that area. At any given point in time, a user may be within range of multiple cells, or none. As a user moves out of range of one cell into the range of another, at some point, the quality of their connection will reduce unacceptably if they do not change to a new cell. The process of doing this is known as a *handover* or *spectrum handoff*.

One common way to model a *spectrum handoff* is to use Markov models. An example of this approach is described by Szálka et al. [46] who introduce a Markov movement-model creator framework (MMCF) designed to construct mobility models for wireless networks using Markovian approaches. They outline a systematic method for defining parameters and constructing an optimised Markov mobility model tailored to specific network topologies and user behaviours. They emphasise accuracy, complexity reduction, and practical usability, demonstrating the utility of MMCF through simulations that compare it with existing mobility models. They argue that their framework is particularly valuable for network operators to implement effective authorisation mechanisms, fraud detection systems, and adaptive network self-configurations, enhancing overall system security and performance.

Markov processes can be effective for simulating the behaviour of multiple actors or *agents* in some scenario, especially when their behaviours could be linked to the occurrence of other events. Shakeel et al. [47] address challenges in providing QoS to multi-class secondary user (SU)s with different delay requirements during *spectrum handoff* in cognitive radio network (CRN)s. They propose a hybrid spectrum access strategy combining interweave and underlay techniques, analysed using a continuous-time Markov chain (CTMC) model. They reveal that the hybrid approach significantly enhances spectrum utilisation, system throughput, and reduces transmission delays, especially for delay-sensitive SU traffic, compared to traditional interweave-only strategies.

One of the strengths of Markov processes is their ability to effectively model relatively simple and generalised systems, especially when the underlying factors are uncertain or poorly understood. Zhao et al. [48] present a novel strategy for dynamic spectrum management (DSM) in cognitive radio network (CRN)s by integrating spectrum prediction, user mobility prediction, and channel selection. Using a high-order hidden bivariate Markov model (H2BMM) for spectrum prediction and a high-order Markov model for mobility forecasting, the proposed method anticipates both channel availability and user movement. The channel selection process incorporates multiple decision factors, including future spectrum availability, to optimise allocation. Their simulated results support that the proposed strategy enhances spectrum utilisation, reduces handoff rates, and improves connection reliability.

Predictive modelling has been applied to channel allocation strategies to anticipate PU behaviour and enhance SU performance. Barnes and Maharaj [22] investigate the use of hidden Markov model (HMM) to predict PU behaviour, analysing the interdependence of spectrum sensing, channel allocation, and SU throughput in CRNs. The proposed proactive channel allocation significantly enhances SU throughput and reduces PU disruption compared to reactive approaches, especially under heavy traffic conditions. Their real-world spectrum measurements validate the simulation results, confirming the accuracy and effectiveness of the predictive modelling approach.

## 2.3.2 Multiagent Modelling

### Summary

*Multiagent modelling plays a central role in this thesis, enabling a shift from traditional, region-based spectrum management to more dynamic systems that consider the behaviour and needs of individual users or agents. Recent research has demonstrated how agents can optimise resource allocation in complex environments, such as satellite communications and device-to-device (D2D) networks, using market-driven schemes and reinforcement learning (RL) techniques. These models emphasise not only efficiency and throughput but also fairness under uncertainty, showcasing the computational challenges and trade-offs involved in real-world applications. Additionally, agent-based simulations of urban mobility highlight how understanding user behaviour at a granular level can enhance the predictive capabilities of spectrum management systems in dynamic environments.*

### Review

Conventional spectrum management approaches are based on allocating spectrum between regions. More advanced approaches may require us to solve the problem by considering the actions and needs of individual users or agents. As well as considering the spectrum needs and the actions of agents, agents can also play a role in managing a network in a more distributed and decentralised way. Li et al. [49] propose a market-driven, agent-based spectrum management scheme to optimise resource allocation in satellite communications. Recognising terrestrial agents as key intermediaries who lease spectrum from satellite providers and resell it to end-users, they introduce a dynamically optimal incentive system based on stochastic processes and contract theory. The goal is to maximise both terrestrial agents' and satellite systems' benefits by considering the satellite's marginal transmission capacity cost, the agent's effort, and market volatility. Their simulated results support the proposed scheme, showing improved profitability for satellite systems and efficient resource utilisation in the presence of changing market conditions and risks.

Often the goal of DSM is described as primarily being a resource allocation problem. The performance of solutions to resource allocation problems can be determined by measuring how effective they were in meeting the needs of the spectrum users. Buermann [50] explore multi-agent resource allocation, emphasising fairness and efficiency under uncertainty, particularly when resource availability is probabilistic rather than deterministic. The concept of *ex-ante envy-freeness* is examined as a fairness criterion where each agent prefers their own allocation to that of any other agent. It is shown that balancing fairness with social welfare optimisation is computationally challenging (strongly NP-hard), with the *price of envy-freeness* scaling with the number of agents. They present an integer programming model to compute optimal allocations when agents' valuation functions are linear and satiable, addressing both theoretical and



practical complexities in resource allocation under uncertainty.

Reinforcement learning (RL) techniques have also been applied to improve dynamic resource allocation in device-to-device (D2D) communication networks. Li et al. [51] propose a deep reinforcement learning (DRL)-based resource allocation algorithm utilising the advantage actor-critic (A2C) method for device-to-device (D2D) communication networks. Traditional resource allocation methods struggle with interference management and adaptive capabilities, leading to suboptimal throughput. The presented multi-agent A2C approach addresses these challenges by dynamically adapting resource allocation decisions through the cooperation of multiple agents. Simulation results show that this approach achieves substantially higher network throughput than baseline methods, including deep Q-networks (DQN) and multi-agent actor-critic (MAAC), demonstrating its effectiveness in D2D resource management.

Many spectrum management problems occur in dense, complex urban areas, so accurately modeling traffic at a low level can provide crucial insights into current and future spectrum demand. Szufel et al. [52] present a multi-agent, discrete-event simulation framework designed to model commuter behaviour and vehicle routing in large-scale transportation systems. They explore the interaction between individual and social knowledge and how these influence traffic flow and congestion. Two simulation approaches (queuing-based and delay-based) are compared to illustrate their impact on traffic analysis outcomes. Applied practically to commuter behaviour in Winnipeg, Canada, the model employs intelligent agents capable of adapting their travel decisions based on personal experience and social information, providing valuable insights into traffic optimisation and infrastructure planning.

### 2.3.3 Abstracted Agent Modelling

#### Summary

*In complex system simulations, it is often impractical to model every component in high detail, making abstraction a valuable technique to balance accuracy and computational efficiency. Various levels of abstraction such as macroscopic, microscopic, submicroscopic, and mesoscopic have been widely adopted in domains like traffic and mobility modelling to represent behaviours at differing granularities. Recent approaches also introduce adaptive abstraction methods, where the level of detail dynamically adjusts during simulation based on factors like agent density or behavioural consistency, significantly improving performance without overly compromising fidelity. These principles, although applied in different problem spaces ranging from blood coagulation to mobile network planning, offer transferable strategies that are relevant for scalable and efficient agent-based modelling in urban and spectrum environments.*

#### Review

It is not always practical or possible to model a system at the highest level of detail. In many common situations, the ability to dynamically adjust the level of detail for different parts of a simulation is crucial. For example, using faster, lower-detail settings may be appropriate in the early stages of exploration to define time ranges, locations, and initial parameters. Once these are established, higher-detail settings can then be used to fine-tune the parameters or to generate the final, more accurate results.

There is arguably a clear logic to this approach and its importance is reflected across the literature in simulation and agent modelling. Szufel et al. [52] describe four abstraction levels for traffic modelling as macroscopic, microscopic, submicroscopic, and mesoscopic. Macroscopic models focus on aggregate traffic flow characteristics like density, speed, and flow, without modelling individual vehicles. Microscopic models simulate individual vehicle-driver interactions, including acceleration, deceleration, and car-following behaviour. Submicroscopic models add finer details, incorporating driver psychology and vehicle performance, such as acceleration curves and responses to traffic signals. Mesoscopic models combine aspects of macroscopic and microscopic approaches, typically using probability distributions for aggregated traffic behaviour while allowing some individual vehicle dynamics. Additionally, hybrid mesoscopic models integrate both micro and macro approaches, applying microscopic modelling in areas of interest while using macroscopic modelling elsewhere for computational efficiency.

Many research problems relate to the challenges of serving the needs of dense populations of people living in large urban areas. A common challenge when solving these problems is finding effective and accurate ways to model the movements and *life patterns* of the population. With newer technologies and higher levels of computation available, many of these problems

are addressed at this level. However, prior to the availability of these technologies, far higher levels of abstraction were required. These same techniques can be applied as a way to apply computational abstractions as well, making this area of research still highly relevant to more modern approaches to urban city modelling. Markoulidakis et al. [53] introduce three mobility models tailored for analysing third-generation mobile telecommunication systems (third generation (3G)). The models are called the *city area model*, *area zone model*, and *street unit model*. They vary in granularity, from broad city zones to detailed street-level user behaviour. Each model facilitates analysis of system design aspects like location management, paging strategies, handovers, and resource allocation in cellular networks. Additionally, they propose an integrated mobility modelling tool, leveraging combined insights from these models to refine predictions about user mobility behaviour, enhancing accuracy and reducing computational complexity in planning and optimising mobile network infrastructure.

Many abstraction approaches in the literature define parts of the modelling rigidly and in advance, whereas Sarraf Shirazi et al. [54] propose an approach that can adjust the level of abstraction at *run-time* based on factors such as density or data resolution criticality. Specifically, they introduce an adaptive abstraction mechanism designed to optimise computational resources in spatial agent-based simulations. Observer agents detect groups of spatial agents clustering together over time and abstract them into single meta-agents. These meta-agents encapsulate both individual behaviours and structural properties. The validity of meta-agents is a trait that indicates whether the conditions for applying the abstraction are met. This validity is periodically checked, and if it is found not to hold due to changes in the dynamics, the meta-agents are decomposed back into individual agents. The approach is demonstrated through a blood coagulation simulation, where it effectively speeds up computation without altering the simulation's overall behaviour. This application is quite different to what the topic of this thesis, but the same principles could be applied to spectrum user modelling. Building on the same concepts, Shirazi et al. [55] also propose a method for reducing computational complexity in agent-based simulations through the abstraction of repeated interaction patterns. By leveraging ML techniques such as neural networks, genetic programming, and clustering, groups of agents exhibiting similar behaviour are abstracted into simpler representations. These abstractions are continually validated and adapted dynamically as the simulation progresses. They demonstrate their approach on biological systems, showing it maintains model flexibility while significantly lowering computational demands.

A similar idea could be applied to radio frequency (RF) propagation modelling where if nothing has changed significantly, the same measurements or calculations could be reusable. Franceschini et al. [56] address adaptive abstraction methods within agent-based simulations, aiming to balance simulation accuracy and computational efficiency. They use traffic flow modelling as a case study to examine various approaches to automatically switching between detailed (micro) and aggregated (macro) abstraction levels. They explore necessary conditions for transition-



ing between these levels, emphasising the importance of clearly defined models, initialisation strategies, and criteria for adaptive switching. By analysing different adaptivity variants, they demonstrate significant improvements in computational performance and increased insights into system dynamics.

## 2.3.4 Life Patterns & Timetables

### Summary

*Agent-based simulations offer a powerful approach to understanding and managing spectrum demand by modelling individual life patterns and timetables across space and time. These models allow planners to assess spectrum allocation effectiveness using agent-level quality of service (QoS) metrics, reflecting both active and passive device usage. Numerous studies have demonstrated the value of agent-based modelling in urban mobility and public policy planning, showing how realistic daily activity schedules enhance the accuracy and relevance of simulations. More recently, large language model (LLM)s have been integrated into such frameworks to generate semantically rich and interpretable activity patterns, enabling simulations that adapt more effectively to dynamic or unforeseen scenarios. Importantly, even in data-scarce environments, both traditional agent-based methods and LLM-augmented approaches can replicate realistic behaviours and support scenario testing, making them versatile tools for complex urban and infrastructure planning challenges.*

### Review

Spectrum management is a large geospatial problem with solutions needing to consider multiple dimensions such as frequency, space, time, and code. Generally, spectrum allocation decisions tend to be made based on the higher level view of patterns such as the density of spectrum demand and overall movement trends between regions. Conversely, one of the simplest ways to know whether the allocations are effective or not is to measure if each user or agent was satisfied with their allocation by quantifying it with criteria such as QoS. The higher level demand data is ultimately generated due to the demands of individual agents. An effective implementation for an agent level simulation could generate the higher level data more accurately and provide a way to assess the quality of allocations.

When modelling a spectrum management scenario, there are additional aspects of the agent life patterns to consider. Specifically, modelling for how spectrum demand varies across space and time generally corresponds to active or passive device usage. Active device usage would be an activity initiated by the agent such as sending a message or taking a call. Passive device usage generally would refer to spectrum usage that occurs as an overhead such as handovers between phone towers as an agent moves across a city, as well as signalling which is required to coordinate communications. Solutions which are able to produce spatial and time series data which includes information about active and passive spectrum usage are extremely relevant to this problem.

Across the literature there are many examples available where researchers have implemented agent-based simulations for a range of different applications. Doraki et al. [57] develop an agent-

based traffic simulation for the city Tehran to evaluate sustainable transport policies. They construct a synthetic population via simulated annealing to mimic daily commuting patterns and test four scenarios, including improved cycling infrastructure, enhanced public transit, and flexible work hours. The model predicts that better bike lanes and transit can cut private car use by up to 46%, while introducing flexible working hours reduces peak-hour traffic volumes by 47%, significantly shrinking travel distances. These findings highlight how policy interventions and timetable shifts in daily routines can alleviate congestion and pollution in urban settings.

Life-pattern modelling has also benefited from data-driven scheduling approaches for agent behaviour simulation. Drchal et al. [58] introduce an activity scheduler to generate realistic daily schedules for agents in urban mobility simulations. Using ML on travel survey data, their scheduler produces sequences of activities (e.g. home, work, shopping) and travel legs that match observed life-pattern distributions. The method is integrated into a multi-modal agent-based model, improving its accuracy in reproducing timing, duration, and mode choice of daily trips. Key findings show that a data-driven approach to timetabling agents' activities provides more realistic peak travel times and mode shares, enhancing the fidelity of the simulated urban mobility scenarios.

For urban city modelling in particular, there is an abundance of papers covering agent modelling. This may be because of how generalisable this strategy is to a wide range of problems. Divasson-J. et al. [59] present a comprehensive review of agent-based models applied to urban mobility. They survey recent studies to identify trends in modelling daily travel behaviour, activity patterns, and transportation scenarios. Their review discusses common methodologies for simulating individual movement (e.g. activity-based simulations, multi-agent traffic models) and how these models incorporate life-pattern data (like diaries and census information) and transit timetables. Key challenges are highlighted, such as calibration of agents' activity schedules and integration of emerging mobility modes. They conclude that agent-based simulations have become invaluable for testing urban transport policies, from congestion pricing to public transit scheduling, while calling for greater use of real-world data to improve scenario realism.

Modelling realistic urban behaviour has also been achieved in data-scarce contexts through synthetic population generation. Li et al. [60] propose a systematic framework to build agent-based mobility models for cities lacking detailed data. Using only open social-demographic information and aggregated mobile phone mobility data, they generate a synthetic population and travel demand for the city Brussels. Their approach ensures the simulated travel patterns (including activity chains and multi-destination tours) closely resemble real-life behaviour despite data scarcity. The resultant MATSim-based model is validated against observed mode shares, trip distances, and traffic counts, showing high fidelity. They demonstrate that even with limited data, realistic daily life patterns and timetables can be modelled, enabling scenario analysis (e.g. external travel demand or infrastructure changes) transferable to other



data-scarce urban regions.

Understanding how people move through and interact with urban environments is essential for designing effective transport systems and public policy. Recent research applies agent-based modelling to simulate realistic human behaviour at scale, offering new tools for studying mobility and planning interventions. Amiri et al. [61] introduce the patterns of life (POL) simulation, an agent-based model designed to generate realistic human daily activity patterns and mobility traces at scale. Grounded in behavioural theories (*Maslow's hierarchy* and *theory of planned behaviour*), the model's agents make decisions about activities (work, socialising, errands) and travel in a city environment. The framework uses open street map (OSM) data for realistic geography and can simulate up to 100k agents over years, producing trillions of location points. Key results demonstrate the model's ability to replicate diverse life schedules and interactions, and its efficient data generation via parallelisation. The POL simulator, released open-source, helps researchers study urban scenarios like public health outbreaks or transport interventions by providing high-fidelity synthetic activity-travel diaries when there is a lack of real data.

Some models integrate daily agent schedules into public transport planning to improve network design under changing demand. Manser et al. [62] extend a city-scale agent-based transport model (MATSim) to optimise public transit network design for Zurich. Their co-evolutionary algorithm iteratively adjusts bus routes, stop locations, and frequencies, while allowing dynamic travel demand responses as agents reroute or change mode in reaction to network changes. Unlike static transit planning methods, this approach integrates full daily activity schedules and mode choices. The optimised scenario provides a sparser network with smaller vehicles but much higher frequency, which increased ridership and reduced subsidies compared to the status quo. They illustrate that incorporating agents' daily timetables and adaptive behaviour into transit planning can reveal unconventional but efficient network designs, particularly relevant for future scenarios with automated vehicles.

Agent simulation research is often based on case studies such as modelling specific cities. Rasca et al. [63] apply an agent-based simulation as a decision-support tool to improve a regional bus line serving small cities in Agder. Adapting a mobility-as-a-service agent model, they simulate nine service scenarios (varying routes, stops, and schedules) to assess impacts on ridership, travel time, and operating cost. Agents represent residents with persona-based travel patterns, enabling realistic responses to each scenario's timetable changes. The model identifies four promising redesigns that substantially increase the number of users. Importantly, one scenario meets the transit agency's cost constraints while maximising usage. Their results demonstrate how incorporating agents' daily activity routines into transit planning can uncover optimal route and scheduling adjustments, providing a powerful tool for enhancing mobility in low-demand areas.

COVID-19 required urban agent modelling to understand virus spread in a complex environment. Manout and Ciari [64] integrate an agent-based travel model with an epidemic simulation to study COVID-19 transmission in Montreal. Using MATSim, they simulate each individual's daily activity schedule (home, work, school, leisure) and associated travel, then feed the close-contact interactions into an infection model (EpiSim). A synthetic population calibrated to Montreal's demographics and travel survey data provides relatively realistic life patterns. Their business-as-usual scenario reveals that regular activities at home, workplaces, and schools are major drivers of virus spread, whereas transit use contributed minimally to infections. Their findings highlight how agent-based mobility models with detailed timetabling of daily routines can evaluate public health interventions (like activity restrictions or staggered schedules) in urban pandemic scenarios. These findings have applications in spectrum management where access must be coordinated between many users.

Understanding how autonomous vehicles reshape daily travel requires models that incorporate realistic scheduling and time constraints. Feng et al. [65] develop an agent-based travel demand model with explicit space-time accessibility constraints to study how private autonomous vehicle (AV)s could change activity-travel patterns. The framework imposes individual travel time budgets (derived from activity diaries) on agents' daily schedules, ensuring their simulated activities and trips respect realistic time constraints. Applied to a Melbourne suburb (Clayton), the model examines scenarios with reduced travel time costs and higher road capacity due to AVs. Contrary to many expectations, results indicate that even with AVs, individuals' activity ranges remain bounded by time. Only modest increases in discretionary trip lengths and congestion were observed. This suggests that human daily scheduling habits (e.g. fixed work hours, limited free time) will continue to constrain mobility.

Detailed modelling of spectrum usage requires advanced knowledge of user behaviours and transport options. By drawing on research that directly addresses transit adoption and ridership modelling, similar approaches could be adapted to represent user spectrum usage with greater accuracy. Zannat et al. [66] integrate a proposed bus rapid transit (BRT) service into an activity-based agent transportation model to forecast ridership in Dhaka, Bangladesh. Calibrating the multi-agent simulator (MATSim) with local travel behaviour (including mode choice models using stated-preference surveys), they simulate how daily travel patterns might shift once the new BRT is introduced. Their scenario analysis varies BRT access modes and fare structures. They demonstrate the power of agent-based timetabled simulations in a data-poor, car-dominated context. By accounting for individuals' activity schedules and mode preferences, planners can better predict usage of a transit innovation and design supporting policies to maximise its adoption.

A recent study explores how large language model (LLM)s can be applied to the generation of





personal mobility trajectories. Wang et al. [67] propose an agent-based framework that uses LLMs to learn habitual activity patterns and daily motivations from semantically rich datasets, such as check-in records, and then use these insights to generate realistic activity schedules. The framework operates in two phases: first identifying self-consistent activity patterns through semantic interpretation, and then inferring motivations to guide activity generation. By incorporating retrieval-augmented strategies, the model produces interpretable outputs that align closely with real-world behaviours, enabling nuanced simulation of daily routines. Evaluation with urban activity data from Tokyo shows that the framework reproduces temporal and spatio-temporal aspects of mobility with strong fidelity, while also adapting to scenarios such as the COVID-19 pandemic. This work demonstrates the potential of LLM-based agent modelling to enhance life-pattern simulations, bridging statistical approaches and behavioural realism in ways highly relevant for urban spectrum demand modelling.

## 2.4 Demonstrations

### 2.4.1 Context

The challenge addressed in this chapter stems from the limitations of static spectrum allocation, which fails to account for the spatial and temporal variations of real-world demand. As outlined in the introduction and approach, achieving efficiency in spectrum usage requires not only allocating spectrum but doing so dynamically in response to changing conditions. Predicting when and where demand will occur is difficult because communication patterns are irregular, user mobility is variable, and the allocation process lacks a clear guide for adjustment. The inability to anticipate demand shifts leads to inefficiencies, motivating the need for modelling techniques that can provide synthetic, adaptable, and scalable demand data suitable for testing spectrum management solutions.

The literature highlights a range of modelling approaches capable of addressing this problem. Markov processes offer lightweight but effective means of simulating mobility and handover events, providing coarse but useful predictions of demand fluctuations. Multiagent modelling advances this further by considering individual actors, enabling decentralised solutions and fairness in allocation. Abstraction techniques address the computational burden by dynamically adjusting levels of detail, while life-pattern and timetable-based approaches create realistic behaviour-driven datasets that better capture the rhythms of urban mobility. Collectively, these methods demonstrate how statistical analysis, behavioural modelling, and computational techniques converge to provide feasible tools for understanding and simulating dynamic spectrum demand.

The demonstrations in this section put these ideas into practice by translating theoretical approaches into practical modelling exercises. The first demonstration applies a region-level Markov chain model weighted by census and mobility data to generate synthetic but scalable demand trends across urban areas. The second demonstration builds on this with an agent-level approach, combining statistical methods and timetable-driven simulations to capture both individual behaviours and aggregated demand patterns. Together, these demonstrations link the ideas from the literature with practical applications in dynamic spectrum management (DSM), showing how abstract models can be turned into workable tools for testing allocation strategies in realistic but manageable settings.

### 2.4.2 Summary

Demonstration 1 outlines a method for generating synthetic time-varying radio spectrum demand data to test spectrum allocation solutions when real-world data is unavailable or unsuitable due to sensitivity, scale, and resolution constraints. The approach uses a Markov chain with initial states weighted by Australian census population data by sub-region, with time-

varying transition weightings based on traffic patterns representing daily movements between home and work. Chosen for its simplicity and compatibility with available population data, the Markov chain method is computationally efficient, well-supported in theory, and capable of representing real or conceptual systems at large geospatial scales. While the model is coarse compared to real-world complexities, it offers valuable insight for conceptual technology evaluation and can use open-source data for broad coverage. Future work could explore more advanced Markov models, greater state complexity, and a wider variety of simulated day types for improved realism.

Demonstration 2 explores generating agent movement data from urban life patterns to support testing and developing spectrum management solutions. Two main simulation methods are considered: a statistical approach using overall trends from a continuous-time Markov chain (CTMC) to guide agents' movements, and a more detailed agent-based modelling approach where agents act according to realistic conditions and personal goals. The statistical method produces large-scale but lightweight simulations, while the detailed method, although more complex, provides a far higher level of detail, making it more suitable to model user spectrum usage. It was simplified through assigning agents attributes and generating realistic timetables with a language model. Both methods are computationally efficient and can scale to represent large populations, producing synthetic datasets in the desired format without privacy concerns of real-world data. However, their simplicity limits realism, and verifying accuracy is challenging since the necessary real-world tests are impractical. Future improvements could include higher-fidelity 3D environments, agent interactions, and detailed traffic simulations, offering significant potential for planning and evaluating spectrum management strategies in the absence of live testing opportunities.

### 2.4.3 Structure

This section presents two main demonstrations of spectrum demand modelling.

1. **Region-level Spectrum Demand Model** — A synthetic spectrum demand generation method using a Markov chain weighted by Australian census data, with transitions modelled according to time-varying mobility patterns. Demonstrates daily demand shifts across regions in an urban environment.
2. **Agent-level Spectrum Demand Model** — A hybrid approach combining statistical and agent-based modelling to simulate individual movement and spectrum usage patterns in an urban environment, using agent attributes, timetables, and routing to generate detailed synthetic datasets.

#### 2.4.4 Demonstration 1: Region-level Spectrum Demand Model

Table 2.1 summarises the structure of Demonstration 1, which focuses on a region-level spectrum demand model. It outlines the main objectives, the modelling methods applied, and the data sources used. Each objective highlights a different aspect of how synthetic demand is generated and varied across time and space, while also indicating how coarse-grained simulations can support spectrum allocation testing. This overview provides a quick reference to the scope and components of the first demonstration.

Demonstration	Objective	Key Method	Data Inputs
1.1	Population movement modelling	Markov chain	Statistical population data
1.2	Time-varying demand generation	Continuous-time Markov chain (CTMC)	Mobility patterns from traffic data (TomTom Adelaide [68])

Table 2.1: Summary of Demonstration 1 – Region-level Spectrum Demand Model

## Demo 1.1: Continuous-Time Markov Chain of a City

One of the highest-level models that can be implemented cheaply to show shifts in spectrum demand is a Markov chain. Rather than looking at the detailed positions of each spectrum user, they can shift between different regions represented in a Markov chain representation of an urban area, over time. The Markov chain below in Figure 2.4.4.1 shows the different states of movement between the regions (2 of each kind). The purpose of this model is to provide a simple representation of spectrum demand to use as an input to basic spectrum management solutions. In some cases, this higher level spectrum demand info may be all that is available so it is applicable even in current generation networks. There are further computational and latency advantages to be able to work with less detailed spectrum demand information.

### Model Definition

- **Regions:** 4 regions divided into 2 types (city and suburb)
- **Users:** 3 user types with shares of the population
  - City workers living in suburbs: 75%
  - Majority suburb inhabitants: 20%
  - Majority city inhabitants: 5%
- **Spectrum:** User demand measured in channel units




Figure 2.4.4.1: Demo 1.1: 2-region Markov chain

### Demo 1.2: User movement modelling

The chart in Figure 2.4.4.2 illustrates the temporal distribution of users across the four regions (CBD 1, CBD 2, Suburb 1, Suburb 2) at key points of the day, based temporally varying population movements derived from real-world traffic data from TomTom for Adelaide [68]. The vertical bars represent the number of users in each region, divided between idle and busy states, with busy users shown in each plot. The four subplots capture user movements across early morning, morning rush, afternoon, and night. In the early period, most users are in the suburbs and idle, whereas during the morning rush a substantial movement into the CBD occurs alongside increased busy usage. By the afternoon, demand peaks in the CBD, reflecting work-related usage, before partially shifting back towards the suburbs in the evening and night. This sequence of transitions highlights the daily migration of users and their changing activity levels, providing a link between population movement and potential network load.




Figure 2.4.4.2: Demo 1.2: Temporal user movement across CBD and suburban regions, showing idle and busy users at four key times of the day



### Demo 1.3: User demand modelling

The chart in Figure 2.4.4.3 builds upon the user movement data by explicitly quantifying network demand in terms of busy users and their required channel units. Each subplot again represents a different time of day, with the left y-axis tracking the number of busy users and the right y-axis showing corresponding demand units. During early hours, demand is low and concentrated in suburban regions. By the morning rush, demand rises sharply in the CBD as workers begin network-intensive activities. The afternoon period shows a sustained high level of demand across both CBD and suburbs, indicating overlapping work and personal usage. At night, although the majority of users have returned to the suburbs, demand remains high due to entertainment-driven activities. These demand patterns demonstrate how user mobility directly drives spectrum load fluctuations across time and space, which is central to effective spectrum management strategies.

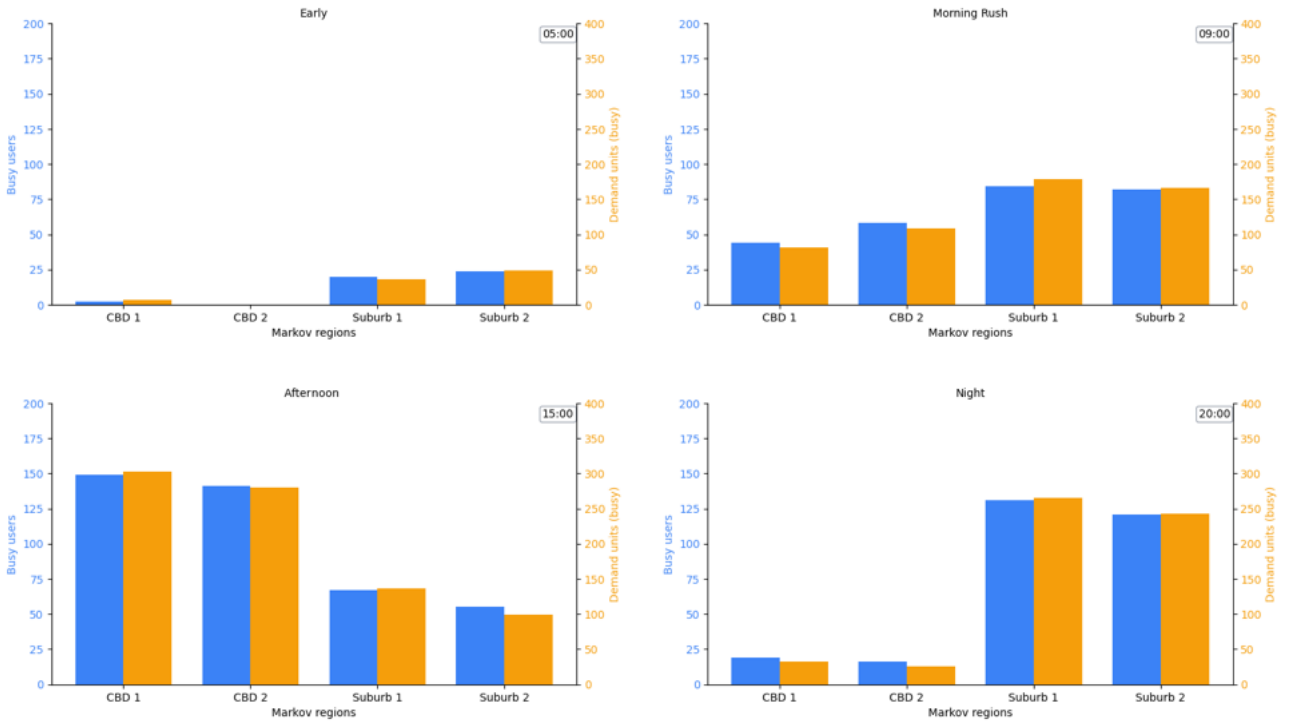


Figure 2.4.4.3: Demo 1.3: Network demand evolution expressed as busy users and required channel units across regions and times of day

### 2.4.5 Demonstration 2: Agent-level Spectrum Demand Model

Table 2.2 presents a summary of Demonstration 2, which develops an agent-level spectrum demand model. It describes the objectives of the demonstration, the key methods used to simulate mobility, navigation, and behavioural patterns, and the data inputs required for each stage. The demonstration combines statistical mobility models, geospatial routing, and timetable-driven agent behaviour to create detailed synthetic datasets. This summary highlights the hybrid nature of the approach and its ability to capture both movement and usage dynamics at the individual level.

Demonstration	Objective	Key Method	Data Inputs
2.1	Statistical mobility simulation	CTMC-based movement patterns	Population statistics, mobility data
2.2	Geospatial navigation modelling (low resolution)	Routing based on geospatial networks	OSM road and building data
2.3	Geospatial navigation modelling (high resolution)	Routing based on geospatial networks	OSM road and building data
2.4	Agent behaviour and spectrum usage modelling	Attribute-based scheduling	Agent profiles and LLM-generated timetables

Table 2.2: Summary of Demonstration 2 – Agent-level Spectrum Demand Model



## Demo 2.1: CTMC-Based Traffic-Weighted Census Population Model

Figure 2.4.5.1 shows a city region defined by geographic bounds sourced from open street map (OSM) [69]. The data includes information about roads, buildings, and other relevant geographic features which enables the simulation to model urban mobility more realistically. The initial and base states of the population density modelling are derived from census data. The Australian Bureau of Statistics (ABS) collects census data at multiple levels, from very fine Mesh Blocks aggregated into Statistical Areas Level 1 (SA1), which hold between 200 and 800 people, and further aggregated into SA2, SA3, SA4 for broader spatial contexts [70]. The time-varying distribution is produced by applying traffic data at different times of the day using the technique of a continuous-time Markov chain (CTMC), where the census-defined statistical areas form the basis of the Markov states. This modelling could be applied to determine appropriate initial positions for agent-based simulation and to statistically decide destinations for agent pathing. At the beginning of the simulation agents could be placed according to population density distribution and destinations generated by a probabilistic model that considers identity, current location, and time of day.




Figure 2.4.5.1: Demo 2.1: Midnight and midday population distribution

## Demo 2.2: Markov Chain-based Agent Simulation at Low Resolution

**Scenario:** Agents represent spectrum users leaving the Adelaide CBD towards the suburbs after a CBD-based event.

**Setup:** Movements are generated by routing agents between Markov states, each mapped to a location. Paths and travel speeds are determined by the Valhalla routing engine [71]. Results are aggregated into a coarse density grid and shown as a plasma coloured heat map to capture spatial density.

**Frame 1 (00:00):** At the initial time (Figure 2.4.5.2), agents are dispersed across the city centre and parklands.

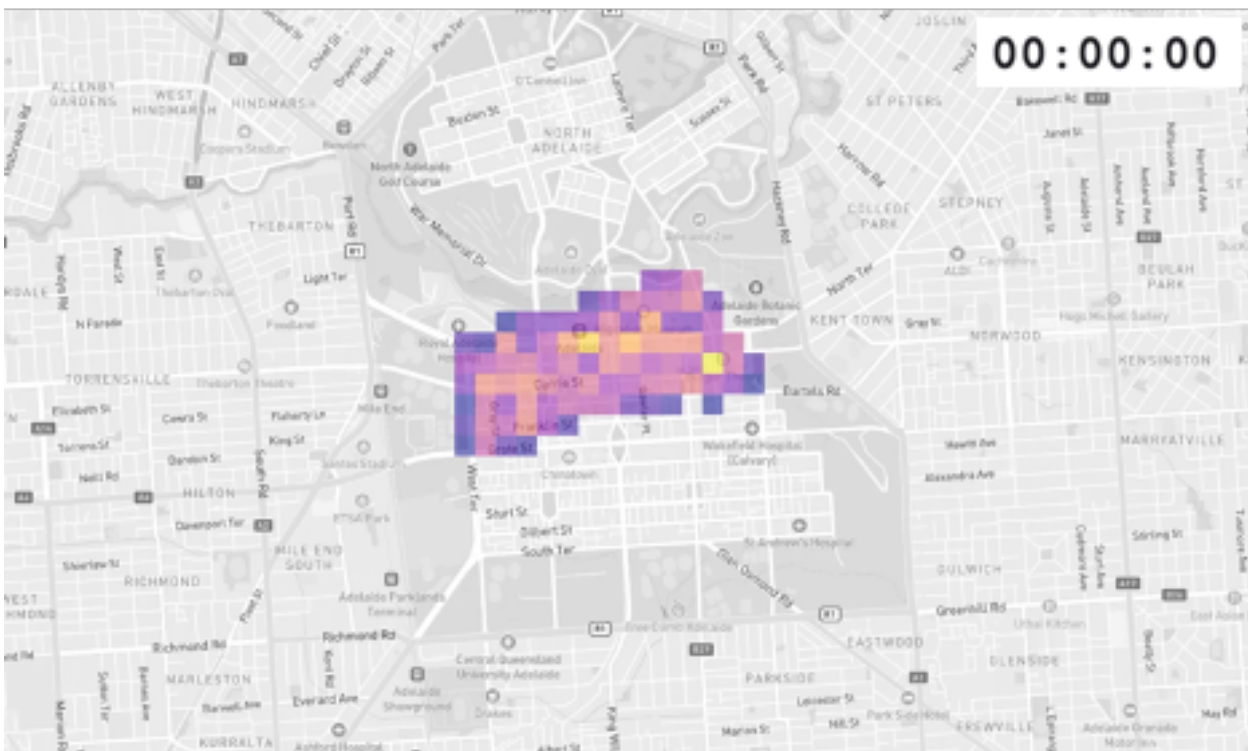


Figure 2.4.5.2: Demo 2.2: Low resolution density grid at initial time 00:00

**Frame 2 (05:00):** Five minutes later (Figure 2.4.5.3), density decreases within the CBD and increases in surrounding parklands, reflecting the outward movement of agents.

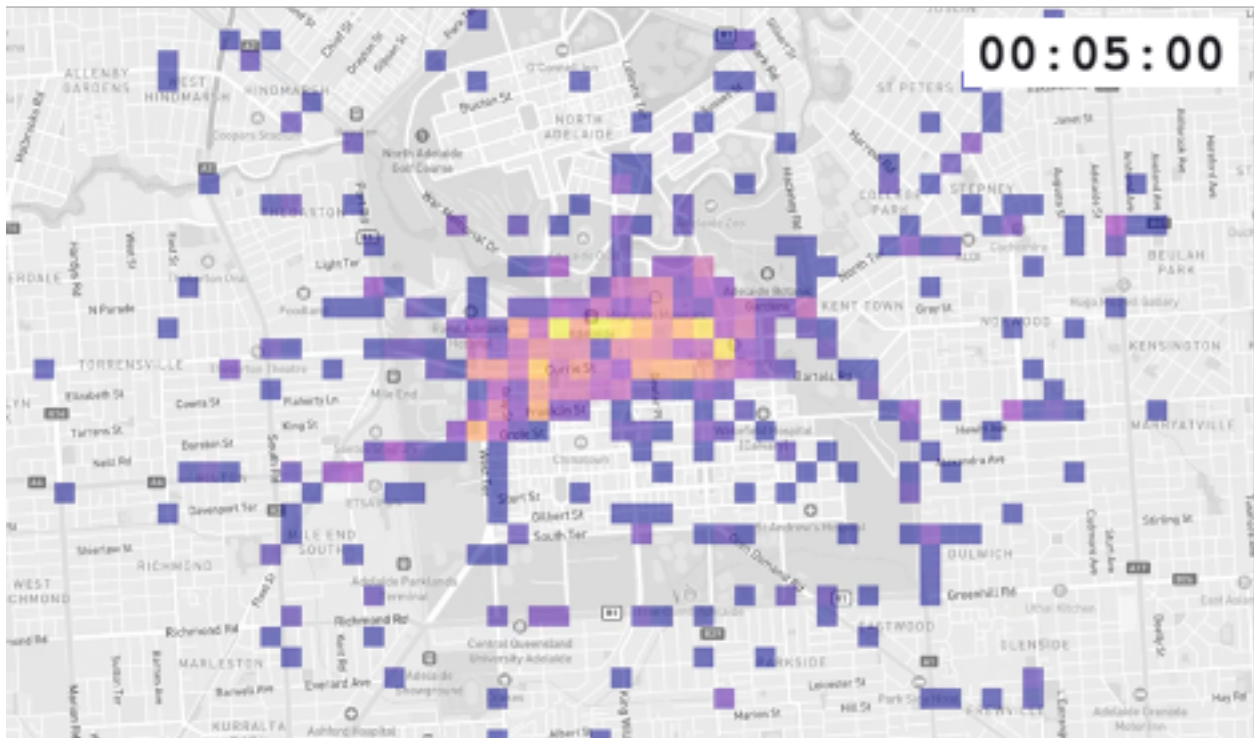


Figure 2.4.5.3: Demo 2.2: Low resolution density grid at 05:00



### Demo 2.3: Markov Chain-based Agent Simulation at High Resolution

**Scenario:** A different post-event movement is simulated which results in a large population shift from North Adelaide into the Adelaide CBD.

**Setup:** A higher resolution grid is used on this occasion to capture finer detail. Movements again follow Markov routing with Valhalla-generated paths. A plasma coloured heat map is once again applied for visualisation of the shifting population density.

**Frame 1 (00:00:00):** At the starting point (Figure 2.4.5.4), agents are relatively evenly distributed across the CBD and North Adelaide. Density is low and widespread.



Figure 2.4.5.4: Demo 2.3: High resolution density grid with agents spread across Adelaide CBD and North Adelaide at 00:00:00



**Frame 2 (00:07:00):** 7 minutes later (Figure 2.4.5.5), density increases within the CBD, showing how this relatively simple Markov state change agent modelling can show changes in population density.



Figure 2.4.5.5: Demo 2.3: High resolution density grid with agents concentrated in the Adelaide CBD at 00:07:00

## Demo 2.4: Schedule-based Agent Simulation

Recent work has demonstrated the potential of applying large language model (LLM)s to agent-based modelling. For example, Wang et al. [67] show how LLMs can learn habitual activity patterns from semantically rich datasets and generate realistic daily schedules for use in mobility and life-pattern simulations.

**Setup:** The process uses prompts to a LLM to generate agent schedules. Inputs are agent attribute mappings, which encode statistical data about the population and environment. These mappings define agent types and their likely behaviours, derived from statistical distributions of the society being modelled.

**Schedule Generation:** The LLM (ChatGPT in the example) generates daily timetables of activities with associated phone usage. Example: Table 2.3 shows an agent’s schedule, linking time, activity, location, and phone usage intensity. Outputs are then reviewed and edited by humans to ensure logical consistency and correctness. This step makes it feasible to scale across many environments without manually constructing schedules.

Time	Activity	Location	Phone Usage
6:00 AM	Morning Routine	Home	Low
7:00 AM	Commute to work (audiobook)	Car	Medium
8:30 AM	Morning Work Focus	Office	Low
12:00 PM	Lunch Break (casual browsing)	Office Cafeteria	High
1:00 PM	Afternoon Work	Office	Low
5:00 PM	Commute back home	Car	Medium
6:30 PM	Evening Hike	Local Trail	Low
8:00 PM	Cooking Dinner	Home	Low
9:00 PM	Watching TV	Home	Medium
10:30 PM	Sleeping	Home	None

Table 2.3: Demo 2.4: Example agent timetable of phone usage



## Spectrum Demand Modelling:

Spectrum demand is derived by combining:

- Agent movements (spatial transitions across home, commute, work, leisure)
- Device usage patterns (low, medium, high usage levels per activity)

Markovian processes can be used to model population-level behaviour. Logical reasoning or real-world data can refine spectrum usage levels. Figure 2.4.5.6 demonstrates continuous phone usage over 24 hours. This method enables realistic synthetic datasets at scale, reducing the need for large manual teams.




Figure 2.4.5.6: Demo 2.4: Continuous phone usage over a full day from midnight to midnight

## 2.5 Discussion

### 2.5.1 Summary

Chapter 2 has four sections: *Introduction*, *Approach*, *Literature*, and *Demonstrations*.

The *Introduction* motivates modelling dynamic spectrum demand across an urban setting. The broad electromagnetic spectrum is distinguished from the subset relevant to telecommunications, and it is argued that practical dynamic spectrum management (DSM) work must be explicit about which parts are in scope and at what level of modelling complexity. The core problem is framed as defining demand well enough to guide dynamic reallocation. Static and command-and-control assignment can fill all the frequency band allocations but still fail on usage efficiency because demand varies by time and place.

The *Approach* section states that conventional spectrum management uses long-term, static allocation, and then sets a proposition and guiding question for the chapter. The proposition is that ideal allocations should track demand across frequency, space, time and code. The research question posed is how to represent that demand so a DSM system can be simulated. Four focus areas are named as the focus for this chapter: Markov chains & processes, multi-agent modelling, abstracted agent modelling, and life patterns & timetables.

The *Literature* section reviews these four areas. The contribution of Markovian methods to mobility and channel availability is described. Multi-agent work is summarised, with allocation treated as a distributed decision problem. Market mechanisms, such as satellite and terrestrial contracting, are discussed. Reinforcement learning (RL) is then examined for device-to-device (D2D) resource control, with fairness and efficiency trade-offs considered. The non-deterministic polynomial-time (NP)-hardness of allocation under uncertainty is also acknowledged. Abstraction strategies are reviewed, showing how switching between macro, micro and meso levels adaptively makes large simulations tractable. *Life-pattern* modelling is then surveyed. Daily agent timetables are discussed, capturing active and passive usage. This enables evaluation of quality of service (QoS) at the user level and supports planning in data-scarce cities.

The *Demonstrations* present two modelling tiers. A region-level continuous-time Markov chain (CTMC) uses Australian census sub-regions and time-varying transition weights guided by traffic rhythms to generate synthetic, daily demand, showing morning inflows to the CBD, midday peaks, and evening returns to the suburbs, suitable for testing allocation logic with lightweight inputs. An agent-level model then produces population density data using two different approaches: (i) CTMC-driven destinations with realistic movement rules, and (ii) a lower-level, hybrid approach where richly attributed agents follow LLM-generated (and human-checked) daily timetables.



## 2.5.2 Key Findings

### Finding 1: The need for dynamic spectrum modelling

Dynamic spectrum modelling is essential because static allocation cannot capture the temporal and spatial variability of spectrum usage. The *Approach* explains that conventional allocation is static, but as demand fluctuates with daily cycles, mobility, and location, it would ideally be handled dynamically. The *Literature* shows how stochastic methods such as Markov processes and agent-based techniques can predict and adapt to these fluctuations by modelling user mobility, handovers, and channel availability. The *Demonstrations* explore these ideas by showing daily demand peaks and shifts reproduced through region-level Markov models and agent-level simulations. This demonstrates that dynamic approaches can represent spectrum demand in a more realistic way.

### Finding 2: The trade-off between efficiency and realism

Spectrum demand models must balance efficiency with realism. To achieve this, the modelling complexity must be carefully matched to the purpose of spectrum management. The literature shows that Markov approaches are computationally efficient but coarse, while agent-based methods capture detailed life patterns and device usage at the cost of high computational load. Abstraction strategies such as mesoscopic and adaptive grouping are presented as ways to preserve fidelity while reducing computational burden. The *Demonstrations* illustrate this trade-off: region-level Markov chains generate computationally efficient but less detailed results, while agent-level models deliver higher behavioural realism and resolution but require simplifications to remain computationally feasible and to overcome data limitations.

### Finding 3: The value of hybrid approaches

Hybrid approaches may provide the strongest practical outcomes because they integrate complementary modelling layers. This is due to the goal of linking demand representation to allocation decisions, requiring insight at both aggregate and individual levels. The literature shows that hybrid methods combining statistical and agent-based simulations improve prediction and scenario testing. The *Demonstrations* explore these ideas further: region-level models capture broad demand shifts across the city, while agent-level simulations represent detailed mobility and usage patterns. The main idea explored is that layered and hybrid strategies may provide the a robust foundation for practical dynamic spectrum demand modelling.

### 2.5.3 Limitations

Although dynamic spectrum modelling provides a more realistic representation of temporal and spatial demand variability, its accuracy remains constrained by the availability and quality of input data. The demonstrations shown for region-level models relied heavily on census and traffic datasets, which may not fully capture transient or irregular usage patterns such as event-driven surges or unusual mobility of users. Agent-level models can include more behavioural realism, but they remain synthetic and are not easily verifiable against real-world data due to privacy and measurement challenges. While dynamic approaches better approximate demand than static allocation, their predictive reliability is bounded by data limitations and the extreme challenge of how to ground results in reality.

The trade-off between efficiency and realism imposes practical limits on how spectrum demand models can be applied at scale. Coarse models, such as Markov chains, offer computational tractability but risk oversimplifying user behaviour and overlooking localised interference dynamics. Conversely, detailed agent-level simulations are computationally intensive and require abstraction or simplification, which can obscure fine-grained effects they aim to capture. The adaptive abstraction strategies discussed in the literature alleviate this to some extent, but transitioning between levels of detail may introduce further complexities. No single modelling approach can simultaneously achieve the highest quality, scalability, and efficiency, constraining their utility for real-time spectrum management. However, in the demonstrations, the hybrid approaches did appear to strike a good balance, especially when the trade-offs are configurable based on the application (speed versus quality).

Hybrid approaches strengthen modelling robustness by combining complementary techniques, but they also inherit the weaknesses of each individual method. Integrating statistical models with agent-based simulations requires careful calibration to avoid mismatches in scale, resolution, or assumptions. The demonstrations highlight this by linking city-wide Markov-driven demand trends with timetable-based agent behaviour. However, such integration relies on synthetic datasets that may propagate biases between them. Furthermore, hybrid frameworks add complexity in implementation and interpretation, demanding greater computational resources and expert or human oversight. While hybrid strategies broaden the applicability of spectrum demand modelling, they remain vulnerable to compounded uncertainty and limited transparency in how aggregate and individual behaviours interact.

### 2.5.4 Future Work

Progress in spectrum demand modelling could require moving beyond reliance on coarse and static datasets. While existing methods demonstrate that synthetic demand can be generated from regional data and traffic data, these inputs are insufficient to reflect transient surges or irregular mobility patterns. Future research should focus on incorporating richer real world data, including event-based demand spikes, high-resolution mobility traces, and a vertical distribution of users within urban and residential environments. Extending models into three-dimensional (3D) space would capture the impact of buildings, elevation, and multipath propagation, supporting the construction of realistic digital twin environments that more accurately replicate urban spectrum usage.

Improving modelling techniques also demands solutions to the trade-off between computational efficiency and behavioural realism. Lightweight statistical methods are computationally efficient, but risk oversimplification, whereas detailed simulations of agent behaviour become impractical at scale. A promising avenue is the development of adaptive abstraction strategies that dynamically adjust the level of detail depending on context, preserving fidelity in dense interference zones while reducing complexity elsewhere. Machine learning (ML) methods could further support real-time decisions about when to apply higher or lower resolution, allowing spectrum demand simulations to run efficiently while still capturing critical fine-grained effects.

Finally, hybrid approaches that combine statistical trends with agent-based modelling offer the most potential for realistic outcomes, but they also introduce challenges of calibration and integration. Future work could refine hybrid frameworks by drawing on insights from urban mobility research, transport modelling, and digital twin systems. Including these models within live data streams, such as sensor networks and anonymised mobility datasets, would enable iterative calibration and validation. Over time, such integration could support adaptive digital twin models that reflect not just average behaviour, but real-world spectrum demand as it evolves across 3D urban environments.

## 2.6 Next Steps

Having examined how to model spectrum demand, the focus of this thesis now turns to how spectrum allocations can be made when the demand is known. Allocation is treated as a constrained resource problem with capacity limits, contiguity requirements and interference protections. Key objectives are throughput, fairness, spatial reuse and reconfiguration cost. The approach sets objectives first, makes constraints explicit, then builds a tunable allocation framework that can trade efficiency and equity. The spatial characteristics of demand are used to guide partitioning so that frequency reuse and interference avoidance may suit the actual RF propagation environment well. Outcomes are assessed with throughput and fairness metrics, and robustness checks account for delay or bias in availability information. The aim is a clear, adjustable method that stays stable under imperfect data.

## Chapter 3

# Spectrum Allocation Mechanisms

*This chapter examines mechanisms for allocating radio spectrum in dynamic spectrum management (DSM) contexts, integrating concepts from resource allocation theory, spatial clustering, fairness metrics, and Voronoi tessellations. It reviews historical and practical approaches to distributing spectrum as a contested, interference-prone resource, highlighting trade-offs between efficiency, fairness, and computational feasibility, and the role of metrics such as Jain's index in quantifying allocation outcomes. A spatial clustering method based on Voronoi tessellation partitioning is explored as a strategy for grouping users and optimising reuse, interference management, and coverage. The chapter presents demonstrations of allocation strategies and Voronoi-driven dynamic partitioning using synthetic mobility data, concluding with recommendations for DSM solutions that emphasise adaptability, fairness, and scalable spatial coordination.*

## 3.1 Introduction

Spectrum allocation mechanisms have evolved over many decades, since spectrum was first recognised as a valuable resource for telecommunications. In the early stages of spectrum management, spectrum was treated more as a technology than a scarce resource [5]. As a consequence, patents and concepts such as intellectual property were applied to restrict others' usage of radio spectrum. Due to numerous issues and incidents, as described by Tripathi [2], governments moved to control the resource, treating it like any other natural resource. Initially, governments prioritised securing access for their own purposes. However, because commercial applications existed before government control, it became clear that spectrum licensing could both generate government revenue and establish a regulated market in which an economically viable telecommunications industry could operate. Over time, technical and regulatory conventions emerged that balance efficiency, fairness, and the prevention of interference, forming the foundation for modern spectrum management.

When it comes to allocating spectrum, there are some common strategies. The first is shared access, in which anyone is allowed to use it, as in the industrial, scientific and medical (ISM) bands that enable globally popular protocols like Wi-Fi and Bluetooth. The second is licensed allocation, where exclusive rights are granted to users following interference analysis. An alternate approach is dynamic allocation, which allocates and deconflicts spectrum in real time, similar to how commercial mobile networks operate. At the primary allocation level [1], spectrum allocation is a resource distribution problem, as access to spectrum must be assigned among competing agents in a way that balances efficiency, fairness, and computational feasibility. Metrics play a critical role here, particularly fairness metrics such as Jain's index, proportional fairness, and utility-based measures, which allow objective evaluation of allocation outcomes. These metrics enable adaptive and responsive dynamic spectrum management (DSM) systems by quantifying whether allocations are equitable across users, a necessity in both static and dynamic allocation scenarios.

Spatial clustering and spatial partitioning are powerful tools for optimising spectrum allocation, as they allow frequency reuse and interference management to be tailored to actual demand and user locations. Clustering techniques, such as k-means or density-based methods, can group spectrum users based on proximity to simplify allocation and maximise reuse. Voronoi tessellations offer an alternative by dividing space into regions around seed points, often based on population density or network topology, creating tessellating polygons that align with coverage areas. In telecommunications, these methods support dynamic and adaptive allocation by creating geographically coherent groups for frequency, time-slot, or code reuse. They are especially relevant in DSM contexts where both centralised and decentralised systems may require optimal spatial division before adapting the process to real-world constraints.

## 3.2 Approach

The goal of this chapter set the foundation for much of the work in the rest of the thesis. As previously stated in the thesis overview, this research started with an industry problem of: *“How to allocate radio spectrum between users with competing needs at the upper allocation level?”*. This chapter covers the main findings made in the process of answering that question. The main limitations encountered for a DSM system arose from the difficulty and technical complexity of defining the spectrum allocation problem itself.

### 3.2.1 Proposition

There are many possible allocations for radio spectrum with numerous criteria available for assessing what aspects of a solution are good and bad.

### 3.2.2 Question

Which allocation mechanisms are suitable for DSM solutions?

### 3.2.3 Topics

**Resource Distribution & Allocation** - Provides some background for spectrum allocation at a higher level when considering access to radio spectrum as a resource.

**Spatial Clustering Techniques** - Common methods to group regions or clusters of users together to make allocations simpler and more feasible.

**Fairness Metrics** - A tangible and measurable way to determine what makes an allocation good or bad.

**Voronoi Tessellations** - An interesting technique often considered in spectrum allocation research to divide spectrum allocations, especially in satellite communications for multibeam assignment.

## 3.3 Literature

### 3.3.1 Resource Distribution & Allocation

#### Summary

*Resource distribution and allocation is a foundational concern in spectrum management, as it involves determining fair and efficient access to the limited and interference-prone resource of radio spectrum. While spectrum itself is not a consumable resource like food, gas, or water, access to it must be carefully managed to avoid signal interference, making its allocation a critical challenge. This section explores both theoretical and practical approaches to resource allocation, including algorithmic fairness, game theory, decentralised optimisation, and economic mechanisms. Through a survey of foundational studies and emerging trends, it highlights the interdisciplinary nature of the field and its relevance to diverse domains such as computing clusters, smart grids, robotics, and next-generation wireless networks.*

#### Review

Resource distribution and allocation is a broad research area concerned with the efficient and fair assignment of limited resources to competing needs. It spans fundamental theoretical frameworks as well as diverse application domains. It is a critical part of this thesis because it is essentially the main goal of spectrum management. It is important to clarify the terminology of words like *resource* in a topic like spectrum management because spectrum itself is not a resource in the same way that an item such as food is considered to be a resource. Access to spectrum is the resource and the reason it has to be distributed is due to the interference or *pollution* Tripathi [2] of the spectrum that occurs when it is accessed.

Resource allocation among multiple agents involves intricate trade-offs between efficiency, fairness, and computational feasibility, requiring a deep understanding of these trade-offs to design effective mechanisms. Chevaleyre et al. [72] provide a foundational overview, discussing preference modelling, social welfare metrics, and algorithmic complexity, while reviewing various allocation procedures, including auctions and negotiations. They introduce multiple languages for expressing preferences and measures for evaluating outcomes, alongside complexity results that support the feasibility of allocation methods. Key application domains such as satellite scheduling and grid computing highlight the practical relevance of this foundational work and motivate ongoing research.

Strategic behaviour in resource allocation requires mechanisms that account for self-interest and potential manipulation by participating agents. Dash et al. [73] advocate integrating game-theoretic principles with computational algorithms to address strategic behaviour in distributed





environments. They stress the importance of incentive-compatible protocols like auctions or pricing schemes that ensure truthfulness and efficiency. They argue that resource allocation algorithms must consider self-interested agents and propose computational mechanism design as a bridge between economic theory and algorithmic implementation. This interdisciplinary approach finds applications in public goods provision, network resource allocation, and task scheduling.

The trade-off between fairness and efficiency is quantitatively analysed by Bertsimas et al. [74], who introduce the *price of fairness* to capture the efficiency loss from imposing equity constraints. Framing the allocation problem as an optimisation task, they explore models such as max-min and proportional fairness, computing the worst-case efficiency losses. Their findings reveal that fairness often comes at a modest cost, making equitable allocations feasible in many practical settings. Their analysis provides a rigorous foundation for balancing social objectives with performance in real-world resource distribution.

Decentralised control mechanisms are essential in scalable resource allocation systems, particularly in networked environments with limited coordination. Kelly et al. [75] present a decentralised approach to network resource allocation by modelling it as a global utility maximisation problem under capacity constraints. They introduce proportional fairness and show how shadow pricing can guide distributed rate control to achieve fair bandwidth sharing.

In large-scale decentralised systems, distributed optimisation plays a key role in enabling agents to make efficient allocation decisions without central coordination. Doostmohammadian et al. [76] examine distributed optimisation algorithms that enable agents to compute resource allocations through local interactions. They address the distributed resource allocation (DRA) problem and explore consensus, primal-dual, and gradient-based methods suited for decentralised computation. They compare linear and non-linear strategies across domains like power grids and robotics, highlighting trade-offs in feasibility, convergence, and robustness. Synthesising theory with practical constraints, they outline current capabilities and open challenges in scaling distributed allocation algorithms in large, dynamic multi-agent systems.

Fair and efficient resource allocation is a central challenge in shared infrastructure systems where competing users' or agents' needs must be balanced against limited capacities and resource availability. Ghodsi et al. [77] address the challenge of resource distribution in heterogeneous computing clusters, introducing dominant resource fairness (DRF) as a generalisation of max-min fairness across multiple resource types. Their fairness mechanism equalises users' dominant resource shares (the resource each user consumes most heavily) to ensure equitable access. They demonstrate that DRF satisfies properties such as sharing incentive, envy-freeness, and strategy-proofness, making it a robust allocation scheme. Implemented in the Mesos cluster manager, DRF improved job completion times and system utilisation compared to traditional

schedulers, significantly influencing data centre scheduling practices.

In smart grid systems, user-centric allocation methods are necessary to ensure efficiency and autonomy in energy consumption. Mohsenian-Rad et al. [78] present a non-cooperative game-theoretic framework for scheduling electricity usage in smart grids, modelling each user's load shifting decisions to minimise personal costs. This approach introduces an autonomous demand-side management algorithm where consumers adjust appliance usage based on pricing signals. Their results support the existence of a Nash equilibrium and design a distributed method for users to iteratively update consumption. Their simulated results show reduced peak demand and overall energy costs, highlighting how incentive-compatible allocation mechanisms benefit both utilities and users.

Effective resource allocation is a central challenge across diverse domains, from disaster relief to robotics and communication networks. Optimisation techniques can be tailored to address both technical and ethical considerations in these complex systems. Pérez-Rodríguez and Holguín-Veras [79] explore welfare-oriented resource allocation in humanitarian logistics, focusing on distributing emergency supplies efficiently after disasters. Their models incorporate deprivation costs, a measure of human suffering from delayed aid, to prioritise timely deliveries. By solving complex allocation and routing problems, they develop heuristic solutions that outperform cost-minimising approaches in reducing societal harm. Case studies reveal that considering deprivation costs leads to fairer, faster aid distribution, bridging logistics optimisation with ethical imperatives.

Optimisation theory plays a crucial role in structuring network protocols to ensure efficient resource allocation across layers. Chiang et al. [80] provide a theoretical foundation for understanding network resource allocation via optimisation decomposition. They show how layered protocols such as congestion control, routing, and scheduling can be derived as distributed solutions to a global utility maximisation problem. Through vertical and horizontal decomposition, they explain how inter-layer and distributed algorithms achieve efficiency and fairness.

Designing effective strategies for secondary user (SU)s to access idle spectrum in multi-channel environments is a central challenge in cognitive radio network (CRN)s. This requires making sequential decisions under uncertainty, often without centralised coordination and based on noisy observations of channel availability. One promising approach is to enable SUs to dynamically access unused spectrum without interfering with primary user (PU)s. Zhao et al. [81] formulate the opportunistic spectrum access problem as a sequential decision problem under uncertainty. They consider a scenario of SUs searching for idle spectrum in a multi-channel environment without a central coordinator. They model each SU's decision process as a partially observable Markov decision process (POMDP), where the state (which channels are free) is not fully known and must be inferred from observations (noisy channel sensing results). They

derive an optimal strategy that dictates which channel a SU should sense and possibly occupy at each time slot to maximise long-term throughput while avoiding interference to PUs.

Managing interference is a central challenge in enabling shared spectrum access (SSA) between licensed and unlicensed users. Etkin et al. [82] take an information-theoretic approach to analyse the capacity of SSA, providing insights into the fundamental limits of interference management when spectrum is used by multiple systems. They study a model with a primary user (PU) (licensed) and a secondary user (SU) (unlicensed) sharing the same band. They introduce the concept of an interference temperature constraint, where the SU must keep its interference to the PU below a specified threshold. By characterising the Shannon capacity region under this constraint, they show how much SU throughput can be achieved without harming the PU's performance. A notable finding is that if the interference threshold is set too conservatively, the SU's capacity becomes very limited. Essentially, underlay spectrum sharing (transmitting below the noise floor) provides small gains unless the SU's link has a very favourable channel. They also examine spectrum etiquette policies and equilibrium power allocations when both users selfishly adjust power. Their theoretical analysis informs spectrum regulators and system designers about the trade-off between allowing secondary transmissions and protecting PU's services.

Dynamic spectrum access (DSA) using auction-based mechanisms has emerged as a compelling approach to improve spectrum utilisation in wireless networks. Zhou et al. [83] apply market mechanisms to spectrum allocation, designing auctions to distribute spectrum rights among users dynamically. They propose an auction framework for a scenario where multiple spectrum owners (sellers) can lease out frequency bands to multiple SUs (buyers) in a geographic region. They develop a strategy-proof double auction called *eBay in the Sky*, which ensures that no participant can benefit by misreporting their true valuation or availability. The mechanism matches buyers to sellers' spectrum bands and determines prices such that the allocation is efficient (maximises total surplus) while guaranteeing truthful bidding.

The integration of cognitive radio (CR) technologies into 5G and beyond networks has become a central focus in addressing the limitations of traditional spectrum allocation methods. Recent research emphasises full spectrum sharing and the application of machine learning (ML) and big data to enable real-time, adaptive spectrum management. Hu et al. [84] represent the next-generation perspective on spectrum allocation, focusing on how CR techniques can be integrated into 5G and beyond wireless networks. They review recent research on full spectrum sharing, where the goal is to enable all spectrum bands to be accessible by any wireless service under proper coordination. They cover developments in spectrum sensing that use big data and ML, which allow networks to predict spectrum availability in real time. They also address advanced resource allocation algorithms for heterogeneous networks, such as combining licensed and unlicensed spectrum and using device-to-device (D2D) communication to



share cellular spectrum. They examine new 5G architectural elements like network slicing and how they interact with spectrum management. For example, one option is to dynamically allocate spectrum to different network slices or services based on demand. Challenges discussed include interference management in ultra-dense networks, security and privacy issues in spectrum marketplaces, and the need for regulatory evolution to embrace dynamic sharing frameworks globally.

### 3.3.2 Spatial Clustering Techniques

#### Summary

*Clustering is a widely used technique across various fields, serving different purposes such as customer segmentation in marketing, object recognition in computer vision, and disease classification in healthcare. Broadly, clustering is valuable for three main reasons: uncovering hidden patterns in unlabelled data, simplifying complex datasets for easier analysis, and enhancing predictive models through feature engineering. In the context of spectrum management, spatial clustering is particularly relevant, as it enables grouping users by location to optimise frequency allocation. This approach allows for more adaptive and efficient spectrum usage compared to conventional, uniform allocation methods like hexagonal reuse patterns.*

#### Review

Clustering is an extremely important area of research across many fields. The reasons for wanting to cluster computationally vary considerably by field. In marketing, it enables customer segmentation and targeted campaigns by grouping customers based on their data [85]. In computer vision, it supports image compression and object recognition by grouping similar pixels. In healthcare and bioinformatics, it helps classify disease subtypes and segment medical images for improved diagnostics [86]. In social network analysis, it reveals community structures and enhance recommendation systems. In security, it facilitates anomaly detection to identify fraud or network breaches.

Spatial clustering is a hot research topic in spectrum management. There are multiple dimensions to consider in spectrum management depending on which aspect of the spectrum allocation and resource allocation process is being considered. Arguably the most obvious domain to cluster is spatially, with the most basic form of this being to group users together based on their similar locations. This kind of clustering can be used to dynamically achieve frequency separation in a way that is optimised for where the spectrum demand is likely to be, rather than using a regular static pattern like a hexagonal reuse pattern. Camino et al. [87] address a complex NP-hard optimisation problem in multibeam satellite communication systems, focusing specifically on designing efficient and irregular beam layouts to meet non-uniform user traffic demands within satellite service constraints. They propose enhancements to existing mixed-integer linear programming (MILP) formulations by incorporating k-means clustering to significantly reduce the computational complexity and number of variables.

### 3.3.3 Fairness Metrics

#### Summary

*Fairness in resource allocation is a critical consideration across various wireless communication systems, particularly as spectrum access become increasingly contested. This section explores key fairness metrics and their applications in dynamic spectrum sharing, high-performance computing, and cognitive radio network (CRN)s. Through a review of recent research, diverse fairness approaches are highlighted from utility-based and proportional fairness to auction-based models each addressing the challenge of balancing efficiency with equitable access among users. These studies demonstrate that integrating fairness into resource management algorithms not only improves user satisfaction and inclusivity but can often be achieved with minimal compromise to overall system performance.*

#### Review

Any decision-making system needs some measure or heuristic to guide it towards better outcomes. For a spectrum management system, the goal is to develop systems that are considered to be *effective*. The problem to solve here is to know what makes one solution more effective than another. If spectrum users are able to provide feedback on the spectrum allocation system or process, one fundamental aspect they may assess is whether they believe the system is *fair*. Indicators and measures of fairness have commonly been in use for a very long time across many fields.

Some of the most well-known fairness indicators are in economics such as the Gini coefficient or index. It has often been used as a way to measure the balance of wealth between societies. In computing, perhaps the most well-known fairness indicator is Jain's fairness index which has been often been used in problems to distribute resources within a computer system. Ikami et al. [88] address fairness in dynamic spectrum sharing between international mobile telecommunications (IMT) and incumbent systems. They introduce multiple fairness indicators and a dynamic allocation algorithm aimed at fairly distributing unused shareable spectrum among mobile network operators based on demand. Simulation results indicate significant fairness and satisfaction improvements while maintaining high spectrum efficiency.

In the context of high-performance computing (HPC), Ngubiri and van Vliet [89] address how to quantify fairness in scheduling parallel jobs on shared computing resources. They begin by observing that traditional performance metrics (throughput, wait time) do not fully capture how fairly different jobs are treated, thus motivating a dedicated fairness metric for schedulers. They propose a new fairness measure tailored to parallel job scheduling, one that accounts for the distribution of wait times or slowdowns across jobs. They critically evaluate this metric by applying it to various scheduling algorithms (like First-Come-First-Served, priority-based

scheduling, etc.), comparing how each algorithm scores in terms of fairness. They find that certain policies that optimise throughput can severely compromise fairness as per the new metric, while more balanced policies improve the fairness score with only a minor impact on efficiency. The real-world applicability is demonstrated through trace-driven simulations. By applying the fairness metric to real HPC workload traces, they show how system administrators might choose a scheduling policy that best balances performance with fairness, ensuring that no subset of jobs or users is consistently disadvantaged in a multi-user computing environment.

Fair resource allocation is a central challenge in wireless communication networks, particularly in systems where multiple users compete for limited spectrum access. In cognitive radio network (CRN)s, this issue is further complicated by the dynamic and opportunistic nature of spectrum usage. Khan et al. [90] cover fairness in the context of CRNs. They define fairness broadly as the equitable allocation of resources like bandwidth, time slots, or power among users, showing common fairness metrics (including Jain's fairness index and Proportional Fairness) for wireless scenarios. They critically review a wide range of CRN protocols and resource management schemes, comparing how each addresses fairness and at what cost to other performance metrics (throughput, spectral efficiency).

In wireless communication systems, ensuring fairness among users is a complex challenge, particularly when users experience diverse and dynamic local spectrum conditions. Traditional fairness metrics often fall short in capturing what the actual experience of using the channels was like. Dianati et al. [91] look to address this problem by introducing a new utility-based fairness index for radio resource allocation. Well-known fairness notions, such as max-min fairness and Jain's index, along with their limitations in wireless contexts, are discussed. For example, Jain's index treats all throughput differences uniformly, which might be inappropriate when users have different channel qualities or requirements. The proposed fairness metric incorporates utility functions to weigh the allocation in a way that reflects diminishing returns so that giving a bit of resource to a starved user improves fairness more than the same bit to an already well-served user. They perform a comparative evaluation, applying the new index and classical indices to a set of allocation scenarios in a wireless network.

In dynamic wireless environments, efficient and fair spectrum allocation remains a central challenge, particularly as cognitive radio network (CRN)s become more widespread. Balancing individual user needs with overall system performance requires novel approaches that go beyond traditional competitive models. Ma and Wang [92] address the challenge of distributing spectrum among secondary user (SU)s such that both efficiency (spectrum utilisation) and fairness are maximised. They frame the problem in interactive (competitive and cooperative) CRNs, introducing a coalition-competition model where users can form coalitions for spectrum sharing. Fairness is defined in terms of each user's spectrum access rate, employing a cooperative game theory approach known as a *differential game* to derive an allocation that is Pareto-efficient

(no one is better off) and fair. Their evaluation comes through theoretical analysis and simulation with their scheme compared against purely competitive allocations (maximising individual throughput) and purely equal-share allocations (fair but underuses spectrum). Their results show an approach that achieves a desirable middle ground of near-optimal spectrum efficiency while each coalition member gets a fair portion of spectrum relative to their needs. Real-world applicability is demonstrated by simulating time-varying spectrum environments. The proposed dynamic allocation adapts to changing conditions (like user mobility and traffic demand) and still upholds fairness. This suggests that the model could be useful in future decentralised spectrum-sharing systems where fairness must be maintained alongside high spectrum utilisation.

As wireless networks evolve toward 6G [93, 94], the equitable distribution of limited spectrum resources becomes increasingly critical. Ensuring fairness in spectrum allocation is essential to support diverse stakeholders and prevent market monopolisation. Khadem et al. [95] explore fairness in the context of spectrum auctions for shared access between incumbents and entrants. They contextualise the problem by noting that traditional spectrum auctions or allocation mechanisms often prioritise efficiency or revenue, potentially sidelining fairness for smaller stakeholders. They propose an auction framework under the enhanced licensed shared access (ELSA) model that explicitly incorporates a fairness criterion into the allocation algorithm. In their design, bids are weighted not just by monetary value but also by a fairness index that ensures a more equitable distribution of spectrum rights among participants (preventing large players from always dominating). They perform a comparative evaluation against a baseline auction (without fairness considerations) and a static sharing scheme. The fairness-aware auction is shown to slightly sacrifice maximum throughput or revenue in exchange for significantly improved fairness as measured by spectrum share distribution.

Efficient and fair allocation of limited wireless resources is a critical challenge in supporting multimedia services over modern communication networks. As user demand for high-quality video and data continues to grow, network providers must balance throughput with equitable service delivery to ensure acceptable quality of service (QoS) for all users. Guan et al. [96] look at radio resource management (RRM) for multimedia services, where both efficiency and fairness are desired for QoS reasons. They introduce a strategy based on *coopetition* a mixture of competition and cooperation among users to allocate bandwidth in an orthogonal frequency-division multiple access (OFDMA) network. Fairness in this work is defined by a utility function that increases with a user's allocated rate but with diminishing returns, aligning with an  $\alpha$ -fairness concept (where  $\alpha$  is chosen to balance total throughput vs. fairness). They formulate the allocation as an optimisation problem that maximises total utility (to address efficiency) subject to fairness and QoS constraints. They critically evaluate their scheme via simulations against purely competitive algorithms (which maximise throughput without fairness) and purely equal allocation, and find that the *coopetition* strategy achieves superior fairness and user satisfaction



while maintaining efficiency.

Covering third generation (3G) wideband code division multiple access (WCDMA) cellular networks, Papavassiliou and Li [97] study how to schedule users in the uplink to maximise total throughput while ensuring fairness among users. The research is motivated by the concept that scheduling can easily favour users with good channel conditions if left unchecked, leading to unfair access for users at cell edge or with poor channels. They define a fairness objective in terms of each user's long-term average throughput and incorporate it alongside throughput maximisation in a single optimisation framework. They propose a scheduling algorithm that uses a weighting factor to strike a balance. Essentially, the algorithm dynamically adjusts priorities so that users who have been underserved get higher priority, achieving a form of proportional fairness over time. They demonstrate that their method achieves better proportional fairness over time compared to a max-throughput scheduler (that ignores fairness) and a round-robin scheduler (that ignores channel quality).

Opportunistic scheduling (serving the user with the currently best channel) is known to boost throughput in wireless networks at the expense of fairness. Gueguen and Baey [98] address the trade-off by proposing an access scheme that maintains high throughput while improving fairness for OFDM-based systems. Looking at 4G wireless systems, the scheme introduces a system of dynamic weights for each user that consider both the user's experienced QoS and past allocation, effectively penalising users who have been served a lot and prioritising those who have been left behind. They define fairness in terms of long-term service equality and use weighted proportional fairness as a guiding principle. Their critical evaluation, through simulation, compares the new scheme to classic proportional fair scheduling and max-rate scheduling. The findings show that with the proposed scheme, the network does not have to choose between high fairness and high throughput. It achieves a balance where the overall throughput remains close to the max-rate case, but the fairness (quantified by Jain's index and outage probabilities per user) is notably better.

Cooperative communication systems rely on the collaboration of multiple nodes to enhance reliability and coverage in wireless networks. However, this collaboration introduces new challenges, particularly in ensuring that all participating nodes are treated fairly during relay selection. Lopez Vicario et al. [99] look at cooperative communications, examining how fairness can be maintained when selecting relay nodes in a wireless network. Relay nodes are intermediate nodes that forward data from a source to a destination, assisting communication when direct transmission is weak. In opportunistic relay selection (ORS), typically the relay that maximises instantaneous performance (highest signal quality) is chosen, which can lead to some nodes always being selected and others rarely, raising fairness concerns. They define a fairness criterion in the context of relay usage (ensuring different users or relays get opportunities to assist) and study the trade-off between system performance (measured by outage probability of the

transmissions) and fairness in relay selection. A unique aspect is the consideration of outdated channel state information (CSI). In practical systems, decisions are made on slightly old channel info, which can inadvertently introduce some fairness (by not always picking the absolute best relay). They provide an analytical evaluation of how varying the selection strategy (from purely max-SNR to more randomised or fair-aware selection) increases outage probability but improves fairness index. The critical finding is that a small sacrifice in optimality (accepting a relay that is not the absolute best) can dramatically improve fairness with only a minor increase in outage probability. This has real-world implications for the design of cooperative protocols as systems can be cheaply modified to reduce excessively greedy relay selection.

In modern wireless communication systems, achieving both efficiency and fairness in spectrum usage is a growing challenge, particularly in environments where multiple users compete for limited resources. Cognitive radio (CR) with its dynamic spectrum access capabilities, presents a promising solution to this issue by enabling intelligent resource allocation strategies. Meylani et al. [100] focus on a modern waveform (low density signature OFDM) for CR and propose a resource allocation algorithm that explicitly optimises a fairness metric. Noting that underlay CR allows SUs to transmit simultaneously with PUs under interference constraints, and fairness among SUs is a concern when resources (subcarriers/power) are allocated, they incorporate Jain's fairness index directly into their allocation algorithm. They modify a baseline iterative loading algorithm (ILRA) to create ILRA-FM (fairness metric), which seeks to maximise spectral efficiency while maintaining a high fairness index across SUs. Their comparative evaluation through simulation shows that the fairness-aware algorithm provides a significantly higher fairness index than the efficiency-only algorithm (ILRA), with only a slight reduction in total throughput.

In wireless sensor networks that employ CR, spectrum must be shared not only fairly but also efficiently due to sensors' energy and resource limitations. Byun and Gil [101] look at strategies for fair spectrum allocation, high spectrum utilisation, priority for urgent data, and minimal spectrum handoffs to save energy. They propose a bi-objective optimisation approach using modified game theory, effectively combining a fairness objective (modelled by weighted proportional fairness using demand-based weights) with an efficiency objective. Their results show that the proposed scheme achieves its design goals, allocating spectrum in a way that approximates weighted proportional fairness, meaning sensors with more urgent data (higher priority) get preference, but among those of similar priority the allocation is fair. It is compared against other allocation schemes in simulations, demonstrating improved fairness (no sensor is persistently excluded over repeated allocation intervals) and reduced unnecessary handoffs (spectrum switching) which in turn preserves sensor energy. This work supports the idea that even in highly constrained networks, fairness metrics can be successfully integrated into resource allocation algorithms to balance performance with equitable access.

### 3.3.4 Voronoi Tessellations

#### Summary

*Voronoi tessellations divide space into regions based on proximity to predefined seed points, with each region containing all points closest to a given seed. These diagrams have widespread applications, especially in telecommunications, where they assist in spatially structuring distributed node networks to improve resource allocation, energy efficiency, and interference management. Recent research demonstrates how Voronoi-based approaches enhance various aspects of 5G and 6G networks, including dynamic power control, device pairing, and infrastructure deployment. Their versatility also extends to traffic flow analysis and UAV optimisation, supporting their value as a foundational tool in modern wireless network planning and operations.*




Figure 3.3.4.1: Voronoi diagram example

#### Review

A Voronoi diagram (example shown in Figure 3.3.4.1) divides space into regions, where each region corresponds to the area closest to one of several specified seed points. Each region, or Voronoi cell, contains all points that are closest to a specific seed compared to any other. These polygons have applications in various fields, including computational geometry, geospatial analysis, and ML. They help in visualising decision boundaries in clustering algorithms and optimising resource distribution problems. Voronoi diagrams can be computed efficiently, but become complex in higher dimensions. In telecommunications in particular, this is quite convenient because distributed node networks is a well-known and popular structure proposed

and in-use for managing and distributing communications across a large area. Essentially, the output of a Voronoi diagram approximately matches a distributed node to the spatial division required for allocating resources and related applications. For this reason, it has become one of the leading approaches to solving this problem across the literature, especially in satellite communications.

Many telecommunications problems look to autonomously create spatial divisions across the region where the communications network operates. A common goal is to distribute power and energy evenly. Su et al. [102] propose a novel approach to optimising multi-tier heterogeneous 5G networks (integrating long-term evolution (LTE) and Wi-Fi) using Voronoi diagrams. They focus on dynamically structuring the radio access layer (RAL) to improve spectral and energy efficiency, effectively manage interference via dynamic frequency reuse, and reduce energy consumption by adaptively switching idle small cells into energy-saving mode. Simulation results demonstrate enhanced performance, particularly in low-load scenarios, highlighting the effectiveness of the Voronoi-based method.

Interference mitigation and mobility management are critical concerns in heterogeneous networks where small cell deployments coexist with macro-cell infrastructure. Xu et al. [103] present a two-step strategy for managing interference and mobility in HetNet small cell networks. The first step uses a Voronoi diagram-based min-max power allocation algorithm to optimise small cell deployment, reducing intra-tier interference. The second employs game theory for dynamic power control to mitigate cross-tier interference, protecting macro-cell served and high-mobility users. Simulation results indicate significant improvements in throughput, outage probability, and energy efficiency compared with sleep mode algorithms, non-cooperative game-based power control without deployment optimisation, and equal power allocation schemes.

Another key goal of spatial divisions in telecommunications is interference reduction or co-ordination. This can be related to spectral energy division but is arguably more useful because it is focused on a more directly practical measure of what allows multiple communication signals to function in the same area, rather than a higher-level goal to distribute energy more evenly. Liu et al. [104] introduce a device-to-device (D2D) pairing scheme leveraging Voronoi diagrams for dedicated cellular-D2D hybrid networks, focusing on reducing co-channel interference. Terminals are paired only if located in adjacent Voronoi polygons, significantly improving coverage probability, spectrum efficiency, and reducing interference compared to conventional methods. Simulation studies validate the efficiency and practical applicability of the proposed scheme.

Dynamic frequency reuse techniques have emerged as a powerful approach to mitigating inter-cell interference in dense wireless networks. Ullah et al. [105] propose a dynamic fractional frequency reuse (FFR) scheme based on Voronoi cell geometry for managing inter-cell interference in orthogonal frequency-division multiple access (OFDMA) cellular networks. Unlike

the traditional hexagonal cell model, the scheme uses a stochastic geometry approach for more realistic modelling. The proposed dynamic frequency allocation significantly outperforms static methods, providing enhanced per-user capacity and capacity density in densely deployed networks.

Voronoi tessellation and diagrams have also been researched extensively for use in modern mobile phone networks. They have been applied for numerous applications such as network traffic modelling, unmanned aerial vehicle (UAV) placement, and measuring the busyness of spectrum channels. Tao et al. [106] introduce an efficient Voronoi diagram-based algorithm for detecting traffic peaks directly from noisy traffic flow data without requiring smoothing techniques. The proposed algorithm computes the prominence of peaks, significantly outperforming existing methods in accuracy, sensitivity, and predictive validity. It enables transportation authorities to optimise traffic network management decisions, enhancing system performance and congestion mitigation.

Codebook design refers to the structured creation of predefined signal patterns, such as beam-forming vectors or pilot sequences, that are used to align transmission and reception in wireless systems. In integrated sensing and communication (ISAC) systems, it is a growing research focus due to its impact on spectral efficiency and detection performance. Cao et al. [107] introduce a multi-resolution hierarchical codebook using Voronoi clustering for cell-free ISAC systems, enhancing channel estimation, target detection, and latency estimation. Compared to traditional centralised methods, the proposed Voronoi-cluster design improves search efficiency and detection accuracy, efficiently mitigating interference and pilot contamination. Extensive simulations validate the scheme's claim of improved spectral efficiency, detection rates, and reduced complexity, making it well-suited for advanced 6G network deployments.

Emerging 6G vehicular networks demand innovative infrastructure strategies to ensure robust, low-latency connectivity across dynamic and wide-ranging environments. Andreou et al. [108] present an innovative approach integrating unmanned aerial vehicle (UAV)s with roadside units (RSU)s to enhance vehicle-to-everything (V2X) connectivity in 6G+ vehicular networks. Utilising a novel Voronoi diagram method with iterative circle expansion, the UAVs' placement is optimised to achieve complete network coverage efficiently. Simulation evaluations demonstrate complete area coverage, highlighting the method's simplicity and effectiveness in optimising infrastructure deployment for intelligent transportation systems (ITS).

## 3.4 Demonstrations

### 3.4.1 Context

The demonstrations in this chapter arise from the underlying problem of how to allocate radio spectrum among competing users in a dynamic spectrum management (DSM) system. As outlined in the introduction and approach, the challenge is not simply one of assigning channels, but of defining the allocation problem itself in a way that is technically rigorous and practically meaningful. Radio spectrum is a scarce, interference-prone medium, and allocation mechanisms must balance fairness, efficiency, and feasibility while adapting to conditions that change in time and space. The demonstrations were designed to address this complexity by creating controlled scenarios in which candidate mechanisms could be tested, compared, and better understood.

The literature shows that spectrum allocation is a cross-disciplinary problem drawing from resource allocation theory, clustering, fairness metrics, and computational geometry. Studies highlight how different fields have approached allocation with methods ranging from auctions and game theory to distributed optimisation and fairness indices. For example, proportional fairness and dominant resource fairness demonstrate how equity and efficiency can be balanced, while clustering and Voronoi-based techniques show the value of spatial partitioning for managing large, mobile populations. This emphasises that allocation problems cannot be solved by a single model but instead require a toolkit of methods adapted to context.

The demonstrations put these concepts into practice by linking theory to practice. Spectrum allocation metrics are tested to quantify fairness under varying conditions and Voronoi tessellations are used to explore spatial allocation strategies. Together, these demonstrations provide synthetic but structured environments to trial methods identified in the literature and to observe their performance against realistic constraints such as user movement, uncertainty, and spatial variability. Their purpose is to bridge the gap between the theoretical models discussed and the practical challenges of DSM.

### 3.4.2 Summary

Demonstration 1 evaluates fairness in spectrum allocation using single-pass, contiguous assignment strategies. Four approaches are compared: greedy, reserve, reserve with cap, and cap. Jain's fairness index is applied to quantify outcomes, revealing that greedy allocation consistently produces the least equitable results, while reserve and cap variants perform more evenly. A second part investigates fairness degradation when actual channel availability diverges from planned availability. Even proportionally fair allocations can fail under such uncertainty, leading to unmet fairness targets or over/under-assignments despite apparent spare capacity.



Demonstration 2 explores spatial spectrum allocation using Voronoi tessellations. In the first case, density-weighted seeds with a minimum separation constraint produce a variable number of clusters tied to covered area. Results show that while dense regions are well handled, sparsely populated areas produce excessively large cells, highlighting the limitations of constant area-per-cluster rules. In the second case, a fixed region with a fixed number of clusters is applied, allowing observation of boundary effects and imbalances in cluster membership over time. Together, these cases illustrate both the strengths and shortcomings of Voronoi methods in DSM.

### 3.4.3 Structure

This section presents three main demonstrations.

1. **Spectrum Allocation Metrics** — Single-pass, contiguous channel assignment strategies (greedy, reserve, capped variants) and their fairness outcomes; analysis of fairness degradation under uncertain channel availability.
2. **Voronoi-Based Clustering** — Spatial partitioning by Voronoi diagrams using density-weighted seeds, shown for (i) variable cluster counts tied to covered area and (ii) a fixed region with a fixed number of clusters.

### 3.4.4 Demonstration 1: Spectrum Allocation Metrics

Table 3.1 summarises the structure of Demonstration 1. It captures spectrum allocation metrics, focusing on comparing different assignment strategies for fairness and analysing the effect of stale availability information. The table links each objective to the corresponding allocation methods such as *greedy*, *reserve*, *reserve + cap* and *proportional fairness*, along with the relevant data sources including simulated user demands and planned versus actual channel availability traces.

Demonstration	Objective	Key Method	Data Inputs
1.1	Compare allocation strategies and fairness	Single-pass, contiguous assignment: <i>greedy</i> , <i>reserve</i> , <i>reserve + cap</i> , <i>cap</i> ; Jain index	Simulated continuous user demands; limited equal-sized channel units
1.2	Assess impact of planned versus actual channel availability	Proportionally fair allocation vs. post-assignment availability error	Planned vs. actual channel availability traces

Table 3.1: Summary of Demonstration 1 – Spectrum Allocation Metrics





## Demo 1.1: Spectrum Allocation Strategies

**Scenario:** A system is considered in which multiple users request access to a shared pool of discrete, equal-sized *channel units* (of an unspecified bandwidth) with randomly generated channel availability. Each user requests a continuous quantity of bandwidth, but allocations must be made using whole channel units. Because the pool is limited and allocations must be *contiguous*, users can receive only a single block of channels during the allocation process. A single left-to-right sweep assigns channels in arrival order, ensuring that each user is considered at most once.

**Setup:** To evaluate allocation behavior and fairness, four simple strategies are compared:

- *Greedy* — a pure first-come, first-served allocation.
- *Reserve* — a greedy allocation that holds a proportional reserve for future arrivals.
- *Reserve with cap* — proportional reserve plus a per-user maximum allocation.
- *Cap* — a strict equal-share per-user cap with no proportional reserving.

For each strategy, the Jain proportional fairness index is computed to assess how evenly resources are distributed. As expected, the *greedy* strategy produces the lowest fairness, whereas the reserve-based and cap-based strategies result in comparatively similar fairness levels.

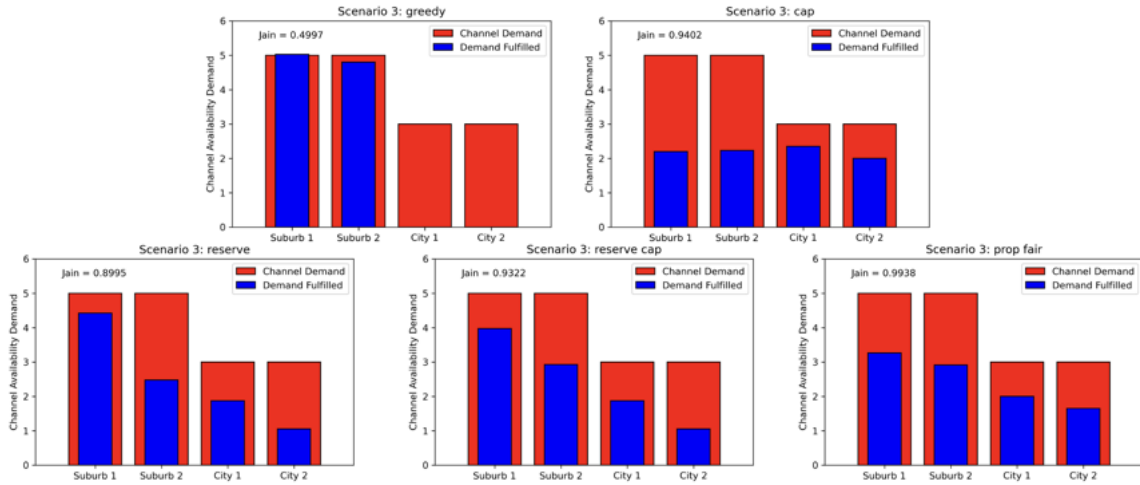


Figure 3.4.4.1: Demo 1.1: Spectrum allocation strategies for channels with Jain fairness scores

## Demo 1.2: Effect of Uncertain Channel Availability

**Scenario:** In Demo 1.1, it was assumed that the measured channel availability was accurate. Under realistic settings, this assumption may not hold. Measurements can be limited by where they were taken, delayed due to latency, affected by the sudden introduction of an interfering device, or based on predictive estimates rather than direct observation. As a result, the allocator may rely on availability information that does not reflect the true state of the system.

**Setup:** The demonstration mirrors Demo 1.1 but now shows the true post-assignment values. Grey bars indicate additions and white bars indicate subtractions relative to the allocator's assumed state, while blue bars show the intended allocations. This reveals how availability errors adversely affect allocation accuracy and fairness.

Deviations between assumed and actual channels availability can be seen to have a significant negative impact on the allocation fairness, with some users being under allocated while others receive more than should be available to them. The contrast between the blue and grey / white bars highlights how availability error can substantially degrade both fairness and resource utilisation.

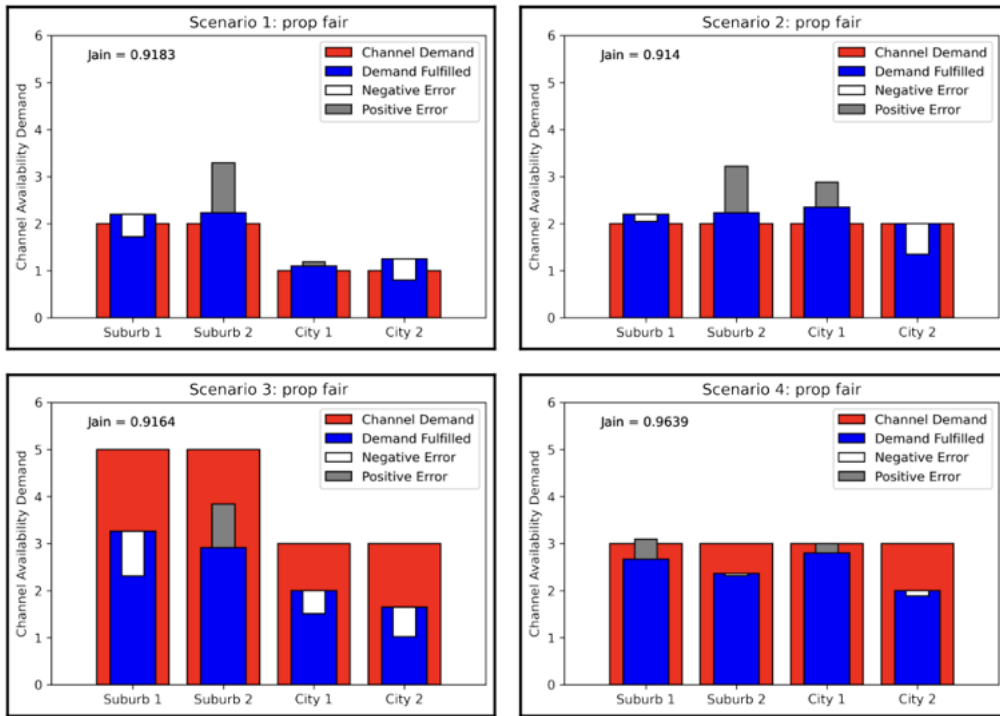


Figure 3.4.4.2: Demo 1.2: Proportionally fair allocation with availability error

### 3.4.5 Demonstration 2: Voronoi-Based Clustering

Table 3.2 summarises the structure of Demonstration 2. It presents Voronoi-based clustering under two regimes: one where cluster counts adapt to area and density, and another with fixed bounds and a fixed cluster count. The table identifies the corresponding methods, such as density-weighted Voronoi seeding and fixed 9-cell Voronoi layouts, applied to population density surfaces and dynamic agent position data.

Demonstration	Objective	Key Method	Data Inputs
2.1	Population density-based dynamic clustering	Voronoi with density-weighted seeds, min-separation and cluster count per covered area	Population density, agent positions over time
2.2	Same as 2.1 but fixed cluster count and showing cluster membership	Same as 2.1 but fixed number of clusters (9), track memberships	Population density

Table 3.2: Summary of Demonstration 2 – Voronoi-Based Clustering

## Demo 2.1: Voronoi Clustering with Dynamic Cluster Counts

**Scenario:** Agents represent spectrum users leaving the Adelaide CBD towards the suburbs after a CBD-based event.

**Setup:** In Figures 3.4.5.1–3.4.5.5, Voronoi diagrams are used to spatially cluster users. Two rules drive the number and size of clusters. First, the number of clusters is determined by dividing the area covered by the users by a chosen constant ( $2\text{km}^2$  in this demo but a minimum of 2 clusters). Voronoi polygons are formed by providing  $n$  centre points, one per polygon. Here, centres are chosen at locations of highest population density while enforcing a minimum separation so seeds do not collapse into the same vicinity.

**Frame 1 (00:00:00).** Two clusters are evenly divided across the area occupied by users at the initial time, reflecting the compact distribution around the CBD (Figure 3.4.5.1).




Figure 3.4.5.1: Demo 2.1: Voronoi clustering at 00:00:00





**Frame 2 (00:02:00).** Two minutes later, most users are driving out of the busy CBD. The area covered by users roughly doubles, so the rule allows four clusters to form (Figure 3.4.5.2).




Figure 3.4.5.2: Demo 2.1: Voronoi clustering at 00:02:00





**Frame 3 (00:04:00).** A further two minutes later, dispersion increases as users continue to leave the CBD. The cluster rule now produces seven clusters, with centres near the densest points while respecting minimum spacing (Figure 3.4.5.3).

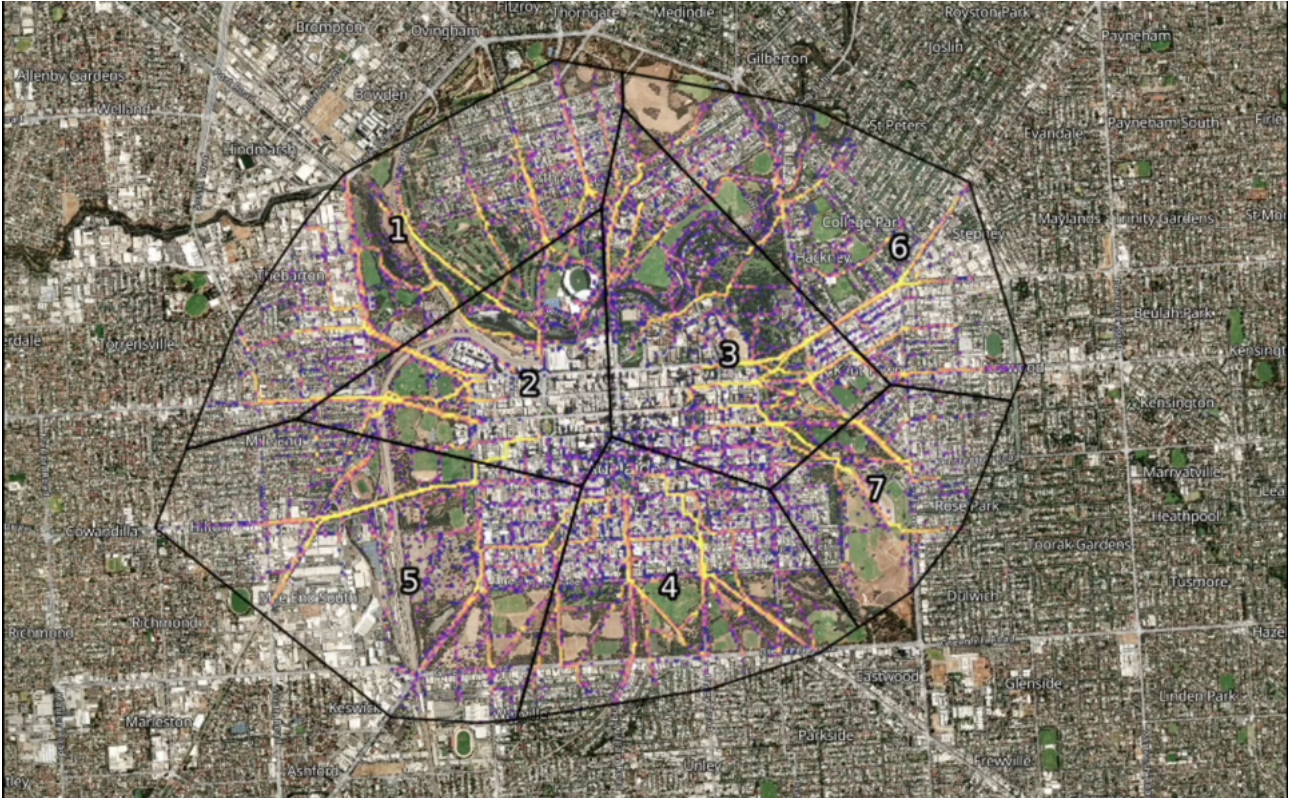


Figure 3.4.5.3: Demo 2.1: Voronoi clustering at 00:04:00





**Frame 4 (00:06:00).** With further outward movement, the covered area grows and the rule permits eleven clusters. Many centres appear on the (east) side where density is highest. However, this creates very large cells in the north, west, and south, revealing a limitation of constant area-per-cluster thresholds (Figure 3.4.5.4).




Figure 3.4.5.4: Demo 2.1: Voronoi clustering at 00:06:00





**Frame 5 (00:08:00).** By eight minutes, many users have reached destinations in the suburbs. The rule now allows sixteen clusters. Imbalances become even more pronounced. For example, the north-western cell (Cluster 1) is much larger than others, again suggesting that more dynamic criteria would produce more desirable partitions (Figure 3.4.5.5).




Figure 3.4.5.5: Demo 2.1: Voronoi clustering at 00:08:00





## Demo 2.2: Voronoi Clustering with Fixed Region and Fixed Cluster Count

**Scenario:** Agents represent spectrum users leaving the Adelaide CBD towards the suburbs after a CBD-based event.

**Setup:** The setup in Figures 3.4.5.6 –3.4.5.9 mirrors the previous example, except both the region bounds and the number of clusters are fixed (to  $k = 9$ ). The top-left panel colours cells to distinguish them; the top-right panel shows population density used for seeding; the bottom panel reports per-cluster membership counts.

**Frame 1 (00:00:00).** Memberships are relatively even at the start, though the difference between the highest (Group 4) and lowest (Group 9) is substantial (roughly 110 member differential), driven by initial density near the CBD (Figure 3.4.5.6).

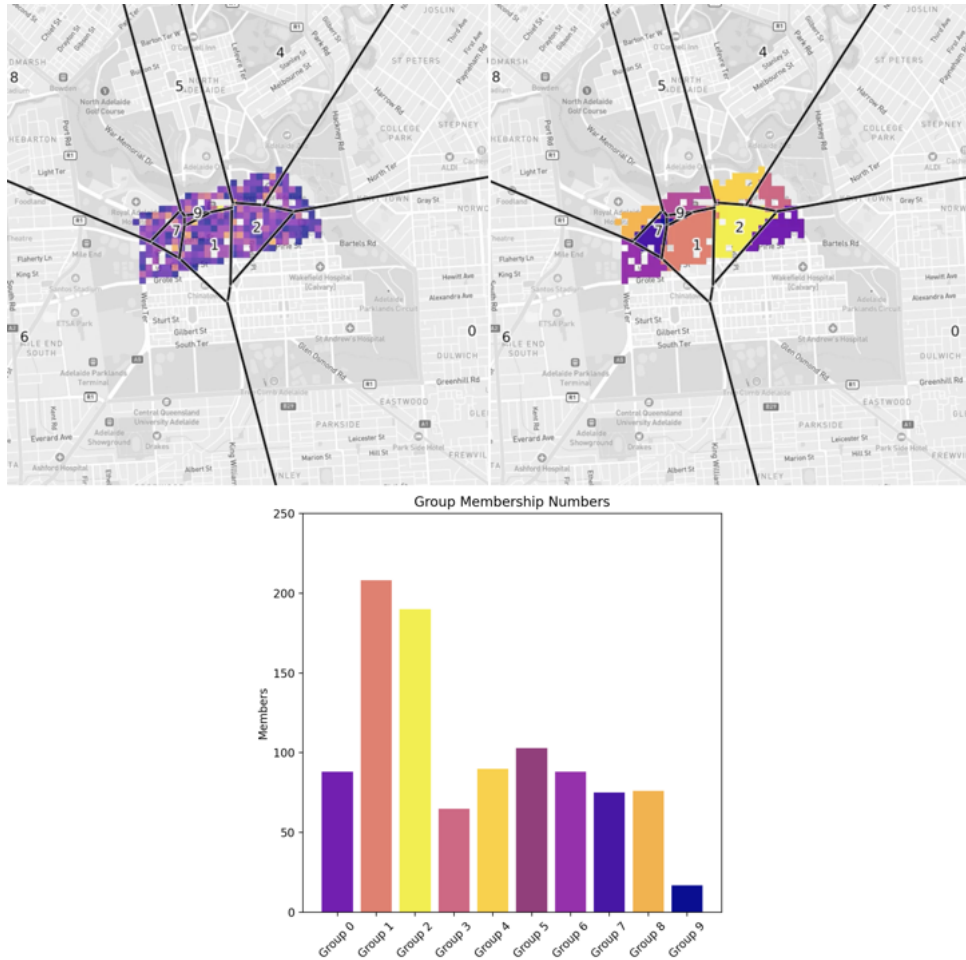


Figure 3.4.5.6: Demo 2.2: Voronoi (fixed region,  $k=9$ ) at 00:00:00



**Frame 2 (00:03:00).** Three minutes later, clusters remain concentrated near the CBD since most users are still central, but membership counts shift as individuals traverse the network and cross cell boundaries (Figure 3.4.5.7).

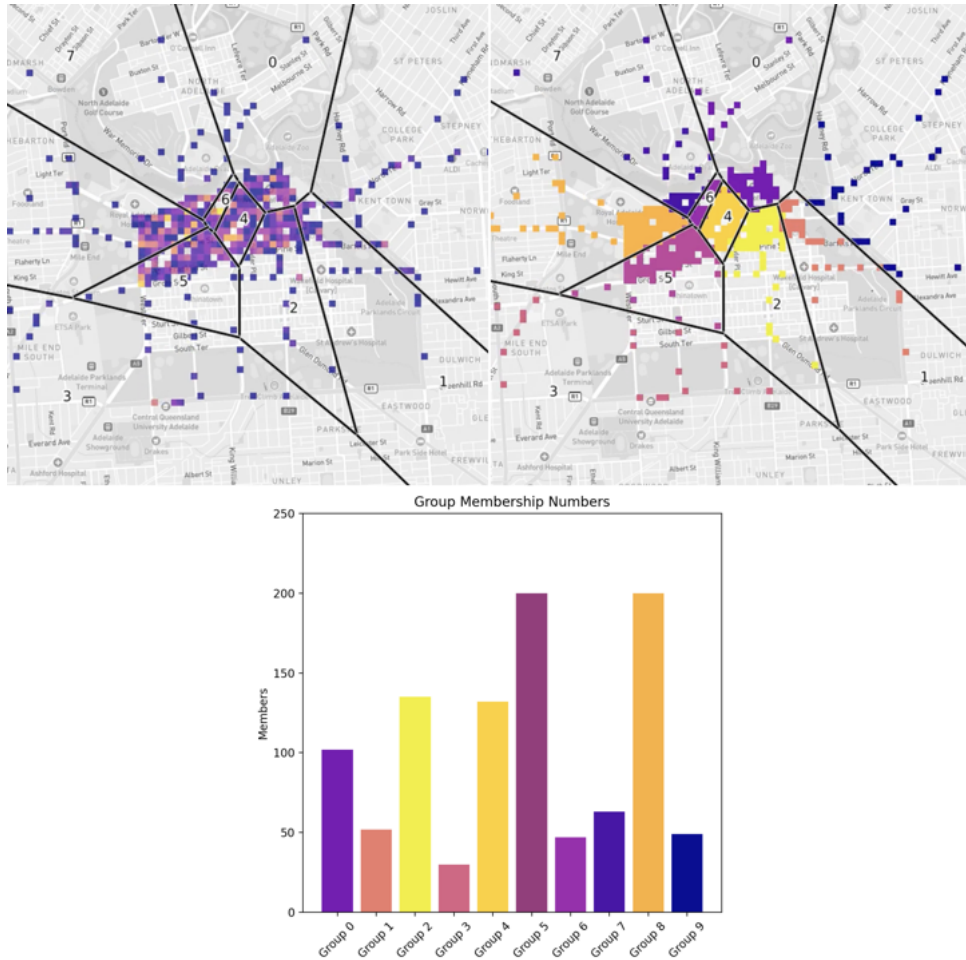


Figure 3.4.5.7: Demo 2.2: Voronoi (fixed region,  $k=9$ ) at 00:03:00



**Frame 3 (00:06:00).** By six minutes, arrangements remain similar around the CBD, but memberships grow in the outer cells (Groups 9, 8, 4, 1) as more users leave the centre (Figure 3.4.5.8).

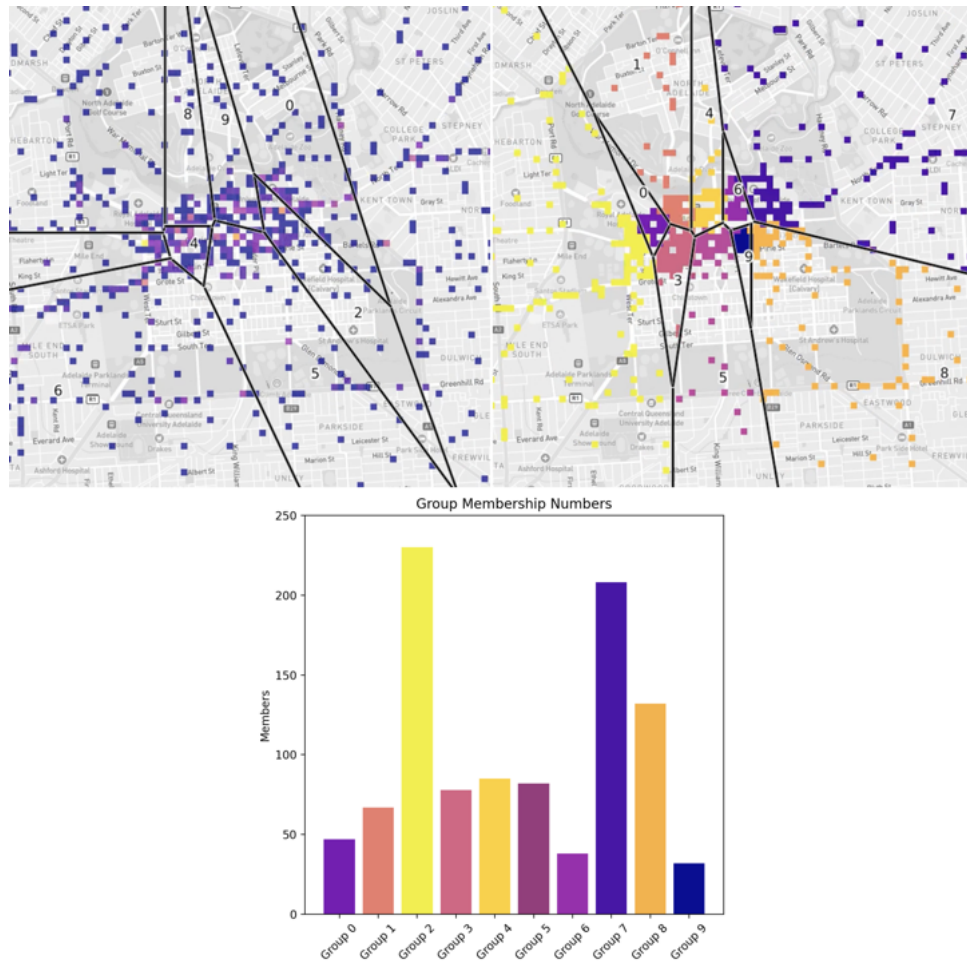


Figure 3.4.5.8: Demo 2.2: Voronoi (fixed region,  $k=9$ ) at 00:06:00



**Frame 4 (00:09:00).** At twelve minutes, dispersion between CBD and suburbs becomes more balanced (slightly favouring the CBD), and cells spread outward accordingly (Figure 3.4.5.10).




Figure 3.4.5.9: Demo 2.2: Voronoi (fixed region,  $k=9$ ) at 00:09:00



**Frame 5 (00:12:00).** At twelve minutes, dispersion between CBD and suburbs becomes more balanced (slightly favouring the CBD), and cells spread outward accordingly (Figure 3.4.5.10).

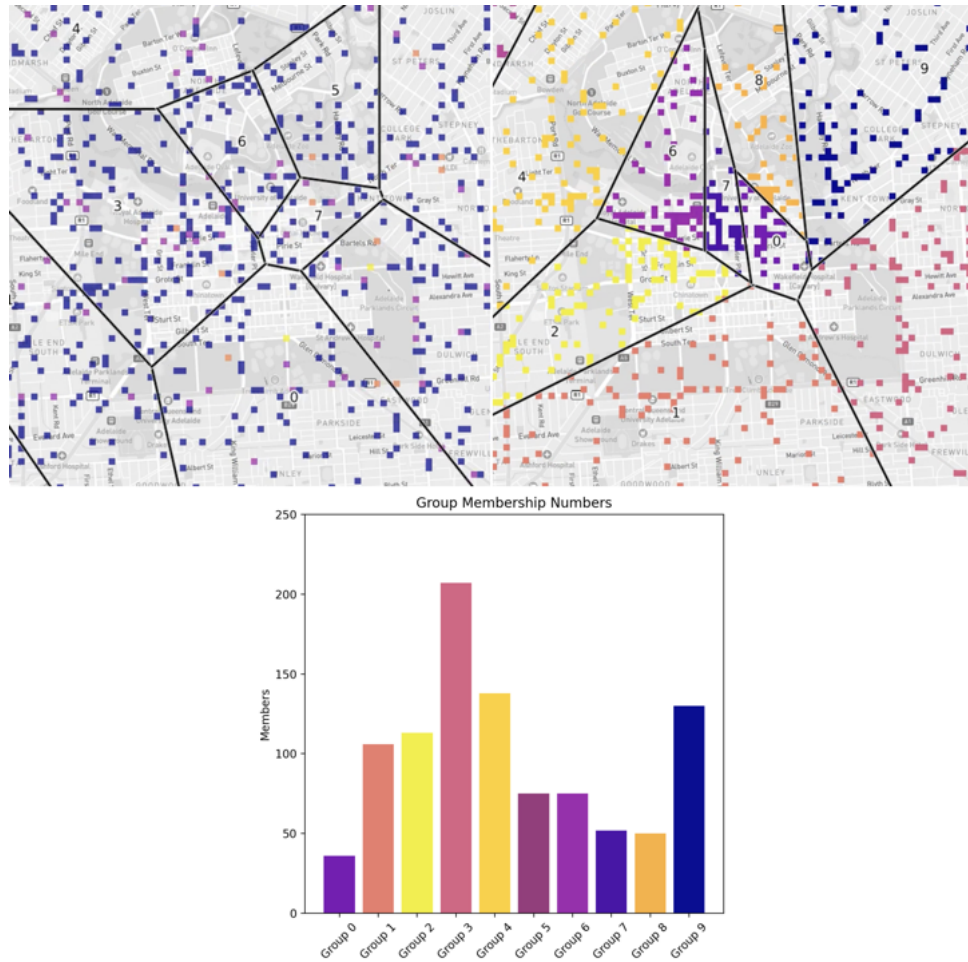


Figure 3.4.5.10: Demo 2.2: Voronoi (fixed region,  $k=9$ ) at 00:12:00

## 3.5 Discussion

### 3.5.1 Summary

Chapter 3 has four sections: *Introduction*, *Approach*, *Literature*, and *Demonstrations*.

In the *Introduction*, an industry problem of allocating radio spectrum among competing users at higher allocation levels is framed. Scope is defined around designing allocation mechanisms for dynamic spectrum management (DSM), and the difficulty of specifying the allocation problem precisely is highlighted. The goal is set as the surfacing of mechanisms that remain practical for future spectrum management systems.

In the *Approach*, the proposition is advanced that many allocations and evaluation criteria are available, and that mechanism choice is determined by problem definition within DSM. The Question is then posed: “which allocation mechanisms are suitable for DSM solutions?”. Emphasis is placed on resource distribution and allocation, spatial clustering techniques, fairness metrics, and Voronoi tessellations.

In the *Literature*, the three focus areas are explored. Spatial clustering for demand-aware grouping and satellite beam layouts is reviewed. Fairness indices such as Jain’s index and utility-based measures are considered. Auction designs that incorporate equity are surveyed. Voronoi-based spatial partitioning for interference management, dynamic reuse, device pairing, and energy efficiency across 5G and beyond is covered. Benefits are identified as scalable sharing incentives and measurable equity, while limits are recognised as arising from modelling assumptions, complexity, and context sensitivity.

In the *Demonstrations*, time-varying demand is simulated using a Markov process initialised from Australian census data, and allocation metrics, clustering health, and Voronoi partitions are evaluated. In Demonstration 1, single-pass contiguous strategies are compared, fairness is quantified, and the degradation of outcomes under stale availability is shown even with proportional allocation. In Demonstration 2, density-weighted Voronoi seeding is applied for dynamic and fixed-k partitions, with boundary effects and constant area rules revealed as limitations, while directions for richer models and adaptive triggers are indicated.

### 3.5.2 Key Findings

#### Finding 1: Fairness Metrics as Allocation Drivers

Fairness metrics provide a reliable basis for evaluating and guiding spectrum allocation. The *Approach* framed the allocation problem as one of balancing competing user demands within constrained resources, identifying fairness as a measurable criterion of effectiveness. The *Literature* highlighted established fairness indices such as Jain's index and proportional fairness, demonstrating their adaptability across wireless and computing contexts while balancing efficiency with equity. The *Demonstrations* validated these measures by applying them to spectrum allocation strategies, where greedy methods produced inequitable outcomes but fairness-aware strategies provided balanced distributions. The implication is that fairness metrics can be formally included in DSM systems to ensure equitable resource access.

#### Finding 2: Spatial Clustering for Spectrum Efficiency

Spatial clustering techniques enable adaptive spectrum reuse and efficient allocation. The *Approach* identified clustering as a central tool for simplifying complex allocation scenarios by grouping users spatially. The *Literature* described k-means, DBSCAN, and other clustering models, showing how they support telecommunications practices such as frequency reuse and traffic management. The implication is that clustering provides a scalable mechanism for managing spectrum in dynamic conditions. By abstracting individual users into spatial groups, it simplifies complex allocation problems while maintaining efficiency. If a resource can only be divided in limited ways (often the case with RF channels), clustering users in this manner could become a practical necessity, and a further benefit is the reduction of overheads in coordination and computation dealing with user clusters rather than potentially thousands of individual users.

#### Finding 3: Voronoi Tessellations for Dynamic Partitioning

Voronoi tessellations offer a practical method to partition spectrum spatially in alignment with user density. The *Approach* noted Voronoi methods as prominent in spectrum allocation research, particularly for satellite and cellular systems. The *Literature* reinforced this by showing their use in interference management, power control, and network deployment, proving their computational efficiency and adaptability in telecommunications. The *Demonstrations* implemented Voronoi-based clustering under both dynamic and fixed-region conditions, revealing how tessellation can track population dispersion and maintain manageable allocations across urban environments. The implication is that Voronoi tessellations provide a mathematically robust and practically effective foundation for dynamic spectrum partitioning in real-world systems.

### 3.5.3 Limitations

While fairness metrics such as Jain’s index and proportional fairness offer a structured means of evaluating allocation outcomes, they introduce limitations when applied in practical DSM contexts. These measures often assume that user demands and channel availability are stable during the allocation process, yet in reality both are highly variable. Demonstration 1 showed that even proportionally fair strategies degrade under inaccurate availability information, producing allocations that appear balanced but fail under actual conditions. This highlights a limitation, as fairness indices capture distributional equity but do not inherently account for uncertainty or dynamics in availability. Furthermore, fairness metrics can oversimplify the trade-off between efficiency and equity, since they treat fairness as an abstracted mathematical property, while practical deployments must also address system throughput, interference avoidance, and operational overhead. In other words, there are so many other factors not considered that it is highly unlikely that the allocations would be as fair as the metrics have indicated, under realistic conditions.

Voronoi tessellations offer a mathematically elegant approach to spatial partitioning, but their limitations emerge in irregular population distributions. Demonstration 2 illustrated that constant area-per-cluster rules can produce excessively large cells in sparsely populated regions, while dense areas remain well captured. Similarly, when fixed-region tessellations were applied, boundary effects and uneven memberships distorted allocations, leaving some cells very full while others remained almost empty. These shortcomings stem from the inherent reliance of Voronoi methods on seed placement and geometric boundaries, which do not fully capture temporal and spatial variation in spectrum demand. Moreover, while computationally efficient, Voronoi tessellations lack built-in mechanisms for incorporating fairness, demand weighting, or interference constraints, requiring additional layers of adjustment. This limits their direct applicability in DSM systems without supplementary models to mitigate imbalance and adapt tessellations to real-world complexity.



### 3.5.4 Future Work

Advancing spectrum allocation could require methods that go beyond static fairness indices by including stochastic and adaptive mechanisms. Future models should incorporate uncertainty directly into allocation algorithms so that fairness outcomes remain valid even under fluctuating channel availability. This could be achieved by integrating fairness metrics with probabilistic demand forecasting and real time monitoring of spectrum conditions. Additionally, integrating richer data sources and extending the set of allocation heuristics beyond Jain's index and proportional fairness to include multi objective utility functions could provide a more comprehensive evaluation of equity, efficiency, and resilience. Such improvements could be validated within a digital twin environment that captures dynamic conditions, enabling systematic comparison of allocation outcomes under realistic scenarios in a simulated environment.

Clustering approaches could evolve toward dynamic and context sensitive models. Rather than relying on fixed memberships or idealised spatial assumptions, adaptive clustering could trigger recalculation based on dispersion thresholds, mobility patterns, and traffic demand. Exploring clustering methods that consider time based stability and varying user densities, such as hierarchical, spectral, or density adaptive approaches, could help reduce the decline in cluster quality over time. The next step would be to include these mechanisms in simulation frameworks that model user movement realistically and quantify trade offs between reclustering frequency, computational overhead, and allocation accuracy. Incorporating clustering outputs into a digital twin could also allow visualisation of cluster health in relation to user locations, movement patterns, and areas of high demand, providing a rich simulated environment for testing robust DSM solutions.

Voronoi based methods present a strong baseline for spatial partitioning, but their limitations point toward hybrid or alternative models. Future research could explore weighted and adaptive Voronoi tessellations where seed placement is responsive to population density, fairness constraints, and interference measurements, rather than static geometric rules. Boundary effects and uneven memberships highlight the need for corrective overlays, such as demand driven adjustments or multi scale tessellations. Alternative partitioning methods, including power diagrams or graph based segmentation, could also be assessed for their ability to capture irregular demand distributions. Using these techniques in digital twin simulations would allow assessment of how spatial partitions evolve with mobility, terrain, and network topology, making it possible to evaluate solutions in a simulated environment.



## 3.6 Next Steps

Analysing the spectrum allocation process in isolation was important and valuable, but its limitations are clear. With a better understanding of allocation mechanisms under ideal conditions, it is time to explore dynamic spectrum management (DSM) under real world conditions. The objective is to coordinate time critical access to radio spectrum across large areas while maintaining fairness, reliability and policy compliance amid mobility and interference. Key constraints include latency, infrastructure limits, geospatial diversity, limited energy and incomplete local information. The general approach is to combine local decision making with hierarchical oversight, represent the environment and reachability in a digital twin, and adapt allocations through rapid feedback and lightweight coordination.



## Chapter 4

# Dynamic Spectrum Allocation in the Real World

*This chapter examines the integration of real-time dynamic spectrum management (DSM) into realistic 3D environments, addressing challenges of large-scale, time-critical radio frequency (RF) coordination under constraints like latency, geospatial diversity, and interference. It reviews decentralised and n-tier spectrum management approaches, mobile ad hoc network (MANET)s, clustering strategies, and isochrone / shortest path tree (SPT) modelling for mobility prediction, drawing on literature that spans distributed control protocols, energy-efficient routing, and advanced clustering methods. Practical demonstrations apply decentralised clustering protocols to simulated urban spectrum demand, showing how localised coordination reduces interference and communication overhead, with visualisations depicting cluster formation and interconnections. The chapter highlights that while these techniques are well-supported in research and suitable for mobile, low-power systems, their application to actual spectrum allocation remains untested, and future evaluation in high-fidelity simulations with realistic RF propagation and human movement modelling is required. This kind of simulation is sometimes referred to as a digital twin.*

## 4.1 Introduction

Integrating real-time dynamic spectrum management (DSM) into a 3D environment is arguably the pinnacle of this field. However, achieving this goal under realistic conditions provides many significant challenges beyond those established in idealised, simplified representations of the problem, often explored in the literature. Realistic modelling addresses these constraints and provides a more accurate foundation for understanding coordination challenges. Across the literature, the term *digital twin* describes a cohesive simulated environment that incorporates a detailed, real-time 3D representation of the geospatial environment [109]. This approach enables a clearer view of how real-world factors influence coordination and creates a reliable platform for comparing idealised and realistic conditions.

Coordinating time-sensitive access to RF spectrum across large areas combines technical and logistical challenges. Large-scale coverage introduces infrastructure limits and latency from the physical distances between communication links. Geospatial diversity complicates transmission behaviour, with dense interference in urban areas and coverage gaps in rural areas. The time-sensitive nature of the problem requires rapid decisions and low-latency communication, placing high demands on computational performance to coordinate many users in real time. These factors are closely linked, with large-scale and geospatial challenges increasing the latency and coordination requirements of a time-critical system. Mobile ad hoc network (MANET)s provide an infrastructureless approach that can adapt quickly to local conditions, reducing reliance on fixed control points and lowering communication delays. A hierarchical management structure allows scalable coordination that addresses local conditions while staying aligned with overall coordination objectives [110, 111, 112].

Under idealised conditions, spectrum users may travel and transmit unrealistically far, producing an unbounded range of possible outcomes. This makes sustained coordinated spectrum access impractical when, under realistic constraints, it may be achievable for limited periods. By modelling reachability, predictive or worst-case planning becomes more practical. Isochrones [113] and shortest path tree (SPT)s [114, 115] are simple ways to represent the area that can be reached within a given range or cost, providing a clear picture of coverage in space. Within a *digital twin*, this becomes a practical tool for spectrum planning.

## 4.2 Approach

### 4.2.1 Propositions

Spectrum management is a large spatial problem with tight temporal constraints.

### 4.2.2 Question

In a realistic setting, which techniques are available to implement real DSM solutions known for requiring significant levels of distributed control and coordination?

### 4.2.3 Topics

**Decentralisation and Hierarchical Spectrum Management** - Provides low- or no-infrastructure solutions for rapid deployment in emergencies, short-term demand, and scalable rollouts. Hierarchical or n-tier solutions allow high-level coordination across large areas while enabling low-latency local decisions.

**Mobile Ad-hoc Network (MANET)** - An area of interest because it is a technology and concept to implement decentralised and n-tier / hierarchical management and control in a real-world setting.

**MANET Clustering** - Is especially critical for DSM-specific use cases of MANETs because the structures that the MANETs can form would need to match the way that spectrum is managed and allocated.

**Isochrones & Shortest Path Trees (SPT)** - A powerful technique with broad applications for accurately modelling and accounting for the potential mobility of spectrum users.

## 4.3 Background

### 4.3.1 MANET Protocols

A common algorithm for MANET formation can be described as follows based on algorithms in [111, 112, 116]. It uses a random start-up, sensing, and self-nomination strategy as shown in Figure 4.3.1.1. The protocol is distributed and only relies on information that can be shared within the signal range for each agent. The overall goal of this algorithm is to determine the optimal number and placement of *cluster heads* which take a leadership role in managing the topology and resource distribution across its own cluster. Depending on the type of MANET, it may also have the role of coordinating with other *cluster heads*, or that role may need to be distributed to edge nodes that can form connections with other clusters.

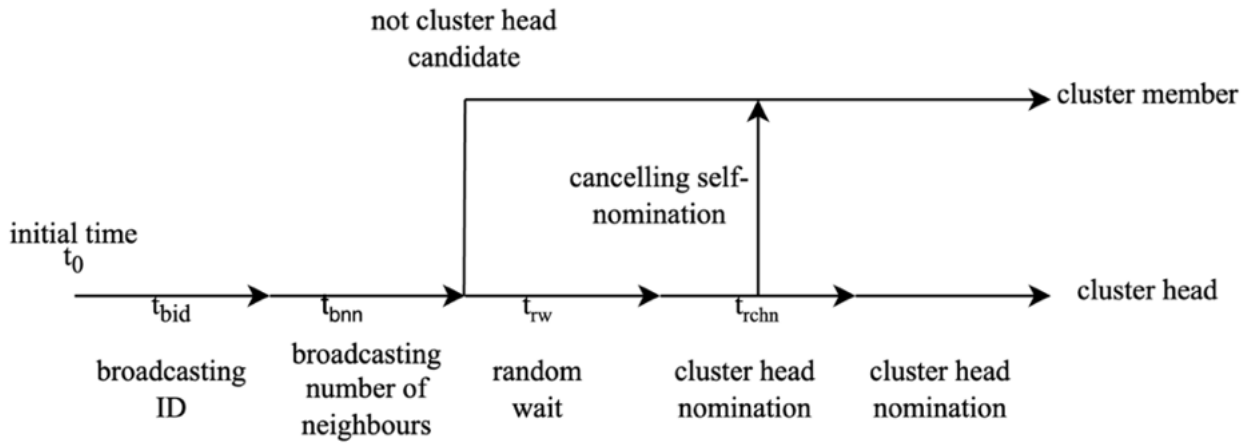


Figure 4.3.1.1: Fundamental mobile ad hoc network (MANET) formation protocol

### 4.3.2 Isochrones & SPT routing Engines

The concept of an isochrone is not new, but has more recently regained popularity, likely due to advancements in computational efficiency and the quality of open-source mapping data such as open street map (OSM) [117]. According to the documentation for the Valhalla open source routing engine, an isochrone can be defined as a line that connects points of equal travel time about a given location, from the Greek roots of *iso* for equal and *chrone* for time [113]. For the demonstrations in this thesis, the related concept of a shortest path tree (SPT) was used instead of an isochrone due to some advantages. When constructed from a routing graph, the only difference between the two is that the *contouring* stage of the isochrone is skipped, leaving the underlying routing graph's reachable path segments. This was preferred as these paths segments can be better representative of the actual reachability. This is because the contouring process of isochrone construction essentially applies an interpolation that can be quite unrealistic, giving poor results where accuracy is needed. An example of a SPT is shown below in Figure 4.3.2.1 and Figure 4.3.2.2.



Figure 4.3.2.1: Shortest path tree (SPT) example for a pedestrian in Adelaide CBD






Figure 4.3.2.2: Shortest path tree (SPT) example for a pedestrian in the Adelaide CBD displayed over satellite imagery





An isochrone can be used to represent contours or regions where a time-varying value is the same over time as a 2D heatmap. The most common time-varying value is travel time for various modes of transport such as walking, cycling, driving etc. The same concept can also be applied to less common scenarios such as the flood distribution isochrone of travel time proposed by Pilgrim [118] where what is shown is the time it would take for flood waters to reach different parts of the map.

There are multiple open-source options available for producing isochrones such as GraphHopper [119], Valhalla [120], and open source routing machine (OSRM) [121].

- **Valhalla** [120] is an open-source routing engine designed to provide flexible and customisable routing solutions for various transportation modes, including driving, walking, cycling, and multimodal options. It utilises open-source data, primarily from open street map (OSM), to deliver dynamic and time-based routing, complete with clear directions and maneuver instructions. Valhalla's architecture is built around tiled, hierarchical data structures, allowing for efficient regional extracts and partial updates [122].
- **GraphHopper** [119] is a fast and memory-efficient open-source routing library and server written in Java. By default, GraphHopper uses OSM data for the road network and incorporates elevation data from the shuttle radar topography mission (SRTM). It supports different algorithms such as Dijkstra's algorithm, A\* search, and contraction hierarchies to enable efficient routing across large graphs [122].
- **Open source routing machine (OSRM)** [121] is a high-performance routing engine specifically designed for road networks. Developed in C++ and optimised for speed, OSRM can compute shortest paths in milliseconds, even on continental-scale road networks. It uses contraction hierarchies, a preprocessing technique that significantly speeds up route calculations by simplifying the underlying graph structure. OSRM supports various routing profiles such as car, bike, and foot, using OSM data to build its routing graphs [122].

Throughout this thesis, Valhalla [120] was chosen because it is simple to setup, and has a very good shortest path tree (SPT) output available. GraphHopper is also very usable for this task as it also has a SPT output available [114] which can be adapted to obtain a similar result.

## 4.4 Literature

### 4.4.1 Decentralisation and Hierarchical Spectrum Management

#### Summary

*Decentralisation reduces decision latency in dynamic spectrum management (DSM), enabling spectrum allocation that adapts to local conditions more quickly than centralised systems. Proven models such as blockchain, cryptocurrencies, and peer-to-peer (P2P) networks highlight its benefits, though risks like the hidden node problem remain when global awareness is absent. Architectures such as spectrum management architecture and protocol (SMAP) demonstrate how decentralisation can be scaled through peer-to-peer exchanges, integration with cloud services, and support for market-based spectrum access while maintaining policy compliance. Hierarchical or n-tier approaches extend these ideas by combining local autonomy with structured coordination, delivering both responsiveness and system stability. Beyond spectrum management, hierarchical optimisation methods such as ant colony clustering in mobile ad hoc network (MANET)s show how multi-level structures can enhance scalability, reduce delay, and improve quality-of-service by balancing local decision-making with inter-tier communication. Together, these designs illustrate how decentralisation and hierarchy jointly provide adaptable, robust, and policy-aligned spectrum management.*

#### Review

Decentralisation is a key concept of dynamic spectrum management (DSM) [123] for numerous reasons with the most significant one being management or decision latency. Some well-known examples of decentralisation being used successfully are blockchain, cryptocurrencies, peer-to-peer (P2P) file sharing, and the Internet. Each of these examples provided extremely popular alternatives to problems that were previously only solved with centralised solutions. The time taken to communicate information about the live state of radio spectrum back to a centralised management could be longer than the duration of the radio communications under some circumstances. By moving the decision-making power closer to where the measurements and data are being taken, the spectrum allocations can match the current state of the spectrum more closely. However, there are some drawbacks to being fully decentralised in this use case. The *hidden node problem* [124, 21] is well-known across wireless networking and it fully applies here too. If the power to allocate radio spectrum is fully-decentralised, multiple spectrum users could be allocated the same spectrum.

Dynamic spectrum management (DSM) is moving from rigid, centralised systems toward more adaptive, distributed approaches. Recent advances in wireless technology, market structures, and regulatory frameworks have highlighted both the need and opportunity for flexible, data-driven spectrum sharing. Karimi et al. [125] introduce spectrum management architecture and protocol (SMAP), a decentralised architecture designed to enhance dynamic spectrum management (DSM) by enabling wireless networks to cooperatively share spectrum resources using an Internet-based control plane. It supports peer exchanges of radio usage data, promoting efficient and fair optimisation of wireless transmission parameters across heterogeneous technologies such as Wi-Fi and LTE. SMAP is structured hierarchically to facilitate scalability, decentralised decision-making, and compliance with local and global spectrum policies. The architecture interfaces with cloud services like regional spectrum aggregators and spectrum access system (SAS), supporting innovative market-based spectrum sharing. Experiments demonstrate SMAP's viability in various scenarios, highlighting its benefits in scalability, robustness, and ability to implement local spectrum-sharing policies effectively.

Hierarchical or n-tier is a related but separate concept to decentralisation. It is a compromise between fully centralised and fully decentralised in which the stability of centralisation is combined with the flexibility and robustness of decentralisation. Such a system could be realised in many ways. One example is that while some decisions could be made *locally* and rapidly, some information could be shared between the different tiers, allowing them to coordinate by having some visibility as to what is occurring in other parts of the system. Sensarma and Majumder [126] propose a hierarchical ant colony optimisation based routing scheme for mobile ad hoc network (MANET)s, addressing the challenge of QoS provisioning in dynamic, infrastructure-less environments where heterogeneous node capacities complicate routing. The method introduces a three-level cluster-based topology where cluster heads are selected using weighted metrics such as energy, mobility, connectivity, and distance, with ants guiding route discovery through pheromone updates and path preference probabilities. The approach considers QoS parameters including delay, bandwidth, residual energy, link expiration time, and hop count for both intra- and inter-cluster routing, using specialised packet types (Route\_Ant, Knave\_Ant, King\_Ant) for discovery and maintenance. Results show that this design improves throughput, reduces delay, lowers routing overhead, and increases scalability compared to flat architectures. The broader contribution lies in demonstrating how clustering with ant-based optimisation can provide robust, energy-aware, and scalable QoS routing in heterogeneous MANET, while future work includes simulations and the integration of new QoS metrics.

## 4.4.2 Mobile Ad-hoc Network (MANET)

### Summary

*Mobile ad hoc network (MANET)s are infrastructure-less, highly dynamic wireless networks composed of mobile nodes with limited energy resources, where connectivity is maintained through distributed clustering and multi-hop communication. Decades of research have proposed numerous strategies to optimise MANET performance, especially for routing, based on metrics such as signal strength, mobility, and energy efficiency. Despite real-world deployments in areas like military communications, MANETs remain challenging to operate due to high overheads and dynamic topologies. Recent studies focus on using techniques such as ML, network coding, and cross-layer optimisation to enhance reliability, energy efficiency, and quality of service (QoS) across various demanding scenarios.*

### Review

Mobile ad hoc network (MANET) formation has been a research topic of interest for decades with numerous methods being proposed and tested in simulated conditions. Summaries of some popular strategies across the literature are reviewed by Agarwal and Motwani [111], Alinci et al. [112], Sharifi and Babamir [116]. MANETs have no infrastructure and a highly dynamic network topology. The nodes detect each other through radio waves and are limited in energy due to their mobility [116]. The way that a MANET is formed is most influenced by the requirements and performance of routing while being limited by the effectiveness of distributed clustering protocols.

In general, the performance of routing is optimised by considering local metrics such as signal strength, distance, connectivity degree, energy efficiency, and the mobility of users [116]. Most proposed MANETs operate using the same fundamental ideas. They generally can be described as a multihop wireless network consisting of mobile nodes with limitations in battery power, spectrum usage is managed collectively rather than by fixed infrastructure. MANET implementations can vary in how much information is collected and shared for self-management with the trade-off being between the quality of the network formation at the expense of increased overheads. Other differences between MANET implementations include how information is shared, either directly from node to node or only through a cluster head. The organisation of clusters can also vary, with all members being either 1-hop or n-hops away from the cluster head [111, 112, 116].

Operating MANETs is known to be a challenging problem across the literature. Although there are already deployable MANETs available in the real-world, mainly for military applications, their usefulness is still quite limited with its use mostly being limited to establishing basic communications in places without existing reliable communications infrastructure. To improve

a MANET's practicality for a wider range of applications, such as commercial telecommunications, solutions are being researched to address its shortcomings. These include lowering the substantial overhead to set up and operate a MANET and improving the performance of its network. Machine learning (ML) in particular stands out as a key approach to consider because its characteristics of being fast, low complexity at runtime, and autonomous complement the needs of a decentralised, quick roll out, low overheads network. Abboud et al. [127] explore applications of ML and deep reinforcement learning (DRL) for MANETs, integrated with software defined network (SDN)s with the goal of optimising energy usage across a network. They address the challenges posed by node mobility and limited resources by balancing the load across nodes, thus preventing rapid energy depletion in specific nodes. By utilising DRL, the proposed approach successfully identifies alternative packet routing paths that avoid nodes with low energy, significantly extending the time until the first node exhausts its energy despite a moderate increase in average hop count.

There are many roles that a MANET can fulfill. One example application of MANETs and arguably the most conventional is to carry communication payloads such as voice, SMS, or data packets. In particular, packet payloads is how most modern networks transmit data such as 4G long-term evolution (LTE) and 5G new radio (NR) so this is an important application for MANETs. However, most conventional networks have static infrastructure unlike a MANET that by definition has mobile infrastructure, or can even be infrastructure-less. This problem is considered to be challenging in this research space and has attracted the attention of a lot of research. For example, Do-Duy et al. [128] investigate packet-level network coding to enhance multicast throughput in integrated MANET-satellite networks, particularly in disaster relief and emergency scenarios. A multicast optimisation problem is formulated to achieve a targeted packet delivery probability while considering node computational constraints. Simulation results demonstrate substantial improvements in achievable throughput compared to conventional routing, particularly benefiting devices with higher computational resources.

In a MANET, new challenges can arise because, in addition to knowing the cost to traverse the network, it is also necessary to determine the locations of other nodes and whether they have moved. Zhang et al. [129] introduce an improved greedy forwarding routing strategy for MANETs designed to mitigate routing holes caused by node movement, energy depletion, and obstacles. The method calculates a metric based on link quality, node energy, and neighbour density to select the optimal next-hop node, incorporating a waiting strategy to resolve temporary routing holes. Experimental evaluations reveal the proposed algorithm significantly reduces energy consumption, enhances packet delivery rates, minimises network delay, and extends network lifetime compared to existing routing protocols such as greedy perimeter stateless routing (GPSR), enhanced message-aware greedy routing (EMGR), and enhanced directional greedy routing (EDGR).

Reliable multicast communication in dynamic spectrum environments requires specialised strategies to account for the unpredictability of spectrum availability. Dong et al. [130] review the specific challenges of multicast communication in multi-channel cognitive radio ad hoc network (CRAHN)s, where spectrum availability dynamically changes. They highlight three primary challenges of time-domain uncertainty, frequency-domain spectrum randomness, and spatial channel variability. They discuss joint multicast routing, spectrum allocation, cross-layer multicast scheduling, and QoS guarantees, emphasising how network coding techniques could significantly improve multicast performance under the inherent uncertainties of CRAHNs.

As wireless ad-hoc networks become more complex and high-capacity, there is growing interest in tightly integrated optimisation strategies that span multiple protocol layers. Kordbacheh et al. [131] present a robust cross-layer optimisation framework combining routing, sub-carrier scheduling, and power allocation in massive MIMO and OFDMA-based wireless ad-hoc networks. Acknowledging practical channel estimation errors and imperfect channel state information (CSI), they propose an iterative convex optimisation approach to maximise throughput under interference, power, and routing constraints. Simulation outcomes indicate that the joint cross-layer design significantly improves network performance, achieving approximately 26.5% higher throughput compared to separate resource allocation strategies.

Mobile ad hoc network (MANET)s have gained significant attention for government and military usage. They provide security benefits due to being capable of adapting to threats to spectrum such as *jamming* attacks which conventional networks generally have no or a minimal capability to respond to. These networks also allow temporary communication systems to be used in critical scenarios such as the response to an emergency that negatively impacted communications (after an earthquake for example). They may also be used for some military purposes in a location without conventional telecommunications infrastructure available. Nicholas and Hoffman [132] address computational complexities involved in dynamic spectrum allocation for military MANETs, emphasising the importance of channel reuse while managing co-channel interference and minimising the frequency of channel changes. Highlighting real-world combat scenarios, they demonstrate that traditional channel assignment algorithms face significant computational difficulties, especially when accounting for cumulative interference from multiple radios. They conclude by identifying inadequacies in existing solutions and recommending research directions to develop more efficient algorithms capable of addressing complex, large-scale scenarios.

Security is a fundamental consideration in MANETs, where decentralised and dynamic structures expose the network to a wide range of threats. Azeez and Otudor [133] propose an access control mechanism integrated into MANET routing protocols to enhance security against malicious attacks. Utilising NS-2 simulation, the model implements a three-stage access control process and evaluates its effectiveness through metrics such as packet delivery ratio and traffic

overhead. Results show that the proposed access control approach provides a good trade-off between security and performance, effectively enhancing the network's resilience to attacks and malicious nodes while maintaining efficient data transmission.

Evaluating the performance of CR-enabled MANETs requires detailed traffic models that reflect the variability and urgency of real-world usage. Gajewski et al. [134] develop a detailed traffic simulation model specifically designed for cognitive-radio-enabled MANETs supporting critical applications like public safety and military communications. They describe the hierarchical node structure, the statistical characteristics of service-specific traffic (voice, video, data), and outline a MATLAB-based traffic generator integrated into OMNET++ simulations. Through realistic modelling of user behaviours, connection randomness, and message prioritisation, this approach provides a comprehensive and accurate simulation tool for evaluating MANET performance under various operational scenarios.

Being able to quantify the performance of a MANET has numerous critical applications, such as simply being able to tell if it is functioning properly. More advanced applications include making real-time adjustments to improve the performance of the network or to respond to spectrum interference challenges or threats. Ahirwar et al. [135] discuss methods for enhancing QoS in MANETs through various meta-heuristic optimisation algorithms, focusing on parameters such as routing, security, energy efficiency, and mobility. They highlight the strengths and weaknesses of different approaches and provide performance analyses of each algorithm. They look into the ongoing challenges in ensuring QoS due to the dynamic nature of MANETs, suggesting directions for future research aimed at improving the reliability, energy efficiency, and security of MANET routing protocols.

Hybrid routing strategies that can adapt to changing mobility patterns are increasingly important for improving the responsiveness and efficiency of MANETs in dynamic environments. Sampath et al. [136] propose the group adaptive hybrid routing algorithm (GAHRA), which intelligently adapts between proactive and reactive routing based on node mobility patterns. Targeting scenarios where nodes exhibit group mobility, GAHRA dynamically adjusts its routing strategy, achieving better network performance through reduced latency and improved adaptability. Simulation results show that GAHRA provides substantial improvements in packet delivery, routing efficiency, and resilience compared to traditional single-mode routing approaches that rely exclusively on either proactive or reactive strategies without the flexibility to adapt to changing network conditions.

### 4.4.3 MANET Clustering

#### Summary

*Clustering is vital for managing mobile ad hoc network (MANET)s, especially in dynamic spectrum management (DSM), by grouping nodes by mobility, energy, or connectivity to optimise performance and resources. A variety of clustering algorithms such as mobility-based, energy-efficient, and weighted methods have been proposed, each with distinct strengths and limitations depending on network conditions. Analytical models and advanced protocols, like multi-hop time reservation using adaptive control for energy efficiency (MH-TRACE) and connected dominating set routing (CRD), further enhance clustering efficiency by reducing computational load and improving key metrics like energy use, throughput, and packet delivery. Survey studies look at the complexity of comparing these approaches and highlight ongoing challenges, driving innovations such as cross-layer frameworks and hybrid clustering models to boost scalability, stability, and QoS in MANETs.*

#### Review

Previously, the literature on mobile ad hoc network (MANET)s covered their usage in proposed spectrum management systems. They have potential for use in dynamic spectrum management (DSM) but may be limited by how well they are formed and maintained, and by device, power, and infrastructure requirements. Clustering methods and protocols play a big role in their success with many different techniques being proposed to find the optimal ways to cluster. Alinci et al. [112] review various clustering algorithms used in MANETs, highlighting four main categories: mobility-based, energy-efficient, connectivity-based, and weighted-based clustering algorithms. Each category selects cluster heads based on different metrics such as mobility, energy, node degree, or combined weight factors. They discuss the advantages and limitations of each approach, emphasising the importance of selecting the appropriate clustering strategy according to the network conditions and desired performance metrics.

Clustering protocols are one of the most important parts of how a MANET functions. Karaoglu et al. [137] develop an analytical model for evaluating the performance of soft clustering protocols, specifically the multi-hop time reservation using adaptive control for energy efficiency (MH-TRACE) protocol in mobile ad hoc network (MANET)s. They highlight how traditional simulation methods for protocol performance evaluation become computationally demanding as network complexity grows. Thus, their analytical model addresses this by quickly estimating critical performance metrics like energy consumption and packet reception rates under various conditions. The model allows optimisation of protocol parameters, enhancing energy efficiency and throughput, and is validated against simulations demonstrating its accuracy and efficiency.

Efficient channel utilisation and robust clustering are essential in cognitive radio (CR) mobile



ad hoc network (MANET)s, where interference management and spectrum agility are critical for maintaining communication quality. Tran et al. [138] propose a connected dominating set (CDS)-based clustering and routing protocol for multi-channel CR MANETs which dynamically selects channels for efficient communication. The protocol consists of a clustering phase to create stable clusters using CDS, followed by a dynamic channel-selection-based routing phase to mitigate interference with primary user (PU)s. Simulation results indicate that the protocol significantly improves the packet delivery ratio, reduces control overhead, lowers delay, and conserves energy compared to traditional routing protocols like ad hoc on-demand distance vector (AODV).

With so many protocols and ways to cluster MANETs available, it becomes challenging to directly compare each without a consistent set of criteria. In the literature, this difficulty has been recognised, as reflected by the quantity of survey papers available to do this comparison. Rahman et al. [139] categorise existing clustering schemes in MANETs based on cluster head selection criteria, thoroughly evaluating their performance concerning QoS metrics. They discuss the trade-offs among different approaches, such as mobility-based or energy-based clustering, identifying significant factors that impact performance, such as scalability and stability. Additionally, they highlight open research challenges and propose solutions, including a cross-layer framework and a hybrid clustering model, to enhance future clustering protocols' efficiency and robustness.

As clustering protocols become more diverse, comparative evaluation frameworks are needed to guide protocol selection under specific performance constraints. Sharifi and Babamir [116] review clustering algorithms in MANETs, categorising them based on their methodologies, and evaluating their performance against critical metrics such as mobility, energy efficiency, and connectivity. They use multi-criterion decision-making methods and the analytical hierarchical process to rank these algorithms. This structured approach identifies strengths and weaknesses of various clustering methods, guiding selection based on specific network requirements and optimisation goals.

Numerous clustering strategies have been proposed to improve the scalability and efficiency of MANETs, each introducing trade-offs that must be carefully considered. Agarwal and Motwani [111] survey various clustering algorithms developed for MANETs, focusing on performance metrics including node ID, mobility, energy, and weighted factors for cluster head election. They categorise and critique methods like Lowest-ID, Max-Min d-cluster, and others, discussing how each method impacts network overheads and scalability. They highlight the challenges faced by different clustering algorithms, including energy efficiency, communication overhead, and stability, offering insights into their suitability under varying network conditions.

An interesting concept for MANET formation is described by Tavli and B. [140]. They exam-



ine the design of energy-efficient real-time data communication protocols specifically tailored for mobile ad hoc network (MANET)s. They introduce the time reservation using adaptive control for energy efficiency (TRACE) family of protocols designed to optimise communication performance across key metrics, including throughput, energy consumption, quality of service (QoS), and network robustness. They cover a wide range of protocol architectures (SH-TRACE, MH-TRACE, NB-TRACE, and MC-TRACE), addressing challenges such as channel access co-ordination, broadcasting and multicasting support, and resilience to channel errors. Extensive analytical modelling and simulation results highlight the effectiveness of the proposed architectures in reducing energy consumption and maintaining QoS. They emphasise the importance of coordinated channel access and cross-layer design in achieving both energy efficiency and high performance in dynamic, infrastructure-less wireless network environments.

#### 4.4.4 Isochrones & Shortest Path Trees (SPTs)

##### Summary

*In the literature, the idea of an ideal isochrone is often used to benchmark transport efficiency by comparing real-world travel conditions to a theoretical model based on constant, uninterrupted speed. Isochrones can also be combined with live traffic data and demographic information to improve business decision-making and location planning. Route optimisation for specific users, such as cyclists, highlights the importance of incorporating individual preferences and safety considerations into travel models. Some frameworks challenge fixed travel-time thresholds by emphasising the need to account for both proximity and diversity of accessible services. Isochrones are useful for identifying gaps in emergency service coverage and can guide infrastructure improvements when traditional administrative boundaries lead to inefficiencies. Alternative methods, such as data-driven modelling and network-based analysis, offer more flexible and detailed approaches to understanding accessibility, equity, and infrastructure resilience.*

##### Review

A common application of isochrones is demonstrated by Śleszyński et al. [141] who propose the concept of an *ideal isochrone*, a theoretical model used as a benchmark to assess the efficiency of transport systems in terms of spatio-temporal accessibility. An ideal isochrone is defined as a perfect circle representing the area reachable in uniform, uninterrupted travel conditions at a constant maximum speed. Real isochrones, derived from actual travel conditions and infrastructure, are compared against this ideal to measure transport efficiency through indicators that evaluate how closely real-world accessibility matches the theoretical optimum. They apply this method to Warsaw's metropolitan area, illustrating how it provides quantitative insights into spatial equity, urban planning, and business decisions. They demonstrate that using ideal isochrones helps identify areas where accessibility is suboptimal, enabling policy-makers to recognise and address inefficiencies, ultimately supporting better transport planning and improved urban development outcomes

Another application of travel isochrones is shown by Efentakis et al. [142] who present a method for calculating catchment areas and reachability analysis using isochrones (areas reachable within certain time intervals) integrated with live traffic data derived from floating car data (FCD). They developed a web-based geomarketing demo utilising live-traffic updates and demographic information from Vienna and Berlin, allowing businesses to assess location strategies dynamically. They highlight that incorporating real-time traffic significantly impacts the accuracy of geomarketing analytics and business decision-making, outperforming static network-based approaches that do not consider traffic fluctuations. The impact of traffic is arguably equally valid for use spectrum management scenarios.

Generally, isochrones for travel are not fine-tuned for more niche applications. Freitas [143] propose an innovative approach to bicycle route optimisation, focusing on the specific needs and preferences of cyclists, which traditional routing algorithms often neglect. The proposed algorithm goes beyond shortest distance or time to incorporate multiple criteria such as gradient, pavement type, urban obstacles, safety, and cyclist comfort. By integrating multi-criteria optimisation, the study demonstrates enhanced cycling experiences, promoting safer and more pleasant rides. Although the methodology addresses important cycling concerns effectively, it also recognises challenges in data collection, computational complexity, and subjectivity in weighting cyclist preferences. This level of fine-tuning for the underlying technique used to generate the isochrones has application in this work for spectrum demand modelling as more details can be considered about the simulated agents' preferences and habits to more accurately model the paths they may take.

The popular *15-minute city* concept aims to provide citizens with essential amenities within a 15-minute walk. It has often been criticised for its uniform temporal threshold not being suitable for every city and demographic. Biraghi et al. [144] introduce the *CityTime* framework as a flexible alternative arguing that proximity should be balanced with the diversity of amenities. They develop a GIS-based analysis that computes isochrones and compares the richness of services around each home location across Dakar, Rio de Janeiro and Milan. Their case studies show that the optimal travel threshold varies across neighbourhoods and cities and that enforcing a strict 15-minute limit can mask inequalities. *CityTime* allows users to customise the mix of services they value, democratising the 15-minute city idea and informing equitable urban strategies. A similar issue could occur in spectrum management problems where systems may be beneficial to those using cars and trains but not to those taking buses as an example.

In any city in the world, emergency medical systems rely on rapid response times. Zhao and Zhou [145] use isochrone coverage measures to evaluate Beijing's pre-hospital emergency facility plan, calculating population and area coverage at 8 to 10 minute thresholds. They find that although more than 90% of the population is covered within 10 minutes, only 78% is served within eight minutes, overlapping near administrative boundaries, creating oversupply in some districts. They recommend adding stations in underserved areas and planning emergency facilities at the city scale rather than within administrative districts. They also note that traffic conditions substantially affect isochrone shapes. Relating this back to spectrum management, emergency medical systems and gaps between coverage of infrastructure are both very relevant. To consider mobility in a spectrum allocation system for those allocated, it is critical to know if the system performs reliably under all conditions that are likely to occur, especially for an emergency medical situation.

Traditional network-based isochrone models require detailed travel cost functions, which are computationally expensive over large regions. Shahabi and Kim [146] develop a data-driven



trajectory generator that creates realistic vehicle trajectories from massive real-world traffic data to generate reachability maps for the Los Angeles metropolitan area. Instead of assigning fixed travel speeds, they interpolate and compute missing traffic data to produce continuous trajectories and then construct isochrones for arbitrary origins and times. The resulting web application allows users to select a location, travel time and view dynamic accessibility maps. Policymakers can use it to evaluate region-wide accessibility and identify disadvantaged areas. This data-driven approach offers a scalable alternative to conventional graph-based methods and supports more responsive transport planning.

Standard isochrone zones assume uniform accessibility within a catchment, which may overlook variations within the zone. To address this, Aydin et al. [147] apply a self-avoiding random walk (SARW) algorithm to educational accessibility in Helsinki and compare it with traditional isochrone-based metrics. Using historical data from 1991–2016, they show that SARW produces detailed node-level accessibility surfaces and captures dynamic changes in hot and cold spots over time. The SARW metric better reflects the street network’s topology and reveals that some neighbourhoods experience declining accessibility despite falling within isochrone boundaries. Their findings suggest that more nuanced metrics, such as SARW, can complement isochrones and provide richer insights into spatial equity.

Electric-vehicle routing poses additional challenges because vehicles must recharge en route. Baum et al. [148] formulate the shortest feasible paths with charging stops problem, extending the constrained shortest path problem with realistic charging models, including varying charging power and battery swapping. They propose CHarge, a combination of algorithmic techniques (contraction hierarchies and ALT etc) that computes optimal electric vehicle (EV) routes even on continental networks. Experiments show CHarge can find high-quality routes in subsecond times and can be adapted into heuristic variants that deliver near-optimal solutions faster.

## 4.5 Demonstrations

### 4.5.1 Context

Earlier discussion outlined the challenges of managing spectrum access under realistic, large-scale, and time-critical conditions, and reviewed decentralisation, n-tier coordination, mobile ad hoc network (MANET) clustering, and shortest path tree (SPT) modelling as candidate solutions. The demonstrations follow directly from this work by providing applied tests of how these mechanisms operate when placed in a constrained and dynamic environment. Their role is to move beyond theory and background, showing how the proposed approaches behave in practice.

The literature review established that decentralisation reduces latency, n-tier structures maintain coherence, and MANETs offer adaptable, infrastructure-light communication. However, it also highlighted that clustering is essential for sustaining performance at scale, and that realistic reachability modelling is required to capture user mobility. These findings position clustering protocols and SPT-based mobility models as critical elements to evaluate, since they directly determine whether decentralised spectrum management can remain both responsive and efficient in realistic conditions. The demonstrations therefore build on this foundation, combining the theoretical insights into applied settings.

The demonstrations implement a simple decentralised clustering protocol, visualise the resulting topology in a simulated city, and extend the analysis by applying SPTs to measure reachability in both urban and rural contexts. Clustering illustrates how spectrum users can be grouped under mobility and distance constraints, reducing the need for global coordination, while SPTs highlight how far users can realistically travel or communicate within fixed budgets of time or cost. Together, these demonstrations show how concepts reviewed earlier in the chapter translate into working examples. They also highlight where the approach remains simplified and where future work, such as detailed 3D propagation and behaviour modelling, would be required to draw stronger conclusions.

### 4.5.2 Summary

These demonstrations explore how decentralised clustering protocols and shortest path tree (SPT) modelling from existing research can be applied to realistic spectrum management scenarios. Clustering helps group nearby spectrum users to reduce interference risk and minimise communication overhead by enabling local coordination through one cluster representative, while decentralisation addresses the limitations of centralised systems in rapidly changing environments by allowing faster, more scalable, and robust decision-making. SPTs complement this by modelling realistic user reachability, showing how far agents can move or communicate within time or cost budgets in both urban and rural environments. Although clustering and SPTs were shown to effectively organise and describe potential spectrum usage, neither has yet been applied to actual spectrum allocation, and the simulations remain limited by simplified movement and life-pattern modelling. Both approaches are lightweight, flexible, and well supported in research, making them suitable for energy-constrained mobile hardware. Future work would focus on testing in more realistic 3D environments with detailed RF propagation, mobility, and behavioural modelling to better evaluate performance.

### 4.5.3 Structure

This section presents three main demonstrations.

1. **Decentralised Clustering Protocols** — A protocol where agents periodically update their neighbour list and share mobility, connectivity, distance, and power metrics to form clusters in a fully decentralised fashion.
2. **Visualisation of Decentralised Clustering Simulation** — Simulation snapshots showing coloured cluster regions, intra-cluster links, inter-cluster connectivity, and cluster heads.
3. **Reachability Modelling with Shortest Path Trees** — Spatial reachability analysis using SPTs to show the extent of movement possible within fixed time budgets across different transport modes and environments.

#### 4.5.4 Demonstration 1: Decentralised Clustering Protocols

Table 4.1 summarises the structure of Demonstration 1, presenting the objective of localised cluster formation, the key method involving periodic neighbour updates, metric exchange and cluster head self nomination, and the data inputs which include agent position data from a routing engine, sensing range, and mobility, connectivity, distance and power metrics.

Demonstration	Objective	Key Method	Data Inputs
1.1	Localised cluster formation	Periodic neighbour updates; metric exchange; cluster-head self-nomination	Agent position data from routing engine; sensing range; mobility, connectivity, distance, power metrics

Table 4.1: Summary of Demonstration 1 – Decentralised Clustering Protocols



## Demo 1.1: Localised Cluster Formation via Metrics

In this demonstration,  $N$  agents are considered that are only aware of their immediate **neighbours**. At regular time intervals  $t_i$ , each agent updates their list of **neighbours**. The clustering metrics **mobility**, **connectivity**, **distance**, and **power** (defined in Table 4.2) are shared between neighbours using lightweight protocols that are assumed to take much less time than  $t_i$ .

The method for clustering to form a MANET follows algorithms described in [111, 112, 116] and uses a random start-up, sensing, and self-nomination strategy (Figure 4.3.1.1 from the *Background*). The process is fully decentralised, relying only on information exchanged within each agent's **sensingDist**. Once elected, a cluster head can coordinate agents outside each other's sensing range by managing their transmissions.

Metric	Description	Agent Attributes	Description
<b>connectivity</b>	Neighbor count	<b>aID</b>	Agent ID
<b>mobility</b>	Agent vector	<b>neighbours</b>	Agents within range
<b>distance</b>	Agent separation	<b>sensingDist</b>	Agent sensing range
<b>power</b>	Signal strength		

Table 4.2: Cluster Metrics

### Using connectivity

1. Each agent broadcasts their **aID**.
2. Each agent calculates their **connectivity** as the number of **aIDs** received.
3. Each agent broadcasts their **connectivity**.
4. **If** their **connectivity** is highest, **then** they self-nominate as cluster head; **else**, they search for a cluster head to join.

### Using mobility

1. **If** an agent's **mobility** exceeds a threshold, **then** they may cancel their cluster-head nomination.

### Using power

1. Each agent may use the received **power** level of a **aID** transmission to choose a preferred cluster head.

### Using distance

1. Each agent transmits their position to establish their **neighbours'** **distance**.

### 4.5.5 Demonstration 2: Visualisation of Decentralised Clustering

Table 4.3 summarises the structure of Demonstration 2. The objective is to illustrate how varying communication distance and link configuration changes the visible clustering topology in the Adelaide CBD. Sub-demonstrations 2.1 and 2.2 show polygon coverage derived from alpha shapes, highlighting the effect of large versus small communication ranges without internal link structures. Sub-demonstrations 2.3 and 2.4 then introduce minimum spanning tree (MST)-based intra-cluster links and inter-cluster connectivity, moving from minimal bridging to a more robust, multi-link topology. Together these frames show how decentralised clustering can scale from broad groupings to resilient inter-cluster networks.

Demonstration	Objective	Key Method	Data Inputs
2.1	Show cluster coverage at large communication distance	Alpha-shape polygons; no MST or interlinks	Agent position data from routing engine; Adelaide CBD map
2.2	Show cluster coverage at small communication distance	Alpha-shape polygons; no MST or interlinks	Agent position data from routing engine; Adelaide CBD map
2.3	Show structured clusters with minimal links	Polygons with MSTs; one inter-cluster link per neighbour	Agent position data from routing engine; Adelaide CBD map
2.4	Show robust inter-cluster connectivity	Polygons with MSTs; multiple inter-cluster links	Agent position data from routing engine; Adelaide CBD map

Table 4.3: Summary of Demonstration 2 – Visualisation of Decentralised Clustering Simulation

## Demo 2.1: Large-Distance Cluster Coverage

### Setup:

- Agent positions dataset generated using the Valhalla routing engine [120].
- Clusters formed using the decentralised cluster formation method from Demo 1.1.
- Communication distance fixed at 10 km.

Figure 4.5.5.1 shows clusters formed when the communication distance is set to a high value. Alpha-shape polygons expand to cover broad areas of the Adelaide CBD because the communication distance is large, resulting in bigger clusters with fewer but larger regions. MSTs and inter-cluster links are not considered here.



Figure 4.5.5.1: Demo 2.1: Cluster coverage at large communication distance

## Demo 2.2: Small-Distance Cluster Coverage

### Setup:

- Agent positions dataset generated using the Valhalla routing engine [120].
- Clusters formed using the decentralised cluster formation method from Demo 1.1.
- Communication distance fixed at 1 km.

The frame in Figure 4.5.5.2 uses a smaller communication distance, so the clusters formed are smaller and more numerous around the CBD. MSTs and inter-cluster links are also not considered here.




Figure 4.5.5.2: Demo 2.2: Cluster coverage at small communication distance

## Demo 2.3: Structured Clusters with Minimal Links

### Setup:

- Agent positions dataset generated using the Valhalla routing engine [120].
- Clusters formed using the decentralised cluster formation method from Demo 1.1.
- Communication distance fixed at 2 km.
- MST produced using the python package SciPy [149].
- One interlink per neighbour for best closest edge node

Figure 4.5.5.3 introduces MST-based intra-cluster structures within each polygon. Each cluster is connected internally with coloured MST links, and exactly one inter-cluster link is drawn between each neighbouring cluster. This produces a minimally connected but stable topology across the CBD.

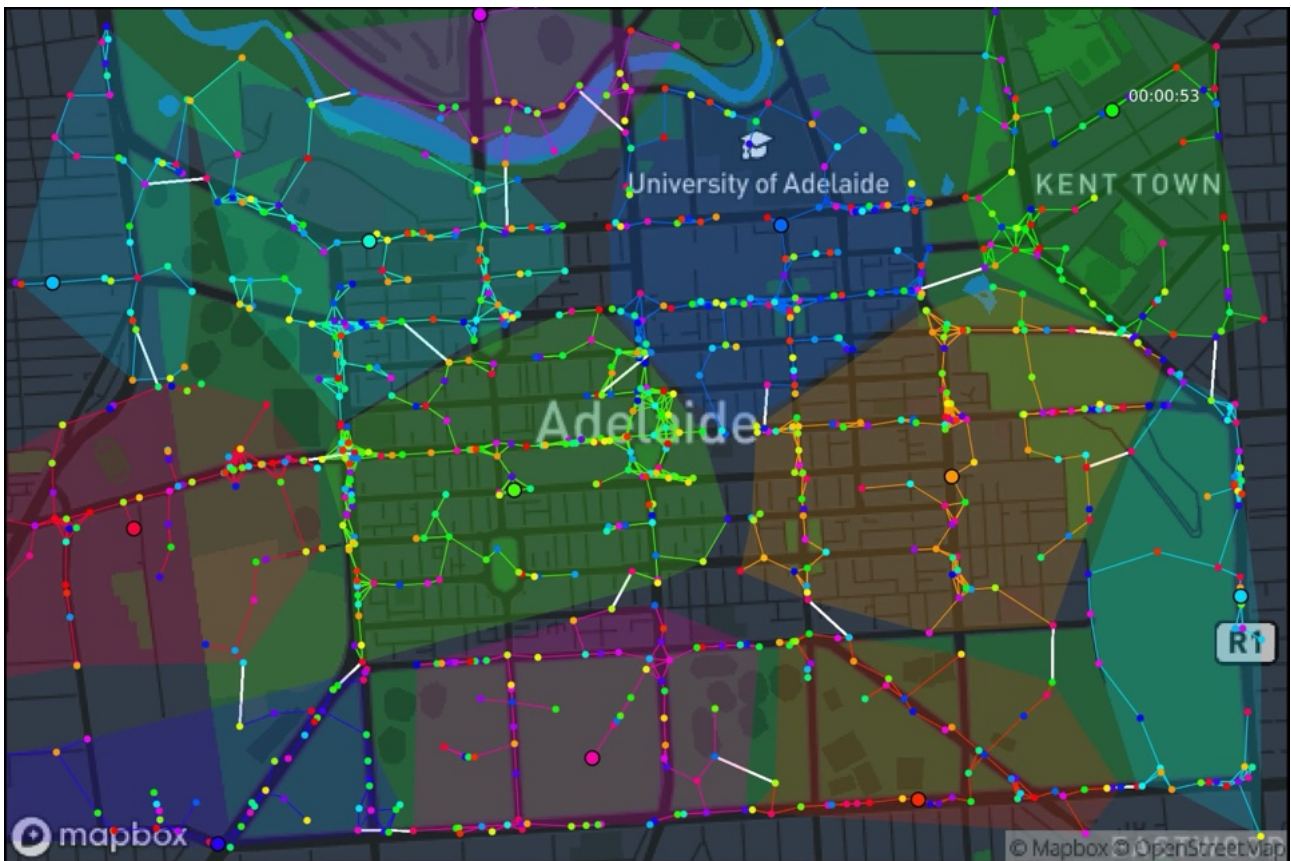


Figure 4.5.5.3: Demo 2.3: Structured clusters with MSTs and minimal inter-cluster links



## Demo 2.4: Robust Inter-Cluster Connectivity

### Setup:

- Agent positions dataset generated using the Valhalla routing engine [120].
- Clusters formed using the decentralised cluster formation method from Demo 1.1.
- Communication distance fixed at 2 km.
- MST produced using the python package SciPy [149].
- Unlimited interlinks per neighbour edge nodes

Finally, Figure 4.5.5.4 extends the previous arrangement by allowing multiple inter-cluster links. Clusters retain MST-based internal structures but are now densely connected to neighbours, producing a more robust and resilient inter-cluster network across the CBD.



Figure 4.5.5.4: Demo 2.4: Robust inter-cluster connectivity with multiple interlinks

### 4.5.6 Demonstration 3: Reachability Modelling with Shortest Path Trees

Table 4.4 summarises the structure of Demonstration 3. The objective is to illustrate how shortest path trees (SPTs) can be used to represent agent reachability in both urban and rural environments under different mobility modes. Each sub-demonstration uses a 10-minute travel budget and overlays SPTs on satellite imagery to highlight the areas reachable within that time frame. By comparing car, walking, and cycling modes across urban and rural settings, the frames show how transport context strongly influences reachability and network coverage.

Demonstration	Objective	Key Method	Data Inputs
3.1	Urban walking reachability	SPT for 10-min walking paths in CBD grid	Path network data; urban satellite imagery
3.2	Urban cycling reachability	SPT for 10-min cycling trips along road network	Road and cycle data; urban satellite imagery
3.3	Urban car reachability	SPT for 10-min car trips over road network	Road network data; urban satellite imagery
3.4	Rural walking reachability	SPT for 10-min walking from small settlement	Path network data; rural satellite imagery
3.5	Rural cycling reachability	SPT for 10-min cycling around rural town	Road and cycle data; rural satellite imagery
3.6	Rural car reachability	SPT for 10-min car trips from rural centre	Road network data; rural satellite imagery

Table 4.4: Summary of Demonstration 3 – Reachability Modelling with Shortest Path Trees





## Demo 3.2: Urban Cycling Reachability

### Setup:

- SPT produced using the Expansion [115] feature in the Valhalla routing engine [120].

Figure 4.5.6.2 presents a 10-minute SPT for cycling in an urban environment. The coverage extends well beyond the CBD grid but remains more constrained than car-based reachability. It demonstrates the level of reach offered by cycling compared to walking and driving.

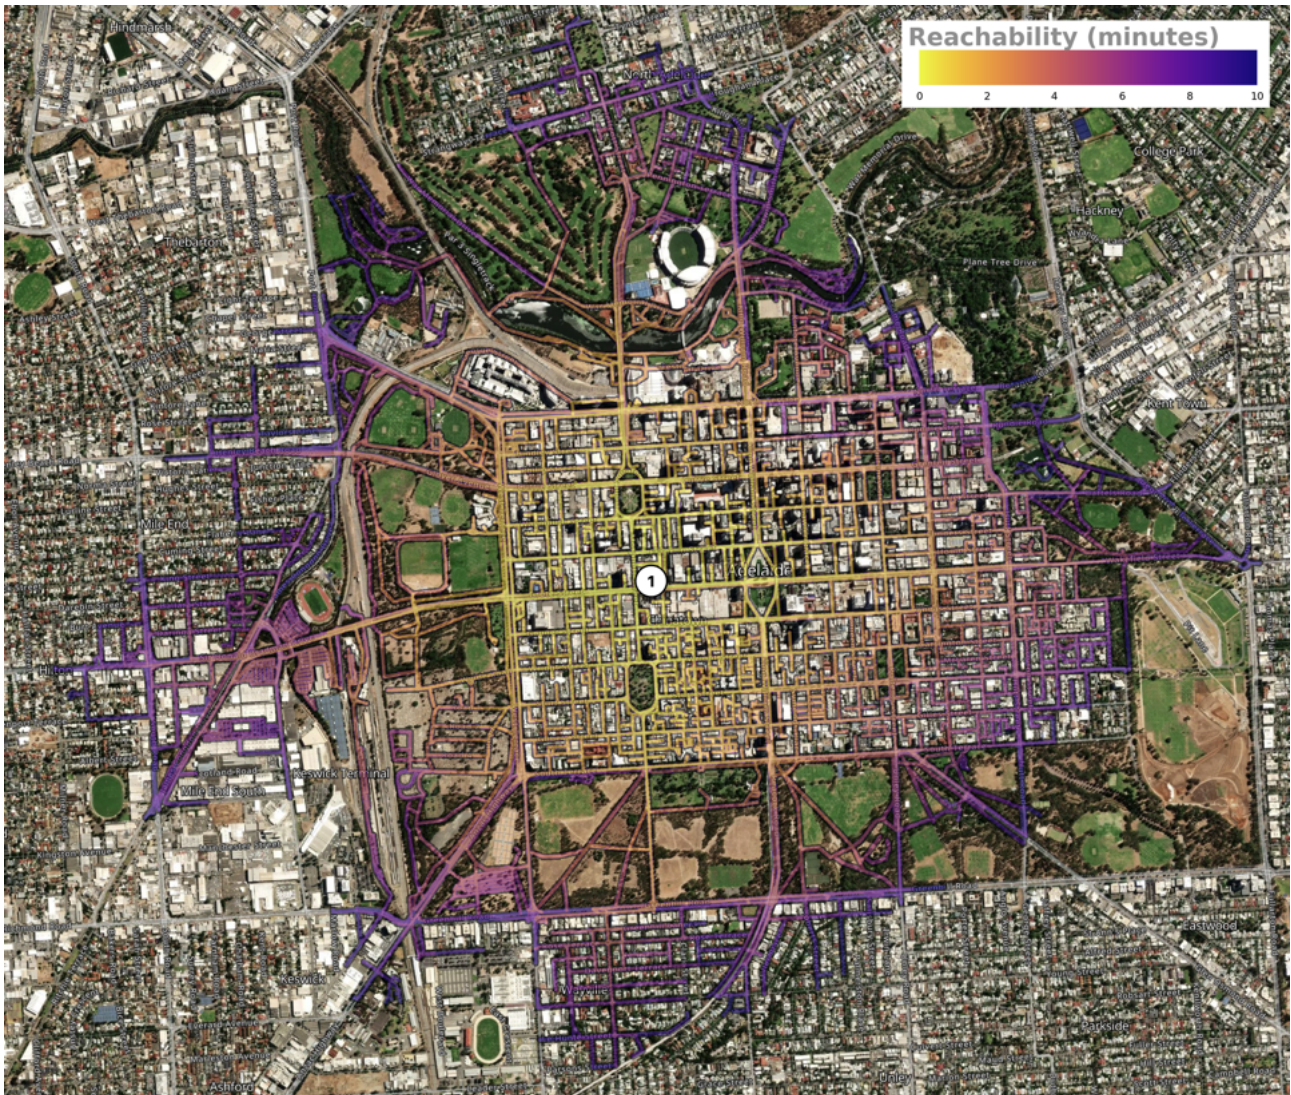


Figure 4.5.6.2: Demo 3.3: Urban cycling reachability within 10 minutes



### Demo 3.3: Urban Car Reachability

#### Setup:

- SPT produced using the Expansion [115] feature in the Valhalla routing engine [120].

Figure 4.5.6.3 shows the SPT reachability of a 10-minute car journey from the Adelaide CBD. The structure highlights the dense and extensive coverage possible within this time, spanning suburbs in multiple directions. The colour gradient illustrates travel cost (in seconds), with yellow areas closer and purple further from the origin.

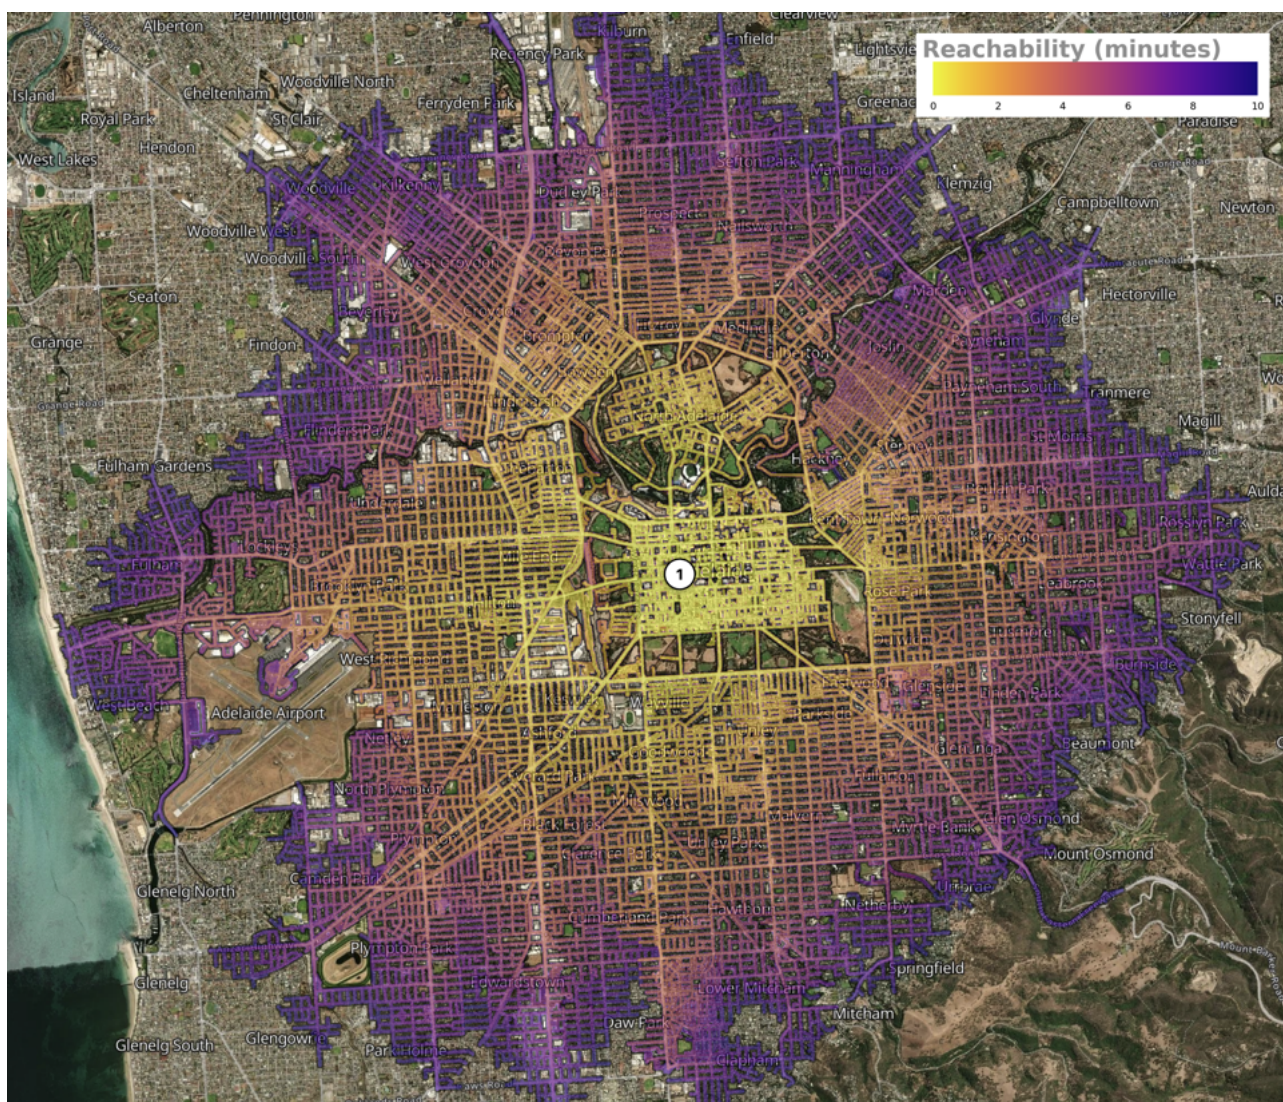


Figure 4.5.6.3: Demo 3.1: Urban car reachability within 10 minutes



### Demo 3.4: Rural Walking Reachability

#### Setup:

- SPT produced using the Expansion [115] feature in the Valhalla routing engine [120].

In Figure 4.5.6.4, the SPT shows a 10-minute walk from a small rural town (Uraidla). Reachability is sparse and limited to the immediate town area, reflecting both slower travel speed and reduced path density compared with urban grids. A key observation is that only the mapped paths are shown, so the area possible reachability would be more extensive than this.



Figure 4.5.6.4: Demo 3.4: Rural walking reachability within 10 minutes



### Demo 3.5: Rural Cycling Reachability

#### Setup:

- SPT produced using the Expansion [115] feature in the Valhalla routing engine [120].

Figure 4.5.6.5 depicts a 10-minute cycling SPT from a rural centre. The coverage extends along connecting roads, capturing both the larger spatial scale possible by cycling and the fragmented structure imposed by sparse rural networks.




Figure 4.5.6.5: Demo 3.5: Rural cycling reachability within 10 minutes



## Demo 3.6: Rural Car Reachability

### Setup:

- SPT produced using the Expansion [115] feature in the Valhalla routing engine [120].

Finally, Figure 4.5.6.6 illustrates rural car reachability within 10 minutes. Coverage spreads widely into the surrounding countryside, with SPT branches extending along arterial routes and highways. This frame shows how motorised transport substantially expands reachability in low-density regions, which are often less walkable due to having less footpaths and more challenging terrain conditions, and private property.

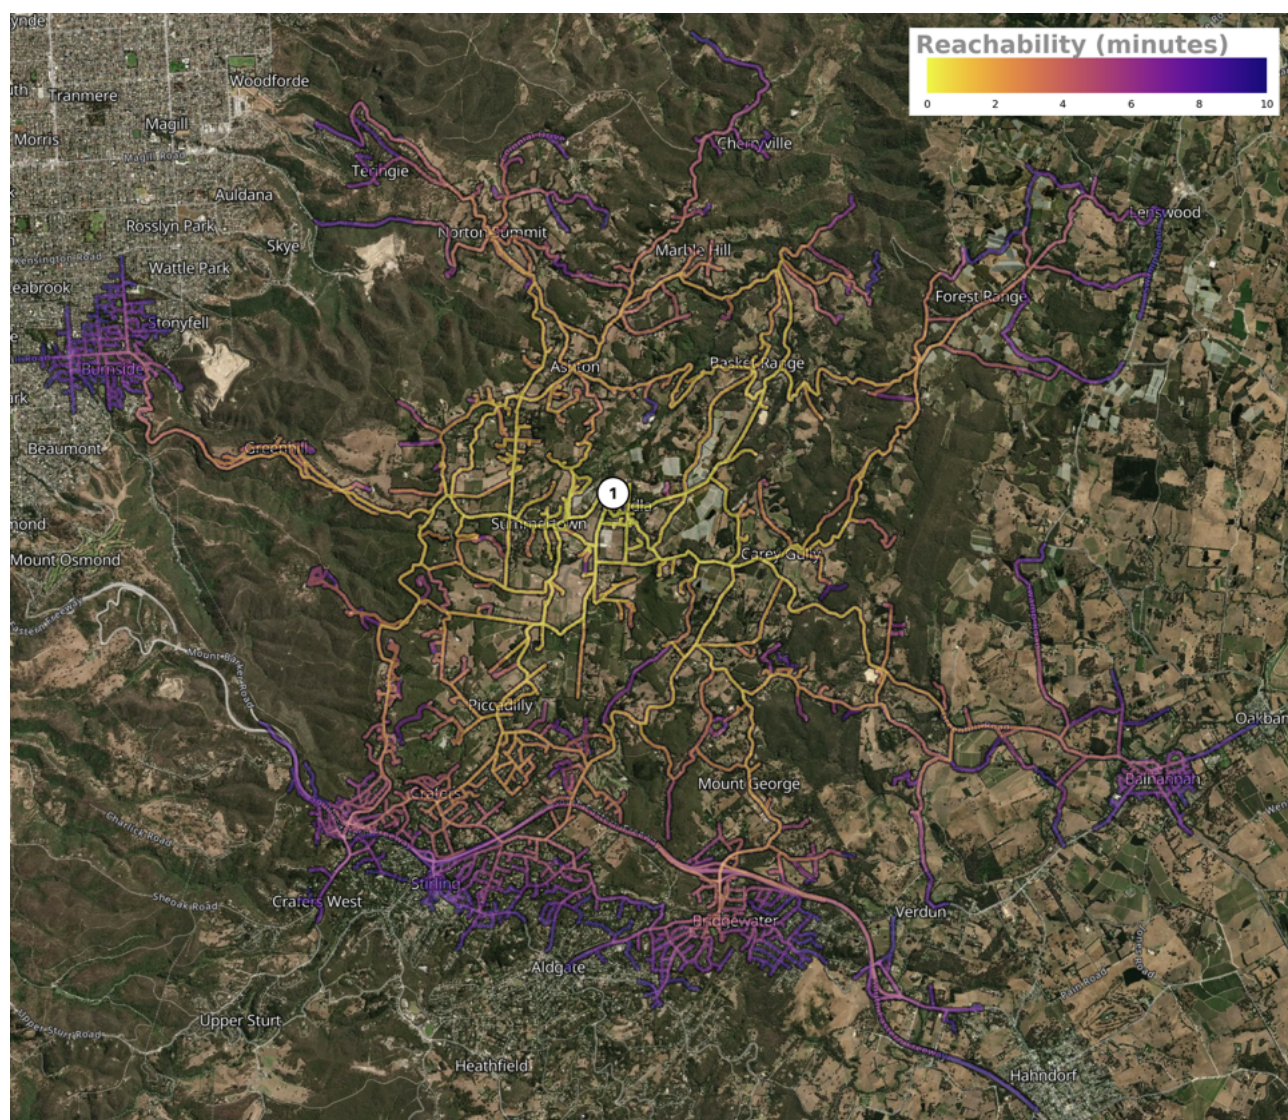


Figure 4.5.6.6: Demo 3.6: Rural car reachability within 10 minutes

## 4.6 Discussion

### 4.6.1 Summary

Chapter 4 has five sections: *Introduction*, *Approach*, *Literature*, *Background*, and *Demonstrations*.

The *Introduction* describes the problem of realistic real-time dynamic spectrum management (DSM) within a 3D geospatial digital twin to coordinate time-sensitive access to radio frequency (RF) spectrum across large areas. Latency, infrastructure limits, and geospatial diversity are identified as constraints, requiring rapid decisions and scalable computation. Mobile ad hoc network (MANET)s and hierarchical control are proposed as mechanisms to address these constraints, while reachability models such as shortest path tree (SPT) are explored to produce realistic mobility modelling and aid spectrum planning.

The *Approach* presents the proposition that spectrum management is a large-scale spatial problem subject to strict temporal constraints. The central question is framed as which techniques can deliver real DSM with meaningful distributed control and coordination. Attention is placed on decentralisation through hierarchical management, MANET formation and clustering, and mobility-aware reachability via isochrones and SPTs.

The *Literature* explores the four focus areas. Decision latency is shown to be reduced by decentralisation, although hidden-node conflicts are introduced, so hybrid hierarchical systems are described as combining local agility with coordinated oversight. Research on MANETs is reviewed across routing, network coding, cross-layer optimisation, and security, with improvements reported in energy use, throughput, and reliability. Clustering surveys and protocols are reported as highlighting trade-offs among stability, overheads, and scalability, while isochrones and SPTs are presented as practical tools for reachability analysis under variable mobility.

The *Background* formalises a distributed cluster-head nomination algorithm for MANET formation using broadcast identifiers, neighbour counts, random waits, and self-nomination to assign cluster heads. Routing engines are reviewed, and Valhalla is selected for its high-quality SPT outputs. SPTs are preferred over isochrones because contouring is identified as introducing unrealistic interpolation. Integration into the digital twin for realistic modelling is enabled by open-source data and tools.

The *Demonstrations* implement a decentralised clustering protocol in which mobility, connectivity, distance, and power metrics are exchanged by agents at regular intervals. Agents are organised into clusters with elected heads that coordinate intra- and inter-cluster links, extending control beyond direct sensing range. In addition, shortest path trees (SPTs) are applied to



model agent reachability within fixed time budgets, illustrating how urban and rural environments, and different transport modes, shape the extent of potential connectivity. The results show both distributed structuring of potential users for DSM through clustering and realistic spatial reachability through SPTs, though allocation is not considered and both mobility and life-pattern modelling remain simplified. Future work is directed toward extending these methods into 3D environments with richer RF propagation, mobility modelling, and behavioural data to evaluate performance and limits under realistic conditions.

## 4.6.2 Key Findings

### **Finding 1: Decentralisation balances local latency with global coordination**

Decentralisation reduces latency in spectrum allocation by distributing decision-making authority to local nodes, aligning with the *Approach*'s framing of spectrum management as a large-scale spatial problem with tight temporal constraints. The *Literature* demonstrates that fully decentralised systems risk hidden node conflicts but n-tier hybrid models balance local adaptability with system-wide coordination, as seen in hierarchical spectrum management protocols. Demonstrations of decentralised clustering protocols support this by showing how clusters can self-organise through lightweight exchanges of mobility, connectivity, and power metrics to maintain coordination under dynamic conditions. This finding suggests that hybrid decentralised designs could be useful for achieving scalable, responsive spectrum management in realistic environments.

### **Finding 2: Clustering enables scalable spectrum management**

Clustering groups of mobile agents based on metrics such as mobility and connectivity, directly addresses the *Approach*'s concern with scalability in real-time DSM solutions. The *Literature* shows that clustering protocols like MH-TRACE and CDS-based routing enhance energy efficiency, throughput, and stability in MANETs, supporting their application in spectrum management. The *Demonstrations* further explore this by visualising decentralised clusters within an urban simulation, where cluster heads coordinate internal communication and inter-cluster links maintain higher-level connectivity. This finding establishes clustering as a potentially useful method for structuring large-scale mobile spectrum networks into manageable, efficient, and stable subsystems.

### **Finding 3: Isochrones and SPTs improve reachability modelling**

Isochrones and shortest path tree (SPT)s provide a means to capture realistic mobility and reachability patterns, directly supporting the *Approach*'s emphasis on modelling user movement for practical DSM planning. The *Literature* shows their effectiveness across transport, emergency response, and infrastructure planning, where these tools highlight gaps, inefficiencies, and equity issues in spatial accessibility. SPTs were adopted for urban simulation because they better represent actual agent reachability than interpolated isochrone contours, enabling more accurate spectrum demand modelling. This finding shows that isochrones and SPTs could be useful for designing DSM frameworks that adapt to real mobility patterns and environmental constraints.



### 4.6.3 Limitations

While decentralisation demonstrably reduces latency, its implementation faces inherent risks such as hidden node conflicts that remain difficult to eliminate in practice. Even with hybrid n-tier systems, the coordination overhead between tiers may introduce new delays that counteract the benefits of localised decision-making. The demonstrations relied on simplified assumptions about neighbour discovery and metric exchange, which do not capture interference, packet loss, or the unpredictability of wireless propagation in dense environments. Furthermore, scalability in real world usage would depend on how effectively distributed nodes can synchronise without centralised visibility, a problem that could become problematic under realistic conditions where devices may vary in power, mobility, and protocol adherence. While decentralisation appears to offer a path to responsiveness, the complexity of balancing local and global control in realistic systems requires further work.

Although clustering was shown to support scalability, the approach depends heavily on stable cluster formation, which is not guaranteed in environments with high user mobility or intermittent connectivity. The literature highlights trade-offs in clustering algorithms between energy efficiency, stability, and communication overhead, and the demonstrations did not extend to conditions where frequent re-clustering would be required. Simulation environments assumed consistent sensing ranges and mobility models that may not translate to the irregular behaviours of real agents. Additionally, cluster head election introduces potential single points of failure, where loss or compromise of a head node could destabilise large sections of the network. Energy and computational constraints on mobile devices further limit the sustainability of clustering protocols at scale, leaving their practical viability for spectrum allocation uncertain.

Isochrones and SPTs provide useful abstractions of mobility but remain approximations that omit important dynamics of real-world movement and radio propagation. The demonstrations assumed deterministic travel networks and uniform agent behaviours, whereas real systems must contend with irregular pedestrian flows, transport delays, and obstacles not captured in the models. Interpolation and simplification in path construction risk misrepresenting edge cases, particularly under congestion or in environments with variable infrastructure quality. Furthermore, while SPTs provide a closer match to actual reachability than interpolated contours, their computational load may become prohibitive when scaled to large populations or fine-grained temporal updates. The absence of 3D propagation effects, interference modelling, and significant behavioural diversity in the current demonstrations restricts the applicability of findings to spectrum allocation in operational settings.

#### 4.6.4 Future Work

Further development could prioritise richer decentralisation strategies integrated with hierarchical management structures. Current implementations rely on simplified neighbour discovery and metric exchange. Future systems could incorporate interference awareness, realistic packet loss behaviour, and adaptive synchronisation across many device types. Incorporating advanced RF propagation models would better reflect real-world diversity in urban and rural settings, allowing for testing under interference-heavy conditions. This would enable hybrid frameworks that balance decentralisation with multi-tier oversight to be trialled under conditions closer to the real world, improving scalability while proving resilience.

Progress will also depend on advancing clustering protocols and MANET implementations under more demanding conditions. Future research could explore stability when agent mobility, energy variation, and irregular participation force frequent re-clustering. Integrating more sophisticated behavioural and life-pattern modelling would be essential to stress-test cluster formation and maintenance in realistic environments. Experimentation with cross-layer optimisation, security-conscious protocols, and adaptive election strategies could be necessary to address vulnerabilities such as single points of failure. These refinements could be evaluated using high-fidelity digital twins that function well for highly realistic agents, enabling assessment of scalability, stability, and resilience.

Finally, mobility modelling beyond SPTs could be explored. Incorporating traffic data, dynamic congestion modelling, and non-deterministic behaviours would improve fidelity, especially in urban environments with variable flows. Alternatives such as stochastic reachability metrics, probabilistic path modelling, or multi-modal trajectory simulations could be tested against SPTs to assess trade-offs in accuracy and computational cost. Realistic evaluation could be achieved by including these models in digital twin platforms with integrated agent movement, RF propagation, and spectrum allocation processes. Demonstrating such systems at city scale would provide a proof of concept for practical DSM frameworks that can adapt to real mobility, interference, and demand patterns.

## 4.7 Next Steps

This chapter introduced real world dynamic spectrum management (DSM) problem definitions and the concept of a digital twin as a possible solution. A key weakness identified is RF propagation modelling under real world conditions, so it is explored in the next chapter. The overarching objective is to enable reliable spectrum allocation and coordination by representing the physical environment and spectrum use within a unified digital model. Constraints include limited and heterogeneous data, variability across urban and rural terrains, material and structural complexity, computational efficiency, and the need for scalability and interpretability. The general approach combines geometry and terrain with physically grounded propagation behaviour, integrating multiple environmental datasets into a consistent three dimensional context and simulating coverage and interference across city and regional scales to guide allocation and protect incumbents.



## Chapter 5

# Modelling Radio Spectrum in 3-Dimensional Space

*This chapter outlines approaches for modelling radio spectrum usage in three-dimensional space, focusing on 3D city modelling, digital elevation models (DEMs) and RF ray tracing to support spectrum allocation in urban and rural settings. It reviews literature on the evolution of 3D modelling from costly LiDAR-based systems to accessible, open-source tools, highlights DEMs' role in terrain-aware propagation modelling, and examines advanced RF ray tracing integrated with machine learning for realistic, scalable simulations. The chapter also introduces Sionna Differentiable RT and radio material modelling for precise electromagnetic analysis. Four demonstrations illustrate practical applications: LOS mapping from elevation data, ray tracing using OSM-derived buildings, high-detail modelling with Google 3D tiles, and terrain mesh-based ray tracing. These examples show trade-offs between resolution, accuracy, data availability, and computational cost, and suggest future improvements like hybrid urban–rural models, enhanced material property data, and integration of multiple data sources for more accurate spectrum management.*

## 5.1 Introduction

In the previous chapter, RF propagation modelling considered only distance, ignoring the physical transmission environment. Under realistic conditions, the radio frequency (RF) propagation environment such as the terrain and buildings all affect how radio signals travel, reflect and attenuate. Without factoring in these influences, a spectrum management *digital twin* cannot reliably represent whether RF transmissions between two or more users would actually be harmful if they were allocated. The approach in this chapter integrates environment-aware RF propagation behaviour into the modelling process to achieve more accurate results across varied conditions. Recent advances in ray tracing engines such as Sionna RT, and the use of terrain-aware simulation, make this possible but require careful selection of input data and processing methods. Work to date has shown that openly available sources such as open street map (OSM) data and digital elevation model (DEM) data can reveal patterns in coverage that are not visible through distance-only methods [150]. The same testing has also identified where basic propagation assumptions fail in realistic terrain and city environments, motivating the use of more advanced modelling techniques presented later.

Urban and rural areas often require different RF modelling methods due to their contrasting physical and electromagnetic characteristics. High-density urban regions benefit from building-aware RF ray tracing that accounts for reflection, diffraction and shadowing caused by streets, buildings and terrain such as trees. Rural environments are better served by models that account for open-space loss and signal blockage from the elevation changes, terrain, and earth curvature. Datasets such as OSM [69], LiDAR-based DEMs [151] and detailed 3D meshes, including sources like Google 3D Tiles [152], supply the structural and topographic detail necessary to support these approaches. These datasets provide the structural and terrain detail needed for accurate environment-specific propagation modelling.

With environmental datasets in place, RF propagation can be modelled using a combination of line of sight (LOS) analysis and RF ray tracing. LOS mapping (often called viewshed analysis [153]) identifies clear transmission paths and is valuable for lower-frequency or rural studies, while ray tracing captures the multipath effects, reflections and obstructions that are critical in urban or irregular terrain. Integration of these methods in tools such as Sionna RT produces coverage maps that reflect both environmental complexity and operational conditions. This combined method can be applied to large regions or detailed street-level scenes without losing the accuracy needed for planning and interference analysis. Outputs from these models make it possible to visualise how variations in the environment cause significant differences in predicted performance. Practical applications of these methods show how different input data sources can be combined to simulate realistic RF propagation in both rural and urban environments.



## 5.2 Approach

### 5.2.1 Proposition

Effective ways to model spectrum usage in 3D space are needed to assess the impact of allocations

### 5.2.2 Questions

**Q1)** What options are available to model usage of radio spectrum in DSM for large urban and rural environments

### 5.2.3 Topics

**3D City Modelling** - Provides two essential options for DSM solution. These are a means to model where RF would cover, and a better way to represent and model users' movements, demand for radio spectrum, and vertical movements (such as up buildings) in a complex urban environment.

**Digital Elevation Model (DEM)** - Plays a similar role where geographical landscapes and terrain is a big factor such as rural areas, or in very hilly urban areas.

**Radio Frequency (RF) Ray Tracing** - Can provide the data to systems to make allocation decisions informed by coverage.

## 5.3 Literature

### 5.3.1 3D City Modelling

#### Summary

*3D city modelling has evolved significantly over the decades and is often used in computer generated imagery (CGI), gaming, and visualisation. It now serves critical functions for governments and industries in planning and analysis. While traditionally reliant on expensive technologies like LiDAR and satellite imagery, recent advances have made detailed global-scale 3D modelling more accessible to individuals through open-source and affordable tools. The integration of GIS-based techniques has further enhanced urban modelling by supporting dynamic, data-rich representations that aid in planning, infrastructure management, and spatial analysis. Recent tools like OpenGERT and novel remote spectrum mapping models demonstrate the growing importance of accurate 3D environmental modelling for applications in communication systems, highlighting a trend toward real-time, automated, and precise urban simulations.*

#### Review

Many global companies and governments rely on the creation and maintenance of digital 3D models for many different purposes and applications. Creating and maintaining the accuracy of these models becomes an increasingly challenging and unsustainable task once the scope expands beyond a small-scale project, such as a single small building. Companies and governments often use technologies such as light detection and ranging (LiDAR) combined with satellite imagery to achieve a level of detail that is considered acceptable for many basic applications. Generally these kind of tasks are not practical to do without being at the size of these companies or governments which have far better access to these technologies than most individuals. More recently however, numerous open-sourced and affordably priced projects have emerged making it far more practical for individuals to get access to 3D modelling data at a global scale with a useful amount of detail in them.

The evolution of digital technologies has significantly transformed how urban environments are modelled and analysed. In particular, the integration of GIS-based techniques has enabled more dynamic and detailed representations of cities, supporting a wide range of planning and analytical applications. Shiode [154] reviews recent advancements in digital 3D modelling of urban environments, highlighting the increasing adoption of geographic information system (GIS)-based techniques to visualise and analyse city structures. They categorise various modelling approaches based on realism, data input methods, and analytical functionality, ranging from basic 2D maps and textured 3D block models to highly detailed computer-aided design (CAD)-based visualisations. They identify different methods of data collection, including ter-



restrial images, aerial photographs, panoramic photography, and light detection and ranging (LiDAR). Applications of these models span urban planning and design, infrastructure management, commercial market analysis, and promotional activities, emphasising the importance of GIS integration for comprehensive spatial analysis and decision-making support. They argue that while standardisation remains elusive due to diverse modelling objectives and methodologies, the continuous evolution of spatial technologies suggests future widespread use of automated, real-time and online 3D urban models.

Accurate environmental modeling is essential for reliable wireless communication simulations, particularly in urban settings. Tadik et al. [155] introduce OpenGERT, an open-source automated geometry extraction tool tailored for accurate ray-tracing propagation modelling, particularly for NVIDIA's Sionna Ray-Tracing (RT) simulations. OpenGERT uses data from OpenStreetMap, Microsoft Global ML Building Footprints, and United States Geological Survey (USGS) elevation datasets to automate the creation of detailed urban environments. They performed extensive analyses in Munich to determine how height, position, and electromagnetic material properties of 3D models affect various RF channel statistics, including path gain, mean excess delay, delay spread, link outage, and Rician K-factor. Their findings reveal that small changes in building heights and positions significantly impact channel metrics, whereas slight variations in electromagnetic material properties have minimal effects. These results bring attention to the critical need for precise environmental modelling to ensure reliable propagation predictions, essential for developing digital spectrum twins and advanced communication networks. Their tool and associated analyses have been open-sourced to encourage further exploration and application.

Accurate estimation of spatial radio power distributions is crucial for effective spectrum management, especially in scenarios where direct access to certain regions is limited. To address this challenge, recent research has focused on innovative modelling and optimisation techniques that enable reliable predictions using sparse and remote measurements. Shen et al. [156] address the challenge of accurately estimating spatial radio power distributions in inaccessible areas using limited measurements from accessible locations. A novel remote compressed spectrum mapping (RCSM) model is proposed, utilising optimised sampling location selection based on maximising matrix determinants to capture spatial correlations efficiently. They apply a computationally efficient greedy optimisation method, significantly reducing computational complexity without sacrificing accuracy. They also incorporate a robust spatial power estimation algorithm using the alternating direction method of multipliers (ADMM). Simulation results confirm that the proposed RCSM algorithm achieves superior estimation accuracy and computational efficiency compared to existing methods, particularly in complex, heterogeneous 3D environments constrained by inaccessible regions.

### 5.3.2 Digital Elevation Model (DEM)

#### Summary

*Digital elevation model (DEM)s are foundational tools in geographic information system (GIS)s, offering 3-dimensional representations of the Earth's surface derived from sources like LiDAR, radar, and photogrammetry. Typically delivered as raster grids (generally in a common format such as GeoTIFF), DEMs support a wide range of applications including slope analysis, watershed modelling, and 3D urban planning. Their quality and utility vary significantly depending on data resolution, accuracy, and generation methods, with research highlighting the need for clear definitions and robust assessment techniques. In the context of this work, DEMs are particularly relevant for spectrum and demand modelling, where terrain data directly influences RF signal propagation and urban activity patterns which ultimately influence demand for radio spectrum.*

#### Review

Digital elevation model (DEM)s are very commonly used when looking at geographic information system (GIS) problems. They are typically a 3D representation of the Earth's surface, based on the elevation of the ground or terrain. They can also include the elevation of trees and buildings, depending on the level of post-processing done to the data afterwards. However, in this case the model is often referred to as a digital surface model (DSM). A DEM is usually provided as a raster grid, generally in a common format such as GeoTIFF, where each cell contains an elevation value. These models are generated using remote sensing technologies such as LiDAR, radar, or photogrammetry. DEMs are critical in analysing slope, aspect, and watershed boundaries in geographic studies. They also serve as foundational layers in 3D simulations, environmental modelling, and infrastructure planning.

Generating a digital elevation model can be approached through several methods, each with its own advantages and limitations. These methods vary based on data sources, processing techniques, and intended applications. Understanding these options is essential for selecting the most appropriate approach for a given project. Guth et al. [157] clarify terminology and definitions surrounding digital elevation model (DEM)s. They discuss distinctions between various DEM-related terms (digital surface versus terrain models) and explain implications of intermediate spheres (biosphere, hydrosphere, cryosphere, anthroposphere) in defining surface models. With so many different ways to produce a DEM (and sources of data), their quality can vary considerably. Polidori and El Hage [158] provide ways to assess the quality of a DEMs. They emphasise the importance of clearly defining nominal terrain surfaces, addressing errors, artifacts, quality criteria, and discuss how these methods apply across various product levels (point clouds, grid surface models, and global DEMs).

Digital elevation model (DEM)s play a critical role in understanding and representing terrain in urban planning and analysis. Various factors, such as resolution, accuracy, and data sources, influence the quality of these models. Understanding how DEMs are generated is key to applying them effectively in real-world scenarios. This includes examining techniques that utilise existing datasets to build detailed, usable models. Scalas et al. [159] discuss the creation and utility of the geometric layer within urban digital twins, focusing on the city of Matera, Italy. They present a method to generate detailed 3D city models using pre-existing public data (open street map (OSM) and aerial DEMs). To demonstrate practical applications, they describe tools for calculating route slopes and generating shadow / light maps for specific times, highlighting potential urban planning and mobility benefits without the need for expensive, custom data collection efforts.

DEMs are of interest because they are used in spectrum modelling and spectrum demand modelling, where terrain can influence patterns of human activity. da Silva et al. [160] compare digital elevation data from the ITU-R database and the shuttle radar topography mission (SRTM) to enhance broadcast coverage and frequency utilisation for digital television. They test different propagation models against measured data, showing how varying resolutions and accuracy of elevation models influence prediction accuracy. Their study highlights the importance of high-resolution terrain data, such as SRTM, for accurate broadcast signal predictions and interference analysis in the VHF and UHF bands.

### 5.3.3 Radio Frequency (RF) Ray Tracing

#### Summary

*Ray tracing of radio frequency (RF) signals has become an essential tool for accurately modelling radio propagation in complex 3D environments. Recent advancements, such as differentiable and dynamic ray tracing methods enable precise and adaptive modelling, particularly in scenarios like vehicle-to-vehicle (V2V) communication, and have been validated through real-world urban data. Integration with machine learning (ML) and simulation platforms has significantly enhanced the efficiency and realism of these models, supporting scalable applications in wireless design, indoor localisation, and digital twins. Furthermore, RF ray tracing is expanding into new domains like mobility and robotics, offering privacy-preserving solutions for high-resolution environmental sensing and autonomous navigation.*

#### Review

Radio frequency (RF) ray tracing for radio propagation modelling is a very important and useful technique for accurately modelling spectrum coverage of a 3D environment. Recent advancements in ray tracing have introduced differentiable and dynamic methods that offer new ways to model radio wave propagation with high precision and resolution. These techniques enable greater flexibility in learning from and optimising environmental parameters, making them valuable for both theoretical analysis and real-world deployment. Eertmans et al. [161] compare differentiable ray tracing (DiffRT) and dynamic ray tracing (DynRT) methods in dynamic scenarios like vehicle-to-vehicle (V2V) communication. They introduce a novel simulation-based metric called the multipath lifetime lap (MLM) that characterises spatial and temporal coherence of radio channels purely from environmental geometry, independently of the frequency used. The comparison employs state-of-the-art tools (3DSCAT and Sionna RT) to validate the method in an urban scenario, demonstrating strong agreement with real-world measurements and clarifying when DynRT or DiffRT is most suitable.

The integration of ML with ray tracing is reshaping how radio propagation modelling is approached. By using AI, these methods dramatically improve computational efficiency and open new avenues for intelligent, adaptive wireless system design. Eertmans et al. [162] propose a generative ML approach integrated into point-to-point ray tracing for radio propagation modelling. This method efficiently identifies valid ray paths among an exponentially large number of possibilities, substantially reducing computational load. Their solution, using reinforcement learning (RL), uniquely generalises across different radio frequency ranges, material properties, and scene geometries. It provides a scalable, computationally efficient alternative that maintains high accuracy without extensive pre-training or datasets.

Datasets tailored to ray tracing and machine learning (ML) integration are essential to devel-

oping realistic models for indoor wireless propagation. Zhang et al. [163] introduce WiSegRT, an extensive dataset specifically created for accurate indoor radio propagation modelling. Using detailed, segmented 3D indoor scenes simulated with differentiable ray tracing (Sionna), WiSegRT provides comprehensive channel impulse response data suitable for training advanced ML models. The dataset supports a wide range of research applications, including accurate indoor localisation, ML-based channel prediction, wireless digital twins, and radio-based object detection. It fills a critical gap in existing radio datasets by incorporating realistic environmental complexity and precise material segmentation.

Ray tracing becomes even more powerful when integrated with broader simulation ecosystems. Combining physical accuracy with network-level simulation tools like ns-3 enhances realism and enables more comprehensive performance evaluations of wireless systems. Zubow et al. [164] introduce Ns3Sionna, integrating the Sionna ray-tracing framework into the ns-3 network simulator to overcome limitations of traditional stochastic channel models. Ns3Sionna provides accurate and physically realistic channel modelling for indoor and outdoor scenarios, offering spatially consistent channel impulse responses. To overcome ray tracing's computational intensity, the system utilises GPU acceleration and intelligent pre-caching strategies, enhancing efficiency for scenarios involving small to medium numbers of mobile nodes.

Digital twins are increasingly used to mirror real-world wireless environments with high fidelity. Ray tracing plays a central role in these platforms, enabling accurate and dynamic replication of complex urban and indoor settings for planning, testing, and optimisation. Noh and Choi [165] propose a high-precision digital twin platform capable of replicating real-world wireless network environments through accurate 3D modelling and ray tracing simulations. Specifically modelling the Suwon Station area in South Korea, the platform integrates microscopic traffic simulation with differentiable ray tracing (using Sionna RT) to produce extensive, high-quality channel datasets. This approach offers a cost-effective, realistic method for network optimisation and experimentation without disrupting real-world service quality.

A critical part of building digital twins for RF simulation involves choosing appropriate ray-launching tools that can handle complex electromagnetic modelling with high fidelity. Extending this idea, Zhu et al. [166] evaluate a range of such software, highlighting trade-offs between accuracy and computational performance in realistic electromagnetic environments. The comparative study evaluates software across increasing levels of complexity (reflections, diffraction, diffuse scattering), highlighting differences in performance, particularly in urban settings. This work serves as an empirical guide for selecting appropriate ray-launching software for digital twin applications, particularly when real-time updates are critical.

Ray tracing is finding novel applications beyond traditional wireless communication, especially in mobility and robotics. From autonomous navigation to mobile network datasets,

these use cases highlight the technique’s versatility and its potential for privacy-preserving, high-resolution environmental sensing. Bastos et al. [167] expand the Raymobtime methodology, previously based on the commercial Wireless InSite software, to integrate with NVIDIA’s open-source Sionna ray tracer. Raymobtime notably supports sophisticated mobility simulations, allowing both transceivers and scattering objects to move. This integration allows efficient generation of datasets suitable for advanced machine learning in 5G and 6G scenarios. Initial results demonstrate Sionna’s performance advantages, although further validation against measured data is necessary.

Privacy-preserving alternatives to visual navigation have emerged through the integration of RF-based sensing and ray tracing technologies. Amatare et al. [168] present a novel testbed integrating differential ray tracing (DRT) using NVIDIA’s Sionna tool for autonomous robot navigation, aiming to address privacy concerns inherent in sensor-based systems. Employing LiDAR-generated data integrated with RF propagation modelling, the testbed facilitates precise navigation in complex indoor environments without exposing sensitive visual data. The proposed approach promises robust, privacy-preserving navigation by using the detailed environmental sensitivity analysis capabilities of DRT.

## 5.4 Background

### 5.4.1 Sionna Differentiable RT

Sionna Differentiable RT is a specialised ray tracing engine for modelling wireless signal propagation with the unique capability of end-to-end differentiability. As part of the open-source Sionna library for communication system simulations, it supports both GPUs for high-performance computation and CPUs when GPU resources are unavailable. Its core purpose is to simulate electromagnetic wave behaviour in complex 3D environments, enabling accurate prediction of channel characteristics in realistic wireless scenarios.

Unlike conventional ray tracers, Sionna RT integrates a physically based renderer (Mitsuba 3) with the Dr.Jit just-in-time compiler, allowing it to trace multipath propagation while computing gradients through the entire simulation pipeline. This makes outputs such as received signal strength (RSS), path delays and frequency responses differentiable with respect to virtually any scene or system parameter. As a result, it supports advanced tasks such as gradient-based optimisation and sensitivity analysis directly on physical channel models.

Sionna RT uses geometry-based ray tracing algorithms to find propagation paths between transmitters and receivers, accounting for effects such as line of sight (LOS) transmission, multi-bounce reflections, diffuse scattering and, in some versions, diffraction. Its realistic material and antenna models incorporate factors such as frequency-dependent material permittivity, antenna radiation patterns and free-space path loss. By summing contributions from all identified paths, the engine produces complete channel impulse or frequency responses for a link.

The differentiable nature of Sionna RT means that it can determine how modifications to any aspect of a wireless scenario, for example a building's material, an antenna's alignment or a receiver's location, would alter the resulting channel response. This capability enables researchers and engineers to directly optimise communication system designs, fine-tune deployment parameters and conduct detailed sensitivity studies, all grounded in physically accurate electromagnetic simulations [169].

### 5.4.2 Radio Materials in Mitsuba 3

A *radio material* is a concept used to categorise and simplify how transmission environments are represented in digital simulations. It contains only the key properties needed to model how RF signals behave in a realistic setting, such as when they are reflected, refracted or diffracted. By capturing the essential electromagnetic characteristics of a surface, the modelling of interactions between RF transmissions and radio materials can be performed with a high degree of accuracy, even when only basic geometric and textural information about objects is available.

In Sionna RT, a radio material is implemented as a class that defines the electromagnetic properties relevant to radio propagation simulations [170]. It is characterised primarily by its relative permittivity ( $\epsilon_r$ ) and conductivity ( $\sigma$ ), which govern how electromagnetic waves are stored, conducted and attenuated in the material. The class may also include parameters for diffuse scattering, such as the scattering coefficient ( $S$ ), cross-polarisation discrimination coefficient ( $K_x$ ) and a scattering pattern function. Together, these parameters enable realistic modelling of reflections, transmissions and scattering within a simulated environment.

- **Relative Permittivity** ( $\epsilon_r$ ): Also known as the dielectric constant, it describes a material's ability to store electrical energy in an electric field.
- **Conductivity** ( $\sigma$ ): Indicates how effectively a material can conduct electric current, influencing how much electromagnetic energy is absorbed during transmission through the material.

These electromagnetic parameters directly determine the reflection, transmission and absorption behaviour of materials when exposed to RF signals. In ray tracing methods for wireless communications, they are used to calculate reflection and diffraction losses with precision [171, 169]. By accurately capturing these physical interactions, Sionna RT supports the realistic simulation and analysis of wireless communication systems in complex environments.



## 5.5 Demonstrations

### 5.5.1 Context

This section explores the problem of modelling spectrum use in environments where traditional allocation processes are insufficient. Urban density, varied terrain, and the rise of complex communication systems mean that simplistic line of sight (LOS) or statistical approaches alone cannot capture the reality of radio propagation. The earlier discussion in the approach section highlighted that improved methods are required to handle reflections, material properties, and vertical dimensions of cities, as well as rural landscapes shaped by terrain.

Research in the literature shows how different modelling techniques contribute to addressing these gaps. Three-dimensional (3D) city models improve the representation of urban RF propagation environments, digital elevation model (DEM)s provide terrain-based accuracy for non-urban areas, and ray tracing methods allow detailed simulation of radio waves interacting with the transmission environment. Recent developments have pushed these methods toward more scalable and automated applications, reinforcing their importance for communication system design and spectrum management.

The demonstrations presented here take these concepts and place them into applied settings. They are not intended as full solutions, but as practical tests that show how the reviewed methods behave when combined with real datasets and simulation tools. This creates a direct link between the challenges identified, the supporting research, and the experimental work used to assess feasibility and performance in context.

### 5.5.2 Summary

Demonstration 1 shows off line of sight (LOS) mapping using elevation data. LOS is a valuable tool for RF modelling, especially for low frequencies and rural environments, as lower frequencies travel farther and pass through objects more effectively, while rural landscapes with varied elevation can greatly influence signal reach. It can offer critical insights in spectrum management by identifying potential interference risks between regions. Advantages include the availability of free, high-resolution elevation data in accessible formats, efficient LOS calculation using raw data, and the ability to reduce computation through downsampling. However, limitations arise from elevation data inaccuracies, reduced usefulness in dense urban or high-frequency scenarios, and the lack of building and vegetation data, which makes 3D modelling preferable in such cases. Future improvements could involve integrating multiple elevation data sources for better resolution, incorporating building data into LOS paths, and verifying elevation data accuracy.

Demonstration 2 shows off ray tracing buildings using OSM Data. Sionna RT combined with the wide coverage of OSM map data offers an effective way to simulate RF propagation in most

mapped urban areas globally by producing ray tracing-based coverage maps through finding the lowest energy path to the ground plane. The accuracy depends on factors such as sampling density, number of reflections, support for propagation effects, and the quality of 3D city models, which are limited by OSM data detail and the assumption of flat ground during processing. Using open-source data enables application in many locations, and the flat-ground approach simplifies experiments, with negligible curvature impact for small urban areas. However, the approach is constrained by incomplete propagation effect modeling, limited material property information, and oversimplified building interiors, often assuming thick concrete. Future improvements may integrate elevation data and Earth curvature for more realistic city models, supported by newer Sionna RT versions that allow 3D mesh measurement surfaces, with real-world validation potentially confirming suitability for spectrum management applications.

Demonstration 3 shows off ray tracing buildings and terrain from Google 3D tiles. Accurate 3D city models can be created using LiDAR data instead of mapping sources like OSM, offering greater detail by capturing varied obstacles such as trees, hills and precise building shapes. When combined with satellite imagery, these models are easier to verify visually, and Google's 3D tiles are a strong example of this capability. The high level of detail in the building, road, park and tree models allows for realistic RF ray tracing and enables coverage mapping overlays that help detect errors like unrealistic signal paths. However, the data is subject to Google API limitations, can be too large to download for extensive areas, is available mainly for major cities and cannot be accessed historically unless previously stored. Processing it is also challenging due to coordinate system differences and the computational load. Despite these constraints, Google's 3D tiles remain the most impressive (in terms of detail and quality) freely available resource for research, providing a benchmark for creating globally scalable, open-source equivalents that could also incorporate historical, updatable and animated scene data.

Demonstration 4 shows off ray tracing 3D meshes from elevation data. In rural RF propagation scenarios, lower frequencies and line of sight (LOS) modelling are often used due to longer propagation distances, but LOS alone fails to account for losses, antenna properties, and other key propagation effects. Ray tracing offers a more comprehensive approach by simulating signal paths as individual rays to capture reflections, refraction, diffraction, and interactions with terrain or vegetation, which can be critical in town centres, hilly landscapes, and forested areas. Using elevation data to create 3D meshes for ray tracing is effective since such data is widely available and can be processed efficiently, even over large rural areas, with the option to downsample for performance and combine with 2D maps for better visualisation. However, limitations include assumptions made between sampled data points, potential inaccuracies from low-resolution data, and the lack of material property information. Future improvements could involve integrating building data to form hybrid urban-rural models and deriving mesh material properties from satellite imagery or machine learning (ML) to improve simulation accuracy.



### 5.5.3 Tools

To facilitate the modelling and visualisation of the realistic scenes with RF propagation coverage maps, additional tools and software were developed specifically for this research. Specifically, the ability to render 3D heatmaps onto the surface of detailed 3D and realistic scenes was a big focus of these tools. By having access to such an intuitive visualisation tool, it became much more feasible to test and explore new spectrum management concepts. The 3D display engine was built in Flutter based on the project by Derocher [172]. In addition to this tool, a python program was developed to produce 3D meshes from digital elevation model (DEM)s, which could then be processed into line of sight (LOS) maps using a viewshed [173] and ray traced using projects like Sionna RT [174]. The Google 3D tiles [152] were processed into a format usable in Sionna RT using the Blender addon Blossm [175].

### 5.5.4 Structure

This section presents four main demonstrations:

1. **LOS Mapping from Elevation Data** — National and state-based elevation data used for visualising LOS coverage at large scales.
2. **Ray Tracing Buildings from OSM Data** — RF path loss simulation using building footprint extrusions.
3. **Ray Tracing Buildings and Terrain from Google 3D Tiles** — High-fidelity modelling using detailed 3D city and stadium models.
4. **Ray Tracing 3D Meshes from Elevation Data** — Terrain-only ray tracing for regional and complex landscapes.

### 5.5.5 Demonstration 1: LOS Mapping from Elevation Data

Table 5.1 summarises the structure of Demonstration 1, outlining the focus on LOS mapping from elevation data. It states objectives that progress from urban to regional and rural contexts, specifies elevation-based visibility as the core method with increasing resolution, and records the elevation grids used as inputs, including SRTM 30m and a VicGrid 10m DEM.

Demonstration	Objective	Key Method	Data Inputs
1.1	Baseline LOS mapping for urban	Elevation-based LOS visibility	30 m SRTM 30m elevation grid
1.2	Regional LOS mapping	Elevation-based LOS visibility	30 m SRTM 30m elevation grid
1.3	Rural LOS mapping	Elevation-based LOS visibility	10 m VicGrid DEM

Table 5.1: Summary of Demonstration 1 – LOS Mapping from Elevation Data

## Demo 1.1: LOS Mapping for Adelaide

### Setup:

- Omnidirectional transmit antenna
- Transmit antenna height: 30 m above ground
- Receiver height: 1.5 m above ground (handheld)
- LOS radius: 50 km
- Elevation data: 30 m resolution SRTM 30m DEM for Australia [176]

The red areas in Figure 5.5.5.1 show LOS for a radius of 50 km from an omnidirectional antenna 30 m above ground, located north-east of the Adelaide CBD. Receivers are assumed to be 1.5 m above ground (typical for handheld devices). The mapping uses SRTM 30m elevation data for Australia [176]. While higher resolution LiDAR data exists locally, its limited coverage reduces applicability for national-scale modelling. This baseline model demonstrates feasibility for use in large-area spectrum management.



Figure 5.5.5.1: LOS Mapping for Adelaide

## Demo 1.2: LOS Mapping for Uluru

### Setup:

- Omnidirectional transmit antenna
- Transmit antenna height: 30 m above ground
- Receiver height: 1.5 m above ground (handheld)
- LOS radius: 50 km
- Elevation data: 30 m resolution SRTM 30m DEM for Australia [176]

Figure 5.5.5.2 shows LOS (light blue) for 50 km from an omnidirectional antenna 30 m high, east of Uluru. Using the same national SRTM 30m dataset, this example illustrates applicability for remote regions. While obvious that Uluru itself blocks LOS, the result visually supports that the model may be accurate, by looking at a simple but distinctive topography.




Figure 5.5.5.2: LOS Mapping for Uluru



### Demo 1.3: LOS Mapping for The Grampians

**Setup:**

- Omnidirectional transmit antenna
- Transmit antenna height: 30 m above ground
- Receiver height: 1.5 m above ground (handheld)
- LOS radius: 50 km
- Elevation data: 10 m resolution VicGrid DEM for Australia [176]

In Figure 5.5.5.3, light magenta areas represent LOS coverage from a 30 m antenna between Halls Gap and Stawell. Higher-resolution (10 m) VicGrid elevation data is used. The rugged terrain with sharp elevation changes dominates RF propagation characteristics (LOS only in this case).




Figure 5.5.5.3: LOS Mapping for the Grampians

### 5.5.6 Demonstration 2: Ray Tracing Buildings from OSM Data

Table 5.2 summarises the structure of Demonstration 2, describing building-aware ray tracing driven by OSM-derived models. It states urban objectives for both full 3D and a 2D view, notes RF ray tracing with reflections as the key method with the second view matching the first, and lists the OSM-based 3D city model as the input.

Demonstration	Objective	Key Method	Data Inputs
2.1	Urban ray tracing 3D	RF ray tracing with reflections	OSM-derived 3D city model
2.2	Urban ray tracing 2D view	RF ray tracing with reflections	OSM-derived 3D city model

Table 5.2: Summary of Demonstration 2 – Ray Tracing Buildings from OSM Data



## Demo 2.1: 3D Ray Tracing in Adelaide CBD

### Setup:

- Transmitter: TR38901 ground antenna pattern, 3.5 GHz / 100 MHz BW facing east
- Data: 3D models derived from OSM [117]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.0 [174]

Figure 5.5.6.1 shows RF path loss for rays from a ground transmitter in central Adelaide. Buildings are modelled as solid concrete, capturing LOS and reflection paths. While LOS remains the dominant part of the coverage, coverage around corners without LOS available is mostly or entirely from reflections due to the high attenuation of the concrete buildings. Yellow shows lower path loss (dB) while the green shows higher path loss (dB).




Figure 5.5.6.1: Adelaide OSM Ray Tracing 3D

## Demo 2.2: 2D Ray Tracing View from Victoria Square

### Setup:

- Transmitter: TR38901 ground antenna pattern, 3.5 GHz / 100 MHz BW facing north
- Data: 3D models derived from OSM [117]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.0 [174]

Figure 5.5.6.2 models RF rays from a 70 m tall transmitter at Victoria Square, facing north as shown by the antenna icon on the map. Coverage extends into North Adelaide with clear LOS in open areas and reflection-based paths into obstructed zones. Again, yellow shows lower path loss (dB) while the green shows higher path loss (dB).




Figure 5.5.6.2: RF ray tracing coverage map based on OSM data for a north-facing antenna at a height of 70m located at Victoria Square, Adelaide



### 5.5.7 Demonstration 3: Ray Tracing Buildings and Terrain from Google 3D Tiles

Table 5.3 summarises the structure of Demonstration 3, focusing on ray tracing with Google 3D Tiles for both buildings and terrain. It sets out objectives that cover urban streetscape RF coverage, stadium RF coverage, and a broader city region RF coverage view. The table notes high-detail mesh ray tracing as the consistent key method applied across all cases, and records the Google 3D Tiles model as the primary data input, with the stadium case explicitly including stadium geometry.

Demonstration	Objective	Key Method	Data Inputs
3.1	Urban streetscape RF coverage	High-detail mesh ray tracing	Google 3D Tiles model
3.2	Stadium RF coverage	High-detail mesh ray tracing	Google 3D Tiles model (stadium geometry)
3.3	City region RF coverage view	High-detail mesh ray tracing	Google 3D Tiles

Table 5.3: Summary of Demonstration 3 – Ray Tracing Buildings and Terrain from Google 3D Tiles

### Demo 3.1: Victoria Square Streetscape Coverage

#### Setup:

- Transmitter: TR38901 ground antenna pattern, 3.5 GHz / 100 MHz BW facing north
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.7.1 shows a 2D and 3D view of RF ray tracing coverage from a Victoria Square transmitter facing north. Strong LOS coverage is seen in open spaces, while narrow streets rely on reflected paths due to building blockage.



Figure 5.5.7.1: (left) 2D bird's-eye view of RF ray traced map of Google 3D tiles data using Sionna RT for a north-facing transmitter located at Victoria Square; (right) Same heatmap overlaid onto the Google 3D tiles data



### Demo 3.2: Adelaide Oval Event Scenario

#### Setup:

- Transmitter: TR38901 ground antenna pattern, 3.5 GHz / 100 MHz BW facing north
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.7.2 illustrates how reflections and gaps in the stadium structure enable coverage into non-LOS areas. This level of detail in the RF propagation modelling could be critical for managing capacity during large events.




Figure 5.5.7.2: (left) 2D bird's-eye view of RF ray traced map of Google 3D tiles data using Sionna RT for a north-facing transmitter located at the Adelaide Oval; (right) Same heatmap overlaid onto the Google 3D tiles data

### Demo 3.3: Regional Coverage from Elevated Transmitter

#### Setup:

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.7.3 shows high-detail mesh ray tracing of RF coverage from a 100 m transmitter. The results highlight how tall buildings shape where the RF coverage is strongest, and where their a dead-zones.




Figure 5.5.7.3: (top) Angled view of Adelaide CBD for reference; (bottom) RF ray traced coverage heatmap using Google 3D tiles and Sionna RT of the same region for an omnidirectional transmitter placed at Victoria Square, Adelaide

### 5.5.8 Demonstration 4: Ray Tracing 3D Meshes from Elevation Data

Table 5.4 summarises the structure of Demonstration 4, presenting terrain-focused ray tracing with meshes derived from elevation data. It states objectives spanning omni and directional antenna use through to obstacle studies and complex mountainous settings, specifies the corresponding ray tracing approaches, and records a common input of DEM-derived 3D meshes across the set.

Demonstration	Objective	Key Method	Data Inputs
4.1	Terrain ray tracing omni	Omni antenna ray tracing	DEM-derived 3D mesh
4.2	Terrain ray tracing directional	Directional antenna with terrain	DEM-derived 3D mesh
4.3	Iconic terrain obstacle study	RF shadowing	DEM-derived 3D mesh
4.4	Complex mountainous terrain	RF coverage in valleys and peaks	DEM-derived 3D mesh

Table 5.4: Summary of Demonstration 4 – Ray Tracing 3D Meshes from Elevation Data

## Demo 4.1: Omnidirectional Terrain Ray Tracing (Adelaide)

### Setup:

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m [152]
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.8.1 shows coverage from an omnidirectional antenna over 30 m DEM terrain mesh. Minimal difference from LOS mapping is observed in the plains; reflections contribute more in hilly areas.




Figure 5.5.8.1: Omnidirectional ray tracing for Adelaide region



## Demo 4.2: Directional Antenna into Hillside

### Setup:

- Transmitter: TR38901 ground antenna pattern, 3.5 GHz / 100 MHz BW facing east
- Data: 3D models derived from SRTM 30m [152]
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.8.2 demonstrates beam shaping with the antenna aimed east toward hills. Back-lobe leakage is visible but at a low power and possibly some reflections back off of the hills too.




Figure 5.5.8.2: Directional ray tracing for Adelaide region towards hillside

### Demo 4.3: Uluru RF Shadowing

#### Setup:

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m [152]
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.8.3 confirms expected blockage behind Uluru and coverage in surrounding flat areas. This is an expected result because the antenna-facing side of the Iconic Uluru landmark (which gets excellent coverage) would provide a significant RF propagation barrier to the areas behind it.




Figure 5.5.8.3: RF Ray tracing for Uluru using Sionna RT for a 3D mesh produced from DEM data

## Demo 4.4: Complex Mountain Terrain (The Grampians)

### Setup:

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m [152]
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 5.5.8.4 shows significant propagation challenges due to rugged peaks and valleys, where LOS is frequently blocked.




Figure 5.5.8.4: Ray tracing for the Grampians using Sionna RT for a 3D mesh produced from DEM data

## 5.6 Discussion

### 5.6.1 Summary

Chapter 5 has five sections: *Introduction*, *Approach*, *Literature*, *Background*, and *Demonstrations*.

The *Introduction* states that effective 3D modelling of spectrum usage is needed to assess allocation impacts within a dynamic spectrum management (DSM). Large urban and rural environments are both considered, with emphasis placed on coverage, demand, and vertical user movement. Reliable radio frequency (RF) planning at city and regional scales is defined as the aim.

The *Approach* proposes combining geometry and terrain with physically based propagation to model usage. Options for modelling radio spectrum in a DSM across large urban and rural environments are examined. Focus is placed on 3D City Modelling for user representation and verticality, digital elevation model (DEM) for terrain influence, and ray tracing for coverage-informed allocation.

The *Literature* explores the three focus areas. 3D city modelling is described as evolving toward accessible, geographic information system (GIS)-integrated pipelines, with evidence from tools such as OpenGERT showing that channel statistics are highly influenced by geometric fidelity. Definitions, quality assessment, and practical construction of DEMs from public data are clarified, with broadcast studies showing that higher resolution terrain leads to improved prediction accuracy. ray tracing research is reported to have advanced, due to differentiable and dynamic methods, machine learning (ML) integration, purpose-built datasets, and its integration with network simulators and digital twins. The need for precise geometry and validated terrain to achieve dependable predictions is made clear.

The *Demonstrations* progress from line of sight (LOS) mapping with national and state elevation grids to RT on building extrusions from open street map (OSM), then to high detail Google 3D Tiles, and finally to terrain-only RT on DEM-derived meshes. Scenarios span urban streetscapes, stadium events, regional coverage, and iconic landforms including Adelaide, Uluru, and the Grampians. Findings show that data resolution, geometric and material realism, sampling density, and supported propagation effects govern outcomes. Limitations include incomplete building and vegetation data, flat-ground assumptions, and platform constraints, with future work on integrating elevation and curvature, richer material models, and measurement-based validation.

## 5.6.2 Key Findings

### Finding 1: Integration of Terrain and Spectrum Modelling

Effective spectrum modelling requires explicit consideration of topography. The *Approach* section frames the problem by identifying digital elevation model (DEM)s as critical for capturing terrain influences on signal propagation in rural and hilly environments. The *Literature* review confirms this by showing that DEM quality, resolution, and processing methods directly affect prediction accuracy, with studies highlighting how high-resolution terrain data improves broadcast and interference modelling. The *Demonstrations* provide show its applicability through line of sight (LOS) mapping around Adelaide, Uluru, and the Grampians, where elevation data was used to reveal shadowing and coverage variation. The implication is that terrain-based DEM integration could be highly useful for dynamic spectrum management (DSM) because their geographical coverage is vast, the information provided is very informative for spectrum planning, and even in urban settings it can be quite critical when there are significant changes in elevation.

### Finding 2: Accuracy Gains from 3D City Models

Accurate modelling of urban RF propagation depends on fine-grained representation of the environment. The *Approach* section frames the problem by pointing to 3D city modelling as a means to capture both user movement and spectrum demand in dense urban settings. The *Literature* supports this with examples of GIS-based modelling and open-source tools, which show that even small variations in building height or placement have measurable impacts on signal quality. The *Demonstrations* explore these concepts through ray tracing exercises in Adelaide, where OSM- and Google 3D-based models revealed reflection paths, stadium coverage, and shadowing from tall buildings. The implication is that it is possible to produce scalable, precise 3D urban models. They could be very useful for making spectrum allocation decisions.

### Finding 3: Ray Tracing could be critical for DSM

Ray tracing provides a robust framework for capturing propagation effects that simpler methods omit. The *Approach* frames the problem by identifying ray tracing as a technique for modelling allocation decisions through reflection, refraction, and diffraction paths. The *Literature* expands on this by showing how advanced ray tracing methods, particularly when integrated with machine learning (ML), deliver efficient and accurate results across dynamic environments. The *Demonstrations* reinforce this with applications using OSM data, Google 3D tiles, and DEM-based meshes, all of which show ray tracing's ability to identify coverage gaps and reflection-based pathways in both urban and rural contexts. The implication is that there is a strong case for ray tracing to be the baseline predictive tool for next-generation spectrum management.

### 5.6.3 Limitations

While DEMs enable explicit consideration of terrain in spectrum modelling, their reliability is constrained by data resolution and coverage. National-scale DEMs, such as WGS84, provide broad applicability but introduce inaccuracies when applied to local-scale analysis where elevation changes are critical. Higher-resolution DEMs, such as VicGrid, improve predictive accuracy but lack consistent availability across regions, creating uneven modelling fidelity. Moreover, DEMs generally exclude vegetation and built structures, meaning that line of sight (LOS) and shadowing effects are represented incompletely in mixed urban-rural environments. The dependence on data from varied sources also introduces inconsistencies in coordinate systems and post-processing methods, further limiting direct integration into DSM, without significant efforts required to provide individual support.

Fine-grained 3D city models enhance the accuracy of RF propagation modelling in dense urban settings, but their applicability is constrained by data quality and scalability. Open-source resources such as OSM offer wide coverage but lack detail in building heights, interior structures, and material properties, resulting in simplified assumptions such as treating all walls as thick concrete. Even high-fidelity data, such as Google 3D Tiles, remain limited to major cities and are subject to licensing restrictions, large data sizes, and computational challenges during processing. Additionally, coordinate mismatches between datasets can reduce accuracy, while the absence of temporal updates restricts the capacity to represent evolving urban environments. These constraints hinder the practicality of maintaining consistent, scalable city models for long-term spectrum allocation. In other words, many of the methods shown in demonstrations are excellent for producing research demonstrations, but would require further work to be applied at scale in an industrial context.

Ray tracing provides high-fidelity modelling of reflection, diffraction, and scattering effects but is computationally intensive, limiting its application at national or regional scales. Despite GPU acceleration and machine learning (ML) integration reducing processing times, trade-offs between accuracy and performance remain, especially when modelling large numbers of mobile nodes or wide-area networks. During demonstrations, numerous optimisations allowed the calculations to be made impressively fast, but in general, the calculations took a few minutes at least depending on the resolution of the result. Ray tracing also depends heavily on input data quality. Incomplete representation of materials, vegetation, or building interiors reduces predictive reliability. Furthermore, current implementations often require simplifying assumptions, such as flat ground or uniform material categories that reduce realism. While ray tracing is powerful for scenario testing, its heavy reliance on data availability and processing capacity could prevent it from being a universally scalable solution for DSM. However, it could be argued that the results were impressive enough to warrant further research on how to achieve the same results at a higher level of efficiency.

#### 5.6.4 Future Work

Further research could advance the use of DEMs by incorporating higher-resolution and multi-source datasets to reduce inaccuracies in terrain representation. Vegetation and built structures, often absent from current elevation models, could be derived from remote sensing or satellite imagery and merged into terrain meshes for more complete propagation environments. Alternative DEM generation and post-processing techniques could also be explored to harmonise coordinate systems and improve interoperability across datasets. These refinements would allow spectrum modelling frameworks to capture complex shadowing and blockage effects in rural and mixed environments with greater reliability.

The development of richer 3D city models would require both broader coverage and deeper fidelity. Open-source data such as OSM offers scale but lacks detail, while high-fidelity options like Google 3D Tiles remain restricted in scope and licensing. Future work could therefore look to hybrid solutions that combine multiple sources, automated building height extraction, and ML methods to infer material properties. Temporal updating and lightweight compression techniques would also be critical to scale such models across evolving urban landscapes without prohibitive storage or computational demands. These improvements could enable more practical use of detailed city models in large-scale spectrum management.

Ray tracing could progress toward more customised and computationally efficient forms. GPU-optimised methods, dynamic ray tracing, and differentiable approaches promise significant performance gains, but continued research would be needed to balance accuracy with runtime, particularly when modelling wide-area or mobile scenarios. Integrating richer radio material models, vegetation, and realistic interiors could increase predictive realism, while integrating network-level simulators would expand its utility in planning and optimisation tasks. Exploration of adaptive sampling, AI-assisted path selection, and realistic digital twin environments could be key to scaling ray tracing into operational spectrum management systems while maintaining a high level of accuracy.

## 5.7 Next Steps

This chapter demonstrates a range of RF propagation modelling techniques. By synthesising this with earlier work on reachability analysis and dynamic spectrum management (DSM) definitions, a more cohesive problem statement can be formed. The next chapter explores a spectrum management digital twin and shows its application to coordinating spectrum access. The central problem is to allocate scarce spectrum across moving ground and aerial users in environments shaped by geography, infrastructure and altitude while meeting regulatory, safety and privacy constraints. Objectives are to anticipate interference in space and time, quantify risk, and coordinate access that preserves service quality and fairness within limited resources. The general approach is to fuse mobility envelopes, environmental context and radio propagation within a digital twin that evaluates feasible user states and allocation options subject to geographic, temporal and policy constraints.



## Chapter 6

# Coordinated Spectrum Access via Digital Twin Simulation

*This chapter introduces the shift from conventional two-dimensional latitude–longitude allocation to advanced three-dimensional (3D) and vertical spectrum management using detailed digital twin modelling of the radio environment. It begins by explaining digital spaces, partial and complete digital twins, and their value in representing real-world environments with high accuracy for coordinated spectrum access. Drones’ spectrum use is examined, noting current reliance on shared industrial, scientific and medical (ISM) bands, regulatory restrictions on licensed spectrum at altitude, and the resulting limitations. The chapter then applies both idealised and realistic models to demonstrate their differing effects, where idealised approaches such as free-space RF propagation and Brownian motion reachability are compared with realistic models incorporating shortest path tree (SPT)s, road network constraints, RF ray tracing, and terrain data to show how environmental and mobility factors alter predicted interference. The demonstrations progress from isolated reachability and RF propagation in urban and rural settings to combined spatiotemporal heatmaps and then to complex multi-user and drone-inclusive scenarios. These include urban and rural driving cases, multi-drone mutual interference mapping, drone flight path coverage summations, and mixed drone–ground coordination with time-resolved signal-to-interference ratio analysis. The results show how three-dimensional mobility-aware modelling in a vertical spectrum context informed by realistic environmental data delivers more accurate, explainable, and computationally viable interference assessments for efficient and resilient spectrum allocation.*

## 6.1 Introduction

The aim of this chapter is to explore new coordinated spectrum access opportunities enabled by the development of a *digital twin* purpose-built for spectrum management. With the capability to simulate protocols and algorithms in advance, a *digital twin* provides a realistic geospatial environment that accurately represents real-world geography, infrastructure, and RF conditions, enabling informed RF planning and allocation decisions before implementation. While the previous chapters built up concepts such as reachability, 3D modelling techniques, RF propagation modelling, and decentralisation, this chapter combines them into a single system, producing what could be referred to as a *spectrum management digital twin*.

A key part of making this coordination effective is combining reachability modelling with RF propagation analysis. Reachability models such as shortest path tree (SPT)s can represent the areas a user could physically move to within a given time, taking into account the real constraints of the road network, terrain, and traffic conditions. RF propagation models, whether simple free-space calculations or detailed ray tracing in complex environments, predict how far and in what directions a transmission could travel. When these two are combined, it becomes possible to calculate specific time-location regions where two or more users could interfere with each other. These regions can be quantified as a measurable risk, enabling spectrum access to be scheduled or adjusted before any harmful interference occurs.

The same principles extend to situations involving multiple spectrum users in shared or adjacent areas. Each user's reachability and coverage footprint can be calculated over time and compared to others to determine when interference is likely. A useful metric for this is the mutual worst signal-to-interference ratio, which measures where all users have sufficient signal strength and can therefore impact each other's operation. In urban environments with dense infrastructure, ray-traced propagation can reveal detailed patterns of interference that simple models cannot, while in rural terrain, elevation and land cover become the dominant factors. This approach allows for more efficient and adaptable allocation strategies that balance service quality with interference control.

Incorporating vertical space into the modelling adds further capability, especially for scenarios involving airborne users. Current spectrum allocation rarely distinguishes between different heights over the same location, yet applications such as drone-to-drone and drone-to-ground communication could benefit significantly from height-aware coordination. By treating planned drone flight paths as moving reachability profiles and combining them with propagation modelling, interference risks can be evaluated well in advance. This enables dynamic allocation that accounts for altitude, building heights, and terrain elevation, allowing multiple aerial and ground systems to operate in the same general area without degrading each other's performance. This vertical perspective becomes increasingly important in modern RF networks



where technologies like beamforming and multiple input multiple output (MIMO) can direct energy precisely, but still require accurate knowledge of where users are in three-dimensional space.

## 6.2 Approach

### 6.2.1 Proposition

- a) Conventional spectrum management allocates in 2D space (latitude and longitude).
- b) Having access to advanced 3D radio spectrum modelling has applications to *deconflict* allocations ahead of time.

### 6.2.2 Questions

- A) How can the possibilities of what an allocated user may do with their access to radio spectrum be considered to *deconflict* allocations to multiple users for an allocation time interval?
- B) How can 3D radio spectrum modelling be applied to manage and coordinate spectrum access for drone-to-drone and drone-to-ground communications?

### 6.2.3 Topics

**Digital Spaces, Partial and Complete Digital Twins** - An important concept because an incredibly in-depth and detailed understanding of the environment the spectrum access is managed in is required to achieve highly coordinated spectrum access.

**Drones' Spectrum Access** - This is relevant to understanding the current ways that drones access spectrum so that a realistic scenario can be produced to implement DSM to provide them with new ways to access spectrum.

## 6.3 Literature

### 6.3.1 Digital Spaces, Partial and Complete Digital Twins

A digital space represents the real world in a computationally efficient and scalable way offering intuitive visualisations and interactivity for users. A digital twin is a related concept that refers to a highly detailed and sophisticated representation of an entity in real-time that is suitable for interactivity with a human user or machine. The concept of a digital twin has become increasingly popularised across society and serves as an example of what technology may look like in the near future. A recent example of this is its advertisement as a prime application for the improvements offered by fifth generation (5G) and sixth generation (6G) networks which promise to provide the higher bandwidths and lower latencies required to achieve digital twins [93, 177]. A fundamental motivation behind digital twins is to allow humans to interact with entities in a virtual environment across great distances in extended reality (XR). For example, this could involve an expert surgeon performing an operation from abroad or an aircraft technician performing maintenance for multiple airports in the region from a single location. Companies such as Meta and Apple have even showcased *avatars* or *personas* respectively that allow users to interact with others and play games such as chess in a more realistic environment [178].

Many of the applications of digital twins in simulation, artificial intelligence (AI), automation and robotics are gaining recognition in the media and in research. Companies like NVIDIA have demonstrated a digital space named Isaac Sim and Isaac Lab [179] where robots can train in a very highly detailed digital twin of a real-world location such as a factory where they could be deployed as shown by Mittal et al. [180]. This application of digital twins does not require human interaction or the use of an XR headset. The representation of the digital space can also be simplified in some way such as a 2D birds-eye view to be used on a 2D screen. This sort of simplification is known as a partial digital twin or a digital shadow. For example, a digital shadow for a city is often represented in 2D with well-known examples such as Google or Apple Maps, providing useful real-time traffic updates on a map that partially represents reality. In telecommunications this is arguably the direction that network planning and design should take, as demonstrated by Borges et al. [177], Testolina et al. [181], Hoydis et al. [182].

The aim of this chapter is to provide both a framework and example for how to make a useful digital twin or digital shadow for simulating telecommunications radio frequency spectrum access between multiple users in a shared space. Prior to building a sophisticated digital space and applying techniques such as AI and machine learning (ML), it is essential to understand how to model both the telecommunications network and its operating environment. Some of the main aspects focussed on are physical separation, natural geography, urban landscapes, frequency characteristics, life patterns, independent frequency usage and network / hive usage, event management and disaster response. For a model to be highly functional and useful, it

needs to be scalable, modular, and computationally efficient to be used quickly and flexibly. A deep understanding of telecommunications theory, modelling, city simulation, and population life patterns is required.

Accurate object localisation is essential for various applications, particularly in environments where traditional vision-based methods face limitations. Recent advancements explore alternative sensing techniques, such as radio frequency (RF) modelling, to enhance detection capabilities in complex or privacy-sensitive settings. Amatare et al. [183] propose a novel approach for real-time object localisation by leveraging RF propagation modelling within digital twins, avoiding the dependency on extensive ML datasets. They utilise Blender and NVIDIA's Sionna RT for generating realistic digital environments, applying ray-tracing techniques to simulate RF signals interacting with objects. By analysing how signals are obstructed or reflected by different objects, the system identifies their locations and shapes effectively, including objects obscured from direct line-of-sight views. Demonstrations conducted in two indoor lab scenarios validated the capability of accurately detecting and locating objects, demonstrating the method's practical applicability, especially in settings where visual-based methods have limitations or privacy concerns.

Looking at a similar topic as the previous paper on object localisation, Amatare et al. [184] introduce RF-Vision, a privacy-preserving framework combining RF propagation and machine learning (ML) within digital twins for object characterisation in indoor settings. They systematically create digital twins using Blender, perform RF propagation simulations using NVIDIA's Sionna Ray-Tracing tool, and use these simulated RF maps to train a ML model to characterise objects based on material properties and shapes. RF-Vision shows strong performance, notably when transmitters are centrally positioned, and excels at detecting objects in non-line-of-sight regions. It provides a viable alternative to traditional visual-based characterisation techniques, emphasising privacy preservation, operational efficiency with smaller datasets, and compatibility with existing RF infrastructure such as WiFi.

Digital twins enhanced with ray tracing have numerous real-world applications across engineering, urban planning, and environmental monitoring. By providing highly accurate simulations of physical environments, they enable predictive modelling and scenario testing with impressive accuracy and detail. This is particularly valuable in safety-critical domains, where understanding the interaction between natural forces and infrastructure is essential. One such application is in geotechnical engineering, where digital twins are used to assess slope stability and predict landslide events. Liu et al. [185] introduce a slope digital twin framework designed to predict rainfall-induced slope instability using historical slope performance and real-time monitoring data. By combining monitoring data (such as rainfall and pore water pressure) and probabilistic analysis, they develop an updated virtual slope model that continuously learns from its physical counterpart, significantly enhancing landslide prediction accuracy. The proposed method was



tested successfully on a real slope in Hong Kong, accurately predicting slope failure during an extreme rainstorm event in 2008. While not a telecommunications application of digital twins, it is interesting to see how the same concept can be applied in other domains. In effect, much of the detail provided in a digital twin could be equally useful across multiple domains, providing further motivation to prioritise the creation and maintenance of them.

### 6.3.2 Drones' Spectrum Access

#### Summary

*Drones are emerging as a compelling use case for more sophisticated and adaptive spectrum management, given their unique operational characteristics and increasing deployment across commercial and civilian sectors. Although drones could theoretically operate in licensed spectrum at altitude with minimal interference to ground users, existing regulations often prohibit such usage. Currently, most commercial drones rely on shared industrial, scientific and medical (ISM) bands like 2.4 GHz and 5.8 GHz. These bands are prone to interference from other devices, limiting performance and reliability, especially for critical applications. As drone usage intensifies with applications like delivery services, advanced spectrum management approaches, including dynamic access and tailored waveform designs, are becoming essential to ensure robust, scalable, and interference-resilient communication systems.*

#### Review

A critical and highly interesting emerging application for advanced and precise spectrum management is drone spectrum usage. In theory, a drone with an antenna with high directivity and good power control could in certain situations make use of licensed spectrum while at altitude with minimal significant RF interference to ground-based spectrum users. Yet, most conventional regulations would make this illegal. Most commercial drones operate on the 2.4 GHz and 5.8 GHz bands (ISM bands), primarily for control signals and video transmission. The 2.4 GHz band is widely used due to its balance between range and data throughput, making it suitable for control commands. The 5.8 GHz band offers higher data rates with reduced latency, which is advantageous for transmitting high-quality video feeds, especially in first-person view (FPV) applications or drone video capture such as what a sport broadcast may use. However, the 5.8 GHz band has less effective range compared to 2.4 GHz. Additionally, some drones, particularly those built by hobbyists, operate on lower frequencies such as 433 MHz and 915 MHz. These lower frequencies can extend control range but may limit video transmission quality due to reduced bandwidth. It's important to note that these frequency bands are shared with other devices, leading to potential interference. For instance, the 2.4 GHz band is also utilised by Wi-Fi networks and Bluetooth devices, which can affect drone operations. To mitigate interference, some advanced drones employ frequency-hopping techniques, dynamically switching frequencies to maintain robust communication links [186, 187, 188].



The choice of spectrum used by commercial drones is a constraint forced by the current state of spectrum management, and by extension the most affordable RF hardware (which is influenced by what can legally be used commercially). In other words, it is not necessarily an optimal technical decision to use these frequency bands for drones. As previously mentioned, interference to drones, especially on shared bands such as the ISM bands can be a problem for drone operation, which is generally not tolerated for critical usages such as government and defence, yet must be tolerated by all other use cases. This provides a motivation for more advanced drone spectrum management. As commercial and non-commercial use cases for drones become more mainstream, the existing spectrum management problems will worsen. For example, it is expected that drone delivery services could become very commonplace for global companies such as Amazon and many others.

Unmanned aerial vehicle (UAV)s are increasingly being deployed across a wide range of civilian and commercial applications, from aerial surveillance to delivery services. This surge in UAV usage places significant demands on existing wireless communication infrastructures. Efficient spectrum management has therefore become a critical enabler for reliable and scalable UAV network operations. Addressing these challenges requires a deep understanding of both the unique characteristics of UAV communications and the limitations of current spectrum allocation frameworks. Jasim et al. [189] provide a comprehensive overview of spectrum management challenges and strategies specifically tailored to unmanned aerial vehicle (UAV) networks. Due to their highly dynamic and adaptive nature, UAV networks require novel spectrum management techniques to facilitate seamless integration and coexistence with existing terrestrial wireless technologies. They categorise existing management schemes into deterministic (resource allocation and access control), opportunistic (spectrum sensing, sharing, and decision-making), and competitive (auction mechanisms). They further outline deployment scenarios, identifies critical mathematical and technological tools like optimisation, ML, and blockchain. They highlight open research challenges such as interference mitigation, spectrum aggregation, and dynamic licensing models, emphasising the need for adaptive and robust solutions to accommodate the rapid expansion and diverse use cases of UAV operations.

Unmanned aerial systems (UAS) are rapidly becoming integral to various civilian applications, ranging from delivery services to infrastructure monitoring. As their usage expands, ensuring reliable and efficient communication becomes a critical challenge. Future wireless networks must adapt to accommodate the unique mobility, scalability, and latency demands of UAS operations. In this context, recent research has begun to focus on innovative spectrum management and waveform design tailored to these emerging requirements. Kakar and Marojevic [190] explore future spectrum management strategies and waveform technologies suitable for the growing number of civilian unmanned aerial systems (UAS). They emphasise the dominance of small and micro unmanned aerial vehicle (UAV)s projected between 2028 and 2032, highlighting their varying communication needs from high-reliability low-throughput con-

trol signalling to high-throughput payload data transmission. They advocate using orthogonal frequency-division multiplexing (OFDM) for its flexibility, spectral efficiency, and compatibility with dynamic spectrum access, considering it a promising candidate for UAV communications. They further emphasise the necessity of dynamic spectrum sharing techniques due to anticipated spectrum congestion and interference, discussing spectrum sharing as a critical approach to meeting future UAS communication requirements efficiently.

As the demand for resilient and efficient communication networks grows, novel strategies for optimising network performance have become critical. One promising avenue involves the use of mobile nodes to address connectivity and capacity challenges in real time. Hunjet [191] explores the strategic use of networked autonomous vehicle (NAV), including unmanned aerial vehicle (UAV), to enhance wireless network performance. They demonstrate that these mobile nodes can be dynamically deployed to target weak points in the network, such as bridges and articulation points to improve connectivity and survivability. They optimise NAV placement and adjust their transmission powers, showing improvements in both network capacity and power efficiency. Through graph-based and optimisation-driven approaches, including particle swarm optimisation (PSO) and cross entropy methods, they illustrate that NAVs/UAVs can reduce interference, support beneficial simultaneous transmissions, and adapt to dynamic environments. These enhancements also contribute to network security by limiting detectability and interception range. They ultimately establish NAVs and UAVs as powerful tools for constructing flexible, high-capacity, and energy-efficient MANETs.

## 6.4 Demonstrations

### 6.4.1 Context

The demonstrations in this chapter address the central problem of spectrum management in increasingly complex and dynamic environments. As outlined in the introduction and approach, conventional spectrum allocation methods are largely two-dimensional and static, failing to account for mobility, altitude, and temporal variability. This limitation becomes critical when multiple users operate simultaneously across shared or adjacent bands, particularly in dense urban environments or when airborne systems such as drones are introduced. By integrating reachability modelling with radio frequency (RF) propagation analysis, the problem shifts from abstract allocation to measurable, risk-aware spectrum planning that accounts for geography, infrastructure, and mobility constraints. Without such methods, interference risk is either underestimated or only addressed reactively, limiting efficiency and scalability of future spectrum use.

The literature review examines this challenge in the context of research on digital twins, telecommunications modelling, and drone spectrum access. Digital twins have been widely adopted across engineering and telecommunications as detailed, interactive representations of physical environments, capable of simulating scenarios with high accuracy. Applications in RF modelling, object localisation, and predictive planning show their value in situations where precision and real-time adaptability are essential. In parallel, drones have emerged as a disruptive case for spectrum access, constrained by current reliance on shared industrial, scientific and medical (ISM) bands yet demanding robust, interference-resilient links for applications such as delivery or surveillance. Studies highlight the need for dynamic spectrum allocation, new waveform designs, and coexistence strategies to integrate drone communications alongside terrestrial networks. These themes converge around the need for scalable, modular, and adaptive models to manage spectrum in both ground and aerial domains.

The demonstrations connect these ideas by providing practical, visual evidence of how digital twin simulations can address the outlined problems. Each demonstration isolates and then combines elements of reachability and RF coverage to illustrate how user mobility and propagation environments jointly determine interference outcomes. Scenarios range from cars moving in structured networks to drones following planned flight paths, showing how co-interference can be mapped, quantified, and anticipated across time and space. By grounding the theoretical foundations from the introduction and literature review in detailed simulations, the demonstrations serve as proof-of-concepts for a spectrum management digital twin. They show how future spectrum allocation frameworks could shift from static, regulation-driven models to dynamic, data-driven systems capable of predicting and mitigating interference before it occurs.

### 6.4.2 Summary

The demonstrations in this final chapter show how a digital twin model can analyse coordinated spectrum access scenarios that would otherwise be too complex to evaluate. The approach is to combine reachability and RF coverage calculations to explore possible interference between spectrum users. A simple starting case is considered where spectrum is allocated by time rather than by area. Users are free to move during their allocation, but the key feature is that they know exactly how long they have access, allowing them to plan. This mobility makes interference harder to predict. Conventional methods might calculate interference for a single location over months or years, but repeating such analyses every 10 minutes across many locations is infeasible without computational support. The digital twin provides a feasible solution, and its potential is demonstrated first by isolating its key functions before combining them to explain results more clearly.

Reachability is first contrasted with Brownian motion before focussing on shortest path tree (SPT)s as the more suitable model. SPTs are demonstrated for both urban and rural environments, and under the assumption that users may walk or drive. The difference between urban and rural cases is stark, shaped by density, traffic, and road networks. RF coverage models based on ray tracing are also demonstrated, showing how propagation varies further due to terrain such as hills or valleys. Free-space conditions are contrasted against ray-traced environments to highlight the importance of propagation effects. These early demonstrations explain how the simulation's building blocks behave alone and in combination. More advanced scenarios then follow, such as multiple mobile transceivers operating close together. By calculating signal-to-interference (SIR) for locations where power levels exceed a threshold of path gain, it becomes possible to map when and where interference occurs. Results are presented as time series heat maps, showing how interference evolves over time.

The same approach applies to future use cases such as drones, which pose challenges due to fast, unconstrained movement. Reachability is less useful here, so drone spectrum management can instead be treated like aircraft managed by air traffic control. By planning and analysing flight paths, RF coverage and interference for drone-to-drone and drone-to-ground access can be evaluated in the same way. With knowledge of all possible interference scenarios across time, spectrum planning can aim to reduce risk rather than eliminate it completely, since total avoidance comes at the cost of efficiency. The SPT model allows weighting of RF coverage by travel time, reflecting the lower likelihood of users straying far from their starting point. More accuracy could come from life patterns or historical movement data, further shrinking areas of probable interference. The outcome is tighter and longer allocations between users, made possible by a digital twin system that reflects live environmental conditions.

### 6.4.3 Tools

The tools first shown in the previous chapter were extended to accomplish more complex spectrum management simulations. Beyond the previous tools for rendering 3D RF ray tracing and line of sight (LOS) heatmaps over a 3D realistic map, new features were added to support animated 3D graphics. These features include combining shortest path tree (SPT)s with RF ray tracing temporally, as well as applying RF ray tracing to drone RF transmissions over time. They also allow combining  $n$  instances of either of these across space and time. This enables results such as the mutual worst signal-to-interference (SIR) shown below in the demonstrations to be calculated and displayed temporally, in 3D space.

### 6.4.4 Structure

This section presents five main demonstrations:

1. **Isolated Reachability Models** – Pure mobility reachability demonstrations including Brownian motion and SPT walking and driving in urban and rural environments.
2. **Isolated RF Propagation Models** – Standalone analysis of free-space and ray tracing RF propagation in both urban and rural environments.
3. **Combined Reachability and RF Propagation** – Integration of mobility models (Brownian, SPT walking/driving) with free-space and ray-tracing propagation in urban and rural settings.
4. **Multi-user Coordinated Spectrum Access** – Vehicle mobility (driving) under free-space and ray-tracing propagation, across urban and rural environments, with time-series visualisation.
5. **Coordinating Spectrum Access with Drones** – RF coverage and interference analysis including urban, rural, and aerial scenarios, as well as multi-drone co-interference and drone-to-ground interactions.

## 6.4.5 Demonstrations 1-3: RF Coverage Reachability

### Summary

A key concept shown in this chapter is the idea that models of spatiotemporal reachability and RF propagation within a highly accurate digital representation of a 3D environment have applications in coordinating spectrum access. The existence of a *digital twin* simulation could provide a way to account for how mobility and RF transmission characteristics impact spectrum allocations under realistic conditions. To explore this kind of *digital twin* simulation, A comparison is made between a basic and a complex model for *reachability* and *RF propagation* for *urban* and *rural* environments, shown across seven demonstrations as per Table 6.1. The reachability and RF coverage models are first presented in isolation, followed by their combination to produce RF coverage reachability heat maps.

Demonstration	Reachability	RF Propagation	Environment
<b>1 – Reachability</b>			
1.1	Brownian random	-	-
1.2 / 1.3	SPT	-	urban / rural
<b>2 – RF propagation</b>			
2.1	-	free-space	-
2.2 / 2.3	-	RF ray tracing	urban / rural
<b>3 – Combined</b>			
3.1	Brownian random	RF ray tracing	urban
3.2 / 3.3	SPT	free-space	urban / rural
3.4 / 3.5	SPT	RF ray tracing	urban / rural

Table 6.1: Combinations of reachability models, RF propagation methods, and environment types

## Demonstration 1: Isolated Reachability Models

Demonstration 1 explores the effect of different reachability models starting from baseline basic models, progressing to more advanced options.

### Demo 1.1: Brownian Random Reachability

Brownian randomness is suitable for a baseline reachability model as it applies no understanding of the real world environment on motion, other than basic movement speed limitations. Figure 6.4.5.1 shows the density of displacement from the origin where 1 million simulated agents make random movements for 10 minutes at 5 m/s from a common starting location. The density pattern is an approximation of a circular Gaussian cutoff.

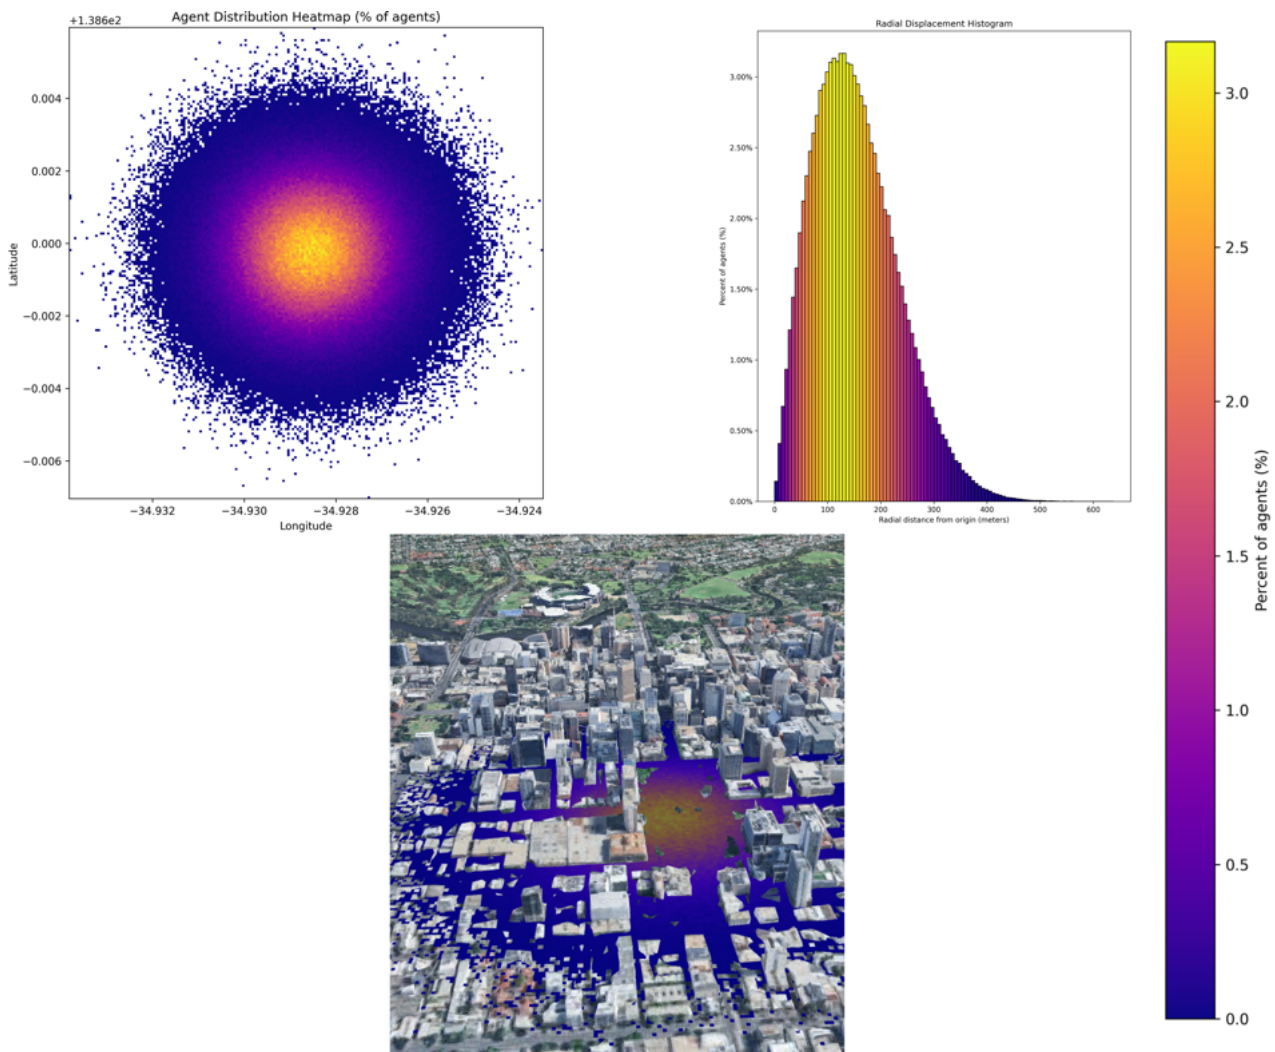


Figure 6.4.5.1: (top left) Bird's-eye view of Brownian motion; (top right) Histogram of radial displacement; (bottom) Overlaid onto a to-scale 3D city model for reference

## Demo 1.2: Isochrone Reachability – Urban

### Setup:

- Data: OSM [117] processed by Valhalla [120]
- Routing Engine: Valhalla [120] using Expansion [115]

Under real world urban conditions, Brownian randomness motion does represent the remote possibility of going very far in the available time as an outlier, but does not consider how the actual real world environment impacts travel. Isochrones and the very related shortest path tree (SPT) concept (shown in Figure 6.4.5.2) provide a way to represent the real world environment's effects on travel. With a precise routing engine, enhanced by traffic data, it can set an outer limit for travel under normal conditions and weight locations by travel time.

The Brownian motion model Figure 6.4.5.1 and walking isochrone on the left of Figure 6.4.5.2 both assume walking speeds of 5 m/s and hence provide similar results. The most notable differences are that 1) the isochrone model shape approximates a diamond shape due to north-south and east-west axes movements generally being shorter paths and 2) the heat distribution is linear for travel time whereas the brownian model approximates a Rayleigh distribution showing much higher density at it's peak. The right side of Figure 6.4.5.1 shows how the reachability changes when able to drive (in free flow traffic conditions). While the brownian randomness model's motion could be increased to car speeds, it would represent the road network very poorly (compared to walking where it is similar) so it would be expected to provide a very different outcome.

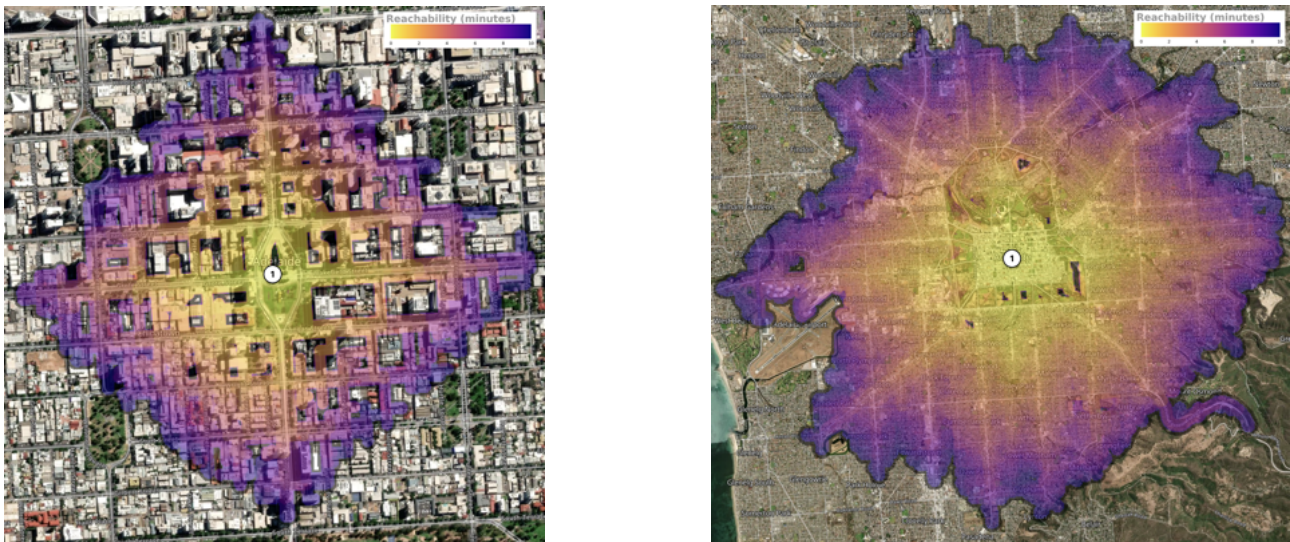


Figure 6.4.5.2: (left) Isochrone for 10 minutes walking in the CBD; (right) Same but driving



### Demo 1.3: Isochrone Reachability – Rural

#### Setup:

- Data: OSM [117] processed by Valhalla [120]
- Routing Engine: Valhalla [120] using Expansion [115]

As seen in the SPT shown in Figure 6.4.5.3, under real world rural conditions, there is a significant divergence from the Brownian random movement model Figure 6.4.5.1 for walking or driving conditions. Routing engines are designed to provide driving directions, so they generally do not provide paths outside of the public roads. This means that there are vastly less destinations included in the rural SPT than the urban ones. It could be argued that this better reflects where most agents would go under real-world conditions, making it an improvement over allowing paths anywhere. However, it is unable to model how locals in particular are very likely leave the main roads to access their properties. Another significant limitation for rural areas is that SPTs for walking in particular will only provide known walking paths. Unlike driving where most roads are covered, these walking paths only cover a fraction of where humans would actually walk. However, for the purposes of getting a general understanding of where spectrum users are likely to go in these areas, it may account for a majority of the travel they are likely to do.





Figure 6.4.5.3: (left) Isochrone for 10 minutes walking in a rural area; (right) Same but driving

## Demonstration 2: Isolated RF Propagation Models

Demonstration 2 explores the effect of different RF propagation models starting from a baseline model before progressing to more advanced options.

### Demo 2.1: RF Propagation in Free-space

**Setup:**

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: Empty 3D model for free-space conditions
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

RF propagation in free-space is a suitable baseline model as it allows isolation of how more complex RF transmission environments impact the path loss (PL) (dB) of the transmitted signals. Figure 6.4.5.4 shows the PL (dB) from a point in the Adelaide CBD if the impact of the buildings and terrain was ignored.




Figure 6.4.5.4: RF propagation in free-space showing path loss (dB) overlaid over the Adelaide CBD for visual reference

## Demo 2.2: RF Raytracing Propagation – Urban

### Setup:

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

RF ray tracing is demonstrated for an urban environment in Figure 6.4.5.5, showing path loss (dB) values for a transmitter located in the Adelaide CBD. It can be seen that the tall city buildings have a big impact on the propagation with most of the transmission power remaining in between the buildings with some leaking out through the gaps.




Figure 6.4.5.5: (top left) Path loss heatmap of Adelaide CBD from ground-based transmitter; (top right) Same scene without heatmap; (bottom) 3D overlay showing signal propagation.



## Demo 2.3: RF Raytracing Propagation – Rural

### Setup:

- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

RF ray tracing is demonstrated for a rural environment in Figure 6.4.5.6. The heat map values for path loss (dB) from the transmitter located at Hall's Gap are very different to the urban RF ray tracing shown in Figure 6.4.5.5, as they are influenced primarily by the landscape rather than by buildings. In general, the RF transmission are free to travel quite far unless there is hilly or mountainous terrain, like in the example Figure 6.4.5.6.




Figure 6.4.5.6: (left) Bird's-eye view of Brownian motion and Free-space; (right) Same data overlaid onto a to-scale 3D city model for reference.

### Demonstration 3 – Combined Reachability and RF Propagation Models

Demonstration 3 explores what occurs when the concept of reachability and RF propagation modelling are combined to create a 3D heatmap. Combinations of the systematically introduced strategies for modelling both are shown to progress from a baseline model to more advanced models.

#### Demo 3.1: Brownian Random – RF Ray Tracing – Urban

##### Setup:

- Movement: Brownian random from initial position
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.5.7 shows the result of combining the Brownian random walking motion from Figure 6.4.5.1 with the urban RF ray tracing from Figure 6.4.5.5. Given that the Brownian random model was approximately a circle, the result is essentially the same as summing many RF traced coverage maps for points sampled from a circle. Although they are sampled from an approximate circle, the city buildings in the gridded structure allow the RF transmissions to propagate north-south and east-west far more easily than diagonally creating a diamond shape. There is also weak coverage shown to the hills through gaps in the buildings. The suburbs' coverage appears weaker than expected, most likely due to limitations of the ray tracing with the flat but jagged light detection and ranging (LiDAR) surface not allowing the *rays* to scatter properly.

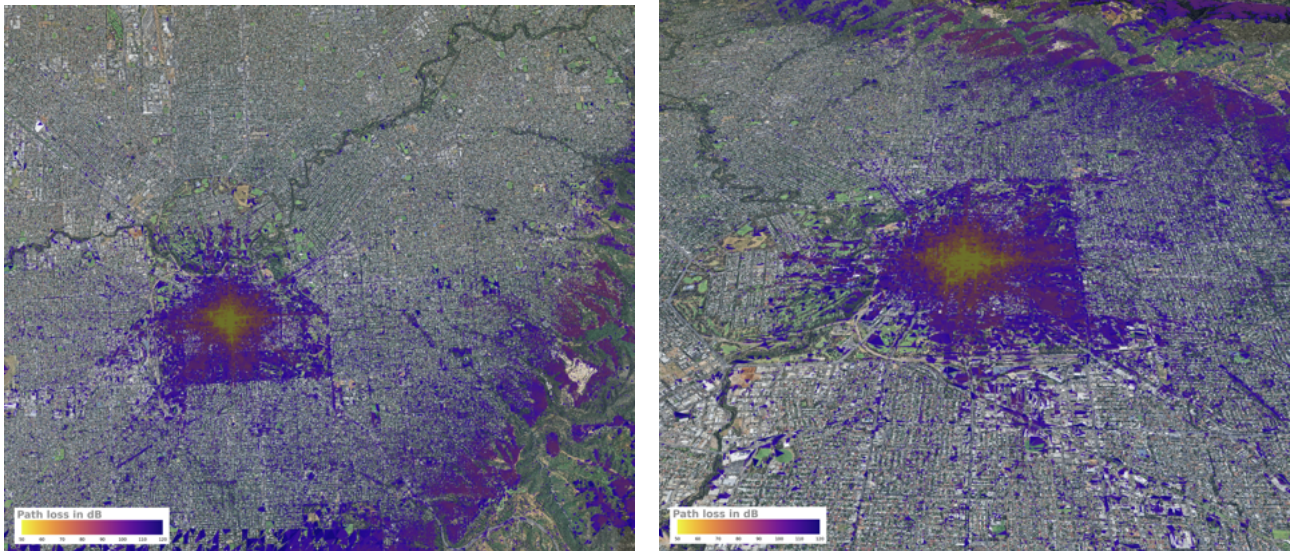


Figure 6.4.5.7: (left) Bird's-eye view of Brownian motion and RF ray tracing path loss (dB); (right) Same data viewed from the west at an angle.

### Demo 3.2: SPT Walking – RF Free-Space – Urban

#### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: Empty 3D model for free-space conditions
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.5.8 shows the reachability of the walking SPT from Figure 6.4.5.2 combined with RF free-space propagation from Figure 6.4.5.4 where the transmission environment effects of the city buildings are ignored. It can be seen that the RF propagation is able to radiate in all directions freely but the strongest power occurs only at the places reachable within 7 minutes of walking.




Figure 6.4.5.8: (top left) Free-Space path loss for a walking SPT at initial time  $t=0$ ; (top right) Same after 7 minutes of SPT walking; (bottom) Bird's-eye view



### Demo 3.3: SPT Driving – RF Free-Space – Rural

#### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: Empty 3D model for free-space conditions
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.5.9 shows the reachability of the driving SPT from Figure 6.4.5.3 combined with RF free-space propagation from Figure 6.4.5.4 where the transmission environment effects of the landscape are ignored (especially mountains around Hall's Gap). It can be seen that the RF propagation is able to radiate in all directions freely but the strongest power occurs only at the places reachable within the 4 minutes of driving.




Figure 6.4.5.9: top left) Free-Space path loss for a driving SPT at Hall's Gap after 30 seconds; (top right) Same after 4 minutes of possible driving; (bottom) Same but after 7 minutes

### Demo 3.4: SPT Walking – RF Ray Tracing – Urban

#### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.5.10 shows the reachability of the walking SPT from Figure 6.4.5.2 combined with RF ray traced propagation from Figure 6.4.5.5 where the transmission environment effects of the city buildings are included at an initial time and 7 minutes later. It can be seen that the dense gridded city centre allows the RF transmissions to propagate north-south and east-west far more easily than diagonally creating a diamond shape. At the initial time surrounded by many tall buildings, the RF transmissions are heavily blocked. Following 7 minutes of possible walking, the strongest power could reach significantly more areas.

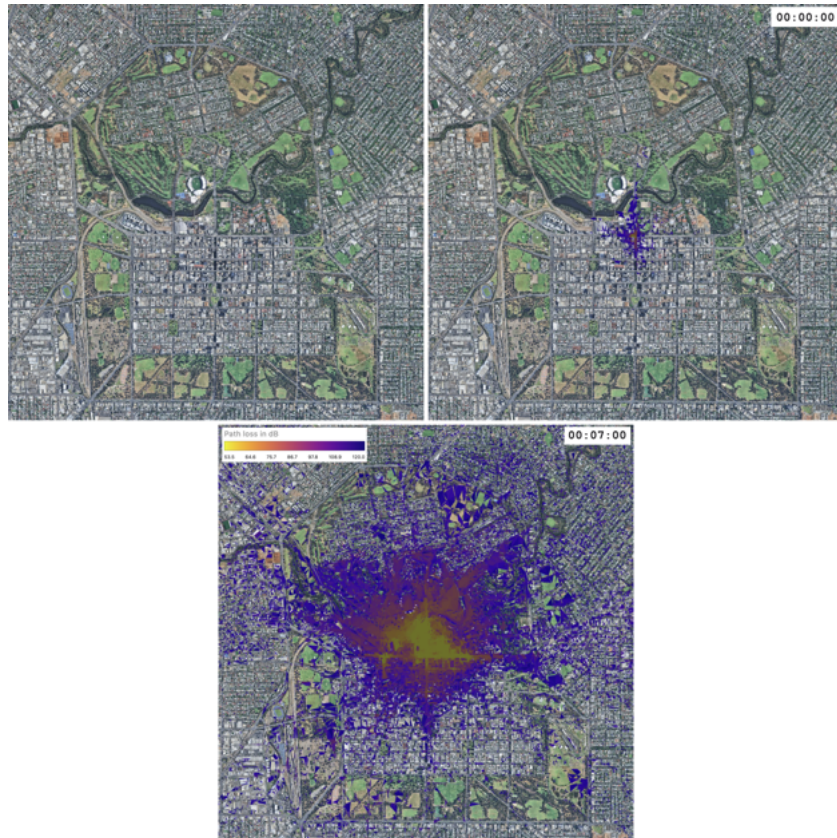


Figure 6.4.5.10: (top left) Bird's-eye view of Adelaide CBD walking SPT combined with ray tracing before the initial time 00:00:00; (top right) Same at 00:00:00; (bottom) Same at 00:07:00





Figure 6.4.5.11 shows the same scene as Figure 6.4.5.10 but from an angled view north west of the transceiver's initial location. From this view point, it is very noticeable that the parklands north of the city centre get significantly more coverage than south due to the wider open spaces allowing nearly free-space propagation.

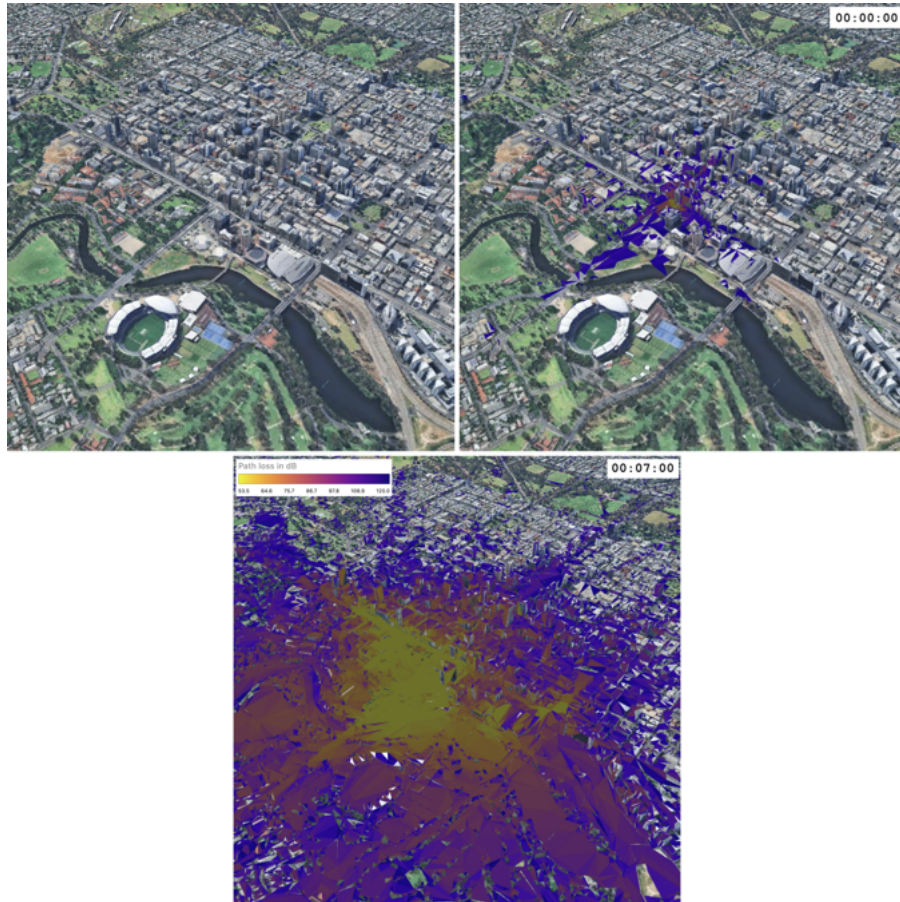


Figure 6.4.5.11: (top left) Closer angled view of Adelaide CBD walking SPT combined with ray tracing before the initial time 00:00:00; (top right) Same at 00:00:00; (bottom) Same at 00:07:00

### Demo 3.5: SPT Driving – RF Ray Tracing – Rural

#### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.5.12 shows a driving SPT, starting from Hall's Gap at four points in time combined with RF ray tracing from sampled points (1 m apart). As time passes, it can be seen that the possible areas of RF interference from that transceiver increase sharply.

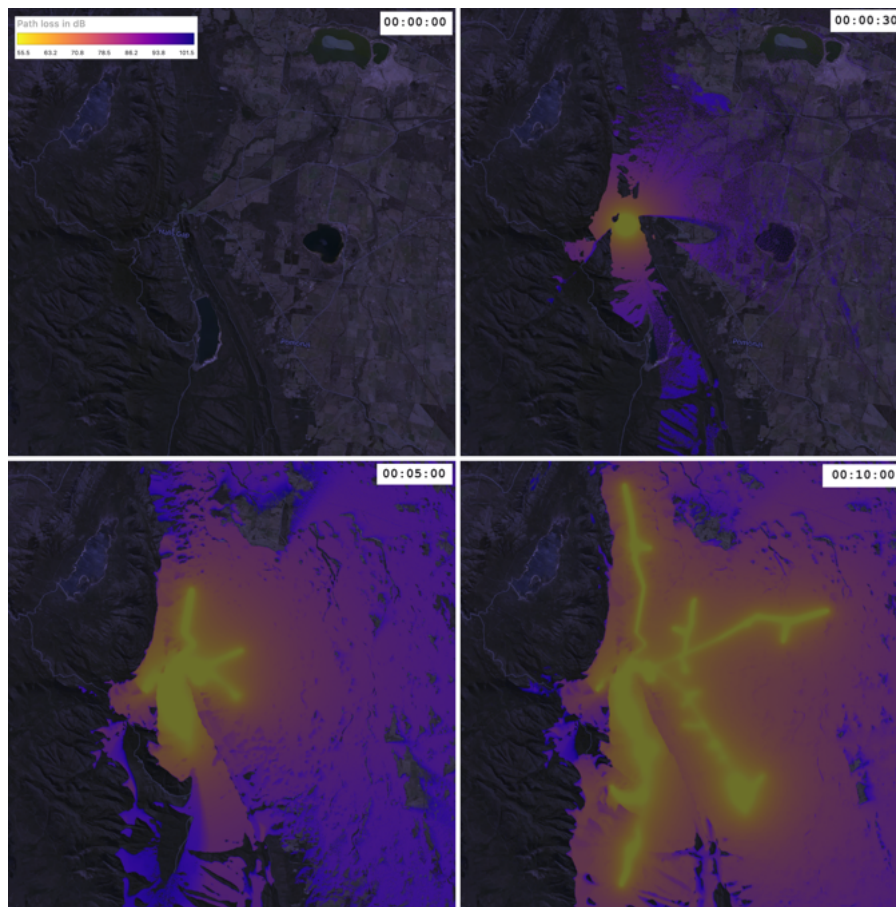


Figure 6.4.5.12: (top left) Bird's-eye view of Hall's Gap showing a driving SPT combined with RF ray tracing at 00:00:00; (top right) Same 30 seconds later at 00:00:30; (bottom left) Same 5 minutes later at 00:05:00; (bottom right) Same 10 minutes later at 00:10:00



Figure 6.4.5.13 shows the same scene and conditions as Figure 6.4.5.12, but from a lower altitude angled view from the south to get a better reference for the elevation of the surrounding mountains of Hall's Gap.




Figure 6.4.5.13: (top left) Close angled view of Hall's Gap showing a driving SPT combined with RF ray tracing at 00:00:00; (top right) Same 30 seconds later at 00:00:30; (bottom left) Same 5 minutes later at 00:05:00; (bottom right) Same 10 minutes later at 00:10:00

### 6.4.6 Demonstration 4: Multi-user Coordinated Spectrum Access

The goal of this demonstration is to show what RF coverage reachability is and how it could be applied to coordinate spectrum access between multiple users. RF coverage reachability is produced by combining a reachability and RF coverage model. The outcome is a time series of heat maps showing where RF transmissions by a given user could reach with increasing allocation time. This information could provide a way to analyse and quantify the possible impact of multiple spectrum allocations to each other over time. In this subsection, shortest path tree (SPT)s are shown as a model of reachability. RF ray tracing is presented as a model of RF propagation. Both are applied under free-space RF propagation assumptions. A 3D mesh produced from Google 3D tiles and regional digital elevation model (DEM) data is used. These are applied to the urban and rural scenarios respectively. The power for the individual spectrum users is measured by path loss (PL) (dB). Possible co-interference between multiple users is measured by calculating the *mutual worst SIR* where two or more users' PL (dB) is above a chosen minimum power threshold. A description of the demonstrations in this subsection is provided neatly in Table 6.2.

Demonstration	Reachability	RF Propagation	Environment
<b>4 – Multi-user</b>			
4.1	SPT driving	RF free-space	urban
4.2	SPT driving	RF free-space	rural
4.3	SPT driving	RF ray tracing	urban
4.4	SPT driving	RF ray tracing	rural

Table 6.2: A description and layout of demonstrations about multi-user coordinated spectrum access



## Demo 4.1: SPT Driving – RF Free-Space – Urban

### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: Empty 3D model for free-space conditions
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.6.1 shows the path loss (dB) values for two transceivers starting on different sides of the Adelaide CBD at an initial time  $t=0$ . Initially, their RF transmissions under free-space conditions radiate in all directions easily, but they are too far from each other to cause signal-to-interference (SIR) at locations where the power of each is above the chosen minimum power threshold.

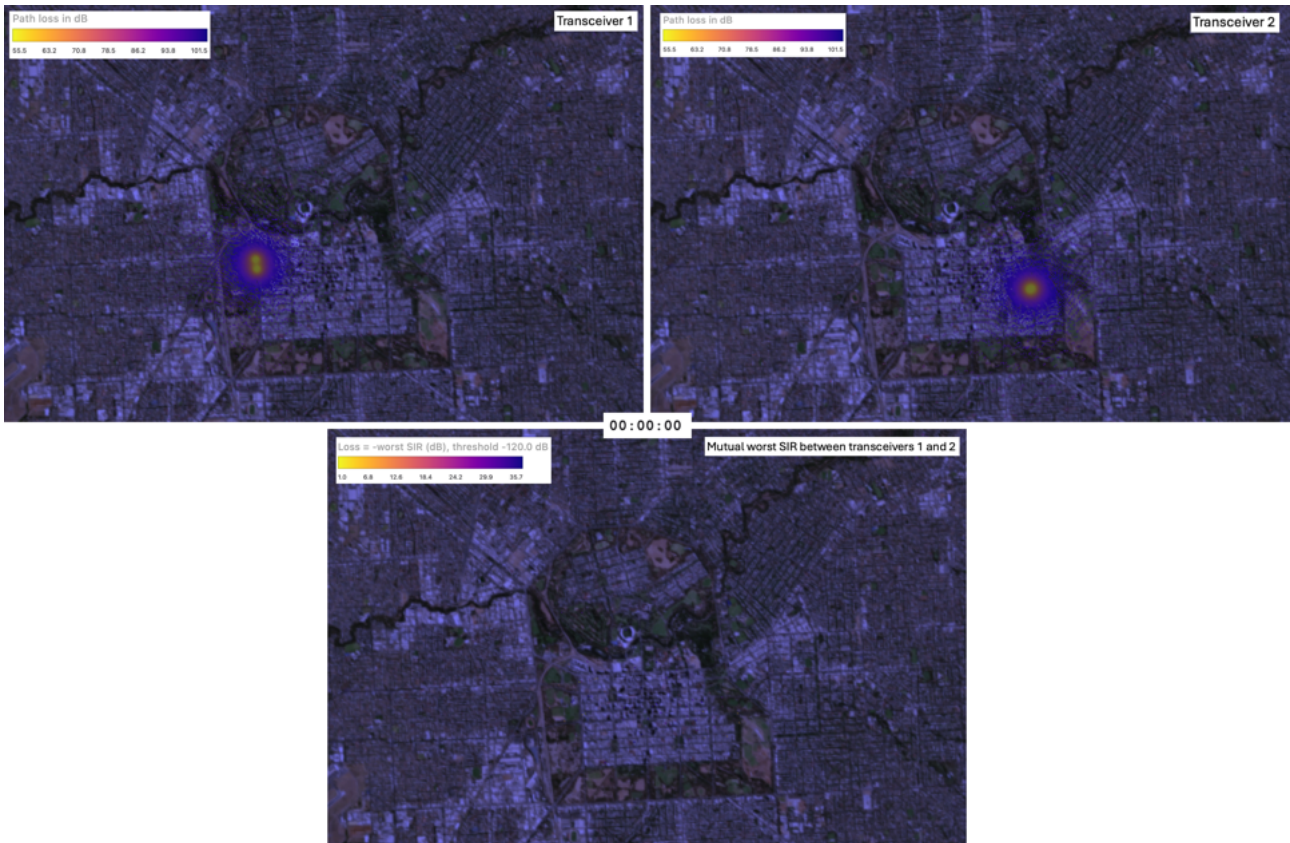


Figure 6.4.6.1: (top left) Transceiver 1's path loss (dB) at an initial time; (top right) Transceiver 1's path loss (dB) at same time; (bottom) Signal-to-interference ratio between the two

Figure 6.4.6.2 shows the path loss (dB) values for two transceivers starting on different sides of the Adelaide CBD at time  $t=1$  minute. After one minute of driving, there are notably more locations where their RF transmissions under free-space conditions can reach and lower mutual SIR values between the two transceivers can be observed at locations where the power of each is above the chosen minimum power threshold.

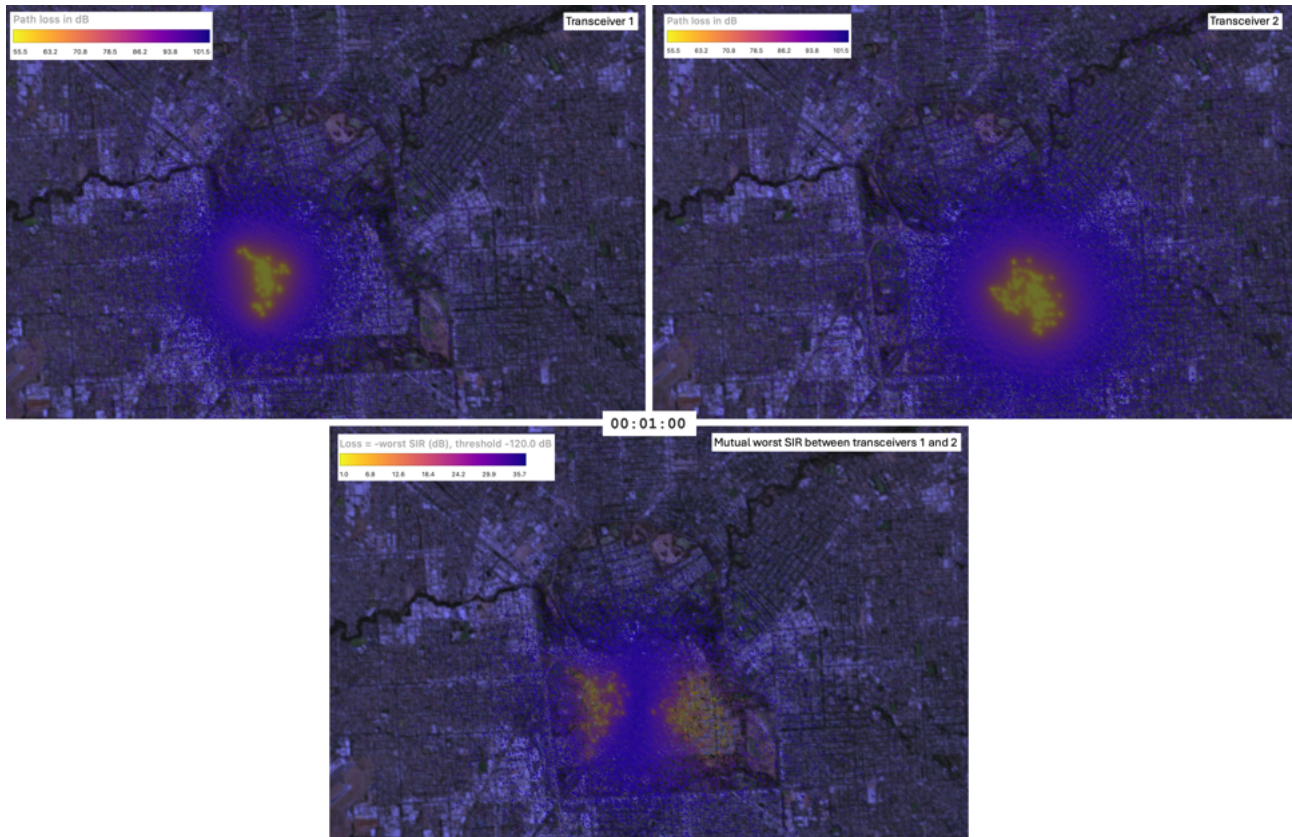


Figure 6.4.6.2: (top left) Transceiver 1's path loss (dB) after 1 minute; (top right) Transceiver 1's path loss (dB) at same time; (bottom) Signal-to-interference ratio between the two

Figure 6.4.6.3 shows the path loss (dB) values for two transceivers starting on different sides of the Adelaide CBD at time  $t=2$  minutes. After two minutes of driving, there are notably more locations where their RF transmissions under free-space conditions can reach and even lower mutual SIR values between the two transceivers can be seen at locations where the power of each is above the chosen minimum power threshold.




Figure 6.4.6.3: (top left) Transceiver 1's path loss (dB) after 2 minutes; (top right) Transceiver 1's path loss (dB) at same time; (bottom) Signal-to-interference ratio between the two



Figure 6.4.6.4 shows the path loss (dB) values for two transceivers starting on different sides of the Adelaide CBD at time  $t=3$  minutes. After three minutes of driving, there are notably more locations where their RF transmissions under free-space conditions can reach and even lower mutual SIR values can be seen now covering the CBD and surrounding inner suburbs between the two transceivers at locations where the power of each is above the chosen minimum power threshold.

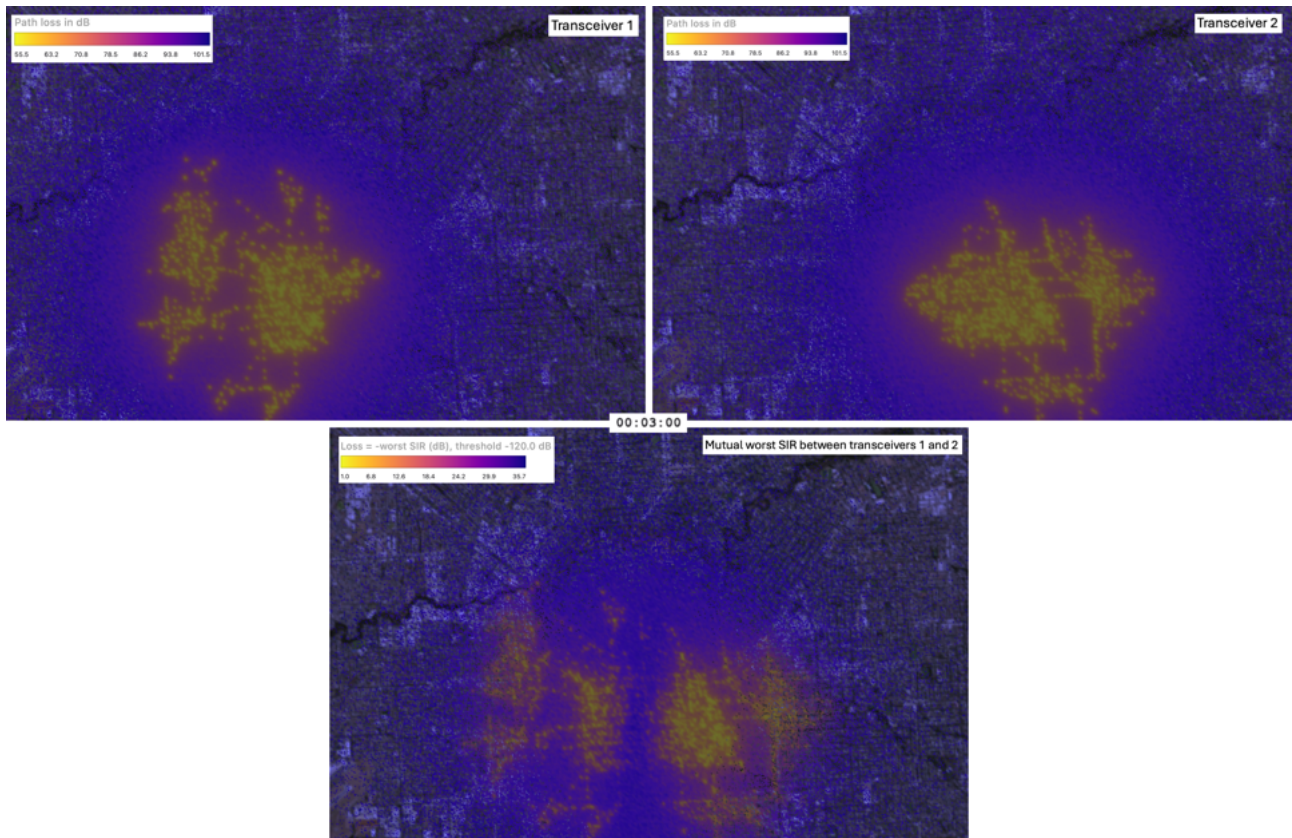


Figure 6.4.6.4: (top left) transceiver 1's path loss (dB) after 3 minutes; (top right) transceiver 1's path loss (dB) at same time; (bottom) signal-to-interference ratio between the two



## Demo 4.2: SPT Driving – RF Free-Space – Rural

### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: Empty 3D model for free-space conditions
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.6.5 shows the path loss (dB) values for two transceivers starting in Hall's Gap and about 10 minutes south of there by car at an initial time  $t=0$ . Initially, their RF transmissions under free-space conditions radiate in all directions easily, but they are too far from each other to cause signal-to-interference (SIR) at locations where the power of each is above the chosen minimum power threshold.




Figure 6.4.6.5: (top left and top right) Hall's Gap at initial time 00:00:00 showing RF free-space SPT path loss for transceiver 1 and 2 respectively; (below) Signal-to-interference ratio for transceiver 1 and 2 at 00:00:00



Figure 6.4.6.6 shows the path loss (dB) values for two transceivers starting in Hall's Gap and about 10 minutes south of there by car after 3 minutes have passed of possible SPT driving. Their RF transmissions under free-space conditions radiate in all directions easily and at this point weakly interfere with each other to cause signal-to-interference (SIR) at locations where the power of each is above the chosen minimum power threshold.




Figure 6.4.6.6: (top left and top right) Hall's Gap after 3 minutes have passed at 00:03:00 showing RF free-space SPT path loss for transceiver 1 and 2 respectively; (below) Signal-to-interference ratio for transceiver 1 and 2 at 00:03:00

Figure 6.4.6.7 shows the path loss (dB) values for two transceivers starting in Hall's Gap and about 10 minutes south of there by car after 5 minutes have passed of possible SPT driving. Their RF transmissions under free-space conditions radiate in all directions easily and interfere with each other to cause signal-to-interference (SIR) at locations where the power of each is above the chosen minimum power threshold.




Figure 6.4.6.7: (top left and top right) Hall's Gap after 5 minutes have passed at 00:05:00 showing RF free-space SPT path loss for transceiver 1 and 2 respectively; (below) Signal-to-interference ratio for transceiver 1 and 2 at 00:05:00



Figure 6.4.6.8 shows the path loss (dB) values for two transceivers starting in Hall's Gap and about 10 minutes south of there by car after 10 minutes have passed of possible SPT driving. Their RF transmissions under free-space conditions radiate in all directions easily and strongly interfere with each other to cause signal-to-interference (SIR) at locations where the power of each is above the chosen minimum power threshold.




Figure 6.4.6.8: (top left and top right) Hall's Gap after 10 minutes have passed at 00:10:00 showing RF free-space SPT path loss for transceiver 1 and 2 respectively; (below) Signal-to-interference ratio for transceiver 1 and 2 at 00:10:00

### Demo 4.3: SPT Driving – RF Ray Tracing – Urban

#### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

Figure 6.4.6.9 shows driving SPTs for two transceivers combined with RF ray tracing path loss (dB) for two transceivers starting at different locations in the Adelaide CBD for the first 1.5 minutes of possible driving in free flow traffic conditions. Only the individual transceivers' path loss (dB) heatmaps are shown here because there are no signal-to-interference (SIR) measurements at this stage because they are too far from each other.

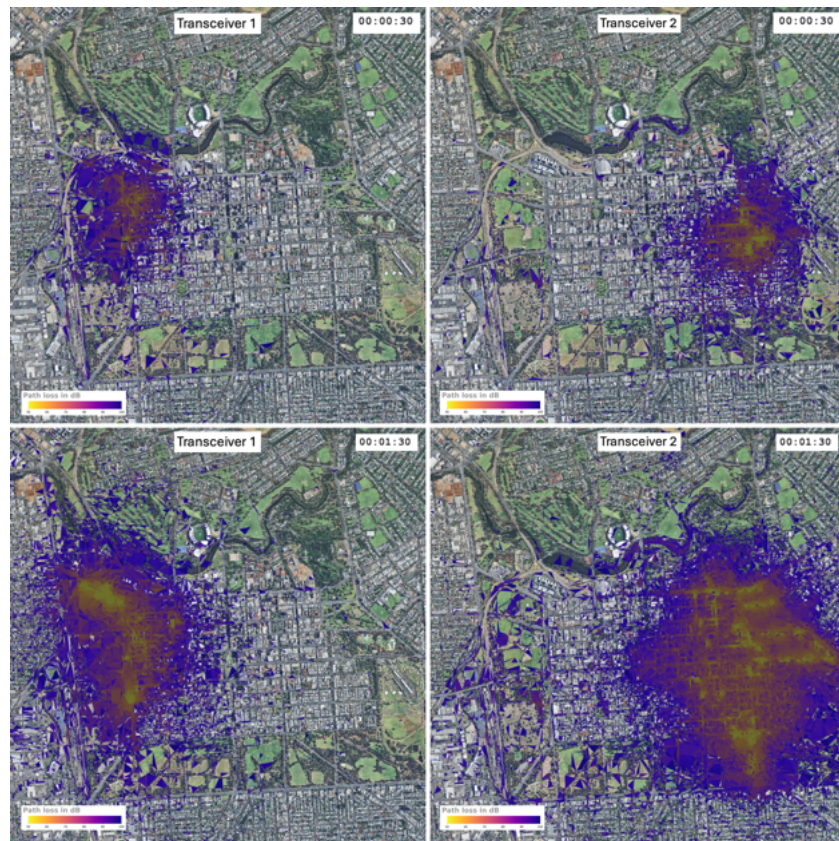


Figure 6.4.6.9: (top left and right) Bird's-eye view of the path loss (dB) for transceivers 1 and 2 combined with a driving SPT after 30 seconds from two initial locations in the Adelaide CBD; (bottom left and right) same for after 1 minute 30 seconds



Figure 6.4.6.10 shows the same scenario as Figure 6.4.6.9 after 2 minutes have passed of possible driving for each transceiver. In the plot below there are now signal-to-interference (SIR) values recorded as each transceiver is contributing power above a chosen threshold to these 3D meshes.

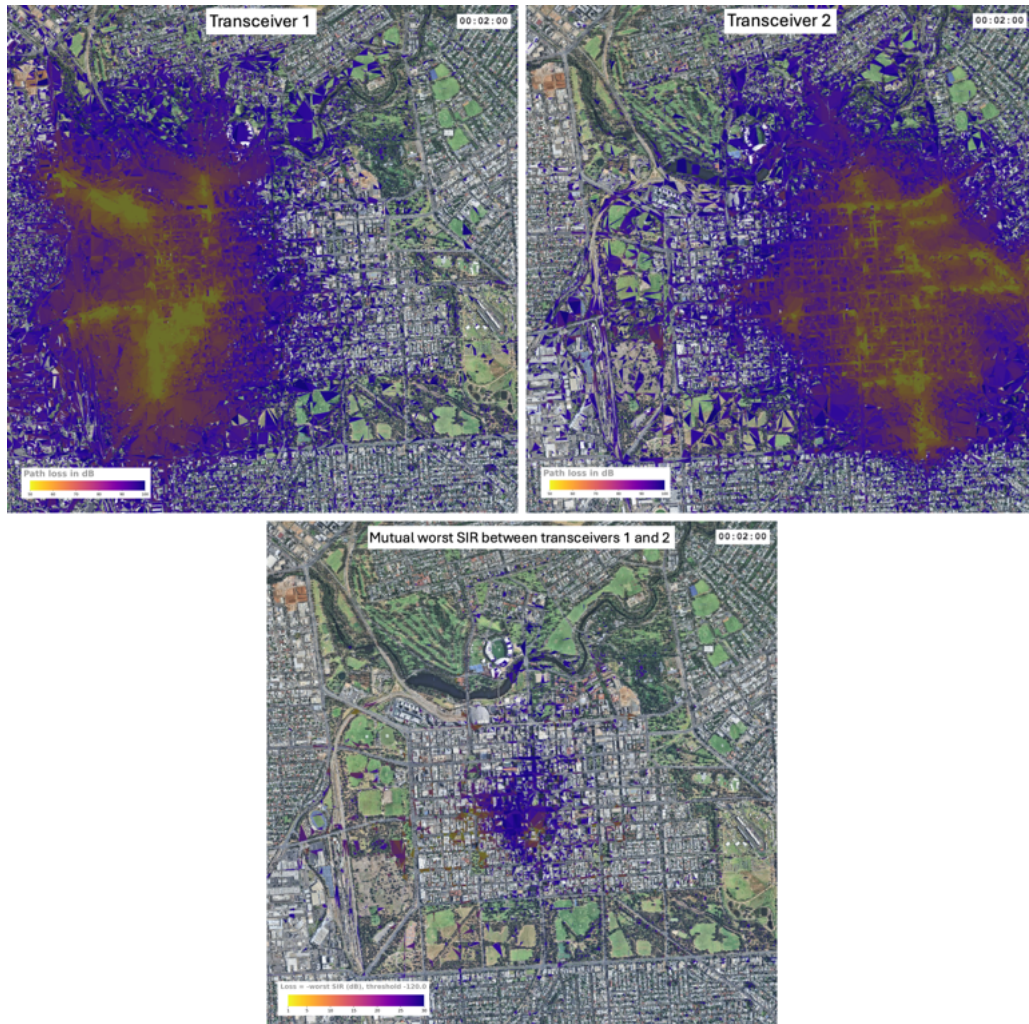


Figure 6.4.6.10: (top left and right) Bird's-eye view of the path loss (dB) for transceivers 1 and 2 combined with a driving SPT after 2 minutes have passed in the Adelaide CBD; (bottom) SIR for the two transceivers

Figure 6.4.6.11 shows the same scenario as Figure 6.4.6.9 and Figure 6.4.6.10 from 2:30 to 4:00 of possible driving in free flow traffic conditions, but only showing the mutual signal-to-interference (SIR) for the two transceivers.

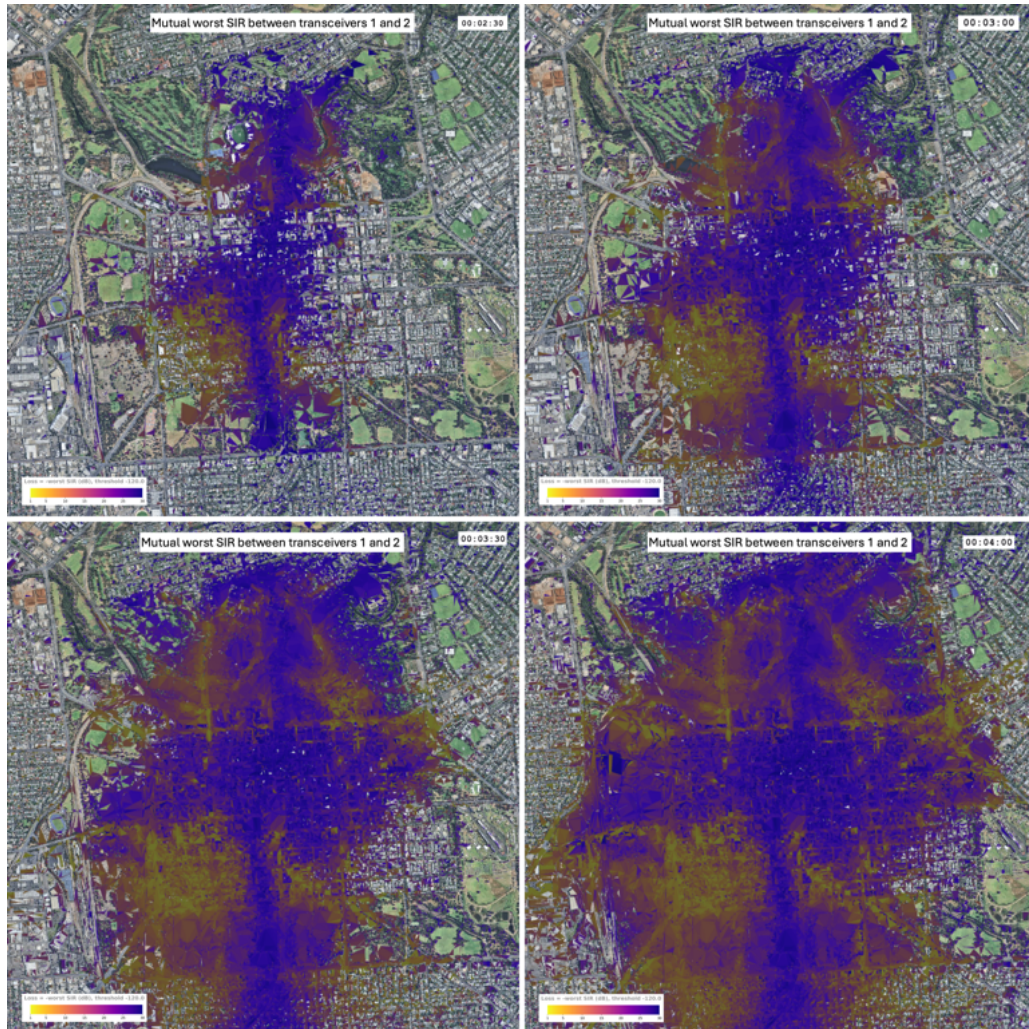


Figure 6.4.6.11: (top left) Bird's-eye view of the SIR between transceivers 1 and 2 after 2 minutes and 30 seconds have passed at 00:02:30; (top right) same at 00:03:00; (bottom left) same at 00:03:30; (bottom right) same at 00:04:00





Figure 6.4.6.12 shows the same as Figure 6.4.6.9 but from a lower altitude angled view from the north-west of the Adelaide CBD.

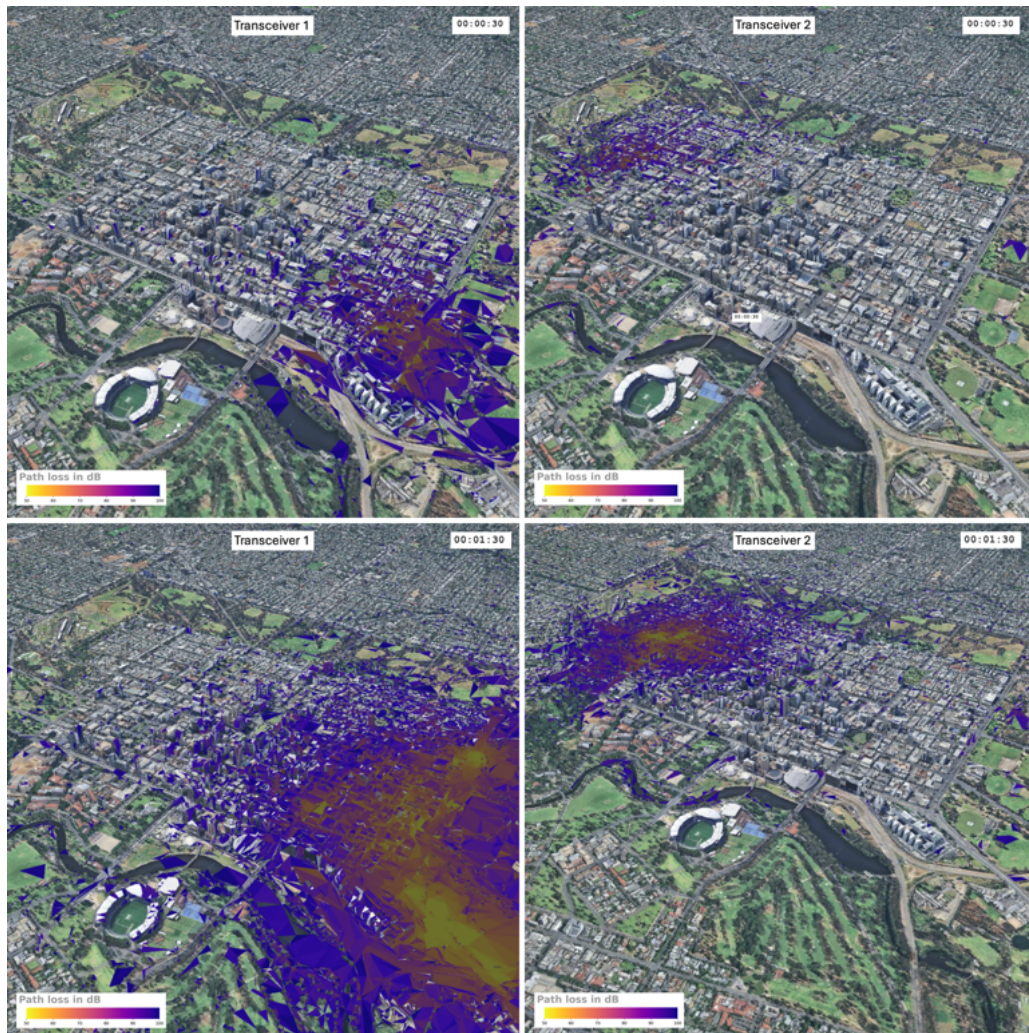


Figure 6.4.6.12: (top left and right) Angled view of the path loss (dB) for transceivers 1 and 2 combined with a driving SPT after 30 seconds from two initial locations in the Adelaide CBD; (bottom left and right) same for after 1 minute 30 seconds





Figure 6.4.6.13 shows the same as Figure 6.4.6.10 but from a lower altitude angled view from the north-west of the Adelaide CBD.

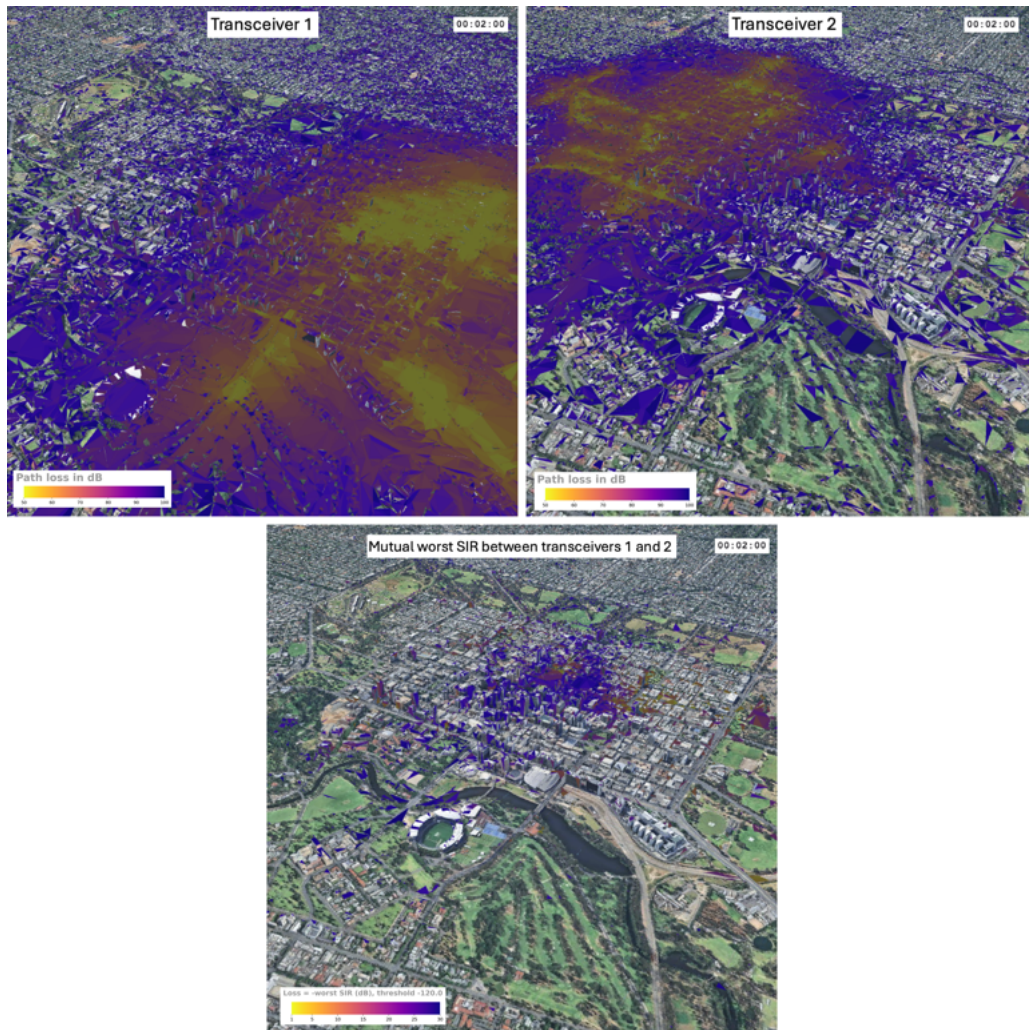


Figure 6.4.6.13: (top left and right) Angled view of the path loss (dB) for transceivers 1 and 2 combined with a driving SPT after 2 minutes have passed in the Adelaide CBD; (bottom) SIR for the two transceivers



Figure 6.4.6.14 shows the same as Figure 6.4.6.11 but from a lower altitude angled view from the north-west of the Adelaide CBD.

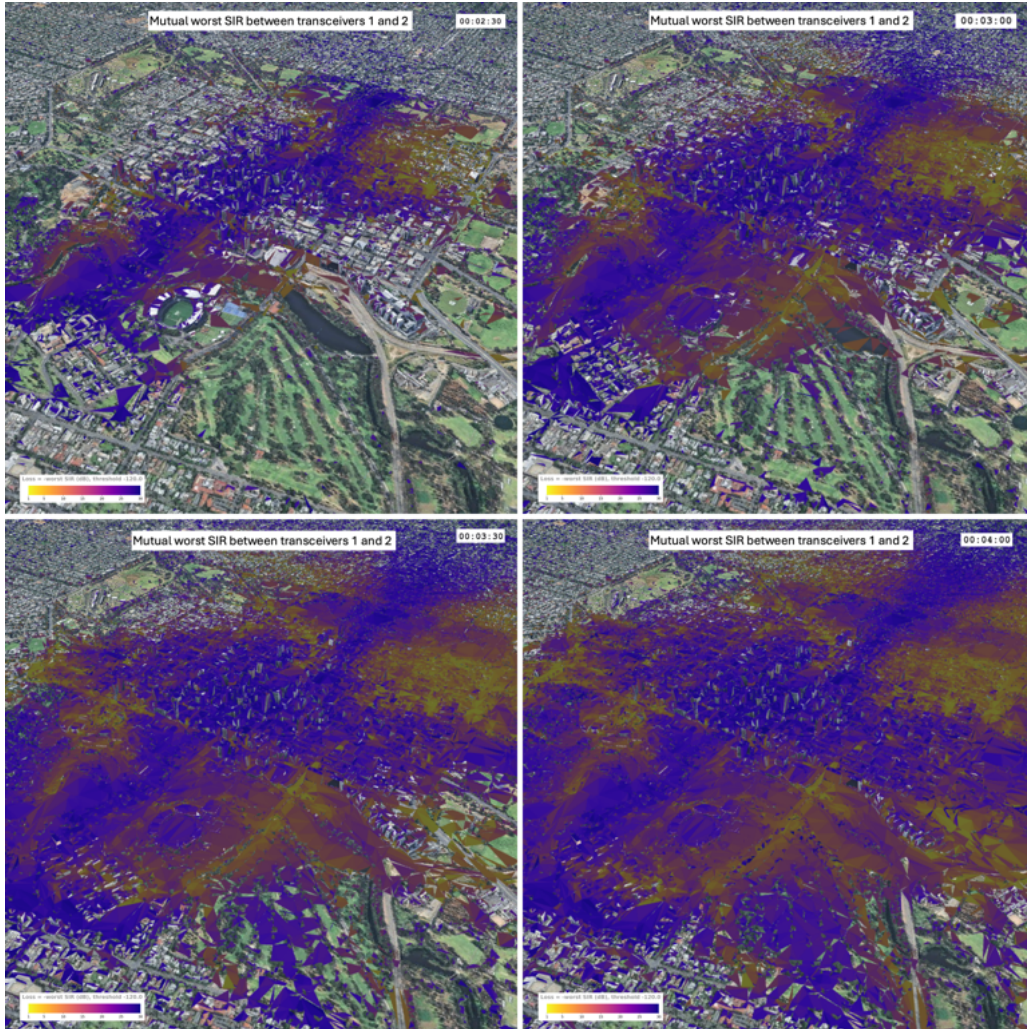


Figure 6.4.6.14: (top left) Angled view of the SIR between transceivers 1 and 2 after 2 minutes and 30 seconds have passed at 00:02:30; (top right) same at 00:03:00; (bottom left) same at 00:03:30; (bottom right) same at 00:04:00



#### Demo 4.4: SPT Driving – RF Ray Tracing – Rural

##### Setup:

- Movement: SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic ground antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The images below Figure 6.4.6.15 show the initial RF ray tracing path loss (dB) from transceivers 1 and 2 at an initial time as well as their mutual worst SIR.

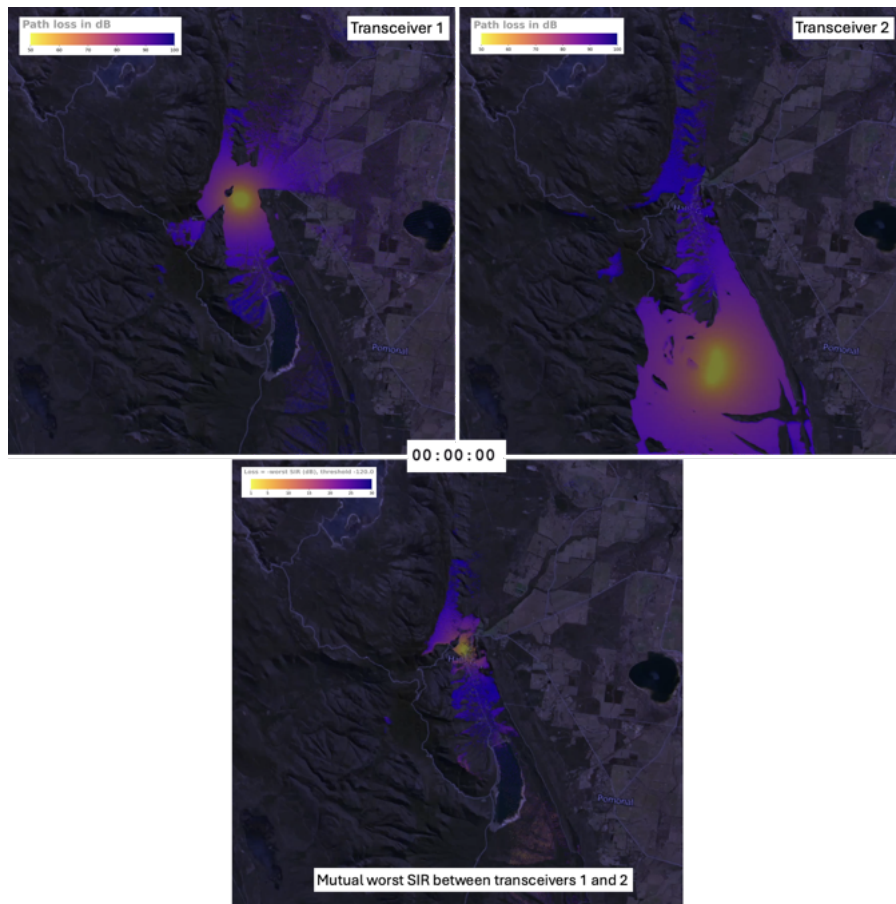


Figure 6.4.6.15: (top left and right) Hall's Gap at initial time 00:00:00 showing RF ray traced SPT path loss (dB) for transceiver 1 and 2; (bottom) Signal-to-interference ratio for transceiver 1 and 2 at 00:00:00



The images below Figure 6.4.6.16 show the RF ray tracing path loss (dB) from transceivers 1 and 2 for their reachable driving locations after 5 minutes as well as their mutual worst SIR.




Figure 6.4.6.16: (top left and right) Hall's Gap after 5 minutes at 00:05:00 showing RF ray traced SPT path loss (dB) for transceiver 1 and 2; (bottom) Signal-to-interference ratio for transceiver 1 and 2 at 00:05:00



The images below Figure 6.4.6.17 show the RF ray tracing path loss (dB) from transceivers 1 and 2 for their reachable driving locations after 10 minutes as well as their mutual worst SIR.




Figure 6.4.6.17: (top left and right) Hall's Gap after 10 minutes at 00:10:00 showing RF ray traced SPT path loss (dB) for transceiver 1 and 2; (bottom) Signal-to-interference ratio for transceiver 1 and 2 at 00:10:00

### 6.4.7 Demonstration 5: Coordinating Spectrum Access with Drones

The goal of this demonstration is to show how drone spectrum planning processes could be integrated with the previously shown system to coordinate multiple ground-based spectrum users using SIR as a measurement for co-interference. To begin, RF coverage modelling is shown to illustrate what it looks like for drones on the ground in an urban environment, a rural environment, and even in the sky from ground-users' RF transmissions. While the reachability of ground-based users is represented using a shortest path tree (SPT) in previous sections, the drone scenario considered is one where a flight path is to be approved ahead of time. Like a SPT, a drone flight path also describes locations that the spectrum user could be at in the future. This means that the RF coverage reachability can also be calculated for the drone flight path as a time-series heatmap of path loss (dB). This allows drone-to-drone co-interference analysis using the pre-existing SIR-based system. Since it is in the exact same format as for ground-based users, it is possible to do the SIR-based co-interference analysis with drones as well. A description of the demonstrations in this subsection is provided neatly in Table 6.3.

Demonstration	Description
<b>5 – Drones' Spectrum</b>	
5.1	RF coverage for drones on the ground for urban areas
5.2	RF coverage for drones on the ground for rural areas
5.3	RF coverage for drones in the sky
5.4	RF drone flight paths
5.5	Multi-drone co-interference
5.6	Drone-to-ground co-interference

Table 6.3: A description and layout of demonstrations about coordinated spectrum access for drones



## Demo 5.1: RF coverage for drones for urban areas

### Setup:

- Transmitter: Isotropic antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The image below Figure 6.4.7.1 shows a RF coverage map produced using ray tracing over Google 3D tiles data of the Adelaide CBD, using Sionna RT. The antenna with an isotropic radiation pattern was placed at a height and location represented by the yellow drone in the image. The purpose of this image was an initial demonstration of what a RF propagation ray traced coverage map for the ground looks like for a drone in the sky. The Adelaide CBD location was chosen as it has interesting RF propagation characteristics. In addition, flying drones in densely populated areas like this is generally restricted in some way, and there is also potentially a lot of RF interference to contend with too.



Figure 6.4.7.1: RF coverage for one drone in the Adelaide CBD above Victoria Square with an omnidirectional antenna



## Demo 5.2: RF coverage for drones for rural areas

### Setup:

- Transmitter: Isotropic antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The images below Figure 6.4.7.2 show the path loss (dB) for a drone near Hall's Gap in various locations viewed from various angles to demonstrate the RF environment transmission properties for a drone in a rural mountainous area.




Figure 6.4.7.2: Drone RF ray traced path loss (dB) around Hall's Gap

### Demo 5.3: RF coverage for drones in the sky

#### Setup:

- Transmitter: Isotropic antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from Google 3D tiles [152]
- Materials: Buildings treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The image below Figure 6.4.7.3 shows a RF coverage map produced for a horizontal reference plane in the sky, based on the result of RF ray tracing of the *rays* produced from an isotropic antenna located in Victoria Square, Adelaide 1.5 metres from the ground. The 3D model data used was processed from Google 3D tiles and the ray tracing engine was Sionna RT. Figure 6.4.7.3 shows where the antenna was placed as a green sphere, and a close up side view of the RF coverage map. A low altitude viewing angle is required to fully see what the coverage is due to challenges in displaying a complex 3D scene like this.




Figure 6.4.7.3: RF Coverage in sky for an isotropic transmitter Victoria Square



The image below Figure 6.4.7.4 shows the same scene as the above image Figure 6.4.7.3, but from afar. By showing it from a greater distance, a better understanding is gained of what an RF coverage map looks like for a dense CBD area with a ground-based transmitter surrounded by skyscrapers.




Figure 6.4.7.4: RF Coverage in sky for an isotropic transmitter Victoria Square from afar



The image below Figure 6.4.7.5 shows the same scene as the above image Figure 6.4.7.3 but from directly below the transmitter. It can be seen that the shape of the centre of Adelaide CBD is essentially printed into the RF radiation transmitted into the sky. This is expected as the LOS as well as the other propagation paths (reflections) have little where else to go and would essentially be funnelled into the sky.




Figure 6.4.7.5: RF Coverage in sky for an isotropic transmitter Victoria Square cross-sectional

## Demo 5.4: RF drone flight paths

### Setup:

- Movement: Example flight path of aerial coordinates
- Transmitter: Isotropic antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The image below Figure 6.4.7.6 shows the sum of all RF path gain values sampled at regular intervals at an altitude of 100 metres (above ground level). The regular values sampled correspond to locations on a drone flight path (conceptual). The resultant map shows all the areas that the drones may cause RF interference to during the flight. While this kind of RF mapping for multiple antennas is often used to determine how they may interfere with each other, they are in this case instead summed together because they represent the RF coverage of the same drone moving over time. The same data could be processed as animation frames to show where the RF interference would be at specific times. By examining the summed version of a drone flight path, it is possible to infer that a spectrum collision or conflict may occur when compared to other known spectrum usage. A time-series analysis can then be performed to confirm this and identify the exact locations and times that require adjustment.




Figure 6.4.7.6: RF Interference calculated for a drone flight plan at 100m from the ground across Adelaide





The image below Figure 6.4.7.7 provides a lower side view of Figure 6.4.7.6 as well as a city-only reference view for a simple comparison of the ray traced RF summed flight path map, and the what the 3D city model looks like from the same angle.




Figure 6.4.7.7: RF Interference calculated for a drone flight plan at 100m from the ground across Adelaide - lower side view

## Demo 5.5: Multi-drone co-interference

### Setup:

- Movement: Example flight path of aerial coordinates
- Transmitter: Isotropic antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The images below in Figure 6.4.7.8 show the mutual worst SIR between 5 drones (increased size for visibility) over mountainous areas near Hall's Gap as a snapshot in time 00:03:00 from 4 angles to help visualise the heatmap in the complex 3D environment.




Figure 6.4.7.8: Close angled view of 5 drone scenario around Hall's Gap showing the mutual worst signal-to-interference ratio all at  $t=3$  minutes shown from 4 different angles



The images below in Figure 6.4.7.9 show the mutual worst SIR between 5 drones (increased size for visibility) over mountainous areas near Hall's Gap across 5 points in time, viewed from a single viewing angle. It can be seen at  $t = 0$  minutes that when there is only one drone to the east (right of image) side of the scene, there is no mutual worst SIR measured there indicating that the signals transitted there are not being interfered with. Whenever two or more drones approach each other such as at  $t = 12$  minutes in the east (right of image), the hottest SIR is near the drones because the other drones' relatively strong signal near them is much weaker than that of the drone closest to it. In between the drones, the SIR is near 1.0 as the path loss (dB) from both is equal.




Figure 6.4.7.9: Far angled view of 5 drone scenario around Hall's Gap showing the mutual worst signal-to-interference ratio at  $t = 0, 4, 8,$  and  $12$  minutes



The images below in Figure 6.4.7.9 show the same scene and conditions as Figure 6.4.7.8, but from bird's-eye view.




Figure 6.4.7.10: Bird's-eye view of 5 drone scenario around Hall's Gap showing the mutual worst signal-to-interference ratio at  $t = 0, 4, 8, 12$  minutes

## Demo 5.6: Drone-to-ground co-interference

### Setup:

- Movement: Example flight path & SPT using Valhalla [120] Expansion [115]
- Transmitter: Isotropic antenna at 3.5 GHz with 100 MHz bandwidth
- Data: 3D models derived from SRTM 30m
- Materials: Terrain treated as solid concrete
- Propagation: Ray Tracing using Sionna RT 1.1 [174]

The images below in Figure 6.4.7.11 show the individual path loss (dB) heat maps for a car and a drone around Hall's Gap after 1 minute has passed. The movements of the car are fully based on reachability while the drone is following a future flight path. The idea is that the mutual worst SIR can be calculated ahead of time to assess the risk of future interference for a given drone flight path.




Figure 6.4.7.11: (top left and right) Bird's-eye view of Hall's Gap showing a drone flying and car's driving SPT combined with RF ray tracing-based path loss (dB) at time  $t = 1$  minute; (bottom) SIR is shown between the car and the drone at time  $t = 1$  minute





The images below in Figure 6.4.7.12 show the individual path loss (dB) heat maps for a car and a drone around Hall's Gap after 3 minutes have passed.




Figure 6.4.7.12: (top left and right) Bird's-eye view of Hall's Gap showing a drone flying and car's driving SPT combined with RF ray tracing-based path loss (dB) at time  $t = 3$  minutes; (bottom) SIR is shown between the car and the drone at time  $t = 3$  minutes



The images below in Figure 6.4.7.13 show the individual path loss (dB) heat maps for a car and a drone around Hall's Gap after 5 minutes have passed.




Figure 6.4.7.13: (top left and right) Bird's-eye view of Hall's Gap showing a drone flying and car's driving SPT combined with RF ray tracing-based path loss (dB) at time  $t = 5$  minutes; (bottom) SIR is shown between the car and the drone at time  $t = 5$  minutes



The images below in Figure 6.4.7.14 show the individual path loss (dB) heat maps for a car and a drone around Hall's Gap after 10 minutes have passed.




Figure 6.4.7.14: (top left and right) Bird's-eye view of Hall's Gap showing a drone flying and car's driving SPT combined with RF ray tracing-based path loss (dB) at time  $t = 10$  minutes; (bottom) SIR is shown between the car and the drone at time  $t = 10$  minutes

## 6.5 Discussion

### 6.5.1 Summary

Chapter 6 has four sections: *Introduction*, *Approach*, *Literature*, and *Demonstrations*.

The *Introduction* establishes a spectrum management digital twin that merges reachability with radio frequency (RF) propagation to predict interference during spectrum planning. A geospatial model of terrain, infrastructure, and conditions is applied to support time and location aware allocation across urban, rural, and vertical spaces. By linking shortest path tree (SPT) based reachability with propagation maps, measurable risk regions are produced to support coordinated access.

The *Approach* proposes integrating geometry and terrain with physically based propagation to model usage in a digital twin. Questions are raised about options for modelling radio spectrum in large urban and rural environments, and about how spectrum users' possible actions can be bounded and drone-to-drone and drone-to-ground links coordinated using SPT reachability and mutual worst SIR. Focus is placed on 3D city modelling for user representation and verticality, digital elevation model (DEM) for terrain influence, and real time coverage-informed allocation guided by propagation and SPT.

The *Literature* explores the two focus areas. Applications that provide propagation modelling for localisation and characterisation while avoiding large visual datasets are highlighted, and cross domain utility in predictive digital twins is shown. For drones, reliance on shared bands that face interference is noted, and emerging strategies in DSM and waveform design are reported as ways to meet reliability and scalability needs. Benefits are identified as predictive accuracy and planning flexibility, with limits imposed by data fidelity, computational cost, and regulatory constraints.

The *Demonstrations* progress from isolated reachability and propagation to their combination, then to multi user coordination and drone scenarios. Urban ray tracing reveals how the shape of a city can cause RF energy to form a diamond grid shape, while rural results are dominated by terrain, and time series heatmaps track evolving mutual worst SIR. Flight paths treated as moving reachability support drone-to-drone and drone-to-ground assessment across altitude. Noted constraints include routing and mesh coverage, with future work on including life patterns, past location data, improved propagation modelling, and apply the findings to implement tighter risk aware scheduling.



## 6.5.2 Key Findings

### Finding 1: Integration of Reachability and Propagation

Combining reachability modelling with RF propagation enables accurate prediction of interference zones. The *Approach* frames this as a step beyond conventional 2D spectrum allocation by explicitly incorporating mobility constraints and environmental factors. The *Literature* situates this within digital twin research, highlighting scalable models that integrate realistic urban landscapes and population movement. *Demonstrations* show how shortest path tree (SPT) models, when combined with ray tracing, reveal interference patterns shaped by city grids and terrain. The implication is that interference risk can now be quantified spatiotemporally, providing applications for real-time or predictive spectrum management.

### Finding 2: Multi-user Spectrum Risk Quantification

Coordinated allocation across multiple users is feasible through mutual worst signal-to-interference ratio (SIR) analysis. The *Approach* identifies this as a means to anticipate allocation conflicts over time rather than reactively. The *Literature* supports the method by categorising advanced UAV spectrum management into deterministic, opportunistic, and competitive frameworks, all of which require quantifiable interference metrics. *Demonstrations* show SPT-driven simulations where co-interference is visualised across time in both urban and rural settings. The implication is that spectrum coordination can be transformed into a predictable and optimisable process rather than an uncertain negotiation.

### Finding 3: Vertical Modelling for Drone Coordination

Vertical spectrum modelling significantly improves management of drone-to-drone and drone-to-ground communications. The *Approach* frames this by asking how 3D spectrum modelling could extend allocation practices to altitude-dependent users. The *Literature* supports this with studies that propose dynamic spectrum access for unmanned aerial vehicle (UAV) networks, emphasising flexibility and interference resilience. *Demonstrations* present simulated results by mapping RF coverage for drones in urban, rural, and aerial environments, and by quantifying mutual worst SIR across drone flight paths. The implication is that height-aware (or vertical) allocation is a necessary condition for scalable UAV spectrum access.

### 6.5.3 Limitations

While the integration of reachability and RF propagation modelling demonstrates improved accuracy in predicting interference zones, several limitations remain. First, the fidelity of the models is highly dependent on the underlying digital elevation model and 3D city mesh resolution. Incomplete or outdated geographic data may reduce predictive accuracy, especially in dynamic urban environments where infrastructure changes rapidly. Second, the computational overhead of combining SPT calculations with detailed ray tracing restricts scalability for large regions or long time intervals. This creates a trade-off between accuracy and processing time, potentially limiting the approach to offline or early planning rather than continuous real-time coordination. Finally, the use of idealised mobility assumptions such as free-flow traffic conditions or fixed walking speeds may not capture real-world conditions well, meaning that interference regions quantified in the simulation may differ from those encountered in reality.

The quantification of spectrum allocation conflicts through mutual worst signal-to-interference ratio (SIR) analysis provides a structured framework for coordination, but it also presents constraints. The method assumes that interference risk can be effectively reduced to a single measurable SIR metric, which may oversimplify scenarios where other factors such as latency sensitivity or waveform diversity influence user experience. Additionally, the current demonstrations treat users as having uniform transmission characteristics, but in real world use cases, different kinds of devices and adaptive protocols may complicate mutual interference estimation. The SIR-based approach also relies on the assumption that spectrum users would disclose honestly or predict their future behaviour accurately. In practice, incomplete knowledge of mobility patterns or strategic misreporting in competitive environments could limit the method's reliability. These challenges highlight that while SIR analysis enhances predictability, it requires supplementary mechanisms for robustness against uncertain user behaviour and devices.

Vertical modelling improves unmanned aerial vehicle (UAV) spectrum coordination but carries its own limitations. Current simulations assume that flight paths are either pre-planned or can be reliably forecast, but real drone operations are often subject to weather, regulatory rerouting, or autonomous adaptation. This uncertainty constrains the applicability of pre-computed interference maps, especially in dense urban airspaces with high drone density. Moreover, the ray tracing models used to evaluate drone-to-drone and drone-to-ground links could face computational limitations when considering so many devices across such a wide area. Another limitation arises from the lack of integration with regulatory frameworks, which may restrict drones from accessing licensed spectrum at altitude, making the scenarios unrealistic in many real world conditions. Such policies were not mentioned or addressed in the demonstrations, and would need to be considered in future work to make a stronger case for the applicability of this work today. While height-aware modelling represents an essential advancement in spectrum management, its adoption depends equally on improvements in predictive modelling, computational optimisation, and a possibility of regulations and policies adapting.

#### 6.5.4 Future Work

Further research could enhance RF propagation and reachability models within the digital twin. The demonstrated integration of SPT modelling with RF propagation established a solid foundation, although accuracy is currently limited by mesh resolution, terrain accuracy, and idealised traffic assumptions. Extending the digital twin with adaptive RF propagation engines that blend ray tracing, empirical datasets, and machine-learned correction factors could improve realism while reducing computational overhead. Life pattern modelling and the incorporation of historic movement records would refine interference risk assessment beyond static or uniform assumptions, producing tighter and more practical spectrum allocations.

Expansion of spectrum risk quantification methods is another priority. While the mutual worst signal-to-interference ratio metric provided a consistent benchmark, its simplification of diverse performance constraints does not capture the full complexity of real-world systems. Future work could incorporate multi-metric evaluation into the digital twin, including latency sensitivity, adaptive waveforms, and expected results and performance for various device types. Spectrum demand models from earlier chapters would be directly linked to allocation experiments, enabling automated testing of decentralised spectrum access coordination. Including competitive and opportunistic allocation strategies in the simulated environment could demonstrate their scalability and robustness under contested conditions.

Vertical spectrum modelling requires deeper research to fully support drone coordination. Current demonstrations assume predictable or pre-planned flight paths, although in reality drone operations are subject to adaptive routing, weather influences, and regulatory constraints. Future research could enhance the digital twin with realistic drone behaviour models, integrating air traffic management logic and simulating high-density UAV operations. Decentralised and hierarchical management concepts such as MANET could then be tested in the simulated environment, demonstrating how airborne and terrestrial systems could negotiate access dynamically. This approach could validate proposed coordination mechanisms while also stress-testing regulatory assumptions, linking technical feasibility with policy adaptability for altitude-aware spectrum management.

## 6.6 Conclusion

This chapter demonstrated a digital twin model applied to a spectrum management problem. While the demonstration shows off a very advanced definition and application of this work, it barely scratches the surface for it could be used for. The results are also highly generalisable, so some additional applications are provided in the next section. Taken together, the chapter shows that integrating reachability with RF propagation provides spatiotemporal risk estimates that enable proactive coordination across urban, rural and vertical spaces. It also establishes a mechanism for multi user allocation through a mutual worst signal to interference ratio that turns coordination from reactive to predictable planning. By incorporating altitude, it shows the value of height aware planning for drone to drone and drone to ground links and motivates vertical spectrum management. Overall the work provides a realistic geospatial basis for allocation decisions before deployment and a template that delivers accurate and explainable assessments suitable for scaling.



# Additional Applications

## Digital Forensics

Digital forensics is the application of scientific methods and technical processes to analyse digital forms of evidence, with telecommunications playing a central role. Mobile devices are deeply integrated into daily life, serving as tools for communication, navigation, and online access, and therefore are almost always carried during activities, including criminal acts. As a result, mobile network data provides a powerful evidentiary source. While call data records reveal which towers were connected and when, they do not directly show exact device locations. This gap drives the use of modelling techniques to reconstruct where a device may have been, combining telecommunications data with external information sources to build a more complete forensic picture.

Where digital twins for spectrum management often project forward, forensics requires historical digital twins that reconstruct the conditions of a telecommunications system at some past moment. This process can be highly complex, as it involves recovering information about legacy networks, devices, and environmental factors that may no longer exist. For example, a forensic digital twin may need to account for the location, configuration, and technology of base stations in an earlier era, shifts in land use such as urban development or deforestation, and historical weather or spectrum usage that could have influenced signal behaviour. Constructing such twins allows analysts to re-simulate the past in a scientifically defensible way, providing courts with reconstructions grounded in telecommunications engineering.

Reachability modelling and RF propagation analysis are key techniques for narrowing possible device locations in forensic contexts. By combining estimated cell coverage areas with models of where a suspect could have travelled within given time constraints, analysts can test alibis or reconstruct possible movements. Unlike conventional spectrum management, forensic analysis may need to consider obsolete network technologies, old spectrum allocations, and the quirks of older devices that no longer operate in today's networks. These added dimensions make forensic digital twins unique. They must integrate technical reconstruction with historical verification, ensuring that conclusions reflect the actual possibilities of the time rather than present-day assumptions. Two real case studies are provided below to demonstrate how concepts and tools from this thesis have been applied in practice.

## Case Study 1 - Cold Case around Picton, New South Wales

Case study 1 [192] takes historical telecommunications analyses to the extreme by examining phone records from 2001, roughly 24 years before the analysis was conducted. The task required reconstructing the telecommunications environment of that time, a process that introduced a range of challenges. It was necessary to determine where the base station was located in 2001, what spectrum and technology were used, how the antennas were configured, and how the local geography and land use had changed since then. Satellite imagery comparisons between 2002 and 2023 (shown on the right of Figure 6.6.0.2) helped identify shifts in tower placement, with tectonic plate movement, poorly aligned older satellite imagery, and low-resolution historical imagery introducing further complications. These steps illustrate the depth of effort required to recover past network conditions in order to generate a reliable historical digital twin.

Beyond tower placement, the case be extended by performing an analysis of the broader environment at the time of the records. This could include seasonal conditions, the weather on the days in question, and the state of the RF environment due to other spectrum users. Line of sight (LOS) RF propagation modelling in 3D coverage simulations (shown on the left of Figure 6.6.0.2) was used to reconstruct how signals would have propagated across the rugged landscape, including the Razorback Range. By incorporating these factors into the forensic twin, it was possible to assess whether a device connecting to the Picton tower in 2001 could plausibly have been at a particular location. This demonstrates how historical digital twins, combined with RF analysis, enable a scientifically rigorous reconstruction of past events.



Figure 6.6.0.1: (left) Line of Sight (LOS) RF analysis for Picton Base Station in 2001 in 3D looking north and down at a height; (right) Satellite imagery comparisons between 2002 and 2023 of the same Picton base station site



## Case Study 2 - Investigation around Leongatha, Victoria

Compared to Case Study 1 in Picton, Case Study 2 [193] is far less historical, going back only a few years from the time the analysis was conducted. This significantly reduced uncertainty, as most details of the base stations involved could be verified directly from contemporary records. The investigation still required careful examination of call data records and the corresponding network coverage, but the reduced temporal gap simplified the process of validating tower locations, technologies, and configurations. By confirming these parameters against current regulator databases and operator information, it was practical to construct a verifiably accurate historical digital twin.

With the historical mobile phone network environment well understood, the analysis focused on applying LOS RF propagation modelling to reconstruct coverage during the times of interest. Reachability analysis was then used to test which movements could realistically have occurred during the time windows in question, offering insights into the suspect's possible presence or absence in specific locations. Although the reconstruction of the forensics digital twin was less complex than the Picton case, the Leongatha investigation highlights the routine utility of digital twins for forensics. Even in relatively recent cases, they provide a structured, defensible method to assess telecommunications evidence, travel constraints, and the assess plausibility of alibis.

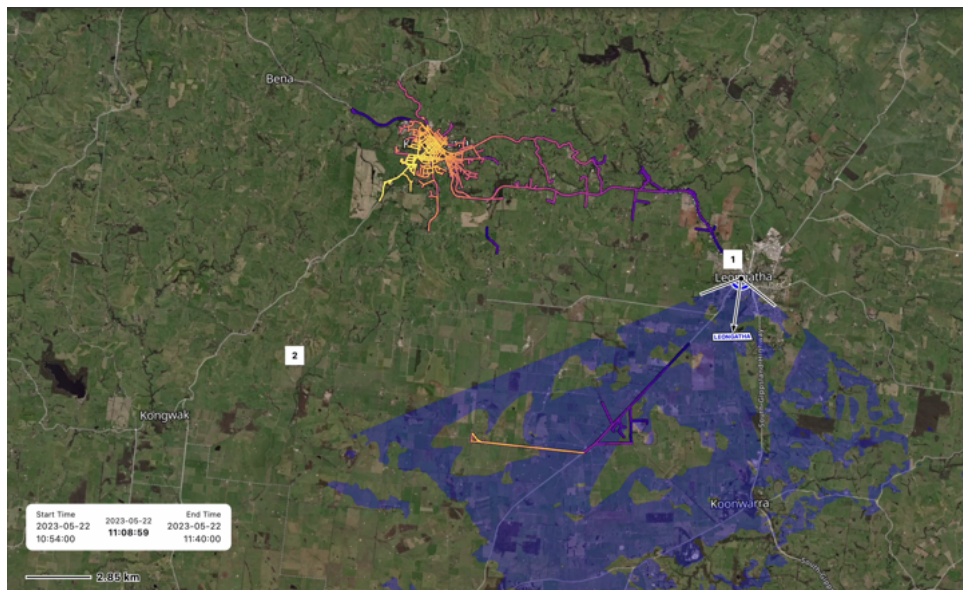


Figure 6.6.0.2: Line of Sight (LOS) RF analysis for time interval of interest - 2023-05-22 11:08:00

## Space Spectrum

Space applications could benefit from RF propagation analysis across varying environments through the atmosphere, in orbit, and in interplanetary links. Digital twins of the RF environment can enable realistic modelling of satellite communication conditions under dynamic factors such as orbital position, beam steering, and atmospheric disturbances. Ray tracing in a digital twin could help assess interference between different constellations, optimise spectrum allocation, and predict reliability for Earth-to-space and space-to-space links. Extending these methods to interplanetary communications would add further value, allowing scenario testing for latency, degradation, and carrier-to-noise ratios that shape mission feasibility.

A key objective in satellite spectrum allocations is to assess whether introducing a new SAT-COM spectrum allocation would interfere with existing ones. This is currently addressed by the International Telecommunications Union (ITU), which ensures that any new allocation maintains a significantly lower carrier-to-noise ratio (CNR) ratio to prevent disruption. However, there are many mathematical and physics factors to consider that are known to vary due to influences such as a varying atmosphere that are difficult to model and account for in real-time. RF Ray tracing has obvious applications to solve this problem, with a demonstration of this provided in Figure 6.6.0.3 where Sionna RT RF ray tracing is applied to a low resolution Earth 3D model from an antenna placed at the location of a conceptual example satellite.




Figure 6.6.0.3: RF ray tracing render

## Maritime Spectrum

Maritime digital twins provide a way to simulate shipping lanes, sea conditions, and vessel movements while incorporating RF propagation and reachability analysis. These models make it possible to forecast coverage, manage interference, and coordinate spectrum use between ships and coastal stations. The clustering of ships, as shown in the left panel of Figure 6.6.0.4, illustrates how groups of vessels can be treated as cooperative spectrum users, while the right panel shows a basic example where an allocation algorithm distributes channels to minimise interference and maintain fair access.

The TRITON: high-speed maritime wireless mesh network project by Zhou et al. [194] used real ship traffic data to study the feasibility of mesh-based ship-to-ship communication in dense shipping lanes. It showed how connectivity could be extended through IEEE 802.16 mesh links, with middleware allowing fallback to satellite links when ships were too sparse or too far off-shore. While TRITON itself was not a digital twin, its methods for modelling vessel density, mesh connectivity, and fallback conditions provide inputs that a maritime digital twin could integrate to explore resilient communication strategies.

Other research has focused on the broader design of maritime communication systems. Surveys of hybrid satellite-terrestrial maritime networks by Wei et al. [195] examined how satellite, coastal, and shipborne systems can be combined to overcome limited shore-based coverage and to adapt to atmospheric and sea-state variations. Similarly, research on dynamic spectrum allocation for maritime cognitive radio systems by Zhang et al. [196] developed a queuing-based algorithm to prioritise high-value services while maintaining throughput for general traffic. These approaches, though not digital twins themselves, provide techniques that can be embedded within a digital twin to test hybrid architectures, spectrum allocation policies, and resilience under realistic maritime conditions.

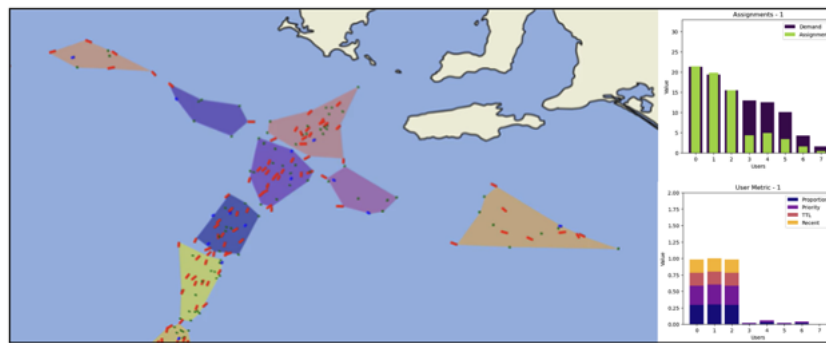


Figure 6.6.0.4: (left) Ship clusters used to coordinate intra-spectrum usage; (right) allocation of spectrum channels between ships using a dynamic spectrum management approach



## Future Work

Future work could first focus on advancing digital twins by integrating richer, more diverse data sources and scaling models to global coverage with high computational power. These ultra realistic simulations could enable deeper insight into large-scale spectrum allocation problems. However, the long-term aim is not centralised heavy computation but lightweight local modelling deployed at the edge. Such decentralised, real-time solutions in the field could allow decisions to be made closer to the point of action, reducing latency and improving responsiveness under dynamic conditions.

Another priority could be to extend decentralisation research, particularly the performance of MANETs under realistic operating conditions. Earlier work identified their promise but lacked a true-to-life testing environment. The digital twin environment developed later in the thesis now makes it possible to simulate devices in realistic propagation conditions, accounting for physical constraints, interference patterns, and device-specific properties. This could open the way to systematically evaluate how decentralised nodes communicate, coordinate, and self-organise in the field, producing insights that are both scalable and transferable to real deployments.

Finally, with realistic models and decentralised environments in place, attention could return to the original spectrum allocation challenge. The question of how best to divide and assign spectrum is now tractable in simulations that account for both regulatory frameworks and environmental constraints. By incorporating real-world propagation, interference, and demand variability into models, candidate allocation strategies could be stress-tested under conditions that mirror practice. This provides a pathway to identifying allocation mechanisms that balance fairness, efficiency, and compliance, while remaining viable for large-scale, dynamic deployment.



# Conclusion

This thesis has explored spectrum allocation solutions to address the growing scarcity of RF spectrum access, driven by the rapid expansion of wireless technologies. The main contribution was to define the allocation problem under realistic conditions and demonstrate how simulated realism can refine problem definitions. By identifying critical subproblems and examining allocation strategies, decentralisation, and clustering methods, the work highlights pathways for overcoming the static nature of spectrum management. A key barrier remains around incumbency and the difficulty of managing access across large areas, but the findings indicate that DSM may be feasible once realism is applied.

The challenge of coordinating users without harmful interference is shown to be complex and time-sensitive. Early approaches such as life pattern modelling, clustering, and hierarchical management strategies provided valuable insights, though they lacked sufficient realism to verify their performance under realistic conditions. The central limitation was the absence of advanced RF propagation models capable of capturing the impact of incumbents and environmental conditions. By introducing the concept of a digital twin, advanced propagation modelling, and reachability analysis, the thesis demonstrated how more sophisticated simulations can provide the foundation for testing coordinated spectrum access in realistic, sandbox-like environments.

The research presented here is an early step in applying digital twins to spectrum management and other large-scale, time-sensitive geospatial problems. It shows how realistic problem definitions can be filtered and tested through scalable simulations. It contributes methods for modelling synthetic demand, evaluating fairness through spatial allocation strategies (such as Voronoi), and testing decentralised clustering and mobility-aware coordination. Finally, it demonstrates how advanced RF propagation in 3D digital twins provides a foundation for interference-aware, multi-user, and vertical spectrum management. Demonstrations of mobility-aware allocation, mutual signal-to-interference (SIR) analysis, and vertical spectrum management illustrate how digital twins can extend spectrum access beyond static two-dimensional planning. Future applications may include forensics, maritime, space spectrum management, and further testing of decentralisation and MANET-based spectrum management within realistic digital twin environments. As artificial intelligence (AI) advances, powered by





the intense global interest arising from the large language model (LLM) technological revolution, its effectiveness in the spectrum management domain may depend on realistic problem representation and evaluation methods. Once these are in place, the path opens towards truly *dynamic* spectrum management.



# Bibliography

- [1] International Telecommunication Union (ITU). Radio regulations, article 5 – frequency allocations. <http://handle.itu.int/11.1002/pub/814b0c44-en>, 2020. Article 5 distinguishes PRIMARY services (names in capitals, with priority) and secondary services (lowercase).
- [2] Madhulika Tripathi. A dynamic spectrum access framework (bring your own spectrum). PhD thesis, Electrical and Electronic Engineering, 2020. URL <http://hdl.handle.net/2440/130173>.
- [3] Laraib Abbas, Umar Shoaib, and Ali Kashif Bashir. Priority based dynamic spectrum management using virtual utility functions in cognitive radio enabled internet of things. Computer Communications, 196:239–248, 2022. ISSN 0140-3664. doi: <https://doi.org/10.1016/j.comcom.2022.10.002>. URL <https://www.sciencedirect.com/science/article/pii/S0140366422003802>.
- [4] Nadine Abbas, Youssef Nasser, and Karim El Ahmad. Recent advances on artificial intelligence and learning techniques in cognitive radio networks. EURASIP Journal on Wireless Communications and Networking, 2015(1):174, 2015. ISSN 1687-1499. doi: [10.1186/s13638-015-0381-7](https://doi.org/10.1186/s13638-015-0381-7). URL <https://doi.org/10.1186/s13638-015-0381-7>.
- [5] J. A. Stine and David Portigal. Spectrum 101: An introduction to spectrum management. page 220, 2004.
- [6] Jon M Peha. Spectrum management policy options. IEEE Communications Surveys, 1(1):2–8, 1998.
- [7] Peter Anker. From spectrum management to spectrum governance. Telecommunications Policy, 41(5-6):486–497, 2017.
- [8] Martin Cave, Christopher Doyle, and William Webb. Essentials of modern spectrum management. Cambridge University Press Cambridge, 2007.
- [9] Z. Luo and S. Zhang. Dynamic spectrum management: Complexity and duality. IEEE Journal of Selected Topics in Signal Processing, 2(1):57–73, 2008. ISSN 1941-0484. doi: [10.1109/JSTSP.2007.914876](https://doi.org/10.1109/JSTSP.2007.914876).

- [10] O. P. Babalola, O. O. Ogundile, and A. A. Owoade. Channel selection and switching methods in dynamic spectrum networks: A survey. TASUED Journal of Pure and Applied Sciences, 2(1):64–82, 2023. URL <http://journals.tasued.edu.ng/index.php/tjopas/article/view/8>.
- [11] Mathias Bohge. Dynamic Resource Allocation in Packet-Oriented Multi-Cell OFDMA Systems. Doctoral thesis, 2010. URL <http://dx.doi.org/10.14279/depositonce-2665>.
- [12] M. Suchanski, R. Matyszekiel, P. Kaniewski, M. Kustra, P. Gajewski, and J. Łopatka. Dynamic spectrum management as an anti-interference method, volume 10418 of Conference on Reconnaissance and Electronic Warfare Systems. SPIE, 2017. URL <https://doi.org/10.1117/12.2269294>.
- [13] A. Kaur and K. Kumar. Imperfect csi based intelligent dynamic spectrum management using cooperative reinforcement learning framework in cognitive radio networks. IEEE Transactions on Mobile Computing, pages 1–1, 2020. ISSN 1558-0660. doi: 10.1109/TMC.2020.3026415.
- [14] J. Kim, G. Casati, A. Pietrabissa, A. Giuseppi, E. Calvanese Strinati, N. Cassiau, G. Noh, H. Chung, I. Kim, M. Thary, J. M. Houssin, F. Pigni, S. Colombero, P. Dal Zotto, L. Raschkowski, and S. Jaeckel. 5g-allstar: An integrated satellite-cellular system for 5g and beyond. In 2020 IEEE Wireless Communications and Networking Conference Workshops (WCNCW), pages 1–6. doi: 10.1109/WCNCW48565.2020.9124751.
- [15] N. Cassiau, I. Kim, E. C. Strinati, G. Noh, A. Pietrabissa, F. Arnal, G. Casati, T. Choi, Y. J. Choi, H. Chung, S. Colombero, P. D. Zotto, E. D. Santis, J. B. Doré, A. Giuseppi, J. M. Houssin, J. Kim, M. Laugeois, F. Pigni, X. Popon, L. Raschkowski, M. Thary, and S. H. Won. 5g-allstar: Beyond 5g satellite-terrestrial multi-connectivity. In 2022 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit), pages 148–153. ISBN 2575-4912. doi: 10.1109/EuCNC/6GSummit54941.2022.9815664.
- [16] Ying-Chang Liang. Dynamic Spectrum Management : From Cognitive Radio to Blockchain and Artificial Intelligence. Springer Nature, 2020. doi: 10.1007/978-981-15-0776-2.
- [17] Joseph Mitola. Cognitive radio an integrated agent architecture for software defined radio. In Cognitive Radio An Integrated Agent Architecture for Software Defined Radio, 2000. URL <https://api.semanticscholar.org/CorpusID:61207536>.
- [18] Xin Tian, Zhi Tian, Khanh Pham, Erik Blasch, and Dan Shen. Jamming/anti-jamming game with a cognitive jammer in space communication, volume 8385 of SPIE Defense, Security, and Sensing. SPIE, 2012. URL <https://doi.org/10.1117/12.918798>.



- [19] Ahmed A. Thabit and Hadi T. Ziboon. Design and implementation of a new detection system based on statistical features for different noisy channels, 2019. URL <https://doi.org/10.1145/3321289.3321294>.
- [20] Yuvaraja Teekaraman, Hariprasath Manoharan, Adam Raja Basha, and Abirami Manoharan. Hybrid optimization algorithms for resource allocation in heterogeneous cognitive radio networks. *Neural Processing Letters*, 55(4):3813–3826, 2023. ISSN 1573-773X. doi: 10.1007/s11063-020-10255-2. URL <https://doi.org/10.1007/s11063-020-10255-2>.
- [21] V. Noel Jeygar Robert and K. Vidya. Genetic algorithm optimized fuzzy decision system for efficient data transmission with deafness avoidance in multihop cognitive radio networks. *Journal of Ambient Intelligence and Humanized Computing*, 14(2): 959–972, 2023. ISSN 1868-5145. doi: 10.1007/s12652-021-03349-9. URL <https://doi.org/10.1007/s12652-021-03349-9>.
- [22] S. D. Barnes and B. T. Maharaj. Prediction based channel allocation performance for cognitive radio. *AEU - International Journal of Electronics and Communications*, 68(4): 336–345, 2014. ISSN 1434-8411. doi: <https://doi.org/10.1016/j.aeue.2013.09.009>. URL <https://www.sciencedirect.com/science/article/pii/S1434841113002604>.
- [23] Hamid Eltom, Kandeepan Sithamparanathan, Rob Evans, Ying Liang, and Branko Risti. Statistical spectrum occupancy prediction for dynamic spectrum access: a classification. *EURASIP Journal on Wireless Communications and Networking*, 2018:29, 2018. doi: 10.1186/s13638-017-1019-8.
- [24] G. Ding, F. Wu, Q. Wu, S. Tang, F. Song, A. V. Vasilakos, and T. A. Tsiftsis. Robust online spectrum prediction with incomplete and corrupted historical observations. *IEEE Transactions on Vehicular Technology*, 66(9):8022–8036, 2017. ISSN 1939-9359. doi: 10.1109/TVT.2017.2693384.
- [25] G. Ding, Y. Jiao, J. Wang, Y. Zou, Q. Wu, Y. Yao, and L. Hanzo. Spectrum inference in cognitive radio networks: Algorithms and applications. *IEEE Communications Surveys & Tutorials*, 20(1):150–182, 2018. ISSN 1553-877X. doi: 10.1109/COMST.2017.2751058.
- [26] X. Li, G. Chen, Y. Xu, X. Wang, and T. Song. Recovering missing values from corrupted historical observations: Approaching the limit of predictability in spectrum prediction tasks. *IEEE Access*, 8:180379–180393, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.3027012.
- [27] X. Li, X. Wang, T. Song, and J. Hu. Robust online prediction of spectrum map with incomplete and corrupted observations. *IEEE Transactions on Mobile Computing*, pages 1–1, 2021. ISSN 1558-0660. doi: 10.1109/TMC.2021.3081715.



- [28] J. Yang and H. Zhao. Enhanced throughput of cognitive radio networks by imperfect spectrum prediction. IEEE Communications Letters, 19(10):1738–1741, 2015. ISSN 1558-2558. doi: 10.1109/LCOMM.2015.2442571.
- [29] J. Zhang, Y. Chen, Y. Liu, and H. Wu. Spectrum knowledge and real-time observing enabled smart spectrum management. IEEE Access, 8:44153–44162, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.2978005.
- [30] J. Zhang, Y. Liu, C. Yong, and L. Cao. Spectrum opportunity exploration based on joint spectrum sensing and the spectrum prediction enabled spectrum database. In 2016 IEEE MTT-S International Wireless Symposium (IWS), pages 1–5. doi: 10.1109/IEEE-IWS.2016.7585415.
- [31] M. Tonnemacher, C. Tarver, V. Chandrasekhar, H. Chen, P. Huang, B. L. Ng, J. Charlie Zhang, J. R. Cavallaro, and J. Camp. Opportunistic channel access using reinforcement learning in tiered cbrs networks. In 2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN), pages 1–10. ISBN 2334-3125. doi: 10.1109/DySPAN.2018.8610474. Covers DSM and RL, but not very relevant.
- [32] International Telecommunications Satellite Organization (ITSO). Overview of the international telecommunication union (itu), 2018. URL [https://itso.int/wp-content/uploads/2018/04/akim\\_Day-1-1-rev1-Overview.pdf](https://itso.int/wp-content/uploads/2018/04/akim_Day-1-1-rev1-Overview.pdf). Accessed: 2025-02-26.
- [33] United Nations Office for Outer Space Affairs (UNOOSA). Itu frequency management, 2022. URL [https://www.unoosa.org/documents/pdf/spacelaw/workshops/SLC2022/presentations/H.\\_Glaude\\_-\\_UNOOSA\\_-\\_ITU\\_frequency\\_management\\_20220511.pdf](https://www.unoosa.org/documents/pdf/spacelaw/workshops/SLC2022/presentations/H._Glaude_-_UNOOSA_-_ITU_frequency_management_20220511.pdf). Accessed: 2025-02-26.
- [34] Maria Kyropoulou, Josué Ortega, and Erel Segal-Halevi. Fair cake-cutting in practice. Games and Economic Behavior, 133:28–49, 2022. ISSN 0899-8256. doi: <https://doi.org/10.1016/j.geb.2022.01.027>. URL <https://www.sciencedirect.com/science/article/pii/S0899825622000331>.
- [35] Australian Communications and Media Authority. Australian spectrum map grid 2012: October 2023 update, October 2023. URL [https://www.acma.gov.au/sites/default/files/2023-10/Australian%20spectrum%20map%20grid%202012\\_October%202023%20update%20FINAL.pdf](https://www.acma.gov.au/sites/default/files/2023-10/Australian%20spectrum%20map%20grid%202012_October%202023%20update%20FINAL.pdf). Accessed from <https://www.acma.gov.au/australian-spectrum-map-grid>.
- [36] Australian Communications and Media Authority. Spectrum auctions, 2023. URL <https://www.acma.gov.au/spectrum-auctions>. Accessed: 2025-05-08.
- [37] Australian Communications and Media Authority. Australian radiofrequency spectrum allocations chart, 2017. URL <https://www.acma.gov.au/sites/default/files/2019->



- 10/Australian%20radiofrequency%20spectrum%20allocations%20chart.pdf. Accessed from <https://www.acma.gov.au/>.
- [38] National Telecommunications and Information Administration. United states frequency allocation chart, January 2016. URL [https://www.ntia.gov/sites/default/files/publications/january\\_2016\\_spectrum\\_wall\\_chart\\_0.pdf](https://www.ntia.gov/sites/default/files/publications/january_2016_spectrum_wall_chart_0.pdf). Accessed from <http://ntia.gov/page/united-states-frequency-allocation-chart>.
- [39] U.S. Department of Commerce, National Telecommunications and Information Administration. Commercial spectrum enhancement act: Annual progress report for 2023. Technical report, National Telecommunications and Information Administration (NTIA), September 2024. URL <https://www.ntia.gov/sites/default/files/2024-10/2023-csea-report.pdf>. Accessed: 2025-05-08.
- [40] Martin Sauter. From GSM to LTE-Advanced Pro and 5G : an introduction to mobile networks and mobile broadband. Wiley, Hoboken, New Jersey, third edition. edition, 2017. ISBN 1-119-34690-8. Includes bibliographical references at the end of each chapters and index.
- [41] Erik Dahlman, Stefan Parkvall, and Johan Skold. 5g Nr: The Next Generation Wireless Access Technology. Elsevier Science & Technology, San Diego, 2020. ISBN 9780128223208.
- [42] Danny W. H. Tham. Carrier to interference (c/i ratio) calculations. Presentation at the World Radiocommunication Seminar 2012, December 2012. International Telecommunication Union, Geneva, Switzerland, 3–7 December 2012.
- [43] Nilesh Khambekar, Chad M. Spooner, and Vipin Chaudhary. Quantified spectrum sharing: Motivation, approach, and benefits. CoRR, abs/1608.07854, 2016. URL <http://arxiv.org/abs/1608.07854>.
- [44] Samsung. 6g spectrum expanding the frontier, 2022.
- [45] Benoît Pierre Freyens and Oleg Yerokhin. Allocative vs. technical spectrum efficiency. Telecommunications Policy, 35(4):291–300, May 2011. ISSN 0308-5961. doi: 10.1016/j.telpol.2011.02.002.
- [46] Tamás Szálka, Péter Fülöp, and Sandor Imre. General mobility modeling and location prediction based on markovian approach constructor framework. Periodica Polytechnica Electrical Engineering, 55:71, 2011. doi: 10.3311/pp.ee.2011-1-2.08.
- [47] Atif Shakeel, Riaz Hussain, Adeel Iqbal, Irfan Latif Khan, Qadeer ul Hasan, and Shahzad Ali Malik. Analysis of efficient spectrum handoff in a multi-class hybrid spectrum access cognitive radio network using markov modelling. Sensors, 19(19):4120, 2019. ISSN 1424-8220. URL <https://www.mdpi.com/1424-8220/19/19/4120>.



- [48] Y. Zhao, Z. Hong, Y. Luo, G. Wang, and L. Pu. Prediction-based spectrum management in cognitive radio networks. IEEE Systems Journal, 12(4):3303–3314, 2018. ISSN 1937-9234. doi: 10.1109/JSYST.2017.2741448.
- [49] F. Li, K. Y. Lam, Z. Sheng, W. Lu, X. Liu, and L. Wang. Agent-based spectrum management scheme in satellite communication systems. IEEE Transactions on Vehicular Technology, 70(3):2877–2881, 2021. ISSN 1939-9359. doi: 10.1109/TVT.2021.3059837.
- [50] Jan Buermann. Efficiency and fairness of resource utilisation under uncertainty: doctoral consortium, 2020. URL <https://eprints.soton.ac.uk/439156/>.
- [51] Xinzhou Li, Guifen Chen, Guowei Wu, Zhiyao Sun, and Guangjiao Chen. Research on multi-agent d2d communication resource allocation algorithm based on a2c. Electronics, 12(2):360, 2023. ISSN 2079-9292. URL <https://www.mdpi.com/2079-9292/12/2/360>.
- [52] Przemysław Szufel, Bartosz Pankratz, Anna Szczurek, Bogumił Kamiński, and Paweł Prałat. Vehicle routing simulation for prediction of commuter’s behaviour. Journal of Advanced Transportation, 2022:1604303, 2022. ISSN 0197-6729. doi: 10.1155/2022/1604303. URL <https://doi.org/10.1155/2022/1604303>.
- [53] John G Markoulidakis, George L Lyberopoulos, Dimitrios F Tsirkas, and Efsthios D Sykas. Mobility modeling in third-generation mobile telecommunications systems. IEEE personal communications, 4(4):41–56, 1997. ISSN 1070-9916.
- [54] Abbas Sarraf Shirazi, Timothy Davison, Sebastian von Mammen, Jörg Denzinger, and Christian Jacob. Adaptive agent abstractions to speed up spatial agent-based simulations. Simulation Modelling Practice and Theory, 40:144–160, 2014. ISSN 1569-190X. doi: <https://doi.org/10.1016/j.simpat.2013.09.001>. URL <https://www.sciencedirect.com/science/article/pii/S1569190X13001330>.
- [55] Abbas Sarraf Shirazi, Sebastian von Mammen, and Christian Jacob. Abstraction of agent interaction processes: Towards large-scale multi-agent models. SIMULATION, 89(4):524–538, 2013. doi: 10.1177/0037549712470733. URL <https://doi.org/10.1177/0037549712470733>.
- [56] Romain Franceschini, Simon Van Mierlo, and Hans Vangheluwe. Towards adaptive abstraction in agent based simulation. In 2019 Winter Simulation Conference (WSC), pages 2725–2736, Dec 2019. doi: 10.1109/WSC40007.2019.9004843.
- [57] Mohammad Erfan Doraki, Akram Avami, Mehrdad Boroushaki, and Zahra Amini. Agent-based modeling for sustainable urban passenger vehicle mobility: A case of tehran. Transportation Research Part D: Transport and Environment, 135:104380, 2024. doi: 10.1016/j.trd.2024.104380.





- [58] Jan Drchal, Michal Čertický, and Michal Jakob. Data-driven activity scheduler for agent-based mobility models. Transportation Research Part C: Emerging Technologies, 98: 370–390, 2019. doi: 10.1016/j.trc.2018.12.002.
- [59] Asier Divasson-J., Ana M. Macarulla, J. Ignacio Garcia, and Cruz E. Borges. Agent-based modeling in urban human mobility: A systematic review. Cities, 158:105697, 2025. doi: 10.1016/j.cities.2024.105697.
- [60] Jingjun Li, Evy Rombaut, and Lieselot Vanhaverbeke. Agent-based digital traffic model generation for regions facing data scarcity using aggregated cellphone data: a case study for brussels. International Journal of Digital Earth, 17(1):2407046, 2024. doi: 10.1080/17538947.2024.2407046.
- [61] Hossein Amiri, Will Kohn, Shiyang Ruan, Joon-Seok Kim, Hamdi Kavak, Andrew Crooks, Dieter Pfoser, Carola Wenk, and Andreas Züfle. The patterns of life human mobility simulation. In Proc. of the 32nd ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (SIGSPATIAL '24), 2024. doi: 10.1145/3678717.3691319.
- [62] Patrick Manser, Henrik Becker, Sebastian Hörl, and Kay W. Axhausen. Designing a large-scale public transport network using agent-based microsimulation. Transportation Research Part A: Policy and Practice, 137:1–15, 2020. doi: 10.1016/j.tra.2020.04.011.
- [63] Sinziana I. Rasca, Bin Hu, Benjamin Biesinger, and Matthias Prandtstetter. Agent-based decision-support model for bus route redesign in networks of small cities and towns: case study of agder, norway. Public Transport, 16:583–618, 2024. doi: 10.1007/s12469-024-00358-7.
- [64] Ouassim Manout and Francesco Ciari. Assessing the role of daily activities and mobility in the spread of covid-19 in montreal with an agent-based approach. Frontiers in Built Environment, 7:654279, 2021. doi: 10.3389/fbuil.2021.654279.
- [65] Lewen Feng, John M. Betts, Liton Kamruzzaman, and Hai L. Vu. Impact of autonomous vehicles on discretionary activities: an agent-based model with space–time accessibility constraints. Transportation, 2024. doi: 10.1007/s11116-024-10502-8.
- [66] Khatun E. Zannat, Janek Laudan, Charisma F. Choudhury, and Stephane Hess. Developing an agent-based microsimulation for predicting the bus rapid transit (brt) demand in developing countries: A case study of dhaka, bangladesh. Transport Policy, 148:92–106, 2024. doi: 10.1016/j.tranpol.2023.12.014.
- [67] Jiawei Wang, Renhe Jiang, Chuang Yang, Zengqing Wu, Makoto Onizuka, Ryosuke Shibasaki, Noboru Koshizuka, and Chuan Xiao. Large language models as urban residents: An llm agent framework for personal mobility generation. In Advances in Neural Information Processing Systems 38 (NeurIPS



- 2024), 2024. URL [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/e142fd2b70f10db2543c64bca1417de8-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/e142fd2b70f10db2543c64bca1417de8-Paper-Conference.pdf).
- [68] TomTom International B.V. Adelaide traffic report – tomtom traffic index. <https://www.tomtom.com/traffic-index/adelaide-traffic/>, 2025. Accessed: 2025-11-19.
- [69] OpenStreetMap contributors. Planet dump retrieved from <https://planet.osm.org>, 2022.
- [70] Australian Bureau of Statistics. Australian statistical geography standard (asgs): Main structure and greater capital city statistical areas, july 2021 – june 2026, 2021. URL <https://www.abs.gov.au/statistics/standards/australian-statistical-geography-standard-asgs-edition-3/jul2021-jun2026/main-structure-and-greater-capital-city-statistical-areas/statistical-area-level-2>. Accessed: 2025-08-20.
- [71] Valhalla contributors. Valhalla, 2022.
- [72] Yann Chevaleyre, Paul E. Dunne, Ulle Endriss, Jérôme Lang, Michel Lemaitre, Nicolas Maudet, Julian Padget, Steve Phelps, Juan A. Rodríguez-Aguilar, and Paulo Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, 2006.
- [73] Rajdeep K. Dash, Nicholas R. Jennings, and David C. Parkes. Computational-mechanism design: A call to arms. *IEEE Intelligent Systems*, 18(6):40–47, 2003.
- [74] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.
- [75] Frank Kelly, Aman Maulloo, and David Tan. Rate control for communication networks: Shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237–252, 1998.
- [76] Mohammadreza Doostmohammadian, Alireza Aghasi, Mohammad Pirani, Ehsan Nekouei, Houman Zarrabi, Reza Keypour, Apostolos I. Rikos, and Karl H. Johansson. Survey of distributed algorithms for resource allocation over multi-agent systems. *Annual Reviews in Control*, 2024. to appear.
- [77] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proc. of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2011.
- [78] A. H. Mohsenian-Rad, Vincent W. S. Wong, Juri Jatskevich, Robert Schober, and Alberto Leon-Garcia. Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid. *IEEE Transactions on Smart Grid*, 1(3):320–331, 2010.



- [79] Noel Pérez-Rodríguez and José Holguín-Veras. Inventory-allocation distribution models for postdisaster humanitarian logistics with explicit consideration of deprivation costs. Transportation Science, 50(4):1261–1285, 2016.
- [80] Mung Chiang, Steven H. Low, A. Robert Calderbank, and John C. Doyle. Layering as optimization decomposition: A mathematical theory of network architectures. Proceedings of the IEEE, 95(1):255–312, 2007.
- [81] Qing Zhao, Lang Tong, Ananthram Swami, and Yunxia Chen. Decentralized cognitive mac for opportunistic spectrum access in ad hoc networks: A pomdp framework. IEEE Journal on Selected Areas in Communications, 25(3):589–600, 2007.
- [82] Raul Etkin, Abhay Parekh, and David Tse. Spectrum sharing for unlicensed bands. IEEE Journal on Selected Areas in Communications, 25(3):517–528, 2007.
- [83] Xueyang Zhou, Sachin Gandhi, Sanjay Suri, and Haitao Zheng. ebay in the sky: Strategy-proof wireless spectrum auctions. In Proc. of the 14th Annual ACM Intl. Conf. on Mobile Computing and Networking (MobiCom), pages 2–13, 2008.
- [84] Feng Hu, Bing Chen, and Kun Zhu. Full spectrum sharing in cognitive radio networks toward 5g: A survey. IEEE Access, 6:15754–15776, 2018.
- [85] Analytics Vidhya. An introduction to clustering and different methods of clustering, November 2016. URL <https://www.analyticsvidhya.com/blog/2016/11/an-introduction-to-clustering-and-different-methods-of-clustering/>. Accessed: February 16, 2025.
- [86] Jonathon Byrd. Data clustering: Intro, methods, applications, November 2023. URL <https://encord.com/blog/>. Accessed: February 16, 2025.
- [87] Jean-Thomas Camino, Christian Artigues, Laurent Houssin, and Stéphane Mourgues. Milp formulation improvement with k-means clustering for the beam layout optimization in multibeam satellite systems. Computers & Industrial Engineering, 158:107228, 2021. ISSN 0360-8352. doi: <https://doi.org/10.1016/j.cie.2021.107228>. URL <https://www.sciencedirect.com/science/article/pii/S0360835221001327>.
- [88] A. Ikami, T. Hayashi, and Y. Amano. Fair spectrum allocation for dynamic spectrum sharing between different radio systems. In 2020 23rd International Symposium on Wireless Personal Multimedia Communications (WPMC), pages 1–6. ISBN 1882-5621. doi: 10.1109/WPMC50192.2020.9309506.
- [89] John Ngubiri and Martin van Vliet. A metric of fairness for parallel job schedulers. Concurrency and Computation: Practice and Experience, 21(12):1525–1546, 2009.

- [90] Ubaid Ullah Khan, Naqqash Dilshad, Mubashir Husain Rehmani, and Tariq Umer. Fairness in cognitive radio networks: Models, measurement methods, applications, and future research directions. Journal of Network and Computer Applications, 73:12–26, 2016.
- [91] Mehrdad Dianati, Xuemin (Sherman) Shen, and Subhas Naik. A new fairness index for radio resource allocation in wireless networks. In Proceedings of IEEE WCNC, pages 1267–1272. IEEE, 2009.
- [92] Zhonggui Ma and Hongbo Wang. Dynamic spectrum allocation with maximum efficiency and fairness in interactive cognitive radio networks. Wireless Personal Communications, 64:439–455, 2012.
- [93] Wanshi Chen, Xingqin Lin, Juho Lee, Antti Toskala, Shu Sun, Carla Fabiana Chiasserini, and Lingjia Liu. 5g-advanced toward 6g: Past, present, and future. IEEE Journal on Selected Areas in Communications, 41(6):1592–1619, June 2023. ISSN 1558-0008. doi: 10.1109/JSAC.2023.3274037.
- [94] Antti Toskala. Trends towards 6g technology. PowerPoint slides, 2024. Unpublished slides.
- [95] Mina Khadem, Masoud Ansarifard, Navid Mokari, Mohammad Javan, Hadi Saeedi, and Eduard A. Jorswieck. Dynamic fairness-aware spectrum auction for enhanced licensed shared access in 6g networks. IEEE Communications Letters, 27(8):1910–1914, 2023.
- [96] Zhangyu Guan, Dongfeng Yuan, and Haixia Zhang. Optimal and fair resource allocation for multiuser wireless multimedia transmissions. EURASIP Journal on Wireless Communications and Networking, 2009:801613, 2009.
- [97] Symeon Papavassiliou and Chengzhou Li. Joint throughput maximization and fair uplink transmission scheduling in cdma systems. EURASIP Journal on Wireless Communications and Networking, 2009:564692, 2009.
- [98] Cédric Gueguen and Sébastien Baey. A fair opportunistic access scheme for multiuser ofdm wireless networks. EURASIP Journal on Wireless Communications and Networking, 2009:726495, 2009.
- [99] Jose Lopez Vicario, Albert Bel, Antoni Morell, and Gonzalo Seco-Granados. Outage probability versus fairness trade-off in opportunistic relay selection with outdated csi. EURASIP Journal on Wireless Communications and Networking, 2009:412837, 2009.
- [100] Linda Meylani, Andik Kurniawan, and M. Sigit Arifianto. Radio resource allocation with the fairness metric for low density signature ofdm in underlay cognitive radio networks. Sensors, 19(8):1921, 2019.



- [101] Sang-Seon Byun and Joon-Min Gil. Fair dynamic spectrum allocation using modified game theory for resource-constrained cognitive wireless sensor networks. Symmetry, 9 (5):73, 2017.
- [102] Jun Su, Mykola Beshley, Krzysztof Przystupa, Orest Kochan, Bohdan Rusyn, Rafał Stanisławski, Oleg Yaremko, Michał Majka, Halyna Beshley, Ivan Demydov, Julia Pyrih, and Ihor Kahalo. 5g multi-tier radio access network planning based on voronoi diagram. Measurement, 192:110814, 2022. ISSN 0263-2241. doi: <https://doi.org/10.1016/j.measurement.2022.110814>. URL <https://www.sciencedirect.com/science/article/pii/S0263224122001117>.
- [103] Xiaodong Xu, Yi Li, Rui Gao, and Xiaofeng Tao. Joint voronoi diagram and game theory-based power control scheme for the hetnet small cell networks. EURASIP Journal on Wireless Communications and Networking, 2014(1):213, Dec 2014. ISSN 1687-1499. doi: 10.1186/1687-1499-2014-213. URL <https://doi.org/10.1186/1687-1499-2014-213>.
- [104] Chun-Peng Liu, Chen-Guang He, Wei-Xiao Meng, and Shuai Han. A design of d2d-pairing scheme on voronoi diagram. In 2017 13th International Wireless Communications and Mobile Computing Conference (IWCMC), pages 202–205, June 2017. doi: 10.1109/IWCMC.2017.7986286.
- [105] Rahat Ullah, Norsheila Fisal, Hashim Safdar, Wajahat Maqbool, Zubair Khalid, and A.S. Khan. Voronoi cell geometry based dynamic fractional frequency reuse for ofdma cellular networks. In 2013 IEEE International Conference on Signal and Image Processing Applications, pages 435–440, Oct 2013. doi: 10.1109/ICSIPA.2013.6708046.
- [106] Rui Tao, Jian Liu, Yuqing Song, Rui Peng, Dali Zhang, and Jiangang Qiao. Detection and optimization of traffic networks based on voronoi diagram. Discrete Dynamics in Nature and Society, 2021(1):5550315, 2021. doi: <https://doi.org/10.1155/2021/5550315>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1155/2021/5550315>.
- [107] Yu Cao, Qi-Yue Yu, Ji-Chong Guo, and Julian Cheng. Voronoi-cluster multi-resolution hierarchical codebook design for cell-free integrated sensing and communication systems. IEEE Transactions on Communications, 72(7):4432–4445, 2024. doi: 10.1109/TCOMM.2024.3369677.
- [108] Andreas Andreou, Constandinos X. Mavromoustakis, Jordi Mongay Batalla, Evangelos K. Markakis, and George Mastorakis. Uav-assisted rsus for v2x connectivity using voronoi diagrams in 6g+ infrastructures. IEEE Transactions on Intelligent Transportation Systems, 24(12):15855–15865, Dec 2023. ISSN 1558-0016. doi: 10.1109/TITS.2023.3273716.
- [109] NVIDIA. Ericsson uses digital twins and nvidia omniverse to plan, optimize future networks, 2024. URL <https://blogs.nvidia.com/blog/ericsson-digital-twins-omniverse/>. Accessed: 2024-09-17.



- [110] Bulent Tavli and Wendi Heinzelman. Mobile Ad Hoc Networks Energy-Efficient Real-Time Data Communications. Springer Netherlands, Dordrecht, 1st 2006. edition, 2006. ISBN 1-280-62512-0. doi: 10.1007/1-4020-4633-2. Includes bibliographical references (p. 221-233) and index.
- [111] Ratish Agarwal and Dr Motwani. Survey of clustering algorithms for manet. arXiv preprint arXiv:0912.2303, 2009.
- [112] Mirjeta Alinci, Spaho Evjola, Lala Argenti, and Kolici Vladi. Clustering algorithms in manets: A review, July , 2015.
- [113] Valhalla Contributors. Isochrone api reference, 2025. URL <https://valhalla.github.io/valhalla/api/isochrone/api-reference/>.
- [114] Peter. Alternative roads to rome, January 19 2016. URL <https://www.graphhopper.com/blog/2016/01/19/alternative-roads-to-rome/>. Accessed: 2025-03-31.
- [115] Valhalla Contributors. Expansion api reference, 2025. URL <https://valhalla.github.io/valhalla/api/expansion/api-reference/>.
- [116] Seyed Ali Sharifi and Seyed Morteza Babamir. Evaluation of clustering algorithms in ad hoc mobile networks. Wireless Personal Communications, 109(4):2147–2186, 2019. ISSN 1572-834X. doi: 10.1007/s11277-019-06673-8. URL <https://doi.org/10.1007/s11277-019-06673-8>.
- [117] Openstreetmap wiki. [https://wiki.openstreetmap.org/wiki/Main\\_Page](https://wiki.openstreetmap.org/wiki/Main_Page). Accessed: 2025-02-27.
- [118] David H. Pilgrim. Isochrones of travel time and distribution of flood storage from a tracer study on a small watershed. Water Resources Research, 13(3):587–595, 1977. doi: <https://doi.org/10.1029/WR013i003p00587>. URL <https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/WR013i003p00587>.
- [119] Graphhopper directions api. <https://www.graphhopper.com/>. Accessed: 2025-02-27.
- [120] Valhalla Contributors. Valhalla routing engine. <https://github.com/valhalla/valhalla>, 2025. Accessed: 2025-02-27.
- [121] Open source routing machine. <http://project-osrm.org/>. Accessed: 2025-02-27.
- [122] İbrahim Sarıçiçek. Routing alternatives with custom spatial data. Medium, January 2024. URL <https://ibrahimsaricicek.medium.com/routing-alternatives-with-custom-spatial-data-dbb2a9175ef2>. Accessed: 2025-03-31.



- [123] Naim Kapucu, Brittany Haupt, Thomas Quint, Mostafizur Rahman, and Murat Yuk-sel. Polycentric governance and decentralized decision-making for pervasive spectrum sharing. *International Journal of Public Administration*, 46(9):659–668, 2023. ISSN 0190-0692. doi: 10.1080/01900692.2021.2010752. URL <https://doi.org/10.1080/01900692.2021.2010752>. doi: 10.1080/01900692.2021.2010752.
- [124] Rodney Martinez Alonso, David Plets, Sofie pollin, Luc Martens, and Wout Joseph. Out-lier detection and spectrum feature extraction based on nearest-neighbors correlation and random forest algorithm. 2023. URL [https://www.techrxiv.org/articles/preprint/Outlier\\_Detection\\_and\\_Spectrum\\_Feature\\_Extraction\\_Based\\_on\\_Nearest-Neighbors\\_Correlation\\_and\\_Random\\_Forest\\_Algorithm/24416629](https://www.techrxiv.org/articles/preprint/Outlier_Detection_and_Spectrum_Feature_Extraction_Based_on_Nearest-Neighbors_Correlation_and_Random_Forest_Algorithm/24416629).
- [125] Parishad Karimi, William Lehr, Ivan Seskar, and Dipankar Raychaudhuri. Smap: A scalable and distributed architecture for dynamic spectrum management. In *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pages 1–10, Oct 2018. doi: 10.1109/DySPAN.2018.8610416.
- [126] Debajit Sensarma and Koushik Majumder. A novel hierarchical ant based qos aware intelligent routing scheme for MANETS. *CoRR*, abs/1312.2062, 2013. URL <http://arxiv.org/abs/1312.2062>.
- [127] Zainab Ali Abbood, Doggu Çagğdass Atilla, and Caggatay Aydin. Enhancement of the performance of manet using machine learning approach based on sdns. *Optik*, 272: 170268, 2023. ISSN 0030-4026. doi: <https://doi.org/10.1016/j.ijleo.2022.170268>. URL <https://www.sciencedirect.com/science/article/pii/S0030402622015261>.
- [128] Tan Do-Duy, Trang Dang Phuoc Hai, Phuc Truong Quang, Loc Nguyen Van Thanh, Phuc Nguyen Van, and Thanh Le Minh. Optimal throughput for multicast over manet-satellite networks. *Journal of Technical Education Science*, (68):74–81, 2022. doi: 10.54644/jte.68.2022.1112. URL <https://jte.edu.vn/index.php/jte/article/view/1112>.
- [129] De-Gan Zhang, Zhao Peng-Zhen, Cui Yu-ya, Chen Lu, Zhang Ting, and Wu Hao. A new method of mobile ad hoc network routing based on greed forwarding improvement strategy. *IEEE Access*, 7:158514–158524, 2019. doi: 10.1109/ACCESS.2019.2950266, ISSN=2169-3536.
- [130] Chao Dong, Yuben Qu, Haipeng Dai, Song Guo, and Wu Qihui. Multicast in multi-channel cognitive radio ad hoc networks: Challenges and research aspects. *Computer Communications*, 132:10–16, 2018. ISSN 0140-3664. doi: <https://doi.org/10.1016/j.comcom.2018.09.003>. URL <https://www.sciencedirect.com/science/article/pii/S0140366418305061>.





- [131] Hamid Kordbacheh, Dalili Oskouei Hamid, and Mokari Nader. Robust cross-layer routing and radio resource allocation in massive multiple antenna and ofdma-based wireless ad-hoc networks. IEEE Access, 7:36527–36539, 2019. doi: 10.1109/ACCESS.2019.2904673, ISSN=2169-3536.
- [132] P. J. Nicholas and K. L. Hoffman. Computational challenges of dynamic channel assignment for military manet. In MILCOM 2015 - 2015 IEEE Military Communications Conference, pages 1150–1157. doi: 10.1109/MILCOM.2015.7357601.
- [133] N. A. Azeez and A. E. Otudor. Modelling and simulating access control in wireless ad-hoc networks. Fountain Journal of Natural and Applied Sciences, 5(2), 2016. doi: 10.53704/fujnas.v5i2.128. URL <https://fountainjournals.com/index.php/FUJNAS/article/view/128>.
- [134] Piotr Gajewski, Lopatka Jerzy, and Łubkowski Piotr. Traffic modelling for mobile ad-hoc networks simulation, Sep. , 2021.
- [135] Gajendra Kumar Ahirwar, Ratish Agarwal, and Anjana Pandey. An extensive review on qos enhancement in manet using meta-heuristic algorithms. Wireless Personal Communications, 131(2):1089–1114, 2023. ISSN 1572-834X. doi: 10.1007/s11277-023-10470-9. URL <https://doi.org/10.1007/s11277-023-10470-9>.
- [136] Shanthini Sampath, P. Punitha, and Karthik Subburathinam. Improvisation of node mobility using cluster routing-based group adaptive in manet. Computer Systems Science and Engineering, 44:2619–2636, 2023. doi: 10.32604/csse.2023.027330.
- [137] Bora Karaoglu, Tolga Numanoglu, and Heinzelman Wendi. Analytical performance of soft clustering protocols. Ad Hoc Networks, 9(4):635–651, 2011. ISSN 1570-8705. doi: <https://doi.org/10.1016/j.adhoc.2010.08.008>. URL <https://www.sciencedirect.com/science/article/pii/S1570870510001046>. Multimedia Ad Hoc and Sensor Networks.
- [138] Thong Nhat Tran, Toan-Van Nguyen, and Beongku An. An efficient connected dominating set clustering based routing protocol with dynamic channel selection in cognitive mobile ad hoc networks. Electronics, 8(11):1332, 2019. ISSN 2079-9292. URL <https://www.mdpi.com/2079-9292/8/11/1332>.
- [139] T. Rahman, I. Ullah, A. U. Rehman, and R. A. Naqvi. Notice of violation of ieee publication principles: Clustering schemes in manets: Performance evaluation, open challenges, and proposed solutions. IEEE Access, 8:25135–25158, 2020. ISSN 2169-3536. doi: 10.1109/ACCESS.2020.2970481.
- [140] B. Tavli and Heinzelman W. B. Energy and spatial reuse efficient network-wide real-time data broadcasting in mobile ad hoc networks. IEEE Transactions on Mobile Computing, 5(10):1297–1312, 2006. doi: 10.1109/TMC.2006.151,ISSN=1558-0660.



- [141] Przemysław Śleszyński, Piotr Olszewski, Tomasz Dybicz, Katarzyna Goch, and Michał A. Niedzielski. The ideal isochrone: Assessing the efficiency of transport systems. Research in Transportation Business & Management, 46:100779, 2023. ISSN 2210-5395. doi: <https://doi.org/10.1016/j.rtbm.2021.100779>. URL <https://www.sciencedirect.com/science/article/pii/S2210539521001620>. Themed Volume on Transport Efficiency.
- [142] Alexandros Efentakis, Nikos Grivas, George Lamprianidis, Georg Magenschab, and Dieter Pfoser. Isochrones, traffic and demographics. In Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, SIGSPATIAL'13, page 548–551, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450325219. doi: 10.1145/2525314.2525325. URL <https://doi.org/10.1145/2525314.2525325>.
- [143] João A. V. Freitas. Route Optimization for Bicycle Users. PhD thesis, 2024. URL <https://www.proquest.com/dissertations-theses/route-optimization-bicycle-users/docview/3143978958/se-2>. Copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2025-03-04.
- [144] Carlo Andrea Biraghi, Ozge Ogut, Tao Dong, and Massimo Tadi. Citytime: A novel model to redefine the 15-minute city globally through urban diversity and proximity. Urban Science, 9(2):36, 2025. doi: 10.3390/urbansci9020036.
- [145] Yuan Zhao and Ying Zhou. Isochrone-based accessibility analysis of pre-hospital emergency medical facilities: A case study of central districts of beijing. ISPRS International Journal of Geo-Information, 13(8):288, 2024. doi: 10.3390/ijgi13080288.
- [146] Cyrus Shahabi and Seon Ho Kim. Evaluating accessibility of los angeles metropolitan area using data-driven time-dependent reachability analysis. Technical report, Institute of Transportation Studies, University of California, Davis, 2023.
- [147] Nazli Yonca Aydin, Emre Yigitbasi, Ylenia Casali, and Bert van Wee. Assessing the evolution of educational accessibility with self-avoiding random walk: Insights from helsinki. Applied Network Science, 8(1):37, 2023. doi: 10.1007/s41109-023-00581-4.
- [148] Moritz Baum, Julian Dibbelt, Andreas Gemsa, Dorothea Wagner, and Tobias Zündorf. Shortest feasible paths with charging stops for battery electric vehicles. In Proceedings of the 23rd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, pages 44:1–44:10, 2015. doi: 10.1145/2820783.2820826.
- [149] Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat,



- Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- [150] Fayçal Aït Aoudia, Jakob Hoydis, Merlin Nimier-David, Sebastian Cammerer, and Alexander Keller. Sionna rt: Technical report. *arXiv preprint arXiv:2504.21719*, 2025. URL <https://arxiv.org/abs/2504.21719>.
- [151] State Government of Victoria. Vicmap elevation 1 m digital elevation model (dem). <https://www.land.vic.gov.au/maps-and-spatial/spatial-data/vicmap-catalogue/vicmap-elevation/1m-dem>, 2025. Accessed: 18 September 2025; Page last updated: 29 July 2025.
- [152] Google. Photorealistic 3d tiles overview. <https://developers.google.com/maps/documentation/tile/3d-tiles-overview>, 2025.
- [153] Yifan Zhang, Mengyu Ma, Jun Li, Anran Yang, Qingren Jia, and Zebang Liu. Efficient and fine-grained viewshed analysis in a three-dimensional urban complex environment. *Environmental Modelling & Software*, 188:106359, 2025. ISSN 1364-8152. doi: <https://doi.org/10.1016/j.envsoft.2025.106359>. URL <https://www.sciencedirect.com/science/article/pii/S136481522500043X>.
- [154] Narushige Shiode. 3d urban models: Recent developments in the digital modelling of urban environments in three-dimensions. *GeoJournal*, 52(3):263–269, Nov 2000. ISSN 1572-9893. doi: 10.1023/A:1014276309416. URL <https://doi.org/10.1023/A:1014276309416>.
- [155] Serhat Tadik, Rajib Bhattacharjea, Johnathan Corgan, David Johnson, Jacobus Van der Merwe, and Gregory D. Durgin. Opengert: Open source automated geometry extraction with geometric and electromagnetic sensitivity analyses for ray-tracing propagation models, 2025. URL <https://arxiv.org/abs/2501.06945>.
- [156] Feng Shen, Guoru Ding, and Qihui Wu. Efficient remote compressed spectrum mapping in 3-d spectrum-heterogeneous environment with inaccessible areas. *IEEE Wireless Communications Letters*, 11(7):1488–1492, July 2022. ISSN 2162-2345. doi: 10.1109/LWC.2022.3175820.
- [157] Peter L. Guth, Adriaan Van Niekerk, Carlos H. Grohmann, Jan-Peter Muller, Laurence Hawker, Igor V. Florinsky, Dean Gesch, Hannes I. Reuter, Virginia Herrera-Cruz, Serge Riazanoff, Carlos López-Vázquez, Claudia C. Carabajal, Clément Albinet, and Peter Strobl. Digital elevation models: Terminology and definitions. *Remote Sensing*, 13(18),



2021. ISSN 2072-4292. doi: 10.3390/rs13183581. URL <https://www.mdpi.com/2072-4292/13/18/3581>.

- [158] Laurent Polidori and Mhamad El Hage. Digital elevation model quality assessment methods: A critical review. *Remote Sensing*, 12(21), 2020. ISSN 2072-4292. doi: 10.3390/rs12213522. URL <https://www.mdpi.com/2072-4292/12/21/3522>.
- [159] Andreas Scalas, Daniela Cabiddu, Michela Mortara, and Michela Spagnuolo. Potential of the geometric layer in urban digital twins. *ISPRS International Journal of Geo-Information*, 11(6), 2022. ISSN 2220-9964. doi: 10.3390/ijgi11060343. URL <https://www.mdpi.com/2220-9964/11/6/343>.
- [160] Marco Aurelio Nunes da Silva, Emanuel Costa, and Markus Liniger. Digital elevation data and their use for improved broadcast coverage and frequency utilization. In *2006 First European Conference on Antennas and Propagation*, pages 1–7, Nov 2006. doi: 10.1109/EUCAP.2006.4584789.
- [161] Jérôme Eertmans, Enrico Maria Vittuci, Vittorio Degli-Esposti, Laurent Jacques, and Claude Oestges. Comparing differentiable and dynamic ray tracing: Introducing the multipath lifetime map, 2024. URL <https://arxiv.org/abs/2410.14535>.
- [162] Jérôme Eertmans, Nicola Di Cicco, Claude Oestges, Laurent Jacques, Enrico M. Vittuci, and Vittorio Degli-Esposti. Towards generative ray path sampling for faster point-to-point ray tracing, 2025. URL <https://arxiv.org/abs/2410.23773>.
- [163] Lihao Zhang, Haijian Sun, Jin Sun, and Rose Qingyang Hu. Wisegrt: Dataset for site-specific indoor radio propagation modeling with 3d segmentation and differentiable ray-tracing: (invited paper). In *2024 International Conference on Computing, Networking and Communications (ICNC)*, pages 744–748, Feb 2024. doi: 10.1109/ICNC59896.2024.10556262.
- [164] Anatolij Zubow, Yannik Pilz, Sascha Rösler, and Falko Dressler. Ns3 meets sionna: Using realistic channels in network simulation, 2024. URL <https://arxiv.org/abs/2412.20524>.
- [165] Yong Jun Noh and Kae Won Choi. High-precision digital twin platform based on ray tracing simulation. In *2024 15th International Conference on Information and Communication Technology Convergence (ICTC)*, pages 1464–1465, Oct 2024. doi: 10.1109/ICTC62082.2024.10827549.
- [166] Michele Zhu, Lorenzo Cazzella, Francesco Linsalata, Maurizio Magarini, Matteo Matteucci, and Umberto Spagnolini. On the digitization of the em environment: A comparison of ray launching solutions. In *2024 IEEE International Mediterranean*



Conference on Communications and Networking (MeditCom), pages 85–90, July 2024. doi: 10.1109/MeditCom61057.2024.10621310.

- [167] Sávio Bastos, Ailton Oliveira, Daniel Suzuki, Luan Gonçalves, Ilan Sousa, and Aldebaro Klautau. Generation of 5g/6g wireless channels using raymobtime with sionna’s ray-tracing. XLI Simpósio Brasileiro de Telecomunicações e Processamento de Sinais, 2023.
- [168] Sunday Amatare, Michelle Samson, and Debashri Roy. Testbed design for robot navigation through differential ray tracing. In 2024 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN), pages 173–174, May 2024. doi: 10.1109/DySPAN60163.2024.10632751.
- [169] Fayçal Aït Aoudia, Jakob Hoydis, Merlin Nimier-David, Sebastian Cammerer, and Alexander Keller. Sionna rt: Technical report. arXiv preprint arXiv:2504.21719, 2025.
- [170] NVIDIA Labs. Sionna rt: Radio material documentation, 2024. URL [https://nvlabs.github.io/sionna/\\_modules/sionna/rt/radio\\_material.html](https://nvlabs.github.io/sionna/_modules/sionna/rt/radio_material.html). Accessed: 2025-02-24.
- [171] MathWorks. Ray tracing propagation model - matlab, 2024. URL <https://www.mathworks.com/help/antenna/ref/rfprop.raytracing.html>. Accessed: 2025-02-24.
- [172] Brandon Derocher. flutter\_scene. [https://github.com/bdero/flutter\\_scene](https://github.com/bdero/flutter_scene), 2021.
- [173] Even Rouault, Frank Warmerdam, Kurt Schwehr, Andrey Kiselev, Howard Butler, Mateusz Loskot, Tamas Szekeres, Etienne Tourigny, Martin Landa, Idan Miara, Ben Elliston, Kumar Chaitanya, Lucian Plesea, Daniel Morissette, Ari Jolma, Nyall Dawson, Daniel Baston, Craig de Stigter, and Hiroshi Miura. Gdal: Geospatial data abstraction library, version 3.11.4, 2025. URL <https://zenodo.org/records/17098582>.
- [174] Jakob Hoydis, Fayçal Ait Aoudia, Sebastian Cammerer, Merlin Nimier-David, Nikolaus Binder, Guillermo Marcus, and Alexander Keller. Sionna: An open-source library for next-generation physical layer research. <https://nvlabs.github.io/sionna/>, 2023. Accessed: 2025-02-27.
- [175] vvoovv. Blossm: addon for blender (openstreetmap, google 3d cities, terrain). <https://github.com/vvoovv/blossm>, 2021. branch release.
- [176] Geoscience Australia. Digital Elevation Data. <https://www.ga.gov.au/scientific-topics/national-location-information/digital-elevation-data>, March 2025. Accessed: 2025-06-24; Page last updated: 5 March 2025.
- [177] João Borges, Felipe Bastos, Ilan Correa, Pedro Batista, and Aldebaro Klautau. Caviar: Co-simulation of 6g communications, 3d scenarios and ai for digital twins. arXiv preprint arXiv:2401.03310, 2024.



- [178] Ruizhi Cheng, Nan Wu, Matteo Varvello, Eugene Chai, Songqing Chen, and Bo Han. A first look at immersive telepresence on apple vision pro. In Proceedings of the 2024 ACM on Internet Measurement Conference, IMC '24, page 555–562. ACM, November 2024. doi: 10.1145/3646547.3689006. URL <http://dx.doi.org/10.1145/3646547.3689006>.
- [179] NVIDIA. Supercharge robotics workflows with ai and simulation using nvidia isaac sim 4.0 and nvidia isaac lab, 2024. URL <https://developer.nvidia.com/blog/supercharge-robotics-workflows-with-ai-and-simulation-using-nvidia-isaac-sim-4-0-and-nvidia-isaac-lab/>. Accessed: 2024-09-17.
- [180] Mayank Mittal, Calvin Yu, Qinxi Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, Ajay Mandlekar, Buck Babich, Gavriel State, Marco Hutter, and Animesh Garg. Orbit: A unified simulation framework for interactive robot learning environments. IEEE Robotics and Automation Letters, 8 (6):3740–3747, June 2023. ISSN 2377-3774. doi: 10.1109/lra.2023.3270034. URL <http://dx.doi.org/10.1109/LRA.2023.3270034>.
- [181] Paolo Testolina, Michele Polese, Pedram Johari, and Tommaso Melodia. Boston twin: the boston digital twin for ray-tracing in 6g networks. In Proceedings of the ACM Multimedia Systems Conference 2024 on ZZZ, MMSys '24. ACM, April 2024. doi: 10.1145/3625468.3652190. URL <http://dx.doi.org/10.1145/3625468.3652190>.
- [182] Jakob Hoydis, Fayçal Aït Aoudia, Sebastian Cammerer, Merlin Nimier-David, Nikolaus Binder, Guillermo Marcus, and Alexander Keller. Sionna rt: Differentiable ray tracing for radio propagation modeling. In 2023 IEEE Globecom Workshops (GC Wkshps), pages 317–321. IEEE, 2023.
- [183] Sunday Amatare, Gaurav Singh, Aavash Kharel, and Debashri Roy. Real-time localization of objects using radio frequency propagation in digital twin. In MILCOM 2024 - 2024 IEEE Military Communications Conference (MILCOM), pages 653–654, Oct 2024. doi: 10.1109/MILCOM61039.2024.10774060.
- [184] Sunday Amatare, Wentao Gao, Aavash Kharel, Raul Shakya, Mohammad Hasibur Rahman, Xiaojun Shang, and Debashri Roy. RF-Vision: Object Characterization using Radio Frequency Propagation in Wireless Digital Twin. SSRN Electronic Journal, Nov 2024. doi: 10.2139/ssrn.5136513. URL <https://ssrn.com/abstract=5136513>.
- [185] Xin Liu, Yu Wang, Raymond C.H. Koo, and Julian S.H. Kwan. Development of a slope digital twin for predicting temporal variation of rainfall-induced slope instability using past slope performance records and monitoring data. Engineering Geology, 308: 106825, 2022. ISSN 0013-7952. doi: <https://doi.org/10.1016/j.enggeo.2022.106825>. URL <https://www.sciencedirect.com/science/article/pii/S0013795222003106>.



- [186] Mark Rutherford. The issues with jamming drone frequencies, 2022. URL <https://d-fendsolutions.com/blog/issues-with-jamming-drone-frequencies/>. Accessed: 2025-03-01.
- [187] Jacob B. Why is 5.8ghz used for fpv and 2.4ghz used for transmitters?, 2020. URL <https://drones.stackexchange.com/questions/810/why-is-5-8ghz-used-for-fpv-and-2-4ghz-use-for-transmitters>. Accessed: 2025-03-01.
- [188] Radio-controlled aircraft. URL [https://en.wikipedia.org/wiki/Radio-controlled\\_aircraft](https://en.wikipedia.org/wiki/Radio-controlled_aircraft). Accessed: 2025-03-01.
- [189] Mohammed A. Jasim, Hazim Shakhatreh, Nazli Siasi, Ahmad H. Sawalmeh, Adel Aldalbahi, and Ala Al-Fuqaha. A survey on spectrum management for unmanned aerial vehicles (uavs). IEEE Access, 10:11443–11499, 2022. ISSN 2169-3536. doi: 10.1109/ACCESS.2021.3138048.
- [190] Jaber Kakar and Vuk Marojevic. Waveform and spectrum management for unmanned aerial systems beyond 2025. In 2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC), pages 1–5, Oct 2017. doi: 10.1109/PIMRC.2017.8292533.
- [191] Robert Anthony Hunjet. Adaptive topologies : improving wireless networks through the use of additional nodes and power control, 2014.
- [192] Patrick Carlyon and Ashlea Hansen. Mushroom cook trial expert’s new cold case alibi claim in rachelle childs investigation. The Advertiser / AdelaideNow, page 10:14AM, June 2025. URL <https://www.adelaidenow.com.au/truecrimeaustralia/dear-rachelle/dear-rachelle-investigation-explosive-evidence-could-blow-open-murder-investigation/news-story/2e4534ac98d84d3c5a3c5d9481f4b0c4>. Accessed: 2025-09-20.
- [193] Sean Fewster. Mushroom cook: Adelaide university research claims new details in erin patterson’s death cap picking trip. The Advertiser / AdelaideNow, July 2025. URL <https://www.adelaidenow.com.au/truecrimeaustralia/police-courts-sa/mushroom-cook-adelaide-university-research-claims-new-details-in-erin-pattersons-death-cap-picking-trip/news-story/a0bce627e7c3cab03d949585a76810fe>. Accessed: 2025-09-20.
- [194] Ming-tuo Zhou, Vinh Dien Hoang, Hiroshi Harada, Jaya Shankar Pathmasuntharam, Haiguang Wang, Peng-yong Kong, Chee-wei Ang, Yu Ge, and Su Wen. Triton: high-speed maritime wireless mesh network. IEEE Wireless Communications, 20(5):134–142, October 2013. ISSN 1558-0687. doi: 10.1109/MWC.2013.6664484.





- [195] Te Wei, Wei Feng, Yunfei Chen, Cheng-Xiang Wang, Ning Ge, and Jianhua Lu. Hybrid satellite-terrestrial communication networks for the maritime internet of things: Key technologies, opportunities, and challenges. IEEE Internet of Things Journal, 8(11): 8910–8934, June 2021. ISSN 2327-4662. doi: 10.1109/JIOT.2021.3056091.
- [196] Jingbo Zhang, Jianyu Yang, Yiyang Zhang, and Shufang Zhang. A dynamic spectrum allocation algorithm for a maritime cognitive radio communication system based on a queuing model. Information, 8(4), 2017. ISSN 2078-2489. doi: 10.3390/info8040119. URL <https://www.mdpi.com/2078-2489/8/4/119>.