Multi-objective site selection and analysis for GSM cellular network planning

A thesis presented

by

Larry Raisanen

to

The Department of Computer Science

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Computer Science

Cardiff University

Cardiff, Wales

September 2005

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Thesis advisor

Dr. Roger Whitaker

Author

Larry Raisanen

Multi-objective site selection and analysis for GSM cellular network planning

Abstract

Although considerable effort has been placed on developing techniques and algorithms to create feasible cell plans, much less effort has been placed on understanding the relationship between variables and objectives. The purpose of this thesis is to improve the body of knowledge aimed at understanding the trade-offs and tensions in the selection of transmission sites and in the configuration of macro-cells for GSM and related FDMA wireless systems. The work begins by using an abstract 2-dimensional (2D) model for area coverage. A multiple objective optimisation framework is developed to optimise the sequential placement and configuration of downlink wireless cells. This is deployed using a range of evolutionary algorithms whose performance is compared. The framework is further tuned via a decoding mechanisms using the best performing evolutionary algorithm. The relationship between primary variables in the 2D model is analysed in detail. To improve realism, the thesis additionally addresses complexities relating to planning in 3-dimensional (3D) environments. A detailed open source static model is developed and the optimisation framework is extended to accommodate the additional model complexities and choices in algorithm design are compared. Finally, sensitivity analysis is performed to determine the relationship between objectives in the 3D model and benchmark solutions are provided.

Contents

List of Figures

List of Tables

Abbreviations

2D	Two dimensional
3D	Three dimensional
ACP	Antenna configuration problem
$a(h_m)$	Correction factor in Hata's empirical model
AI	Angle of incidence
AIM	Angle of incidence matrix
α	Pair-wise cell overlap constraint expressed as value between 0-1
APP	Antenna placement problem
Ave.	Average
az	Azimuth
β	Multi-wise cell overlap constraint expressed as value between 0-1
C	Cell
CBS	Candidate base station, candidate site (location)
CBS'	Cell plan (based on all CBS in π with non-zero power)
Cove.	Cover
CPP	Cell planning problem (both APP and ACP)
f	Frequency
G	Gains (associated with transmission)
G2PS	Generational two phase strategy
G3PS	Generational three phase strategy
GA	Genetic algorithm
Gens.	Number of generations
GSD	Greedy sequential decoder

GSM	Global system for mobile communications
h_b	Height of base station
HDIAG	Horizontal antenna diagram
h_m	Height of mobile receiver
L	Losses (associated with transmission)
LD	Large directive antenna
MOA	Multiple objective algorithm
MOP	Multiple objective problem
MCO	Multi-wise cell overlap (constraint)
MCO-P	Multi-wise cell overlap (constraint)–iterating by-power
MCO-S	Multi-wise cell overlap (constraint)–iterating by-site
MOGA	Multiple objective genetic algorithm
NSGA	Non-dominated sorting genetic algorithm
OD	Omni-directional antenna
Over.	Overlap
p	Power
P^A	Power of transmission
P^A_{STP}	Power of transmission at a given STP from antenna A
PCO	Pair-wise cell overlap (constraint)
PCO-P	Pair-wise cell overlap (constraint)–iterating by-power
PCO-S	Pair-wise cell overlap (constraint)–iterating by-site
PESA	Pareto envelope-based selection algorithm
π	Permutation (associated with decoder)

PL	Propagation loss
PLM	Propagation loss matrix
PMCO	Pair-wise and multi-wise cell overlap (constraint)
PMCO-P	Pair-wise and multi-wise cell overlap (constraint)–iterating by-power
PMCO-S	Pair-wise and multi-wise cell overlap (constraint)–iterating by-site
R	Distance
RTP	Reception test point
S	Spread (as a measure of solution distribution)
SD	Standard deviation (duplicate reference below)
SD	Small directive antenna (duplicate reference above)
SE	Simulation environment
SEAMO	Simple evolutionary algorithm for multi-objective optimization
Sols.	Number of solutions
SOP	Single objective problem
SPEA	Strength Pareto evolutionary algorithm
S_{PFA}	Set based on a Pareto front approximation
S_q	Service threshold
STP	Service test points
Traf.	Traffic
TSCA	T_{MAX} site configuration algorithm
TTP	Traffic test points
VDIAG	Vertical antenna diagram
W	Working area

N	Total number

Acknowledgments

Firstly, I would like to thank my wife Sarah, who provided the emotional support I needed to complete this work. Without her, this effort would probably have terminated prematurely.

Secondly, I would like to thank my supervisor, Dr. Whitaker, for helping me secure an opportunity to work on a PhD. Also, for providing considered and detailed comments to improve papers and this thesis along the way.

Thirdly, I would like to thank Cardiff University and the EPSRC for funding me in one way or another for three years.

Fourthly, I would like to thank all those people involved in the scientific endeavor as a means to forward science and aid humanity, and never abuse it as a means to seek profit or exclude others.

Finally, I need to give thanks to some of the people who, by virtue of providing information freely, helped this work take shape:

- I would like to thank everyone who helped create the Colt Distribution of Open Source Libraries for High Performance Scientific and Technical Computing in Java. This library saved me a lot of time, by providing high quality java classes for performing a host of tasks. I wish I would have found it and used it sooner.

- I would like to thank everyone who helped create the JISC funded Landmap Project: a joint project between MIMAS and University College London Dept. of Geomatic Engineering and other project partners. The project provided a high resolution Digital Elevation Model for the British Isles which was used in creating data sets.

- I would like to thank James Gosling, and everyone at Java Sun, for creating the Java programming language. I would also like to thank all the contributors to the Java forums for providing useful and considered responses to my queries in such a timely manner.

- I would like to thank everyone at the Apache Software Foundation, which provides support for the Apache community of open-source software projects. I used Ant, Forrest, and the Jakarta POI to save some time while running the hundreds of intensive trials provided herein.

- I would like to thank everyone at LaTeX who provided such a fine high-quality typesetting system. The production of the many research and technical papers (along with this very thesis) would not have been so easy without this sensible word-processing alternative.

Without all of you, this work would not have been possible.

Dedicated to my wife, Sarah.

Chapter 1

Introduction

Mobile communication systems help meet the growing needs of people who are on the move, and who want reliable access to telecommunication services wherever they are. In Europe, countries decided, as part of the general ideology behind the European Union, to develop a pan-European mobile phone standard in 1982 to allow users to roam throughout Europe with mobile voice and data services. This was termed the global system for mobile communications (GSM) when finalized in 1991. Today, the GSM Association reports that GSM is the fastest growing communications technology of all time, with the billionth user connected in the first quarter of 2004, which is only 12 years after the commercial launch of the first GSM networks. They also report that GSM accounts for 75% of the world's digital mobile market and 74% of the world's wireless market, with operators in more than 210 countries and territories. This makes GSM not only the European standard but the global standard for second generation mobile telephony. This tremendous growth has made good and efficient mobile network design highly desirable.

This thesis is devoted to improving the body of knowledge aimed at understanding the relationships, trade-offs, and tensions between variables and objectives during the configuration of macro-cells for GSM and related FDMA wireless systems. Together with a detailed approach to achieving automatic site placement, operators will be empowered to make better choices during network planning, which will aid them in meeting the demands of subscribers more efficiently. Section 1.1 sets the background to the problem of mobile network design and the motivation for resolving it. Section 1.2 provides the research scope and objectives. Finally, Section 1.3 delineates the structure of the thesis.

1.1 Background and Motivation

To allow subscribers the freedom to roam anywhere within a service area adequate received signal strength needs to be made widely available. As the roaming provision and number of subscribers increases, the density of sites needed to meet demand increases. The proliferation of cellular wireless services for mobile communication has led to the *antenna placement problem* (APP).

For cellular wireless systems, mobile communication is facilitated by base stations which have an appropriate spatial distribution. The area of service coverage from a single antenna at a base station constitutes a *cell*, which is a region where the radiated signal power from the serving antenna is of sufficient strength to be received by subscribers. As the power of transmitted signals must be restricted, multiple cells are required to provide wide area coverage. The collection of cells across the service area constitutes a *cellplan* or *network*.

The APP involves selecting base station site locations from a set of candidates, which are dispersed irregularly and which achieve varying qualities of signal propagation due to the surrounding terrain, buildings, antenna heights, etc. Selected sites must be configured to provide adequate service coverage and capacity while adhering to constraints involving regions which can be served by more than one antenna. Such constraints are imposed to ensure that the potential for interference is controlled while providing regions in the network for call *handover*, which are necessary for seamless call transfer between cells. Areas covered by more than one cell, termed *overlap* or $multi - coverage$, must be carefully controlled both to maintain network operation and minimize the total commitment to infrastructure.

Control over the area covered by a cell is gained first by controlling the spatial density of cells, and secondarily through antenna configuration. Antennas at selected sites need to be configured to meet service demands. This is an adjunct to the APP and is known as the *antenna configuration problem* (ACP). The difficulty is there are a wide range of potential configurations at each site, starting with the type and number of antennas and ending with choosing the best transmission powers, azimuths, and tilts. Efficiently configuring antennas is not a simple task. Not only are there many potential combinations of settings, but changing the configuration at one site can affect other sites. For example, if power is reduced in one cell to reduce its traffic load, another cell might pick up the unloaded traffic, and then itself become *overloaded.* This makes it difficult to perform local optimization, as it inevitably affects the global picture. This is why the computational complexity of this problem is so high, because if one makes even a small local change, measurements throughout the network need

to be taken if certain hard constraints, for instance, are being followed.

Finally, it is highly desirable that the APP and ACP be achieved as quickly as possible, particularly with the growing trends toward dynamic structuring, and using the fewest number of sites possible in order to keep costs low, as commissioning sites represents an operator's dominant cost. In short, the overall goal of cell planning is to select a sufficient subset of base station sites at the appropriate spatial density from a set of candidate sites to meet network design objectives, and to configure settings at base stations to meet subscribers' needs and minimize an operator's financial outlay. Discovering the models, settings, and objectives requires forging links with partners in industry who can benefit from the work academics provide.

Due to the fact that local changes affect the global landscape, and the resultant interdependencies between multiple and often conflicting objectives, resolving the APP and ACP is a circular problem: Ideally, to select the best sites one needs to know how sites are configured, but to configure sites one needs to know which sites are selected. Therefore, if the relationships, trade-offs, and tensions between variables can be clarified, the key to unlocking this circular problem may be found and the problem reduced to its most basic components. For example, it is possible that if there are five objectives to meet, investigators are trying to actively meet all five objectives. However, if there is a strong relationship between one of these objectives and the remaining four, it is not infeasible to imagine that instead of attempting to design a network based on five objectives, one could be designing based on one or two.

The number of factors involved in solving the APP and ACP means that auto-

matic software for designing cell plans has become increasingly common [85]. We classify software as *automatic* if the associated computer program has autonomy in the selection of base station locations and the configuration of antennas. Because of the high demand for GSM services, and the pressing need of operators to find the lowest cost base station infrastructure to reliably meet the growing needs of subscribers, *automatic cell planning* which can optimise a cellular network configuration and expedite the engineering process by eliminating manual interventions, decision making, and judgments, is highly desirable.

The underlying algorithm in the software is required to tackle an NP-hard [54] optimization problem with 2^n solutions for the APP given n candidate sites or $(m+1)^n$ for the ACP (with m configurations possible at any site). Consequently, heuristic and meta-heuristic techniques have become increasingly popular for solving the APP.

1.2 Research Purpose and Scope

The aim of the research is (i) to produce a methodology to effectively quantify the relationship between pertinent variables in cell planning and (ii) to deploy the methodologies in a variety of cell planning scenarios to clarify their impact.

The approach starts in the 2D model as a greedy sequential decoder and becomes a more complex site configuration algorithm in the 3D model. By model we mean the information necessary to create and use a data set capable of simulating a static environment for GSM cell planning. Although we do not seek to concurrently resolve the more well-studied channel assignment problem during this approach, every effort is taken to minimize multi-coverage in the 2D model and co-channel interference in

the 3D model. The reason for optimizing a sequential approach to cell planning is to create a flexible model which can be used to quickly configure a network, as the current trend is toward dynamic dimensioning, or which could be easily adapted to find a very refined solution, as might be the case for an operator using static dimensioning.

To use the approach a simulation environment for investigating the APP and ACP which can accurately measure important operator objectives is designed. To this end, we use an available 2D model for some theoretical investigations, and then engineer a more realistic 3D model capable of measuring all pertinent network design criteria. This is an important step as there is currently no standardized, realistic GSM model used in the literature which is made publicly available. By publicly available we mean that the information is available for free in the public domain for personal and academic use.

The approach is then used to discover the fundamental tensions, conflicts, and relationships between interrelated network design criteria, particulary as 100% service coverage is approached. This is important as operators commonly wish to meet multiple design criteria (such as high service coverage, high system capacity, high roaming provision–call handover, low cost, low interference) simultaneously. As the goal is to satisfy multiple objectives simultaneously, this is termed a multi-objective optimization problem (MOP).

It has become increasingly common to resolve MOPs using genetic or evolutionary algorithms, as they are well-suited to this task. This thesis focuses primarily on resolving two fundamental cell planning objectives: providing the required *service*

coverage at the lowest possible *financial cost.* These two conflicting objectives always exist when setting up cellular network services, as adding base stations to improve service coverage inherently increases the cost of the network. Despite the importance of finding an optimal trade-off between these two objectives, we are not aware of any studies directly addressing this issue. This may occur because the focus tends to be on producing one best cell plan, rather than a range of cell plans.

The cell plans we produce are the first step in establishing a cost effective operational network. In the 2D model, the tension between cost and coverage is resolved without detailed *dimensioning* at individual cells [21]. During the 2D approach, only the power of transmission is altered. In the 3D model, individual cells undergo dimensioning. This second stage may involve adjusting additional variables at selected base stations such as tilt and azimuth, in order to ensure the traffic load does not exceed 43 erlangs, which is a setting commonly employed (e.g., [89, 81, 40]). Additionally, multiple directed co-sited antenna may be invoked at this stage to increase the capacity for multiplexing, given knowledge of anticipated traffic patterns. Known as sectorization, this is common operational practice in mobile telephony, whereby using multiple co-sited antenna is generally far cheaper than commissioning a new site.

To optimally resolve the competition between service coverage and financial cost, we introduce a multiple objective optimization framework that does not require *a priori* knowledge of the relative importance of service coverage versus cost. This is achieved by providing a range of alternative site selections which approximate the best possible trade-offs (i.e., Pareto front) between cost and coverage. This means

that unlike the current convention for cell planning, which generally seeks to generate a single cell plan given information on the relative importance of objectives, a radio engineer will be able to choose from a range of alternative cell plans, given visual and detailed information regarding each. This is particularly beneficial when there is a non-linear relationship or unknown dependency between the objective functions as for the general APP. This is also useful given the demarcation between the needs of operators who wish to entice subscribers with high quality service levels, versus low-cost providers who, to keep overheads low, may simply seek the lowest cost per unit of coverage.

As far as we are aware, this method has only been considered for cell planning in [57], where a genetic algorithm was developed specifically for the APP. However, the total network cost was not considered. Unlike [57], the framework we propose considers financial cost and is flexible because it is possible to 'plug-in' any multi-objective optimization algorithm (MOA) which seeks to approximate a Pareto front. This flexibility is achieved by making the cell plan representation independent from the task of the MOA, which, in the case of an integer-coded permutation, is to find optimal orderings of candidate site locations to optimise two objectives. Exploring a search space in this manner is common practice in many discrete optimization problems and has also been successfully applied to partitioning problems, the minimum span frequency assignment problem, the travelling salesman problem, and set covering problems [43].

1.3 Thesis Outline

The main contributions following the introduction in Chapter 1 are the following:

- In Chapter 2: **A review of the literature is undertaken.** We start be reviewing GSM cell planning (including past approaches, simulation environments, and cell plan quality measurement). We show that the most realistic GSM simulations are not made publicly available, that most approaches to planning take a long time to resolve and do not tackle the circular problem of cell planning, and that the fundamental tensions and conflicts are not investigated. We report that many previous approaches have successfully used genetic algorithms.

  We proceed to discuss the use of different algorithms for cell planning (including sequential and greedy algorithms, genetic algorithms, and multi-objective genetic algorithms–MOGAs). After this, the parameters of MOGAs are delineated, including representation, fitness evaluation, selection, recombination and mutation, search space, and parameter tuning.

  Finally, we cover the metrics used to evaluate non-dominated solutions, including domination, pareto-optimality, weak-domination, and the set cover metric. This provides the necessary background for evaluating multiple non-dominated cell plans found in a number of testing situations presented.

- In Chapter 3: **A 2D simulation model and multi-objective cell planning strategy are introduced and tested.** The 2D cell planning model is a simplified network model appropriate for theoretical investigations. Within

Chapter 3, it is used to investigate the suitability of four multi-objective algorithms to find an optimal ordering of candidate sites, where the order of sites is used by a greedy decoder to determine a cell plan. The problems consider area coverage and cost. Results are expected to indicate the most effective multi-objective algorithm, which is carried forward and used throughout the thesis.

- In Chapter 4: **The 2D framework is tuned.** Six decoders, which differ by how they control cell overlap and how they select the next candidate site, are tested on their ability to optimise coverage and cost by controlling cell density. Results indicate the most effective decoder.

- In Chapter 5: **2D cell planning is analyzed.** The rate at which infrastructure cost changes is fundamental in determining the amount of coverage, and related capital investment, which the operator should employ to maximize profit. This clarifies the relationship between coverage and cost and suggests why investigators and operators might not wish to simply maximize coverage.

- In Chapter 6: **A 3D simulation model is engineered, design objectives defined and single site coverage tested.** This is undertaken to create a realistic wireless simulation environment and define network components capable of allowing the measurement of all pertinent design objectives. The goal of the 3D model is to meet or exceed the quality of all other available models and be the first whose method of creation and contents are made publicly available. This assures results presented herein can be challenged and extended in future

scientific works.

- In Chapter 7: **A 3D cell planning strategy is extended from the 2D model.** Changes in the model (in particular, traffic load and increased site complexity) lead to a pre-processing phase which considers each site locally, and in the absence of all others. This phase is termed TSCA, and it is introduced to the cell planning strategy to configure cells before optimizing on coverage and cost. At this time, a binary-representation model is also added to compare and contrast to the standard integer-representation used earlier.

- In Chapter 8: **Pertinent 3D variables and parameters are tuned and tested.** Results indicate a trade-off for many parameter settings (e.g., GA population size) between solution quality and speed of execution, with different settings being more appropriate depending upon the application. These are discussed and algorithms for 3D planning are tuned.

- In Chapter 9: **Tests to compare 3D cell planning strategies and objective relations are completed.** Results distinguish between cell planning approaches, and clarify the correlations between design objectives, and generally indicate if using optimized settings in the cell planning strategy is capable of meeting all pertinent network design objectives.

Finally, conclusions are drawn and presented in Chapter 10, along with directions for future work.

Chapter 2

Review of Literature

2.1 GSM Cell Planning

GSM cell planning is a complicated task. First, developing a GSM simulation environment which is appropriate and practical for solving real world problems is difficult. Three main contending models with varying strengths and weaknesses are discussed in Section 2.1.1. A common fault in each case that affords a high degree of realism is that the topographical and traffic data used is not made publicly available, nor is how this data incorporated or used made transparent. This is a severe limitation as it eliminates the possibility for other researches to challenge findings or beat benchmarks solutions. We address this when developing our own model in Chapter 6.

Secondly, an approach to cell planning needs to be developed to tackle the NP-hard cell planning problem (CPP), which may include aspects of both the antenna placement (APP) and configuration (ACP) problems. Consequently, heuristic and

meta-heuristic techniques have become increasingly popular for solving the CPP. A number of contending algorithms with varying strengths and weaknesses are discussed in Section 2.1.2, and the ones selected for use in this thesis are discussed in more detail in the Section 2.2. Selecting a well-researched algorithm for use in the optimization process allows the primary focus of this thesis to remain on developing new optimization methodologies. This allows attention to be placed on tailoring and applying these methodologies for purposes of this application.

Finally, comparing the performance of multiple objective algorithms is problematic because a set of solutions rather than a single solution is obtained. Several alternative approaches to this are discussed in Section 2.3.

2.1.1 Simulation environments

Models for simulating wireless coverage range from the abstract (e.g., graph theory models) to the detailed (e.g., models based on specific areas with known traffic and topology). Models in-between these two extremes either seek means for increasing the realism of an abstract model, or reducing the complexity of detailed models in order to increase computational speed. Thus, each simulation model represents a trade-off involving the level of detail and computational speed. The three main models for simulating wireless networks include the demand node concept, disk graph models, and discrete test point or cell models.

Demand node concept

The demand node concept was first introduced in [44], and has been used in several studies since (e.g., [5, 24, 83, 77]). The core concept is that a demand node represents the centre of an area where traffic is being generated by users. The main advantage of this is that by combining the traffic from a small geographical area into a single point (or demand node), the number of computations required during cell planning is reduced, but in the process realism is also reduced. For example, in [83] only 288 demand nodes were required for an area 9 km^2, whereas in a discrete model with test points uniformly distributed every 200m there would be 2025 'nodes'. There are much fewer demand nodes in effect because each subsumes a number of test points. However, the process of amalgamating a number of test points into one demand node has the same effect as any type of compression (e.g., CD music to MP3): the resolution or clarity is diminished. In short, fine resolution (e.g., nodes every 50m) affords more realism than low resolution (e.g., nodes every 500m).

Investigators using the demand node model to date have allowed freedom in selecting candidate sites. This makes the possibility of finding a suitable site in any given situation much more likely than in models using a restricted number of potential sites. For example, [83] distributed candidate sites uniformly every 100m, which results in 8,100 candidate sites. This translates to a candidate site density per km^2 of 100.00 versus, for example, 0.24 in the highest candidate site density test scenario used in [63]. The strongest argument against using uniform candidate site distribution in simulation models is that it is unlikely many network operators are granted such freedom.

An advantage of this approach is that certain constraints can be formulated to ensure only cells which meet the constraints are allowed during cell planning. For example, one might only consider cells which cover less than a certain number of demand nodes, and thereby ensure that cells cannot become overloaded. However, as we will see, this advantage is also present in other models, and therefore such advantages have more to do with the planning algorithms employed, and less to do with the simulation model. Although not a limitation of this simulation per se, investigators using this model have not investigated cell planning expansion scenarios.

Disk graph model

The application of disk graphs to cellular system design was first used by [28] for use in resolving the related channel assignment problem. The intersection disk graph model (with non-uniform traffic distribution) has been used subsequently by [37, 38, 39], and the unit disk graph (with uniform traffic distribution) by [53]. The strength of the approach used in [37], for example, is that it allows consideration of many pertinent network design objectives. This means that one can try to resolve CPP and the channel assignment problem simultaneously. The computational time using this approach is also low.

However, a major weakness of this approach is that potential transmitters are often assigned only one out of a range of configurations, which inevitably restricts the total number of possible designs. An additional and significant weakness is that the disk approach assumes idealised propagation modeling (i.e., identical cell shape). While the size of cells can be altered based on non-uniform traffic distribution ([38]),

the shape of the cell is always the same, which tends to be circular. Another problem is that the hierarchical optimization method often employed with the disk graph model (e.g., see [37]) allows each pixel to be a potential base station site (or disk), and therefore suffers from the same problems as mentioned for the demand node in terms of allowing uniform candidate site distribution. On the other hand, there does not seem to be any reason why a limited set of sites could not be used. It appears this is a design choice made by the investigators and not necessarily a limitation of the model.

Test point and cell models

The test point model is perhaps most clearly delineated in [68] and [67], but had appeared in similar form in earlier works (e.g., [31]). In this model, the working area or simulation area W is discretized into test points at Cartesian coordinates (i.e., x,y) at a given resolution. For example, test points may be available every 100m over a 10 km^2 area. Test points are used in order to allow signal strength measurement (see RTP below) in areas which an operator wishes to service (see STP below) because of traffic demand by customers (see TTP below). The three sets of test points distributed in W are reception test points (RTP) where signal reception can be tested, service test points (STP) where signal quality must exceed a signal quality threshold to be useable by a customer, and traffic test points (TTP) where a given amount of traffic (in erlangs) from customers using their mobile devices is associated. In this model, the TTP $\subseteq$ STP $\subseteq$ RTP.

This model has been used by the current author (e.g., [62, 64, 63]) and the

ARNO (*Algorithms for Radio Network Optimization, IT Project 23243*) partners (e.g., [10, 81, 89]). It has also been arguably adapted by [29] who extended the model with a typical urban area map, altitude data, and hypothetical traffic data, although it is not made evident how all this data was specifically used to enhance the simulation. This idea appeared also appeared earlier in [31], who instead of test points, divided the simulation area into cells and incorporated data regarding topology, traffic distribution, and population, although, again, how and to what extent this data was used was not made clear.

The advantage of this approach is that it allows the measurement of all network objectives (e.g., coverage and capacity). Also, although candidate site locations (CBS) are defined by their Cartesian coordinates, they are not uniformly distributed. CBS are also restricted, unlike either of the previous two models, which more realistically reflects an operator's planning environment. Like the disk graph model, investigators here have also considered cell planning expansion scenarios (e.g., [40]).

A disadvantage is that investigations using this approach have had the highest computational time, as investigators tend to use high resolution (e.g., test points every 200m) to increase realism. However, a trade-off could be made to increase speed by, for example, using a lower resolution (e.g., test points every 400m). However, this would be at the expense of realism.

2.1.2 Past approaches to the cell planning problem

The first published paper on optimizing antenna placement dates back to 1994 [6]. Since then a large number of approaches and scenarios have appeared in the

literature. Although exact approaches are only feasible for relatively small test problems, they have been applied in a range of papers [58, 59, 60, 84]. However, in some papers, they have been relaxed or selectively applied. Sequential, or greedy, algorithms have been less well used and then predominantly for comparison purposes [58, 5]. Deterministic heuristic algorithms have also been proposed by a range of authors [14, 42, 76, 25, 77, 58, 59, 24, 55, 89]. Frequently, these approaches exploit observations (e.g., density of base station locations) about the CPP and incorporate them to enhance the performance of the cell plans obtained. However, meta-heuristic algorithms based on simulated annealing [1], tabu search, and genetic algorithms [33] are far more popular.

Simulated annealing has been adopted for the CPP in [2, 6, 54, 3, 40], and tabu search for the CPP in [29, 52, 81, 4]. Both these techniques operate by ranking solutions using a cost function. Given a solution, small changes are made to create a neighborhood of solutions from the current solution. The meta-heuristic then guides the acceptance of new solutions available in the neighborhood. The advantage of this approach is that it has the ability to escape from local minima in the search space (regarding the cost function), thereby improving performance. Differences in the application of these approaches involve how the cell planning problem is modelled, the formulation of rules to create neighborhoods, and the cost function used to rank solutions.

Genetic algorithms (GAs) have also become increasingly popular for the CPP [52, 50, 29, 38, 13, 58, 57, 69]. With the exception of [57, 36], these approaches predominantly seek to optimise a single function (or a linear combination of multiple

objective functions) to create a population of high quality solutions. These algorithms mimic evolution and natural selection through fitness assignment, selection, recombination, and mutation on a population of solutions. For a genetic algorithm to succeed, a suitable representation of the problem needs to be used. While the most popular representation for GAs in the literature is a binary string, several authors in cell planning have introduced the notion of real valued representations which provide information regarding the number of CBS and their locations (e.g., [29]).

2.2 Algorithms

2.2.1 Greedy sequential algorithms

Greedy algorithms are used to produce a good solution quickly by selecting a *myopically* best addition until a feasible solution is constructed. While this process cannot guarantee an optimal solution is found, it does limit the number of steps required, and in this way, tends to be faster than alternatives (e.g., genetic algorithms). A typical means for minimizing the number of steps is to follow a sequential pattern (e.g., by next position in a list). An example use of the greedy sequential algorithm (GSA) is in bin packing problems, where the goal it to place n items into bins of common capacity, C. The problem is to minimize C given the number of bins is m. A possible GSA to solve this problem would be to order the items by size (smallest first) and then put the next item (in sequential order) into the bin with the smallest size. In this way, all items are placed into bins in n iterations, after which the bin with the largest size defines C. While the disadvantage of GSA is that alternative

algorithms could find a better solution, it is unlikely they could do it in fewer steps. Thus, the advantage of GSA is that it requires less computational effort.

In cell planning, GSA works by commissioning antennas or sites based on a given ordering, not allowing antenna reconfiguration or decommissioning. The quality of the solutions found will depend largely upon the criteria used to configure antennas initially and the order in which candidate sites are considered. GSA have been used to resolve the related graph based problems and channel assignment [9, 8] and during cell planning to meet coverage and interference constraints [77].

2.2.2 Genetic algorithms

Genetic algorithms are stochastic search algorithms inspired by Darwin's principles of natural selection, by which the most adapted members of the population survive long enough to reproduce and thereby evolve the general population toward higher levels of fitness, which first appeared in 1859 [16]. It was over a century later (1962) before the first work on genetic algorithms appeared in a series of papers on adaptive systems theory by J.D. Bagley, who in 1967 coined the phrase 'genetic algorithm' and wrote its first application [7]. Eight years later (1975), the first seminal theoretical book was published by John Holland, often considered the father of genetic algorithms [33]. However, only over the last decade have GAs been successfully adapted to solve multiple objective problems. An excellent overview of this area is given in [17].

The general principle behind GAs is to breed a new population through a process of selection, recombination, and mutation. This occurs over a number of generations

to try to improve the performance of the population. The expectation is that desirable characteristics in solutions from one generation will combine to produce better solutions for the next generation. While these algorithms can be relatively simple in structure, they can be powerful mechanisms to search for improvements in a given population.

There are five procedures involved in setting up a basic genetic algorithm: representation, fitness evaluation, selection, genetic modification (crossover and mutation), and replacement. The order of these is as appears in Algorithm 1. As with all machine learning problems it is worth tuning the GA's parameters, such as mutation rate and recombination probability to find appropriate settings for the problem class being worked on. Theoretical upper and lower bounds for these parameters can help guide the initial selection.

Representation

The problem to be solved must first be represented by chromosomes, in such a way that an objective evaluation can take place of the solution the chromosome represents. Chromosomes are typically represented as simple bit-strings, integer-strings, data, or instructions, although there is no limit to the variety of data structures which could be employed. Different representation schemes can lead to differences in terms of solution quality and computational time.

Fitness evaluation

After the representational model has been decided, an initial parent population is generated. The parent population may be randomly determined or a heuristic could

Algorithm 1 Basic genetic algorithm

```
1: Choose appropriate representation
2: Initialize parent population of size n
3: Evaluate each parents's fitness
4: repeat
5:     Create offspring population of size n through reproduction
6:     for i = 1, ..., n do
7:         Select best-ranking parents to reproduce
8:         Mate pairs according to given scheme
9:         Apply crossover operator
10:        Apply mutation operator
11:    end for
12:    Evaluate each offspring's fitness
13:    Replace less fit parents with offspring
14: until Terminating condition
```

be used to form more suitable possible solutions. During each successive generation, each individual is evaluated, and a 'goodness' or fitness value is returned by a fitness function based on desirable objective measures. The population is then ranked such that those having better fitness (representing better solutions to the problem) are at the top of the list for use in generating offspring.

Selection

Generating offspring is done by applying any or all of the potential genetic operators: selection, crossover (or recombination), and mutation. A pair of parents (with those near the top more likely to be chosen) are then selected for breeding. In other words, the chance for being selected is directly proportional to fitness. This is commonly done using roulette wheel selection or tournament selection. The main disadvantage of each is that they can result in premature convergence: roulette wheel selection because it continually favours the stronger candidate, and tournament selection because it has a strong tendency to select those with better fitness.

Crossover

Following selection, crossover (or recombination) is performed upon the two chosen parents. The probability of applying crossover (i.e., that two selected parents will breed) commonly varies from between 0.6 and 1.0. To test, a random number between 0 and 1 is generated; if it falls under the crossover threshold, mating occurs, else the parents may be brought into the next generation unchanged. In either case, two 'new' offspring are added to the next generation's population.

The way that crossover occurs can vary by the type of representation employed.

In any case, the chromosomes of the parents are mixed in some way during crossover, typically by swapping portions of the chromosome. This process is repeated with different parents until there are an appropriate number of candidate solutions in the next generation's population.

In single point crossover, one crossover point is selected. The binary string, for example, from the beginning of the chromosome to the crossover point is copied from the first parent, and the rest is copied from the other parent. For example,

10101 011 + 11011 **111** => **10101-111**, 11011-011

In two point crossover, two crossover points are selected. The binary string from the beginning of the chromosome to the first crossover point is copied from the first parent, the part from the first to the second crossover point is copied from the other parent and the rest is copied from the first parent again. For example,

10 0010 **01** + 11 **0001** 11 => **10-0001-01**, 11-0010-11

In multi-point crossover, as many crossover points as desired are selected, within reason given the size of the chromosome. For instance, if the chromosome length was 5 it would be impossible to have 6 crossover points. For example, a triple point crossover scheme might look as follows:

11 0000 **100** 10 + 01 **1010** 110 **00** => **11-1010-100-00**, 01-0000-110-10

The well-known cycle crossover [75] can be used with integer strings where each member can only appear once. This ensures that each index position in the resulting child is occupied by a value occupying the same index position in one of his parents. As an example, suppose we start with Parents A and B and attempt to form Offspring C:

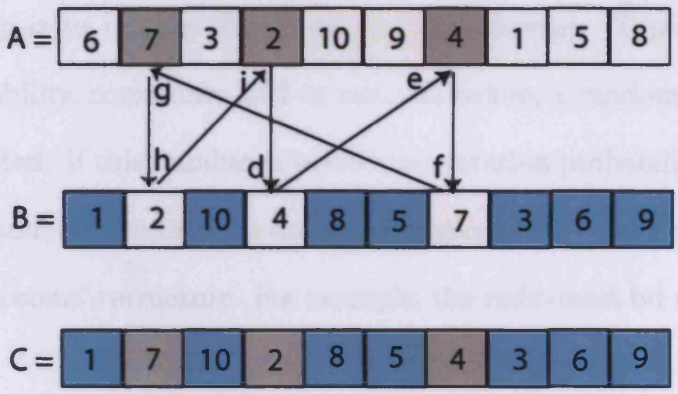

Figure 2.1: Cycle crossover example.

First, an index position from Parent A is randomly selected. For example, position 4 (at i) might have been selected, which contains the value 2 in Parent A. We fill in position 4 with the value 2 in Offspring C, and proceed until we find value 2 in Parent B as follows:

Looking at Parent B, we see that the corresponding value at position 4 is the value 4 (at d), which in Parent A occupies the 7th position (at e). The 7th position contains the value 4, which is then added to our Offspring C. We continue to search in the same fashion until we find our starting value 2:

Looking at Parent B again, we see that the corresponding value at position 7 is the value 7 (at f), which in Parent A occupies the 2nd position (at g). The 2nd

position contains the value 7, which is then added to our Offspring C.

As the corresponding 2nd position in Parent B is the value 2 (at h), we stop, and then fill in the remaining index positions with the values from Parent B.

Mutation

The next step is to mutate the newly created offspring. Typically this is done with a low probability, commonly 0.01 or less. As before, a random number between 0 and 1 is generated. If this number is below the mutation probability, the offspring's chromosome is randomly mutated in some way, typically by simply randomly altering bits in the chromosome's structure. For example, the right-most bit might be 'flipped'

binary: 01001-**0** => 01001-**1**

or, in the case of an integer permutation, two positions can be 'swapped'.

integer: 1-**2**-34-**5** => 1-**5**-34-**2**

Replacement

These processes result in the next generation's genetic pool of chromosomes, which now contains both parents and their offspring. These are likely to be different from the previous generation's genetic pool. However, most GAs only allow the same number of parents to survive to subsequent generations. As we now have double the genetic material, decisions have to be made regarding which individuals survive. This can be done in a variety of ways. For example, offspring may replace their parents

or the least fit individuals in the population, and this could be done either unconditionally, probabilistically, or only if the offspring are better than the individuals they are meant to replace (i.e., elitism). Alternatively, all solutions could be pooled together and then compete for inclusion in the parent population. For example, the most diverse solutions might be saved, or the ones with the highest fitness, or some combination. Generally speaking, if an appropriate replacement strategy is being used the average degree of fitness will increase each generation until a maximum or optimal level is obtained, at which point the algorithm will terminate. However, any other termination criteria could be set. For example, the GA could terminate once a solution deemed 'good enough' was found.

Multi-objective algorithms

Multiple objective problems (MOP) are problems that involve attempting to optimise several competing objectives simultaneously. This is distinguished from single-objective problems (SOP) in which only one objective is optimised, as in the simple GA mentioned previously (Algorithm 1). When solving a SOP, the search focuses on providing one best solution, whereas when solving a MOP the search attempts to capture a range of near-optimal solutions, which recognizes the fact that a single solution may not exist that optimizes all competing objectives simultaneously. This range of near-optimal solutions is called a Pareto set, and often represents trade-offs among competing objectives.

Multi-objective GAs are particularly well suited to manipulating many parameters simultaneously, as different population members can be searching different areas of

objective space simultaneously. This *parallelism*, of searching in several direction at the same time, is desirable as the CPP is difficult to formulate in terms of a single value to be minimized or maximized. Rather, multiple conflicting objectives are involved, where trade-offs must occur. By using parallelism, GAs produce many, equally good solutions to the same problem on multiple objectives. An operator can then select one of these candidates to use based on which objective(s) are deemed most important. This gives the operator a choice which is not possible using a simple GA. There are now so many different multi-objective algorithms (MOA) that a complete review would be a paper in-and-of-itself; the interested reader is referred to [94, 17] for reviews.

2.3 Metrics

Regardless of which optimization technique is adopted to solve the CPP, it is necessary to resolve the conflict between competing multiple objectives, although this is often done implicitly rather than explicitly. However, comparing the performance of multiple objective algorithms is problematic because a set of solutions rather than a single solution is obtained. Although several alternatives have been proposed (e.g., [19, 46, 20, 23, 30, 48, 82, 79, 72, 80, 92, 90, 93]) no approach proves more prevalent. We, therefore, adopt a number of metrics to use depending on the specific test employed. These metrics will be important in establishing a good MOA for purposes of this investigation.

2.3.1 Domination and Pareto optimality

In order to compare solutions based on multiple objective measures, it is necessary to cover the concepts of *non-domination*, *Pareto optimality*, and the *Pareto front*.

Definition 1 *(Non-domination & Pareto optimality) Let* $o_1, o_2, \ldots, o_n$ *be objective functions which are to be maximized. Let* S *be the set of obtained solutions.* $s \in S$ *is dominated by* $t \in S$ *(denoted* $t \succ s$*)if* $\exists j \in \{1, \ldots, n\}$ *such that* $o_j(t) > o_j(s)$ *and* $\forall i$, $1 \leq i \leq n$, $o_i(t) \geq o_i(s)$. *A non-dominated solution is said to be Pareto optimal.*

Pareto optimal cell plans are non-dominated in the sense that it is not possible to improve the value of any objective without simultaneously degrading the quality of one or more of the other objectives. The set of all possible Pareto optimal solutions in the entire search space is called the *Pareto front*. In Figure 2.2, a hypothetical Pareto front is indicated for the objectives of cost and coverage. The most desirable cell plan in the Pareto front depends on which objective is most important. However, in the absence of such a relative ranking of objectives, solutions from the Pareto front must be regarded as equivalent. During subsequent tests, our approach is to generate a set of alternative solutions (i.e., cell plans) which approximate the Pareto front.

Cell plans which lie on the Pareto front represent the best trade-off between coverage and cost. Due to search space size, it is only feasible to approximate subsections of the true Pareto front for a non-trivial problem scenario. To clearly establish the relationship between the objectives, it is desirable to achieve a close approximation to the Pareto front. To help achieve this, sets of approximate solutions can be united to form a larger set; non-dominated solutions extracted from this set can then result in a

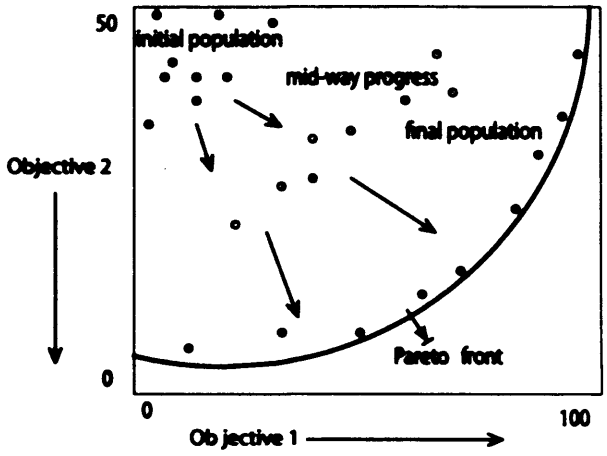

Figure 2.2: Progress towards Pareto front of objective 1 (to be minimized) and objective 2 (to be maximized).

closer approximation of the Pareto front. We term this set S_{PFA}, for Pareto front approximation. An example of this is provided in Figure 2.3, in which 6 approximation sets are combined to form S_{PFA}.

Despite the potential strength of using Pareto optimality within the context of cell planning in this way, it has not been addressed adequately in the literature. This may be partially due to the fact that there are a number of alternative strategies available, which can also handle multiple objective problems. These strategies are:

1. Combine all objectives into a single scalar value, typically as a weighted sum, and optimise the scalar value.

2. Solve for the objectives in a hierarchical fashion, optimizing for a first objective then, if there is more than one solution, optimise these solution(s) for a second objective, and repeat.

3. Obtain a set of alternative, non-dominated solutions, each of which must be

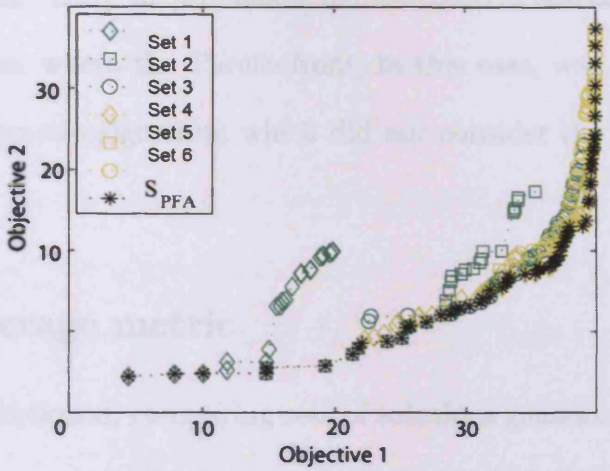

Figure 2.3: Extracting non-dominated solutions to form S_{PFA}

considered equivalent in the absence of further information regarding the relative importance of each of these objectives.

Each approach involves exploring the search space of all possible cell plans to find one or more suitable solutions. Approach one is by far the most popular approach in the literature (e.g., [6, 73, 13, 87, 77, 35, 81, 4, 5, 29, 24, 86, 55, 56]). The biggest problem with this approach is that setting the relative weights of different components in the cost function may lead to inappropriate favoring or penalizing of different objectives. Approach two may be combined with approach one, as in [40, 52, 89, 69], which may involve changing the objective function at different points in the search in a phased or staged manner. This approach effectively prioritizes different single optimization objectives *a priori* and therefore has similar problems

to the first approach. Only in [57] has a multi-objective search been implemented using approach three, where the Pareto front, in this case, was approximated using a problem specific genetic algorithm which did not consider the financial cost of the cell plan.

2.3.2 Set coverage metric

As previously mentioned, comparing sets of solutions generated by different MOA is problematic because a set of solutions is obtained. Although several alternatives have been proposed (e.g., [23, 82, 93, 92]) no single approach proves more prevalent. We adopt the approach first given in [93], and which tested favorably against other metrics in a recent comprehensive study [92], to calculate a *set coverage* metric. This involves the concept of *weak* domination.

Definition 2 *(Weak Domination) Let* $o_1, o_2, \ldots, o_n$ *be objective functions which are to be maximized. Let* S *be the set of obtained solutions.* $s \in S$ *is weakly dominated by* $t \in S$ *(denoted* $t \succeq s$*) if* $\forall i$, $1 \leq i \leq n$, $o_i(t) \geq o_i(s)$, *or* $t \succ s$.

Definition 3 *(Set Coverage) Let* S_A *and* S_B *be two sets of solutions. The set coverage metric of set* S_A *with respect to* S_B *is the percentage of solutions in* S_B *which are weakly dominated by at least one solution from* S_A.

When comparing solutions set S_A with S_B, the higher the set coverage metric the greater the superiority of the solutions in S_A over S_B. These metrics are important for assessing relative cell plan quality from different potential decoding approaches. In our application, a high set coverage metric of S_A over S_B indicates the superior

efficiency of the cell plans in S_A over S_B, in the sense that overall, the cell plans in S_A are frequently achieving higher coverage for the same cost or the same coverage level at lower cost than those in S_B.

In other words, Solution A *weakly dominates* solution B if A and B have the same performance across all objectives or A dominates B. For two sets of solutions S_A and S_B, the *set coverage metric* of set S_A with respect to S_B is the percentage of solutions in S_B which are weakly dominated by at least one solution from S_A. The higher the set coverage metric obtained, the greater the superiority of S_A over S_B.

2.3.3 Measure of solution distribution

To measure the distribution, or spread, of solutions along the Pareto front, a metric proposed in [72] has been implemented. The spacing measure is based on the range of values for d_i, which is the distance (in terms of solution space) between the i^{th} element of the solution set and its nearest neighbour. The average of d_i values for a solution set of size n is denoted $\overline{d}$. Then the measure of spread is defined as:

$$S = \sqrt{\frac{1}{(n-1)} \sum_{i=1}^{n} (\overline{d} - d_i)^2} \tag{2.1}$$

for the n members in the final population. Note that $S = 0$ indicates all members of the Pareto front are spaced equidistantly in the solution space.

2.3.4 Convergence

The ability of an algorithm to rapidly converge to the final solution set is desirable, as long as it converges to the global optimum. We use the term convergence to

refer, not to an unchanging population in terms of chromosomal structure, but to the population's overall fitness not improving for an extended period of time. This indicates that the GA has stopped finding better solutions, but does not guarantee that no more changes would take place if the GA kept running. It is an advisory measure for when it seems appropriate to stop. That is, if there is no change to the overall fitness for a number of generations (e.g., 100), the algorithm should terminate.

2.3.5 Speed of execution

The time it takes for a given approach to find a solution, or set of solutions, is important in measuring the trade-off often found between solution quality and the speed of execution. For example, a fast algorithm may provide a solution with a quality of 85% in 30 minutes, but a complex algorithm a quality of 97% in 300 minutes. Depending on whether quality or speed is more desirable, either approach could be preferable. Of course, if the speed of execution is not measured, trade-off comparisons and choices are not possible.

Chapter 3

2D Model and Cell Planning Strategy

The purpose of this chapter is to introduce a starting point for cell planning, in the form of an abstract distance-based 2D model which addresses the fundamental issue of area coverage taking into account the primary issue in site configuration (i.e., transmission power). A general strategy for multiple objective optimisation of primary metrics is developed and explained in Section 3.2. This strategy is deployed using four state-of-the-art multiple objective genetic algorithms. The best performing algorithm is then determined and carried forward for further tuning (Chapter 4) and in the deployment of more granular modeling (Chapter 6).

3.1 2D Model for Simulating Wireless Coverage

The general model for simulating wireless coverage we adopt (Figure 3.1) focuses on simulated down-link received signal strength. This is the standard approach used in frequency division communication systems such as the current European GSM mobile telephone system [85]. This model was first used in [68] and [67]. In this model, the working area or simulation area W is discretized into test points at Cartesian coordinates (i.e., x,y) at a given resolution. For example, test points may be available uniformly every 100m. The two sets of test points distributed in W are reception test points (RTP) where signal reception could be tested and service test points (STP) where signal quality must exceed a signal quality threshold (denoted S_q) to facilitate wireless communication. In this model STP $\subseteq$ RTP. Also within W lies a set of randomly placed candidate base station locations transmitting at a given power (denoted p). Each commissioned site (i.e., where $p > 0$) has an associated cost. The *cell* (denoted C) formed at a base station site corresponds to the set of STP *covered*, where the signal received is higher than S_q.

More formally, the following sets form the input to our formulation of the APP:

- A set of candidate base station site locations CBS $= \{CBS_1, \ldots, CBS_{n_{CBS}}\}$, where each CBS_i has an associated commissioning cost $\$(CBS_i)$.
- A list of possible transmission powers $p_0, p_1, p_2, \ldots, p_k$ in ascending order of magnitude. Zero power is denoted by p_0.
- A set of reception test points RTP $= \{RTP_1, \ldots, RTP_{n_{RTP}}\}$, at which signal reception quality is measured.

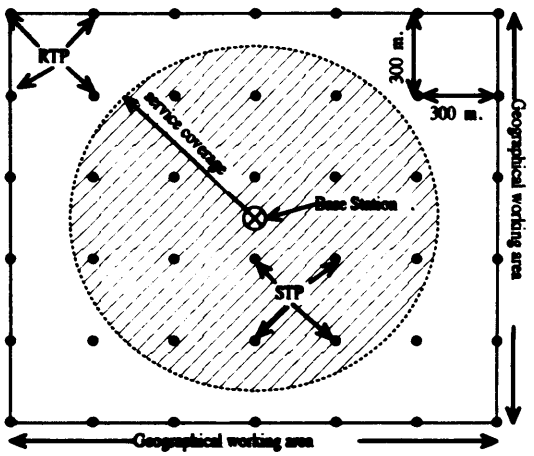

Figure 3.1: Working area for simulating wireless coverage

- A set of service test points STP = $\{STP_1, \ldots, STP_{n_{STP}}\}$, where a signal must be received above a minimum specified service threshold, S_q, to ensure a required quality of service.

For purposes of candidate sites, we assume that each base station is operating a single omni-directional antenna with an isotropic radiation pattern. As traffic, terrain, and shadowing are not considered the radiation patterns are idealised and circular, similar to those found in disk graph models.

3.1.1 Propagation, estimating signal quality

A given service test point STP is said to be *covered* by antenna A if the received signal strength from A, denoted P^A_{STP}, is greater than S_q, which is set to -90 dBm. As the signal strength attenuates with distance as a function of the *path loss* between source and receiver, additional losses and gains may also affect the received signal

strength. We calculate the received signal strength as:

$$P^A_{STP} = P^A - PL - L + G$$

where P^A is the power at which A is transmitting, PL is the path loss experienced between A and STP, L is the aggregation of other losses experienced, and G is the aggregation of gains experienced. For experimental purposes, we assume $G = L$. For each combination of A and STP, PL may be recorded in the field or estimated using a free space path loss or empirical model. In the absence of data from the field, we employ a path loss model based on a set of comprehensive measurements done for Japanese environments in 1968 by Okumura [88], with the derived curves later transformed into parametric formulas by Hata [32]. We adopt Hata's urban empirical model proposed for small-medium cities:

$$PL = 69.55{+}26.16\ \log(f){-}13.82\ \log(h_b){-}a(h_m){+}(44.9{-}6.55\ \log(h_b)){*}\log(R) \quad (3.1)$$

where:

$$a(h_m) = (1.1 *\ \log(f) - 0.7) * h_m - (1.56 *\ \log(f) - 0.8) \quad (3.2)$$

given particular values for the variables frequency (f), base station height (h_b), mobile receiver height (h_m), and the distance (R) in kilometers from each base station to each RTP. The default settings are f=800 MHz, h_b = 31 meters, and h_m = 1.5 meters unless otherwise noted. While Hata's model was used, other propagation models could have equally been adopted.

3.1.2 Assessing cell plan quality

The objectives we are concerned with relate to financial cost and area coverage of service test points. These measures are competing and are the multiple objectives that will be considered during optimization. More formally, the *cost* of a cell plan CBS', which consists of all sites with a non-zero power setting, denoted $cost_{CBS'}$, is defined as:

$$cost_{CBS'} = \sum_{CBS_i \in CBS'} \$(CBS_i),$$

and the *service coverage* of a cell plan, denoted $cover_{CBS'}$, is the proportion of STP covered in CBS', where a given covered STP is denoted STP_{cov}, to the total number of STP in the working area, W:

$$cover_{CBS'} = \frac{\sum_{STP_{cov} \in CBS'} 1}{n_{STP}}.$$

3.2 Strategy for 2D Multi-objective Cell Planning

The intent of the strategy for multi-objective cell planning for assessing the tensions and trade-offs in cell planning is to find a range of high quality cell plans in terms of service coverage and financial cost. By producing a range of near optimal cell plans the trade-offs between objectives can be assessed, and the need for weighting objectives common in many approaches is avoided. Figure 3.2 is a visual aid to the multi-objective cell planning strategy, termed a generational two phase strategy ($G2PS$), where the first phase employs a greedy sequential algorithm to create a cell plan from a given ordering of sites and the second phase a multi-objective algorithm

(MOA) to find optimal orderings of candidate sites.

This approach is introduced after considering the way in which the problem may be solved manually, by sequentially commissioning sites based on a simple heuristic, in our case, represented by the decoder. The candidate base station location CBS_i is referred to as the i^{th} base station. We use the permutation π of the potential base station locations to represent cell plans. Under the permutation π, the i^{th} base station location is denoted $\pi(i)$. We introduce a decoder which translates a permutation π into a cell plan, which is in effect, the set of rules that may be applied in the manual planning process. The decoder is a greedy sequential algorithm for creating a cell plan, and the ordering of sites in π determines which sites are commissioned, and at what power level. A **2D cell plan** is then defined as the allocation of a power setting to each candidate base station site in π, where sites allocated a zero power setting are not commissioned.

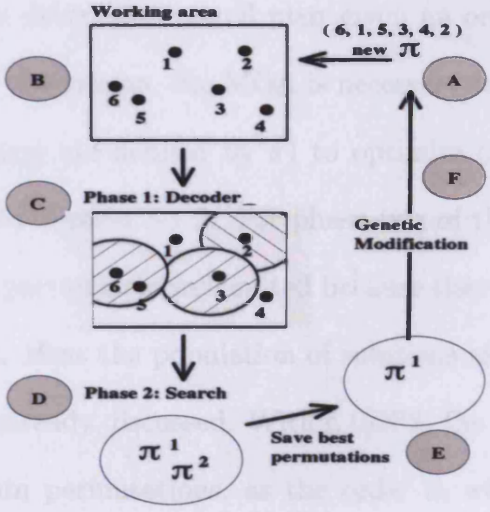

Figure 3.2: Multi-objective cell planning strategy

In phase one of G2PS, a greedy sequential decoder (GSD) is used to control cell density during the formation of a given cell plan. This is an important issue for mobile networks as it can greatly impact service coverage levels, system capacity, interference, and channel assignment. The GSD controls areas of cell overlap by setting criteria for whether a given power allocation at a candidate site is permissible. A power allocation is permissible only if a multi-coverage, or overlap, constraint is met. By controlling overlap, the GSD can eliminate poor cell plans from consideration. The general approach of decoding optimized permutations is well established in many applications of genetic algorithms, such as for the travelling salesman problem (e.g., [12]) and bin-packing problems (e.g., [66]). As we wish to compare the performance of MOA here, a common decoder is used throughout. However, as there is much freedom in the design of decoders, a range of decoders is investigated later (Chapter 4).

As the decoder only determines a cell plan given an ordering of sites, it cannot improve cell plans. For this reason, the MOA is necessary to operate on a population of potential site orderings (as defined by π) to optimise cell plans with respect to the objectives defined in Section 3.1.2. For phase two of the cell planning strategy, genetic approaches are particularly well suited because they maintain and optimise a population of solutions. Here the population of solutions are represented as integer-permutations, or π, as already discussed. Within G2PS, the MOA is used to optimise the order of sites within permutations, as the order in which sites are considered for power allocation determines the cell plan formed. In other words, the MOA is a means for performing a meta-heuristic search (Section 3.4) to find optimal candidate

site orderings. This occurs over a number of generations. In short, the following two components define our framework:

1. *Decoder:*

   The sequential configuration of a permutation of candidate base station sites, where the ordering determines the configuration made.

2. *Optimization of permutations:*

   The application of a multi-objective optimization approach to produce a collection of candidate site orderings, which when decoded approximate the optimal trade-offs between competing objectives.

The way in which these two components operate together can be seen in Figure 3.2: The decoder accepts the site permutation from A, which represents the locations visible at B. It decodes the permutation, which results in a cell plan visible at C. Information regarding each cell plan is then passed to phase 2 at D. At D the multi-objective algorithm receives information from the decoder regarding the objective values obtained by each cell plan. It assigns fitness values to each solution, and saves the n most fit permutations, depicted at E. At F the solutions from E compete in binary tournaments to become parents which undergo cross-over and mutation to form the new permutations which are represented at A. The process continues until the final generation, when the non-dominated solutions at E are saved.

3.3 Decoder

We introduce a decoder which translates a permutation π into a cell plan. This approach mimics the way in which the problem might be solved manually. The decoder is effectively a greedy, sequential algorithm for creating a cell plan, which is dependent on the order of inspection for commissioning potential sites occurring in π. The decoder acts by limiting the permissible amount of overlap, or multi-coverage, that occurs between cells.

More formally, the subset of service test points covered by a particular antenna A is the *cell* served by A, denoted c_A. Note that cells served by different antennae are not necessarily disjoint since an STP can potentially be covered by more than one antenna. Such an STP is referred to as an *overlap STP*. An overlap STP which is contained in more than two distinct cells can also be referred to as a *soft* handover STP. For a cell c_A, the subset of overlap STP is denoted o_A. For c_A, the overlap percentage is defined as

$$\frac{|o_A|}{|c_A|} \times 100.$$

Controlling the size and distribution of overlap, or handover, regions is crucial for both operational and financial reasons. Handover regions are a prerequisite for seamless call transfer between cells for mobile users. However, if very large handover regions are permitted, there is a greater potential for *interference* due to strong signals being received from multiple sources. In frequency division multiple access systems, large handover regions increase the need for large channel separation between adjacent cells in the frequency assignment problem. Large handover regions may also adversely affect the cost of the network by increasing the total number of base stations required

to cover a given area.

The decoder adds cells iteratively to create a cell plan CBS' as follows:

- Initially $CBS' = \emptyset$.
- Potential sites $\pi(1), \pi(2), \ldots, \pi(n)$ are inspected (in the order induced by π) for possible selection.
- At iteration j $(1 \leq j \leq n)$, $\pi(j)$ is considered for addition to the set CBS'.

  For a given power setting, handover between $c_{\pi(j)}$ and CBS' is *feasible* if the hand-over percentage for $c_{\pi(j)}$ is less than the maximum permitted.

  The largest power setting, denoted p_{max}, is identified from the list $p_0, p_1, p_2, \ldots, p_k$ such that handover is feasible between $c_{\pi(j)}$ and CBS'.

  If $p_{max} \neq p_0$, then $\pi(j)$ is added to CBS', and the transmission power of $\pi(j)$ is recorded as p_{max}. Otherwise $\pi(j)$ is not added to CBS'.

A number of observations can be made regarding this approach. Firstly, the approach is greedy in the sense that once a base station location is added to CBS' at power p_{max}, the base station cannot be removed from the cell plan CBS' nor can its transmission power be adjusted. Secondly, for a particular list of potential site locations, characteristics (e.g., cost and coverage) of the resultant cell plan CBS' is entirely dependent on the order (i.e., permutation π) in which the base stations are considered for selection. It is our aim to find the best permutations, which lead to Pareto optimal cell plans, using genetic algorithms.

Despite its simplicity, the decoder is potentially a powerful tool as it has the following strengths:

- Guarantees feasible cell plans (i.e., without an inordinate amount of overlap between cells) given a list of candidate site locations.
- Provides freedom in determining the cell density in networks by allowing the planner to alter the maximum level of permissible overlap.
- Creates a range of networks (when guided by a multi-objective algorithm) in terms of service coverage and cost to cater for low-cost and high quality service providers.
- Grants the planner freedom to 'plug-in' the best available multi-objective algorithm without changing cell representation or decoder specification.

3.4 MOGA: Selection

There are now many potential multiple objective optimization approaches [17]. However, fundamentally, the overall performance of any optimization approach adopted will be governed by the decoder specification, and its ability to translate an abstract ordering of candidate base station sites into a highly efficient cell plan (e.g., high coverage relative to total cost of operational base stations). While a decoder creates a cell plan given a permutation π, the decoder is not capable of using information from the evaluation of the objective functions to *improve* cell plans. For this purpose, a multi-objective genetic algorithm is used to find optimal orderings of candidate sites subject to optimizing two or more conflicting objectives.

In this section, we are concerned with determining the best multi-objective algorithm to find optimal orderings. There are now so many different MOA that testing

each would not be sensible. Thus, the interested reader is referred to [94, 17] for reviews, and we focus here on comparing the performance of three popular GA: PESA [15], NSGA-II [18], and SPEA-2 [91], which have been shown to be superior to older multi-objective algorithms, e.g., MOGA [22], VEGA [71], NSGA [74], SPEA [94], NPGA [34], PAES [47], and one less known one: SEAMO [78], as it promised to be a simpler and faster approach for comparison purposes.

3.4.1 SPEA-II

The Strength Pareto Evolutionary Algorithm version II (SPEA-2) is an enhancement of that originally proposed in [94], and is described in detail in [91]. SPEA-2 has been used in numerous studies (e.g., [11, 51]) where good performance, in comparison to other MOAs, has been reported. In SPEA-2, the most fit individuals from the union of archive and child populations are determined by computing a fitness value for each solution which is the sum of two parts. The first part is a *raw fitness* value based on how many solutions it dominates, and the second is a *density estimate* based on the its proximity to other solutions in objective space. At each generation, the most fit n solutions are saved to the archive, and genetic operators are applied to form a new child population. This process is repeated until termination.

More specifically, SPEA-2 maintains two fixed sized populations each of size n: a main population P_t and an archive population A_t, which stores the most fit individuals in each generation. The algorithm proceeds with the multi-set union, $P_t \cup A_t$, denoted U_t, at iteration t. For each member i from U_t, a *strength value*, denoted S_i, is defined.

This represents the number of members from U_t which i dominates.

$$S_i = |\{j : j \in U_t \wedge i \succ j\}| \tag{3.3}$$

The raw fitness R_i is then calculated as the sum of the strengths of the solutions j which dominate i in U_t.

$$R_i = \sum_{j \in U_t, j \succ i} S_j$$

Following this, a density estimate is computed to distinguish individuals with the same raw fitness and measure the dispersion of solutions. The density estimation D_i is a decreasing function of the distance to the nearest neighbour in objective space. For a complete explanation, please see [91]. Adding D_i to the raw fitness value R_i yields the fitness value F_i, which is computed for each member i from U_t.

Using these functions, an individual with a fitness less than 1 is non-dominated. In all cases, the smaller the raw fitness value, the stronger the individual in relation to others, and the smaller the density estimate, the less crowded the individual is in objective space.

The algorithm proceeds by copying all the non-dominated individuals in U_t to the archive population at the next generation A_{t+1}. If the size of A_{t+1} exceeds n, then a truncation operator is applied, which removes members with the largest density estimate (i.e., most crowded) iteratively with ties being broken by considering the second smallest distance and so on. If the size of A_{t+1} is less than n, it is filled with the best of the dominated solutions (based on fitness values F_i) from U_t. Once the archive population A_{t+1} is at size n for generation $t + 1$, binary tournaments (using fitness F_i for selection) with replacement are performed on members of A_{t+1} to fill the mating pool with n members. Recombination and mutation operators are

then applied to the mating pool in order to form set P_{t+1}. At this point we have obtained all population members at generation $t + 1$, and the process is repeated until termination criteria are met. We summarize the pseudo-code in Figure 3.3.

Generate initial population P_0 and empty archive A_0 of size n

Set $t = 0$, $T = \sharp$ of generations

while $t < T$ **do**

$U_t = P_t \cup A_t$

Calculate fitness F_i, $\forall i \in U_t$

Copy all with fitness < 1 to A_{t+1}

if $|A_{t+1}| < n$ **then**

Fill A_{t+1} with dominated members of U_t in increasing order of fitness

else if $|A_{t+1}| > n$ **then**

Reduce A_{t+1} by means of the truncation operator

end if

Create mating pool of size n by performing binary tournaments with replacement on A_{t+1}

Create P_{t+1} of size n by applying recombination and mutation to the mating pool

$t = t + 1$;

end while

Figure 3.3: The SPEA-2 procedure

3.4.2 NSGA-II

NSGA-II is a fast elitist non-dominated sorting genetic algorithm ([18] for a full description), which has been well studied (e.g., [46, 20]). NSGA-II is similar to SPEA-2, but uses slightly different mechanisms. For example, in NSGA-II the most fit individuals from the union of archive and child populations are determined by a *ranking* mechanism (or *crowded comparison operator*) composed of two parts. The first part 'peels' away layers of non-dominated fronts, and ranks solutions in earlier fronts as better. The second part computes a dispersion measure, the *crowding distance*, to determine how close a solution's nearest neighbors are, with larger distances being better. At each generation, the best n solutions with regard to these two measures are saved to the archive, and genetic operators applied to form a new child population. This process is repeated until termination.

More specifically, NSGA-II maintains two populations, the archive population A_t, which stores the most fit individuals in each generation, and a child population C_t. The child population C_t is formed from A_t by performing selection, crossover and mutation. As in SPEA-2, a simple binary tournament with replacement is used to find parents. In the binary tournament, two population members, i and j, are chosen at random (uniformly) from A_t. Member i wins the tournament if i dominates j, and i becomes a parent. A second parent is selected similarly and both are put forward for crossover and mutation to create an offspring for C_t. This is repeated until C_t is size n. Note that each population A_t and C_t remains fixed at size n. The algorithm proceeds with the multi-set union $A_t \cup C_t$, denoted U_t. Then, n members are selected from U_t to form the new parent population for the next generation A_{t+1}.

Before this occurs, U_t is partitioned into a number of sets called *fronts*, which are constructed iteratively. Front 1, F_1, consists of the non-dominated members from U_t. Front 2, F_2, consists of the non-dominated members from the set $U_t - F_1$. In general, front i, F_i, consists of the non-dominated members from the set $U_t - (F_1 \cup F_2 \cup \cdots \cup F_{i-1})$. Fronts are constructed in this manner until U_t is empty. The authors of NSGA-II describe an efficient algorithm (fast non-dominated sort) to partition U_t into fronts ([18]).

Fronts are then used to form A_{t+1}. If $|F_1| = n$, then the new parent population, A_{t+1}, is set to contain the members of F_1. Otherwise, $F_1, F_2, \ldots$ are added to A_{t+1} in order (front one followed by front two and so on). The process stops with front F_j being added to A_{t+1} when $F_1 \cup F_2 \cup \cdots \cup F_j \leq n$ and $F_1 \cup F_2 \cup \cdots \cup F_j \cup F_{j+1} > n$. If $|F_1 \cup F_2 \cup \cdots \cup F_j| < n$ then solutions have to be chosen from F_{j+1} for inclusion in A_{t+1} to ensure that population A_{t+1} is size n. For each solution i in F_{j+1}, a crowding distance cd_{ik} is calculated based on each of the multiple objectives, o_k. The solutions in F_{j+1} are ordered with respect to ascending magnitude of their value in o_k. The first and last members from the front are assigned infinity as their crowding distance with respect to o_k. For all other members of the front, cd_{ik} is calculated as the sum of the absolute difference between the objective function values of i and its adjacent solutions in the front. The overall crowding distance for a solution i, $CrwD_i$, is defined as:

$$CrwD_i = \sum_{\forall o_k} cd_{ik}$$

The greater the crowding distance the greater the isolation of that particular solution in objective space. Solutions with the highest crowding values are added to A_{t+1} until

$|P_{t+1}| = n$. At this point, the procedures described are re-iterated until a termination condition has been satisfied. We summarize the pseudo-code in Figure 3.4.

3.4.3 PESA

The Pareto Envelope-based Selection Algorithm, PESA, is described in [15]. It uses different mechanisms than SPEA-2 and NSGA-II. The main differences are that its archive population is not of fixed size and only allows non-dominated solutions to be members–which is a more limited set than the previous two GA allowed. If the archive ever exceeds n solutions, a *squeeze factor* is calculated for all members of the archive. The squeeze factor is the total number of members in the same sub-region of a *hyper-grid* (which partitions the search space into sub-regions [49]). The higher the squeeze factor, the more local neighbors a solution has. Random members from the grid region with the highest squeeze factor are then removed until the size of the archive is reduced to n. Genetic operators are then applied to archive members to form a new child population. This process is repeated until termination.

More specifically, at generation t, PESA operates by generating an internal population (IP_t) and an empty external population (EP) archive, each of size n. EP is repeatedly updated as the search progresses and used to create new population members for the internal population at subsequent generations[1]. The objective function values are calculated for each member in IP_t, and non-dominated members are identified. The non-dominated members are then added to the current, existing EP_t, based on the following rules. Firstly, if member i of IP_t is not dominated by any

[1]Unlike the other algorithms presented, the size of EP changes dynamically at each generation, as only non-dominated solutions are stored.

Generate initial population A_0 of size n

Rank and sort A_0 based on non-domination level

Apply selection, recombination and mutation to create a child population C_0 of size n

Set $t = 0$, $T = \sharp$ of generations

while $t < T$ **do**

$U_t = A_t \cup C_t$

Partition U_t into fronts $F_1, F_2, \ldots$

Set $A_{t+1} = \emptyset$, $i = 1$.

while $|A_{t+1}| \leq n$ **do**

Calculate crowding distance in F_i

if $|F_i| + |A_{t+1}| \leq n$ **then**

$A_{t+1} = A_{t+1} \cup F_i$

else if $|F_i| + |A_{t+1}| > n$ **then**

Decreasing sort of F_i members by crowding distance

$A_{t+1} = A_{t+1}+$ the first $(n - |A_{t+1}|))$ elements of F_i.

end if

$i = i + 1$

end while

Calculate crowded comparison operator $\forall\ i \in A_{t+1}$

Create C_{t+1} of size n by applying recombination and mutation to parents selected via binary tournaments on A_{t+1}.

$t = t + 1$;

end while

Figure 3.4: The NSGA-II procedure

member of EP, it is added to EP. Secondly, once i has entered the archive EP, solutions in EP which it dominates are deleted. After all non-dominated members of IP_t have been considered for inclusion in EP_t, it is possible that EP_t is over-full (i.e., $|EP| > n$), in which case a member(s) of EP needs to be deleted. At this point, a hyper-grid is imposed to partition the search space into sub-regions [49]. This takes place to assess the density of solutions throughout the search space.

Following this a 'squeeze factor' is calculated for all members of EP. For a particular individual $i \in EP$, this is the total number of members of EP (inclusive of i) in the same sub-region of the hyper-grid. Therefore, the higher the squeeze factor, the more local neighbors (in terms of the search space) a solution has. A random member from the grid region with the highest squeeze factor (as long as it does not harbor a most extreme objective function value) is then removed, with ties being broken randomly, until the size of EP is reduced to n. Once $|EP| = n$, IP_{t+1} is generated from EP through one of two operators as defined in Figure 3.5.

3.4.4 SEAMO

Finally, the Simple Evolutionary Algorithm for Multi-objective Optimization, known as SEAMO, has performed particularly well on the benchmark test knapsack optimization problem [78]. The main difference between SEAMO and the other algorithms is that it is steady-state and has only one population (of constant size n) to maintain. The main advantage of SEAMO is the simple approach it uses to dispose of all selection mechanisms based on fitness or rank. Instead, the search progresses based on three simple rules:

1. Parents can only be replaced by their own offspring (diversity preservation).

2. Duplicates in the population are deleted (diversity preservation).

3. Offspring can only replace parents if superior (elitism).

Genetic operators are applied to each parent in turn to form a new child, which is considered for substitution into the parent population based on the three rules. This process is repeated until termination.

More specifically, to begin a parent population of size n is created [2]. Then, for each objective function, the *best-so-far* value from the parent population is recorded. New solutions called *children* are formed one-by-one, evaluated, and tested against their parents immediately before the next child is generated, evaluated, and tested. This involves each solution in the parent population serving as the first parent once, and a second parent being chosen at random (uniformly). Crossover and mutation is then applied to form an offspring. If the offspring's objective value improves on any *best-so-far*, then it replaces one of the parents (starting with first, then second), as long as the parent itself does not harbor any other *best-so-far*. If the offspring dominates one of the parents, then the offspring replaces a parent (staring with first parent, then the second), unless the offspring is a duplicate, in which case the offspring is deleted. This process is repeatedly applied to update the parent population, until a termination condition is satisfied.

[2]Unlike the other algorithms presented, SEAMO has only one population (of constant size n) to maintain.

3.4.5 Recombination and mutation

Each of the algorithms considered has a specific method for selecting parents. SPEA-2 bases selection on fitness, NSGA-II on the crowded comparison operator, PESA on non-dominated members of its archive set, and SEAMO uniformly. However, common recombination and mutation operators have been applied throughout despite the fact that this may have unintentionally favoured one algorithm over another. However, within the scope of this work it was not possible to test for this potentiality. Cycle crossover [75] has been used as the recombination operator and the mutation operator involves the simple transposition of candidate base station locations in a randomly selected pair of positions. This was governed by a mutation rate (set to 1%) to restrict the frequency of mutation.

3.5 MOGA: Performance

The performance of the algorithms have been compared using a wide range of synthesised test problems, each of which has been randomly generated. Each test problem gives the location and cost of candidate sites, with test points uniformly distributed in the working area every 300m. The cost of sites were randomly determined values between 1 to 100 to make the selection of sites more difficult for the algorithm than if realistic estimates (which would be closer in value) were used. Test problems are classified in two ways: the size of area in which they are positioned and the density of candidate sites, as documented in Table 3.1. For example, the working area of 45×45 with a 0.12 site density carries 244 randomly positioned candidate

base station locations and 22,500 test points.

Region size km^2	Density of sites per km^2		
	0.03	0.06	0.12
15×15	7	14	28
30×30	27	54	108
45×45	61	122	244

Table 3.1: Number of candidate sites in nine problem classes defined by region size and density

Power Setting	dBW	Watts
p_1	30	1000
p_2	27	501
p_3	24	251
p_4	21	125
p_5	18	63

Table 3.2: Power settings used in tests

Combining the size of regions and the density of sites leads to a total of nine test problem classes, as indicated in Table 3.1. For each test problem class, we produce five incidences on which each algorithm is tested. This means that average algorithm performance is estimated and compared using five problem instances from each of nine classes, with four different algorithms, leading to a total of 180 experiments. All problem instances are available at: `http:// www.raisanen.co.uk/ datasets.html`.

To encourage a fair comparison between algorithms, common parameter settings have been adopted for each experiment. Five non-zero power settings (displayed in units of dBW and Watts) have been used, as specified in Table 3.2. Unless otherwise specified, a population size of 100 is adopted using 500 generations. Additionally, the same random starting populations have been used for each problem class.

We consider the performance of each GA in four ways: (1) the average performance (in terms of the objective values of members in the final population) compared to other GAs across all test problems using the set coverage metric, with diagrams to show obtained Pareto fronts and cell plans, (2) the average measure of population spread, (3) the speed of convergence to solutions in the final population, and (4) the average speed of execution.

3.5.1 Average performance across test problems

In terms of the average performance of each GA compared to other GAs across all test problems using the set coverage metric, it was found that NSGA-II achieves the best performance by weakly dominating an average of 91.15% of solutions obtained by other algorithms in terms of service coverage and cost, closely followed by SPEA-2 (89.67%), then PESA (64.01%) and finally SEAMO (59.53%). See Table 3.3 for details. However, the dominance of NSGA-II is hidden by these averages, as it held the highest score in 9 out of 10 problem instances. The second closest was SPEA-2 which held the highest score in only 3 out of 10 problem instances.

In Figure 3.7 we plot the Pareto fronts (i.e., non-dominated solutions from the final population) achieved by each algorithm on the large region problem at each density. Despite the differences in relative algorithm performance, the Pareto fronts obtained are closely clustered in real terms. Generally, the plots show that that as candidate site density increases, solutions with a higher level of coverage are achievable. Also, in the most dense problem instances, lower cost solutions with higher coverage are achievable due to more freedom in site selection and cost.

In Figure 3.8, we display an example of cell plans with the highest coverage for the large region problem at each density, with density increasing left to right.

3.5.2 Measure of solution distribution

To measure the distribution, or spread, of solutions along the Pareto front, a spacing measure proposed in [72] has been implemented (Section 2.3.3).

It was found that PESA performed the best on this measure, with an average spacing value of 19.75, followed by SPEA-2 at 20.91, NSGA-II at 21.46, and SEAMO at 26.84 (Table 3.4 for details). It is little surprise that the algorithm which is specifically designed to encourage spacing in objective space, PESA, performed the best, and that the algorithm with no direct measure to control dispersion, SEAMO, performed the worst.

3.5.3 Convergence

The ability of an algorithm to rapidly converge to the final solution set is desirable. This has been investigated for each algorithm using the largest problem class with the highest site density. Each algorithm has been applied for 1500 generations, and intermediate populations (produced every 250 generations) have been compared against each other, using the set coverage metric. The results indicate that PESA and SEAMO converge quickly. For example, at generation 750 dominating 88.88% and 71.43% of solutions in the generation 1000 respectively. SPEA-2 and NSGA-II converge more slowly dominating 22.22% and 33.33% respectively (Table 3.6 for details). However, considering that SPEA-2 and NSGA-II outperform PESA and SEAMO in

terms of final solutions (Section 3.5.1), this may indicate they are better at improving solutions over time, despite the evidence from this specific testing instance.

3.5.4 Speed of execution

The average execution times varied from an average of 7.4 seconds to complete the smallest problem, to 3006.7 seconds (roughly 50 minutes) to complete the largest, with SEAMO performing the fastest marginally (Table 3.5 for details). All recorded times were obtained using a Pentium IV 1.8 GHz processor, 256 megabytes of RAM, Java JDK 1.3.1-04, and Windows XP Professional. It is suspected that obtaining only marginal differences in execution time was due to a bottleneck incurred by the computationally intensive decoder.

3.6 MOGA: Conclusions

In this chapter we have have introduced a general framework for applying multiple objective genetic algorithms to the antenna placement problem. The key aspect of this framework is a decoder which uses an ordering of candidate site locations to construct a cell plan. Subsequently, the performance of four GAs to find an optimal ordering of potential site locations were compared and evaluated using a range of test problems classified by size and density.

We found that all the algorithms considered find closely comparable solutions, in real terms. However, there are differences. NGSA-II and SPEA-2 have very similar performance throughout, confirming the findings in [91] concerning the performance of these algorithms. Meanwhile, PESA generally obtains slightly lower quality sets

of solutions, but has the best performance in terms of distribution of solutions and speed of convergence. Both the advantage and disadvantage of SEAMO lie in its simplicity. This makes the algorithm conceptually elegant, easy to implement, and fast to run, but this simplicity appears to impede the overall quality and distribution of the solutions obtained, with SEAMO obtaining the lowest performance measures in these areas. On balance, we consider NGSA-II to be the strongest performing algorithm for purposes of cell planning when using the general framework proposed. This is mainly based on the consistent comparative quality of the solutions obtained. NSGA-II is carried forward for subsequent implementation in all experiments within this thesis.

Generate initial internal population IP_0 of size n

$EP = \emptyset$

Set probability P_c for tournament selection

Set $t = 0$, $T = \sharp$ of generations

while $t < T$ **do**

Determine non-dominated solutions in IP_t and try to add them to EP

while $|EP| > n$ **do**

Calculate the squeeze factor for each member of EP

Remove random member from EP with the highest squeeze factor

end while

Set $IP_{t+1} = \emptyset$

while $|IP_{t+1}| < n$ **do**

Generate a random number $x \in (0, 1]$

if $x \leq P_c$ **then**

Perform binary tournaments to select two parents, and apply recombination and mutation to form a child

Put the child in IP_{t+1}

else if $x > P_c$ **then**

Perform a binary tournament to select one parent and mutate to form a child

Put the child in IP_{t+1}

end if

end while

t=t+1;

end while

Figure 3.5: The PESA procedure

Generate initial parent population P_0 of size n

Record the *best-so-far* for each objective function

Set $t = 0$, $T = \sharp$ of generations

while $t < T$ **do**

  for $\forall i \in P_t$ **do**

    i is the first parent

    Select a second parent uniformly randomly

    Apply crossover and mutation to produce offspring $\bar{i}$

    Evaluate $\bar{i}$'s objective vector

    if $\bar{i}$'s objective vector improves on any *best-so-far* **then**

      $\bar{i}$ replaces one of the parents and *best-so-far* is updated

    else if $\bar{i}$ dominates one of the parents **then**

      $\bar{i}$ replaces the parent it dominated (unless $\bar{i}$ is a duplicate, then it is deleted).

    end if

  end for

  $t = t + 1$;

end while

Figure 3.6: The SEAMO procedure

Set coverage metrics											
Algorithm		Problem Instances (km^2– number of candidate sites)									
S_A	S_B	15-7	15-14	15-28	30-27	30-54	30-108	45-61	45-122	45-244	Ave.
SEAMO	SPEA-2	100.00	92.00	86.83	91.34	43.23	19.03	39.17	7.85	18.00	**55.27**
	NSGA-II	100.00	92.00	84.33	92.80	39.81	13.71	37.15	9.20	11.41	**53.38**
	PESA	100.00	96.00	97.50	96.80	60.29	47.74	63.32	38.05	29.90	**69.96**
	Ave.	**100.00**	**93.33**	**85.56**	**93.64**	**47.78**	**26.83**	**46.55**	**18.37**	**19.77**	***59.53***
SPEA-2	SEAMO	100.00	100.00	100.00	100.00	95.38	94.55	93.74	92.11	82.64	**95.38**
	NSGA-II	100.00	100.00	97.50	100.00	74.80	69.67	86.56	52.04	49.97	**81.17**
	PESA	100.00	100.00	100.00	100.00	86.47	86.75	97.65	83.98	77.21	**92.45**
	Ave.	**100.00**	**100.00**	**99.17**	**100.00**	**85.55**	**83.65**	**92.65**	**76.04**	**69.94**	***89.67***
NSGA-II	SEAMO	100.00	100.00	100.00	100.00	94.83	98.18	93.74	84.33	81.79	**94.76**
	SPEA-2	100.00	100.00	100.00	96.36	75.34	80.96	92.36	65.91	47.14	**84.23**
	PESA	100.00	100.00	100.00	100.00	91.67	88.99	100.00	87.22	82.22	**94.46**
	Ave.	**100.00**	**100.00**	**100.00**	**98.79**	**87.28**	**89.38**	**95.37**	**79.16**	**70.38**	***91.15***
PESA	SEAMO	100.00	90.00	70.89	98.33	45.43	80.93	50.57	53.22	65.21	**72.73**
	SPEA-2	100.00	86.00	67.33	93.01	31.29	75.53	35.80	31.28	22.28	**60.28**
	NSGA-II	100.00	86.00	64.83	94.46	35.35	76.13	35.12	26.91	12.31	**59.01**
	Ave.	**100.00**	**87.33**	**67.69**	**95.27**	**37.36**	**77.53**	**40.50**	**37.14**	**33.27**	***64.01***

Table 3.3: The ave. set coverage values obtained in each problem class, for all pairwise comparisons of algorithms

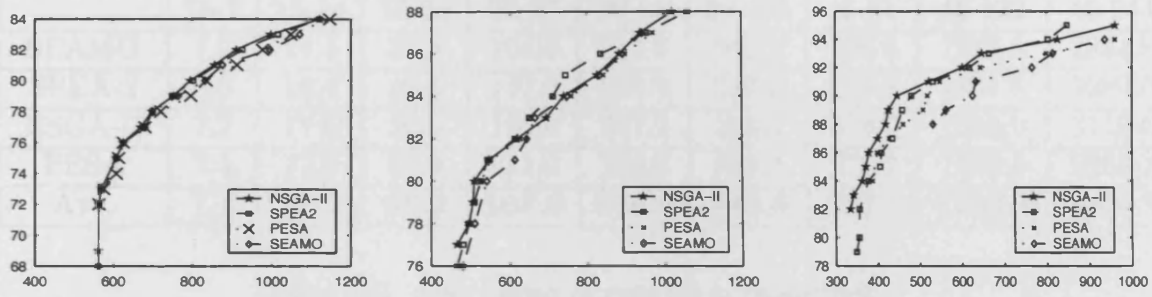

Figure 3.7: Pareto fronts (coverage verses cost) for large problem size instances (v4:45x45) with 61, 122, and 244 candidate sites

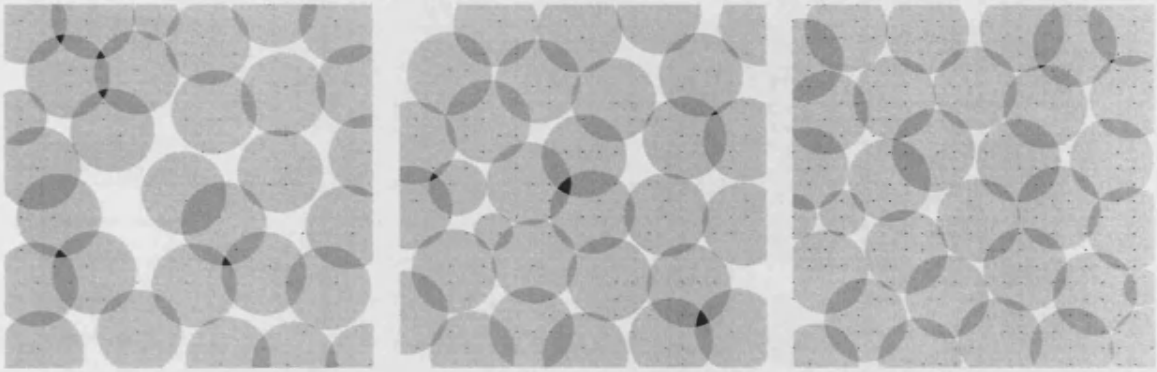

Figure 3.8: Example cell plans with the highest coverage at each density for the large region problem

Algorithm	Problem Instances (km^2– number of candidate sites)									
	15-7	15-14	15-28	30-27	30-54	30-108	45-61	45-122	45-244	**Ave**
SEAMO	3.94	19.63	13.07	14.32	39.92	53.68	25.24	38.07	33.71	**26.84**
SPEA-2	3.94	21.31	13.90	14.94	13.83	15.37	29.35	33.36	42.17	**20.91**
NSGA-II	3.94	21.31	15.02	18.05	25.63	14.84	27.34	30.16	36.83	**21.46**
PESA	3.94	19.29	8.50	14.09	21.53	29.78	18.57	23.81	38.26	**19.75**

Table 3.4: Ave. spacing values by algorithm for each test problem class

Algorithm	Problem Instances (km^2– number of candidate sites)								
	15-7	15-14	15-28	30-27	30-54	30-108	45-61	45-122	45-244
SEAMO	7.2	17.4	37.8	104.0	242.8	542.2	526.6	1272.8	2854.8
SPEA-2	7.8	18.4	40.2	107.0	250.8	553.0	543.8	1305.6	3089.0
NSGA-II	7.2	17.6	39.0	106.0	247.2	548.6	539.6	1300.0	3119.6
PESA	7.4	17.2	38.8	111.0	246.6	555.2	578.4	1293.8	2963.2
Ave.	**7.4**	**17.7**	**39.0**	**107.0**	**246.9**	**549.8**	**547.1**	**1293.1**	**3006.7**

Table 3.5: Ave. speed of execution in seconds

Algorithm	Generation forming S_A	Generation forming S_B					
		250	500	750	1000	1250	1500
SEAMO	250		12.50	0.00	0.00	0.00	0.00
	500	100.00		28.57	0.00	0.00	0.00
	750	100.00	100.00		71.43	16.67	14.29
	1000	100.00	100.00	100.00		33.33	28.57
	1250	100.00	100.00	100.00	100.00		71.43
	1500	100.00	100.00	100.00	100.00	100.00	
SPEA-2	250		0.00	0.00	0.00	0.00	0.00
	500	100.00		27.27	11.11	9.09	9.09
	750	100.00	100.00		22.22	18.18	9.09
	1000	100.00	100.00	100.00		54.55	36.36
	1250	100.00	100.00	100.00	100.00		45.45
	1500	100.00	100.00	100.00	100.00	100.00	
NSGA-II	250		0.00	0.00	0.00	0.00	0.00
	500	100.00		9.09	8.33	33.33	33.33
	750	100.00	100.00		33.33	33.33	33.33
	1000	100.00	100.00	100.00		75.00	58.33
	1250	100.00	100.00	100.00	100.00		58.33
	1500	100.00	100.00	100.00	100.00	100.00	
PESA	250		0.00	0.00	0.00	0.00	0.00
	500	100.00		55.55	55.55	44.44	44.44
	750	100.00	100.00		88.88	77.77	66.67
	1000	100.00	100.00	100.00		77.77	66.67
	1250	100.00	100.00	100.00	100.00		88.89
	1500	100.00	100.00	100.00	100.00	100.00	

Table 3.6: Comparison of intermediate populations for each algorithm, using the set coverage metric for a total of 1500 generations

Chapter 4

2D Cell Plan Decoders

In Chapter 3, the strategy of multi-objective cell planning was delineated. Phase one of the G2PS approach in Chapter 3 used a decoder to determine cell plans given a candidate site permutation π. As with the choice of GA, the specification of the decoder represents a degree of freedom within the multiple objective optimisation framework. The purpose of this chapter is to determine if changes in the specification of the decoder (i.e., sequential algorithm) lead to changes in the quality of the cell plans produced. Six alternative decoders are tested in this chapter. These are explained in Section 4.1 and tested in Section 4.2 on their ability to control cell density. Throughout, NSGA-II is applied, and primary objectives of coverage and cost are considered. Following this, conclusion on the decoders are drawn (Section 4.3).

4.1 Decoders

A decoder produces a cell plan given a permutation of candidate site locations, where each candidate site location is identified by an integer. Given n candidate sites, there are $n!$ possible orderings. The order in which candidate sites are considered for power allocation affects the characteristics of the resultant cell plan. We denote a given ordering of candidate sites π. There are two issues that a decoder must address: (i) the criteria for determining whether a given power setting is permissible (Section 4.1.1) and (ii) a method for selecting the next candidate site for consideration (Section 4.1.2). These factors determine the cell density and therefore the cost and coverage characteristics of the cell plans produced.

4.1.1 Criteria for power allocation

In order to produce efficient cell plans, it is crucial to control cell density and overlap. If cells are too dense they will overlap and duplicate service to too many STP. If excessive this leads to the deployment of more base stations than is necessary for the level of service coverage required, which unduly raises cost. Additionally, high levels of overlap between cells can lead to difficulties when allocating channels to base stations under FDMA protocols such as GSM. STPs receiving service coverage from multiple base stations require wider channel separation for receiving equipment to make a distinction between signals. This constitutes a separate and well studied optimization problem known as the *channel assignment problem* beyond the current investigation.

While high levels of cell overlap is not desirable, some overlap between cells is

tolerated for two reasons. Firstly, cell overlap permits the possibility of seamless *call − handover* for mobile users, where a call can be transferred from one base station to the next without service interruption. Secondly, as base station coverage areas are often irregularly shaped, allowing some cell overlap facilitates cell plans with higher service coverage.

One way to achieve the correct balance of cell overlap is to impose constraints on the degree of overlap allowed between a cell being considered for addition to a cell plan and those already commissioned. Below we define three types of constraint methods for controlling cell overlap in this way. Note that as cells are defined as a set of points (i.e., STPs covered by a given base station), cell size is the cardinality of the set rather than a geographical area. This permits the consideration of cells of irregular shapes, such as when using propagation data from the field.

Pair-wise cell overlap constraint (PCO)

The first constraint restricts cell overlap on a pair-wise basis. Like all constraints, it is checked each time the allocation of a power setting to a site is considered, and is imposed with respect to the set of cells already allocated a power (denoted W). As cells can be of different sizes, this constraint is imposed relative to the percentage of overlap in the smaller cell to avoid situations where a larger cell could completely envelope the smaller without violating the constraint. This can be seen in Figure 4.1, where if 5 were commissioned first, 4 could still be commissioned after. We term this a *pair-wise cell overlap constraint (PCO)* and formally define it as follows:

Definition 4 *Under the* ***pair-wise cell overlap constraint (PCO)****, the*

allocation of p_i to form cell C_i is permissible if and only if the number of STP within C_i exceeds 0 and

$$|C_i \cap C_j| \leq \frac{\alpha}{100} \times min\{|C_i|, |C_j|\}, \forall C_j \in W.$$

This means that pairwise cell overlap is restricted to at most α of the smallest cell in the pair, where α can be set between 0% and 100% inclusive. In other words, if α is exceeded between the cell under consideration and any other cell currently commissioned, the cell C_i is not added at that power level.

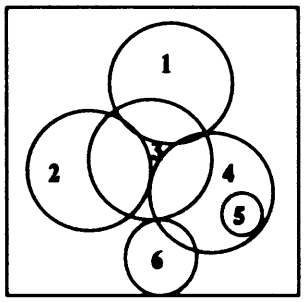

Figure 4.1: Potential cell planning problems

Multi-wise cell overlap constraint (MCO)

The second constraint restricts cell overlap on a multi-wise basis. This constraint was devised to ensure that the proportion of STP covered more than once, *within* the cell being considered, does not exceed a set level. It helps avoid a situation which could occur in PCO, where several cells commissioned in close proximity to one another effectively envelope the 'central' cell without violating the PCO constraint. This can be seen in Figure 4.1, where if 3 were commissioned first under some settings

1, 2, and 4 could still be commissioned after. MCO stops this more frequently because if cells 1, 2, or 4 overlap slightly with another cell (not shown), the MCO constraint is likely to be exceeded, whereas the PCO constraint would not. We term this a *multi-wise cell overlap constraint (MCO)* and formally define it as follows, where M_j is the subset of STP within C_i which are covered more than once by adjacent cells:

Definition 5 *Under the **multi-wise cell overlap constraint (MCO)** the allocation of p_i to form cell C_i is permissible if and only if the number of STP within C_i exceeds 0 and*

$$\frac{|M_j|}{|C_i|} \times 100 \leq \beta$$

This means that cell overlap is restricted to the percentage of overlap that occurs *within* the cell being added (denoted β) as a result of its intersection with cells in W, where β can be set between 0% and 100% inclusive. It is worth observing that while the PCO constraint will hold true at any stage in the creation of the cell plan, the MCO constraint is only guaranteed true at the point at which the current cell is being added.

The potential weakness of MCO is that if a loose constraint (i.e, high β value) is used, similar problems to that discussed for PCO can occur, albeit not as frequently. Alternatively, if a tight constraint (i.e., low β value) is used, favorable opportunities for higher cell power allocation can be missed. This can be seen in Figure 4.1, where if cells 3 and 6 were commissioned, it is possible that the small amount of extra overlap caused by the intersection of 4 and 6 could mean that 4 is not commissioned at the higher power level (as depicted).

Pair-wise and multi-wise cell overlap constraint (PMCO)

The final constraint to consider is a combination of PCO and MCO, termed PMCO. This constraint was developed to curb the likelihood of incurring the potential weaknesses found in PCO and MCO. For example, it has the same power as MCO to more frequently avoid the problem in PCO where cells envelope a 'central' cell. Additionally, it has an enhanced ability to avoid the problem experienced by MCO of not commissioning at a more favourable power setting due to a small amount of extra overlap. It achieves this by having a higher tolerance (higher β setting) than MCO. However, this more loose β setting does not have the potentially disastrous consequences it would have for MCO as it is protected by the α constraint as well. We term this a *pair-wise and multi-wise cell overlap constraint (PMCO)* and formally define it as follows:

Definition 6 *Under* ***pair-wise and multi-wise cell overlap constraint (PMCO)*** *the allocation of p_i to form cell C_i is permissible if and only if the PCO and MCO constraints are not violated.*

Thus, under this approach, not only is the pairwise cell overlap restricted to at most α of the smallest cell in each comparison pair, but also to the overlap that occurs *within* the cell being added.

4.1.2 Selecting next candidate site

We also define two methods for ordering the iteration through the list of candidate sites, π, when applying a method for cell overlap constraint (Section 4.1.1). This

involves progressing either primarily by-site or by-power.

By-Site (-S)

Under the *by-site* approach, sites are iterated over (in the order dictated by π) only once. However, when considered, if the current power setting violates an overlap constraint, a lower power setting is attempted, up to as many times as there are power settings. Once a site is allocated a power setting, the next site in the permutation is considered. The highest power setting (p_k) is allocated first, followed by progressively lower power settings until p_0 (off). The procedure is detailed in Algorithm 2.

Algorithm 2 The *by-site* procedure for selecting the next site in a given ordering

1:

2: Let $L_1, L_2, \ldots, L_{nbs}$ denote an ordering of candidate base station locations, where there must be at least one base station

3: Let $p_0, p_1, \ldots, p_k$ denote the available power settings, where there must be at least one power setting

4: **for** $j = 1$ to nbs **do**

5: Set $i = k$.

6: **while** allocation of p_i to L_j is not permissible (by *PCO*, *MCO* or *PMCO*) and i > 0 **do**

7: $i = i - 1$

8: **end while**

9: **end for**

By-Power (-P)

Under the *by-power* approach, sites are iterated over (in the order dictated by π) and considered for allocation at the highest power setting first. If allocated a power setting, the site is no longer visited. While at least one site remains without an allocated power setting, the power setting is reduced, and the sites which are not yet commissioned are considered again in the order dictated by π. The procedure is detailed in Algorithm 3.

Algorithm 3 The *by-power* procedure for selecting the next site in a given ordering

1:
2: Let $L_1, L_2, \ldots, L_{nbs}$ denote an ordering of candidate base station locations, where there must be at least one base station
3: Let $p_0, p_1, \ldots, p_k$ denote the available power settings, where there must be at least one power setting
4: Set $i = k$.
5: **while** not all candidate sites have been allocated a power setting **do**
6: **for** $j = 1$ to nbs **do**
7: **if** L_j has no power setting allocated **then**
8: **if** allocation of p_i to L_j is permissible (by *PCO*, *MCO* or *PMCO*) and i > 0 **then**
9: L_j is allocated power setting p_i
10: **end if**
11: **end if**
12: **end for**
13: $i = i - 1$;
14: **end while**

Both the by-site and by-power approaches were designed to mimic different approaches a radio engineer might take when attempting to allocate power settings manually. Given there are three methods for cell overlap constraint and two methods for selecting the next candidate site, six decoders are possible: PCO-P, PCO-S, MCO-P, MCO-S, PMCO-P, PMCO-S.

4.1.3 Example of each decoder's action

To appreciate the different approach each decoder takes, consider the following cell planning scenario: There are seven candidate sites, each of which can transmit at zero power(i.e., off), one power, or two power. Dotted lines are used to depict potential cells, and solid lines are used to depict cells which have been formed at power level one or two.

Example of PCO-S, MCO-S, and PMCO-S

In Figure 4.2 the ordering of π is 1, 2, 3, 4, 5, 6, 7, although any ordering of sites could have been adopted. In this figure, we consider PCO-S, MCO-S, and PMCO-S. Taking $\alpha = 30\%$, PCO-S (left) begins by adding cell 1 at full power, as no other cells are commissioned. Cell 2 is also selected, as the overlap between 2 and 1 is less than 30%. Cell 3 is also selected, as the overlap between (3 and 1) and (3 and 2) is less than 30%. Similarly, in this scenario, cells 4, 5, 6, and 7 are also selected at full power.

To demonstrate MCO-S (middle), we use $\beta = 30\%$. MCO-S begins like PCO-S by adding cells 1 and 2. However, at cell 3 the combined overlap with cells 1 and 2 exceeds 30% at full power. At half power, the overlap with cells 1 and 2 is reduced

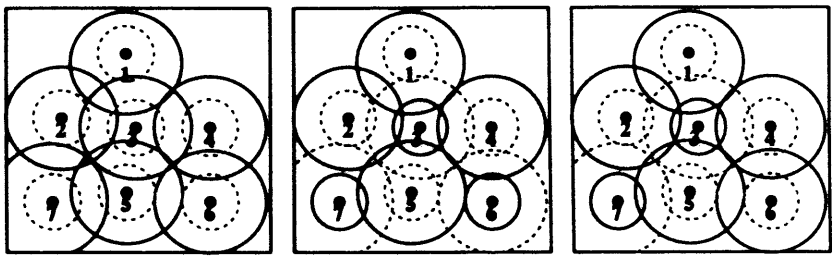

Figure 4.2: Using by-site: PCO-S (left), MCO-S (middle), PMCO-S (right)

considerably, and cell 3 is commissioned. Cells 4 and 5 are added at full power, as beta never exceeds 30%. Conversely, cells 6 and 7 can only be commissioned at half power, as overlap exceeds β at full power.

Finally, PMCO-S (right) is considered with $\alpha = 30\%$, $\beta = 50\%$. PMCO-S proceeds similarly to MCO-S until cell 6. At cell 6, the pair-wise overlap does not exceed α = 30% with cells 4 and 5, nor does the β exceed 50%–whereas for MCO-S the β did exceed the more restrictive 30% constraint; so PMCO commissions 6 at full power. For cell 7, it is assumed that the combined overlap with cells 2 and 5 does exceed 50% in this case, so it can only be commissioned at half power.

Example of PCO-P, MCO-P, and PMCO-P

In Figure 4.3, we demonstrate the by-power method using the ordering of π: 3, 6, 7, 1, 5, 2, 4. Taking α = 30%, PCO-P (left) turns on cells 3, 6, 7, 1, 5, 2, and 4 at full power as the constraint is never exceeded, and therefore has no cells to consider at half power. When compared with PCO-S, the same sites are turned on despite the ordering of π has changed as well as the method for determining the next candidate site. A situation which would create a distinction using PCO-P and PCO-S can be seen in Figure 4.1, where if the α setting were 0% and cell 4 were commissioned

before cell 6, cell 6 could not be commissioned. However, if proceeding by power, cell 4 would not have been commissioned at the low power before cell 6 was commissioned at the higher power.

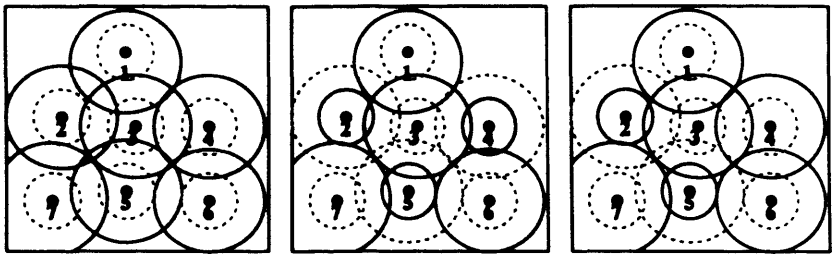

Figure 4.3: Using by-power: PCO-P (left), MCO-P (middle), PMCO-P (right)

Unlike the comparison between PCO-P and PCO-S, MCO-P selects different cells from MCO-S due to the ordering of π. To demonstrate MCO-P (middle), we use $\beta = 30\%$. MCO-P begins by adding cells 3, 6, 7 and 1 at full power, as the β constraint is never exceeded. Cells 5, 2, and 4 are not selected at full power, as β exceeds 30% in each case when the combined overlap with cells already commissioned is considered. MCO-P then lowers the power level, and reconsiders cells 5, 2, and 4, which are now permissible.

Finally, PMCO-P (right) is considered with $\alpha = 30\%$, $\beta = 50\%$. It proceeds similarly to MCO-P in commissioning cells 3, 6, 7, and 1. Cells 5 and 2 are not selected at full power, as the β constraint would be exceeded. However, unlike PCO-P, cell 4 is commissioned at full power, as β does not exceed 50%. PMCO-P then lowers the power level, and reconsiders cells 5 and 2, which are now permissible to be selected.

A number of observations can be made regarding this approach. Firstly, the approach is greedy in the sense that once a base station location is added to CBS' at power p_{max}, the base station cannot be removed from the cell plan CBS' nor can its transmission power be adjusted. Secondly, for a particular list of potential site locations, characteristics (e.g., cost and coverage) of the resultant cell plan CBS' is entirely dependent on the order (i.e., permutation π) in which the base stations are considered for selection. It is our aim to find the best permutations, which lead to the best approximations of cell plans on the Pareto front between the chosen objectives.

4.2 Decoders: Performance

We rigorously investigate each decoder within $G2PS$ for multi-objective cell planning. The model we adopt considers power allocation at base station sites, which is the most important configuration variable. Similarly, the most fundamental objectives are considered when comparing decoder performance. These are the total cost of operational base stations and total area coverage provided. Six general decoders are introduced and compared, taking into account a range of parameter settings for each. Using a sample of randomized test problem scenarios and a general simulated model for wireless propagation, comprehensive empirical evidence is provided from which detailed conclusions are drawn.

The performance of each decoder (using the $G2PS$ framework) was tested using a wide range of synthesised test problems. Test problems are classified by the density of candidate sites in a 30 km^2 working area, as documented in Figure 4.4. Five densities are used. For each test problem class, we produced five randomly generated instances,

denoted version 1 (v1), v2, v3, v4 and v5. This means that the total number of problem instances is 25, which was seen as a sufficient number of testing scenarios to differentiate between the performance of the decoders. Admittedly, the choice is also somewhat arbitrary given the exploratory nature of this work. All problem instances are available at: `http:// www.raisanen.co.uk/ datasets.html`. To help maintain a fair comparison between each decoder, common parameter settings for NSGA-II were adopted for each test–Figure 4.5.

Region size km^2	Density of sites per km^2				
	0.03	0.06	0.12	0.18	0.24
30×30	27	54	108	162	216

Figure 4.4: Number of candidate sites in five problem classes defined by region size and density

Parameter	Description	Setting
...for Hata formula		
f	frequency	800 Mhz
h_b	base station height	31 metres
h_m	receiver height	1.5 metres
S_q	Service threshold	-90 dBm
...for decoders		
α(PCO)	0%, 10%, 20%, 30%, 40%, 50%	
β (MCO)	0%, 10%, 20%, 30%, 40%, 50%	
β (PMCO)	α +10%, α +20%, α +30%, α +40%, α +50%	
p_i(in dBW)	30, 27, 24, 21, 18, 0 (off)	
p_i(in Watts)	1000, 501, 251 125, 63, 0 (off)	
...for NSGA-II		
n	number of generations	500
N	population size	50
—	mutation rate	0.01
—	cross-over rate	1.00

Figure 4.5: Summary of all parameters used in testing

To determine which decoder performed the best, we tested each decoder on every

test problem instance using the following constraint settings:

- PCO-P and PCO-S using all alpha settings
- MCO-P and MCO-S using all beta settings
- PMCO-P and PMCO-S using all pairs of alpha, beta (where beta > alpha)

This leads to 1850 tests (i.e.,(5 problem classes * 5 problem instances * 4(6) for PCO-S, PCO-P, MCO-S, and MCO-P) + (25 * 2(25) for PMCO-S and PMCO-P)). For each test problem instance, we pooled solutions from each decoder run at each constraint setting and then removed dominated solutions to form set S_{PFA} (Section 3.1). In this way, S_{PFA} is our best estimate of the true Pareto front, and as such, a concrete benchmark against which all decoders settings can be compared on a relative basis using the set coverage metric. That is, for each decoder at each constraint setting the set coverage metric is calculated with respect to S_{PFA}, which represents the frequency with which the decoder produces non-dominated solutions as compared to the solutions from S_{PFA}. The average performance was then measured over all test problems, with maximum (Max), minimum (Min), and average (Ave.) scores provided.

4.2.1 Performance of PCO-P vs. PCO-S

The average performance for PCO-P and PCO-S for each problem class is given in Figure 4.6. In Figures 4.7 and 4.8, their respective maximum, minimum, and average performances (along with standard deviation) are given by constraint setting. On average, PCO-S with $\alpha = 50\%$ performs the best, on average weakly dominating

8.52% of S_{PFA} across all test problems, followed by PCO-P with $\alpha = 40\%$ at 8.07%. Notably, taken together the two aforementioned decoders achieved the best average performance across all problem classes. This provides some evidence that less restrictive alpha constraint settings are better for optimization purposes. When comparing the difference between PCO-P and PCO-S, Figures 4.7 and 4.8 show that PCO-P is superior overall, achieving an average S_{PFA} domination of 5.98% as compared to 4.57% from PCO-S. The average performance of the PCO based decoders is 5.28%.

Decoder	Test problem class					Ave.
α	27 v1-5	54 v1-5	108 v1-5	162 v1-5	216v1-5	
PCO-P 00	13.47	5.94	2.72	6.94	5.43	6.90
PCO-S 00	11.98	2.42	0.25	0.49	0.00	3.03
PCO-P 10	3.25	5.67	5.64	5.33	1.09	4.20
PCO-S 10	3.24	1.73	0.00	0.00	1.30	1.25
PCO-P 20	10.68	4.46	3.83	2.68	3.17	4.97
PCO-S 20	10.68	2.82	1.08	0.99	0.00	3.12
PCO-P 30	7.64	3.68	3.21	0.53	6.30	4.27
PCO-S 30	8.43	3.35	5.81	0.25	2.85	4.14
PCO-P 40	4.66	15.38	2.61	9.93	7.80	8.07
PCO-S 40	4.66	15.07	6.27	8.39	2.37	7.35
PCO-P 50	18.01	9.94	4.50	2.60	2.28	7.46
PCO-S 50	19.32	10.70	8.49	2.90	1.21	8.52

Figure 4.6: Ave. PCO weak domination of S_{PFA} by problem class

Stat	*alpha* level						Ave.
	α=0%	α=10%	α=20%	α=30%	α=40%	α=50%	
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	25.00	11.43	21.57	22.37	28.28	34.69	23.89
Ave.	6.90	4.20	4.97	4.27	8.07	7.46	5.98

Figure 4.7: Analysis by α values of the weak-domination of PCO-P over S_{PFA}

Stat	*alpha* level						Ave.
	α=0%	α=10%	α=20%	α=30%	α=40%	α=50%	
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	20.83	12.24	21.57	26.32	31.31	34.69	24.49
Ave.	3.03	1.25	3.12	4.14	7.35	8.52	4.57

Figure 4.8: Analysis by α values of the weak-domination of PCO-S over S_{PFA}

4.2.2 Performance of MCO-P vs. MCO-S

The average performance for MCO-P and MCO-S for each problem class is given in Figure 4.9. In Figures 4.10 and 4.11, their respective maximum, minimum, and average performances (along with standard deviation) are given by constraint setting. On average, MCO-P with $\beta = 50\%$ performs the best, on average weakly dominating 9.52% of S_{PFA} across all test problems, followed by MCO-S with $\beta = 50\%$ at 9.42%. Unlike the best two PCO based decoders, together these decoders did not achieve the best average performance across all problem classes. The best two MCO based decoders only performed well in the two lowest density situations. In higher density situations, decoders with stricter constraint settings performed better. This suggests β constraint settings need to be tailored based on the candidate site density to optimise performance. When comparing the difference between MCO-P and MCO-S, Figures 4.10 and 4.11 show that MCO-P is superior overall, achieving an average S_{PFA} domination of 6.18% as compared to 4.05% from MCO-S. The average performance of the MCO based decoders is 5.12%.

Decoder	Test problem class					Ave.
β	27 v1-5	54 v1-5	108 v1-5	162 v1-5	216v1-5	
MCO-P 00	13.47	5.94	2.72	6.94	5.43	6.90
MCO-S 00	11.98	2.42	0.25	0.49	0.00	3.03
MCO-P 10	3.25	5.43	5.07	1.91	0.77	3.29
MCO-S 10	3.24	2.74	0.40	0.00	2.85	1.85
MCO-P 20	13.51	6.52	3.30	8.52	1.37	6.64
MCO-S 20	14.76	2.71	0.17	0.00	0.49	3.63
MCO-P 30	13.97	4.28	2.53	3.00	4.26	5.61
MCO-S 30	15.54	1.54	0.00	0.00	0.28	3.47
MCO-P 40	8.85	5.19	1.28	7.22	3.15	5.14
MCO-S 40	10.00	3.43	0.20	0.76	0.00	2.88
MCO-P 50	32.04	7.39	4.44	1.69	2.06	9.52
MCO-S 50	36.94	7.13	1.25	1.01	0.78	9.42

Figure 4.9: Ave. MCO weak domination of S_{PFA} by problem class

Stat	*beta* level						Ave.
	β=0%	β=10%	β=20%	β=30%	β=40%	β=50%	
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	25.00	12.62	27.45	27.63	18.37	46.94	26.34
Ave.	6.90	3.29	6.64	5.61	5.14	9.52	6.18

Figure 4.10: Analysis by β values of the weak-domination of MCO-P over S_{PFA}

Stat	*beta* level						Ave.
	β=0%	β=10%	β=20%	β=30%	β=40%	β=50%	
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	20.83	12.24	27.45	31.58	20.41	46.94	26.58
Ave.	3.03	1.85	3.63	3.47	2.88	9.42	4.05

Figure 4.11: Analysis by β values of the weak-domination of MCO-S over S_{PFA}

4.2.3 Performance of PMCO-P vs. PMCO-S

The average performance for PMCO-P and PMCO-S for each problem class is given in Figures 4.12 and 4.13. In Figures 4.14 and 4.15, their respective maximum,

minimum, and average performances (along with standard deviation) are given by constraint setting. On average, PMCO-P (α=40%,β=50%) performs the best, on average weakly dominating 11.05% of S_{PFA} across all test problems, followed by PMCO-S(α=50%,β=60%) at 8.52%. Like the best two PCO based decoders, these decoders performed well (albeit not the best) across all problem classes, and they perform consistently well during optimization with a less strict α constraint. Unlike MCO, the β constraint usually performed best at (α + 10%) or (α + 20%); thus, it does not need to be tailored to candidate site density, probably because PMCO uses a combination of constraints. When comparing the difference between PMCO-P and PMCO-S, Figures 4.10 and 4.11 show that PMCO-P is superior overall, achieving an average S_{PFA} domination of 6.67% as compared to 4.59% from PMCO-S. The average performance of the PMCO based decoders is 5.63%.

4.2.4 Comparison across all decoders

Looking at all the test results, the best performing constraint method was PMCO, with an average weak domination of S_{PFA} of 5.63% across all constraint settings. Also, the by-power method for selecting the next candidate site outperformed by-site in each paired comparison. The best performing decoder overall was PMCO-P (α=40%,β=50%), which on average weakly dominated 11.05% of S_{PFA} across all test problems. This supports the conjecture given in Section 4.1.3 that this decoder avoids inefficient configurations within a cell plan.

Results also suggest that when there is low site density, a less strict constraint is preferable when optimizing cover and cost, probably as with few degrees of freedom,

Decoder	Test problem class					Ave.
α,β	27 v1-5	54 v1-5	108 v1-5	162 v1-5	216v1-5	
PMCO-P 10,20	3.25	5.67	5.06	2.61	6.01	4.52
PMCO-P 10,30	3.25	5.67	5.64	5.33	1.09	4.20
PMCO-P 10,40	3.25	5.67	5.64	5.33	1.09	4.20
PMCO-P 10,50	3.25	5.67	5.64	5.33	1.09	4.20
PMCO-P 10,60	3.25	5.67	5.64	5.33	1.09	4.20
PMCO-P 20,30	10.68	4.51	12.14	1.97	8.25	7.51
PMCO-P 20,40	10.68	4.46	4.97	5.67	3.68	5.89
PMCO-P 20,50	10.68	4.46	3.83	2.68	3.17	4.97
PMCO-P 20,60	10.68	4.46	3.83	2.68	3.17	4.97
PMCO-P 20,70	10.68	4.46	3.83	2.68	3.17	4.97
PMCO-P 30,40	7.64	2.44	5.02	5.59	2.06	4.55
PMCO-P 30,50	7.64	3.86	9.60	1.49	11.72	6.86
PMCO-P 30,60	7.64	3.68	3.21	0.53	6.30	4.27
PMCO-P 30,70	7.64	3.68	3.21	0.53	6.30	4.27
PMCO-P 30,80	7.64	3.68	3.21	0.53	6.30	4.27
PMCO-P 40,50	5.88	14.62	9.98	17.07	7.67	11.05
PMCO-P 40,60	4.66	15.38	2.61	9.93	7.80	8.07
PMCO-P 40,70	4.66	15.38	2.61	9.93	7.80	8.07
PMCO-P 40,80	4.66	15.38	2.61	9.93	7.80	8.07
PMCO-P 40,90	4.66	15.38	2.61	9.93	7.80	8.07
PMCO-P 50,60	18.01	9.94	4.50	2.60	2.28	7.46
PMCO-P 50,70	18.01	9.94	4.50	2.60	2.28	7.46
PMCO-P 50,80	18.01	9.94	4.50	2.60	2.28	7.46
PMCO-P 50,90	18.01	9.94	4.50	2.60	2.28	7.46
PMCO-P 50,100	18.01	9.94	4.50	2.60	2.28	7.46

Figure 4.12: Ave. PMCO-P weak domination of S_{PFA} by problem class

coverage is maximized by being able to commission cells at full power when given the chance. However, as candidate site density increases, and the likelihood of cell overlap increases, a stricter setting is helpful. This may be due to the fact that with increased degrees of freedom, a better solution with less overlap is more likely to be found, with fewer sites ultimately being necessary to maximize coverage.

Decoder α,β	Test problem class					Ave.
	27 v1-5	54 v1-5	108 v1-5	162 v1-5	216v1-5	
PMCO-S 10,20	3.24	1.73	0.00	0.00	0.88	1.17
PMCO-S 10,30	3.24	1.73	0.00	0.00	1.30	1.25
PMCO-S 10,40	3.24	1.73	0.00	0.00	1.30	1.25
PMCO-S 10,50	3.24	1.73	0.00	0.00	1.30	1.25
PMCO-S 10,60	3.24	1.73	0.00	0.00	1.30	1.25
PMCO-S 20,30	10.68	3.18	0.61	0.99	0.78	3.25
PMCO-S 20,40	10.68	2.82	1.02	0.99	0.21	3.15
PMCO-S 20,50	10.68	2.82	1.08	0.99	0.00	3.12
PMCO-S 20,60	10.68	2.82	1.08	0.99	0.00	3.12
PMCO-S 20,70	10.68	2.82	1.08	0.99	0.00	3.12
PMCO-S 30,40	8.43	3.11	0.33	0.56	2.08	2.90
PMCO-S 30,50	8.43	2.90	0.39	0.28	0.29	2.46
PMCO-S 30,60	8.43	3.35	5.81	0.25	2.85	4.14
PMCO-S 30,70	8.43	3.35	5.81	0.25	2.85	4.14
PMCO-S 30,80	8.43	3.35	5.81	0.25	2.85	4.14
PMCO-S 40,50	5.88	12.01	4.48	2.66	7.88	6.59
PMCO-S 40,60	4.66	15.07	6.27	8.39	2.37	7.35
PMCO-S 40,70	4.66	15.07	6.27	8.39	2.37	7.35
PMCO-S 40,80	4.66	15.07	6.27	8.39	2.37	7.35
PMCO-S 40,90	4.66	15.07	6.27	8.39	2.37	7.35
PMCO-S 50,60	19.32	10.70	8.49	2.90	1.21	8.52
PMCO-S 50,70	19.32	10.70	8.49	2.90	1.21	8.52
PMCO-S 50,80	19.32	10.70	8.49	2.90	1.21	8.52
PMCO-S 50,90	19.32	10.70	8.49	2.90	1.21	8.52
PMCO-S 50,100	19.32	10.70	8.49	2.90	1.21	8.52

Figure 4.13: Ave. PMCO-S weak domination of S_{PFA} by problem class

Stat	*alpha*, *beta* level						Ave.
	0,0	10,20	20,30	30,50	40,50	50,60	
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	25.00	14.08	21.57	35.21	30.99	34.69	26.92
Ave.	6.90	4.52	3.25	6.86	11.05	7.46	6.67

Figure 4.14: Analysis by selected α and β values of the weak-domination of PMCO-P over S_{PFA}

Stat	*alpha*, *beta* level						Ave.
	0,0	10,30	20,30	30,60	40,60	50,60	
Min	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Max	20.83	12.24	21.57	26.32	31.31	34.69	24.49
Ave.	3.03	1.25	3.25	4.14	7.35	8.52	4.59

Figure 4.15: Analysis by selected α and β values of the weak-domination of PMCO-S over S_{PFA}

4.3 Decoders: Conclusions

In this chapter, we have considered the effects of changing the decoding algorithm within the optimisation framework. The framework optimizes the ordering of candidate sites, which are transformed into cell plans using a decoder. The decoder algorithm is fundamental as it governs the performance of cell plans found, independent of the technique used to optimise the ordering of candidate sites.

Six decoders have been introduced and evaluated. These decoders use different criteria and approaches to mimic the way in which a radio engineer could approach the commissioning of base stations, on a manual basis. The crucial issues considered are those of power allocation to create cells, and the overlap between cells, given an irregular dispersion of candidate sites. A total of 1850 experiments have been performed using a range of randomized test problems with different density characteristics. For each test problem, the solutions from all decoders were aggregated and the non-dominated members were used to form a best approximation to the Pareto front between coverage and cost. Against this, the performance of each decoder, at each setting, was compared using set-wise weak domination. On averaging over all test problems, the results indicate that PMCO-P is the highest performance approach overall. The results also show that use of the by-power ordering method is consistently

more effectively than by-site. Identifying the decoder with the highest performance means that we can now more accurately assess the trade-off between coverage and cost for the cell planning problem.

Chapter 5

2D Cell Infrastructure Efficiency

In this chapter, we look at ways of finding the characteristics of high performance cell plans under the abstract 2D model. High performance in this context means attaining a high coverage level for relatively low total base station cost. To this end, we introduce the concept of *base station infrastructure efficiency.* This relates network cost to the spatial availability of received signal strength. Our interest is in finding the cell density at which infrastructure efficiency is maximised. Note that this is distinct from finding an optimal level of coverage or optimal level of expenditure. Infrastructure efficiency gives an assessment of a cellular network *independent* of the particular level of coverage or particular level of expenditure.

Additionally, to look at the effects of changing the level of expenditure in the network we introduce the concept of *marginal infrastructure cost of coverage.* This is the increase in cost for one additional unit of coverage. From an economic viewpoint, the rate at which the infrastructure cost of changes is fundamental in determining the *amount* of coverage, and therefore capital investment, which the operator should opti-

mally sustain for profit maximisation. Marginal cost measures the diminishing return from increased expenditure, which is valuable information for network operators.

We wish to observe common trends in the behaviour of this function to quantify the potential diminishing return (in terms of service coverage) for additional infrastructure expenditure. Our aims in this regard are two-fold:

1. Firstly we wish to analyse the effects of *cell density* on the infrastructure cost of the network. High cell density permits higher levels of area coverage but at greater financial cost. We seek to establish the most cost efficient levels for inter-cell overlap, and determine the level of coverage which can be achieved in these circumstances (Section 5.3).

2. Secondly we wish to analyse the effects of *increasing infrastructure expenditure* on service coverage. We seek to determine the extent to which additional investment becomes progressively less effective in facilitating service coverage. This analysis can be used to inform operators regarding the best level of service coverage (Section 5.4).

We consider the implications of cell density on channel re-use across a network. As cell density increases, it becomes progressively more difficult to satisfy signal-to-interference requirements for service area coverage. We quantify the increase in the required span of channels induced by increased cell density (Section 5.5). After this, conclusion on infrastructure efficiency are drawn (Section 5.6). And finally, we analyse the basic 2D model to determine its overall strengths and limitations, and provide motivation for developing a more realistic 3D model (Section 5.7).

5.1 Cost and Profit Literature

In contrast to the wide variety of techniques and models for cell planning, only a few papers have analysed the cost and profit aspects of cellular networks. A frequent underlying premise is that a network is already in existence, and optimal configuration is required. For example in [26], system parameters (e.g., cell size, channel allocation) are optimised to maximise profit. The emphasis in this chapter is configuration to optimise revenue rather than costs, which is appropriate given that site selection is assumed to be fixed. A range of five potential base station types, with different costs and cell radii are available. A Poisson process is used to model traffic and a range of experiments are performed, which determine the optimal cell type and radius for revenue maximisation. However, uniform, regular location of base stations is assumed. In [27] this approach is extended to include the issue of mobility and overlap between cells.

Other papers have mainly considered costs rather than revenue. Sarnecki [70] briefly discusses cost implications of micro-cells. In [45] Katz compares the cost differences between analogue and digital micro-systems. In [65] an engineering cost model has been proposed but the focus was from the regulators perspective rather than the operators. Beyond these papers, cost implications are generally limited to detailed estimates of different components which are incorporated into the objectives to produce a single network design. There is an apparent lack of cost analysis for service coverage from 'green-field' cellular network planning.

5.2 Assessing Infrastructure Efficiency

To assess infrastructure efficiency, we introduce a new measure called the unit cost of coverage (Definition 7), which determines the cost of providing service coverage on a per test-point basis.

Definition 7 *(Unit Cost of Coverage) Let* $\text{cost}_{CBS'}$ *be the total cost of the cell plan and* $\text{cover}_{CBS'}$ *be the number of STP covered. Then the unit infrastructure cost of coverage is:*

$$\frac{\text{cost}_{CBS'}}{\text{cover}_{CBS'}}$$

Note that this metric directly measures the infrastructure efficiency of any given cell plan, with minimisation of this objective preferential. Additionally, it can be applied to directly compare cell plans which have both different total costs and different levels of service coverage. This metric is related to the concept of *domination*, which is a well established concept in multiple objective optimisation used to compare different solutions under multiple objectives. If desired, total cost can also be re-stated as maximisation (rather than minimisation) objective by maximising $-cost_{CBS'}$.

The relationship between unit cost of coverage and Pareto optimality is established by the following Theorem, where TC = $cost_{CBS'}$ and SC = $cover_{CBS'}$:

Theorem 1 *If cell plan s weakly dominates cell plan t then s has lower or equal unit cost of coverage to that of t.*

Proof Let

$$u_s = \frac{TC_s}{SC_s}, \quad u_t = \frac{TC_t}{SC_t}$$

be the unit cost of coverage for s and t respectively. Since s weakly dominates t, $TC_s \leq TC_t$ and $SC_s \geq SC_t$. Hence $u_s \leq u_t$, as required.

Corollary 1 *The highest infrastructure efficient cell plans (i.e., with minimal unit cost of coverage) are Pareto optimal.*

Proof Follows directly from Theorem 1.

Corollary 1 is important important because it reaffirms that we need to search for and consider Pareto optimal cell plans. This can be approximated by finding the best possible sets of mutually non-dominated cell plans, with different levels of service coverage and total cost. To achieve this, the G2PS multiple objective optimisation framework is employed.

5.3 Effect of Cell Overlap on Infrastructure Efficiency

To apply the G2PS strategy, randomly generated test problems are used, classified by size of region and density of candidate base station sites–Table 5.1. This has led to nine problem classes. The maximum cost of a site is permitted (at random) to be twice the minimum cost of a site. This has been applied to mimic the site cost variations which are frequently seen in operational planning scenarios. For each of the nine problem classes, five randomly generated problems have been created (versions $v1, v2, v3, v4, v5$) leading to 45 specific problems overall. For all problem instances,

service test points are located on a regular grid with 100m intervals. Problems are available at `http://www.raisanen.co.uk/datasets.html`.

During all testing in this chapter, we apply the G2PS strategy (Chapter 3.2) using the PCO-P decoder (Chapter 4.1) with omni-directional antenna, where received signal strength is determined by Hata's model (as in Chapter 3.1). Although PMCO-P was found to be the highest overall performing decoder, PCO-P is used to control cell density because pair-wise cell overlap (PCO) constraints are conceptually easier to relate to the issue of cell density. Also, in contrast to PCO, the multi-wise cell overlap constraint (Definition 6) is local rather than global. Therefore it is only imposed when a site is commissioned and may be legally violated later in the algorithm when subsequent sites are added. The 'by-power' approach to next candidate site selection is applied here because this was found to be superior overall to the 'by-site' approach in Section 4.2.4.

Within PCO-P, the parameter α continues to be central to controlling infrastructure efficiency, because it dictates how closely packed cells (and therefore base stations) are. Increasing α will potentially increase total cost of the network (because more base stations can occur) and also increase the potential for coverage. The optimal setting for α, with a view to maximising infrastructure efficiency, represents a useful 'best possible' guide for network planners. The parameters for the model are stated in Table 5.2, and power settings applied are defined in Table 5.3.

Region size km^2	Density of sites per km^2		
	0.03	0.06	0.12
15×15	7	14	28
30×30	27	54	108
45×45	61	122	244

Table 5.1: Number of candidate sites in nine problem classes defined by region size and density

Parameter	Description	Setting
f	frequency	800 Mhz
h_m	receiver height	1.5 metres
n_{gen}	number of generations	500
n	population size	50
—	mutation rate	0.01
—	cross-over rate	1.00
S_q	service threshold	-90 dBm
k	number of power settings	5

Table 5.2: Settings for propagation model and NGSA-II

Power Setting	dBW	Watts
p_5	30	1000
p_4	27	501
p_3	24	251
p_2	21	125
p_1	18	63

Table 5.3: Power settings used in tests

5.3.1 Results and analysis

We aim to find the effects of varying inter-cell overlap (controlled by α) on infrastructure efficiency. Applying the optimisation framework from Chapter 3.2, a set of non-dominated solutions have been found to approximate the Pareto front for each of the 45 test problems, using 13 alternative values of α (0 to 60% in 5% steps). This

leads to a total of 585 test instances, from which conclusions have been drawn. For each non-dominated solution found, the unit cost of coverage has been calculated and assessed. This has been analysed across problem instances by calculating the mean and standard deviation in each problem class, as shown in Tables 5.4 and 5.5.

The results in Table 5.4 show that the highest level of infrastructure efficiency is achieved with a modest amount of cellular overlap. An overlap of approximately 10% dominates as that providing the best infrastructure efficiency values. This is particularly the case for the larger and denser test problems. The profile of the average unit cost of coverage, as maximum permitted cell overlap (α) increases, is shown in Figures 5.1, 5.2 and 5.3. These functions become smoother for the larger, increasingly dense test problem instances, and show an interesting general trend. As α increases, the unit cost of coverage initially falls, up until α reaches a value of approximately 10%. Beyond this value, the unit cost of coverage progressively increases, in a non-linear fashion. Thus, for the best infrastructure efficiency 10% is best, although from a practical standpoint it may not be sufficient for meeting an operator's coverage or handover requirements.

This behaviour can be explained by the freedom for cell site selection. When α is zero (i.e., cell overlap is not permitted) the choices for site selection are most restricted. This has two effects. Firstly it may not be possible to take advantage of a cheapest subset of sites, particularly if their spatial locations would necessarily induce overlap. Secondly, when commissioning cells so they do not overlap, the cells are likely to be smaller (i.e., lower power). This results in cells serving a small number of STP, which will adversely affect the unit cost of coverage. When a high level of

inter-cell overlap is permitted, the progressive increase in unit cost of coverage occurs for two reasons. Firstly, inefficiency is introduced by a high proportion of STP being contained in multiple cells. Secondly, the increased freedom for commissioning cells means that high levels of coverage can be achieved using progressively more expensive sites. In Figures 5.4, 5.5 and 5.6, the average coverage achieved in each problem class is shown. For each geographical region, note that as candidate site density increases, the achievable coverage also increases. To achieve high levels of coverage (e.g., 90% and above), either candidate site density should ideally be increased, or less favourably, maximum permitted inter-cell overlap should be increased. The former alternative is more desirable since it means that the required level of coverage can be achieved using the most efficient level for α.

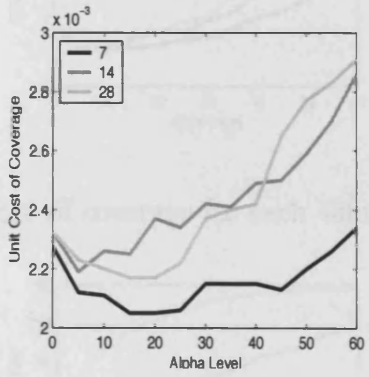

Figure 5.1: Ave. unit cost of coverage for each 15km × 15km problem class

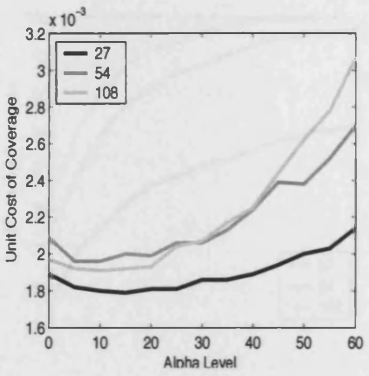

Figure 5.2: Ave. unit cost of coverage for each 30km × 30km problem class

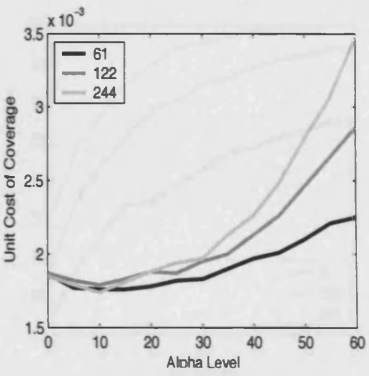

Figure 5.3: Ave. unit cost of coverage for each 45km × 45km problem class

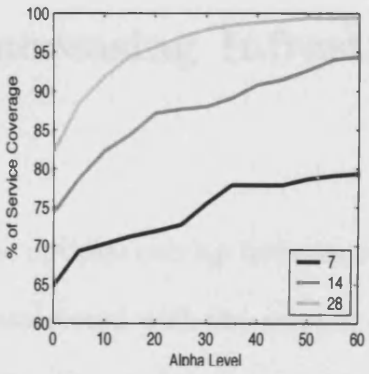

Figure 5.4: Ave. coverage as a function of maximum permitted inter-cell overlap (α) for each 15km × 15km problem class

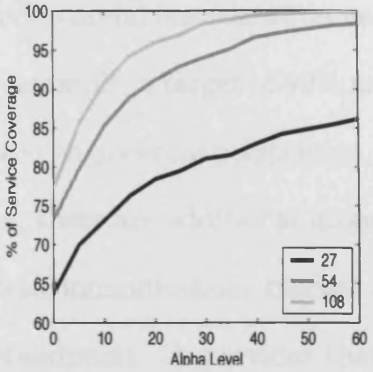

Figure 5.5: Ave. coverage as a function of maximum permitted inter-cell overlap (α) for each 30km $\times$ 30km problem class

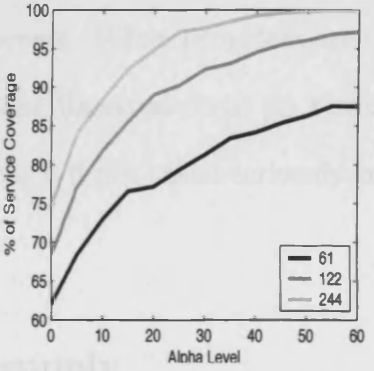

Figure 5.6: Ave. coverage as a function of maximum permitted inter-cell overlap (α) for each 45km $\times$ 45km problem class

5.4 Effects of Increasing Infrastructure Expenditure

As well as determining the optimal overlap between cells, it is important to understand the cost implications associated with the *amount* of service coverage provided, under idealised conditions. In this section, we consider the cost change due to increased spatial service provision. Effectively, we show how the law of diminishing returns, in terms of service coverage, applies to cellular communication networks.

Thus, when planning a network, careful consideration needs to be given to the target level of service coverage. For example, a target of 99% area coverage may represent a plausible objective from a radio engineering perspective, but this may well be a poor choice economically. However, there are additional intangible issues to consider. For example, to even enter the telecommunications market a certain level of coverage is necessary from a practical standpoint. A provider that does not provide excellent wide-area coverage is likely to lose customers quickly. In addition, there may be benefits to being the 'best' service provider in terms of coverage for marketing purposes and the recruitment of customers. When providers are considering cost implications they need to consider both the financial costs (in terms of the cell plan efficiency) and these intangible costs, which if not taken seriously could limit their viability in a competitive market.

5.4.1 Demand and supply

Economics captures idealised behaviour under particular assumptions. We apply economic supply theory to the spatial availability of mobile services, which we consider to be the network operators output. Under perfect economic behaviour, the network operator will ideally wish to know the answer to two fundamental questions: how much will it cost to produce this output and how much revenue will be earned by selling it. By examining how costs and revenues change with output, the operator can maximise profits. Crucial concepts in this regard are marginal cost and marginal revenue. *Marginal cost* is the increase in total cost when output is increased by one unit. *Marginal revenue* is the increase in total revenue when output is increased by

one unit. The important observation is that as long as marginal revenue is greater than marginal cost, an *additional* unit of output will increase total profit. Similarly, if marginal revenue is less than marginal cost, an *additional* unit of output will decrease total profit. Consequently we can use marginal cost and marginal revenue to determine the output level which maximises total profit (i.e., when marginal revenue equals marginal cost). At this point, profit maximisation is achieved. In our scenario, output represents spatial availability of the carrier signal (e.,g., 1% coverage of total STPs equals 1 unit).

In this chapter, we simplify this economic analysis by focussing on marginal cost. This is because marginal revenue data is unavailable, difficult to accurately determine and is also commercially sensitive. In contrast, it is possible to more accurately assess marginal cost of service provision as it is dependent on the principles of radio engineering as described in this thesis. Beyond formal economic analysis, estimated marginal cost plays an important role. From the operators perspective, it is useful when deciding the appropriate *level* of capital investment, even in the absence of information concerning marginal revenue. Analysis of marginal cost determines the point at which investment has a reduced effect in providing coverage, and additional service coverage becomes undesirable. In practice, the operator will have to make a judgement on when this point occurs, in the absence of marginal revenue. Service coverage levels with high marginal cost indicate the increased rate of change in the trade-off between service coverage and cost. A hypothetical optimal trade-off between service coverage and cost is shown in Figure 2.2 (where cost is objective 2, and coverage is objective 1), where the progressive increase in gradient represents the

reduced return, in terms of service coverage, for additional investment. This concept is the cornerstone of economic supply theory and motivates our investigation of marginal cost for the cellular communication network. This has not been addressed in the literature previously.

5.4.2 Determining MC from the Pareto front

We define the *marginal cost of service coverage* at service coverage level s to be the minimal increase in total cost if service coverage is to be increased from s to $s+1$. Due to the density of STPs in a given region, we measure service coverage in percentage units of the total number of STPs. We develop a technique to estimate marginal cost for *specific* irregular cell planning scenarios, such as the test problems in Section 5.3. This involves two components. Firstly, the Pareto front between service coverage and total cost must be approximated, by obtaining a range of non-dominated solutions through the application of the G2PS optimisation framework. Secondly, marginal cost is derived by obtaining a numerical estimate of the gradient using specific points from the approximated Pareto front (Equation 5.1).

Given a range of non-dominated cell plans $CBS'_1, CBS'_2, \ldots, CBS'_k$, we identify each by the pair (s_i, c_i), where s_i is the percentage service coverage and c_i is the total cost of CBS'_i. We assume that the non-dominated cell plans are indexed by increasing service coverage level, that is $s_1 < s_2 < \cdots < s_k$. We estimate the marginal cost of service coverage using a subset S of non-dominated of cell plans from $CBS'_1, CBS'_2, \ldots, CBS'_k$. Graphically, the selected subset of non-dominated cell plans represent extreme points on the Pareto front, identified by forming the convex

hull of points when plotted in the service coverage and cost plane, as shown in Figure 5.7. Members of S are identified as follows. Cell plan CBS'_1 is included by definition. The lowest index cell plan (i.e. lowest cost) in S is denoted by index l. Each cell plan from $CBS'_1, CBS'_2, \ldots, CBS'_k$ with an index greater than l (i.e. higher cost) is considered for inclusion in S. The cell plan CBS'_m from $CBS'_1, CBS'_2, \ldots, CBS'_k$, where $m > l$, is identified such that

$$mc_l = \frac{c_m - c_l}{s_m - s_l} \tag{5.1}$$

is minimal. CBS'_m is added to S and the process is repeated until CBS'_k is included in S. Visualisation of this process is shown in Figure 5.7. The minimal value for mc_l defines our estimate of marginal cost at coverage level s_l. This is a best possible estimate of marginal cost at output level s_l in the sense that mc_l is a lowest estimated increase in cost for an additional unit (1 percent) of service coverage, given the non-dominated solutions found–Figure 5.7.

5.4.3 Results

Using the parameter settings in Table 5.2 and the optimal setting for α, we derive the approximated Pareto front and resultant marginal cost curve for each of the medium and large test problems. These are displayed in Figures 5.8 and 5.9. Note that the approximated Pareto fronts become increasingly convergent and smooth as candidate site density increases. This is because higher density test problems are less distinguishable and consequently the characteristics of the resultant cell plans, in terms of cost and service coverage, are similar. Additionally, the profile of the Pareto fronts is remarkably similar, in exhibiting a progressive and increasing gradient which

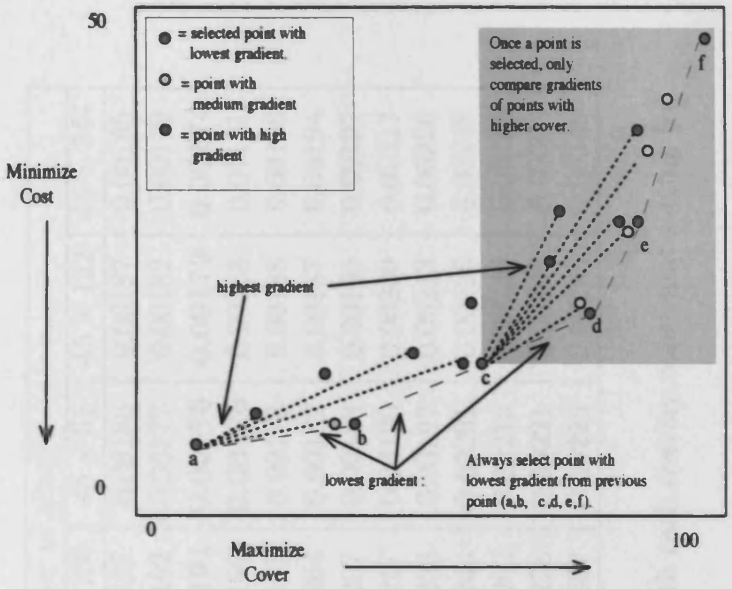

Figure 5.7: Selection of extreme points for estimating the Pareto front

is captured by the corresponding marginal cost curve. In most cases, a significant increase in marginal cost occurs, which is frequently characterised by a sharp terminal peak. This quantifies the difficultly in covering additional service test points without incurring multi-covered STP, which induces significant increase in cost from increased cell density. The start of the sharp terminal peak in marginal cost provides a useful critical point in terms of service coverage level. Coverage beyond this point needs to be carefully considered given the increased marginal cost of coverage, which is equivalent to a significantly diminished return in service coverage for every further unit of capital expenditure.

α	Test Problem Class ($km^2 \times$ number of sites)								
	15 × 7	15 × 14	15 × 28	30 × 27	30 × 54	30 × 108	45 × 61	45 × 122	45 × 244
0%	0.00228	0.00232	0.00232	0.00189	0.00209	0.00197	0.00186	0.00187	0.00185
5%	0.00212	**0.00219**	0.00223	0.00182	**0.00196**	0.00192	0.00177	0.00182	0.00179
10%	0.00211	0.00226	0.00220	0.00180	**0.00196**	**0.00191**	**0.00176**	**0.00179**	**0.00174**
15%	**0.00205**	0.00225	**0.00217**	**0.00179**	0.00200	0.00192	**0.00176**	0.00183	0.00181
20%	**0.00205**	0.00237	**0.00217**	0.00181	0.00199	0.00193	0.00178	0.00188	0.00188
25%	0.00206	0.00234	0.00222	0.00181	0.00206	0.00204	0.00182	0.00187	0.00194
30%	0.00215	0.00242	0.00237	0.00186	0.00206	0.00207	0.00183	0.00195	0.00197
35%	0.00215	0.00241	0.00240	0.00186	0.00213	0.00217	0.00190	0.00200	0.00213
40%	0.00215	0.00249	0.00242	0.00189	0.00224	0.00225	0.00197	0.00213	0.00226
45%	0.00213	0.00250	0.00265	0.00194	0.00239	0.00244	0.00201	0.00226	0.00248
50%	0.00220	0.00259	0.00277	0.00200	0.00238	0.00262	0.00210	0.00246	0.00278
55%	0.00226	0.00270	0.00283	0.00203	0.00252	0.00278	0.00221	0.00266	0.00306
60%	0.00234	0.00286	0.00291	0.00214	0.00270	0.00307	0.00225	0.00286	0.00349

Table 5.4: Ave. unit cost of coverage for each test problem class with each overlap constraint setting (α)

α	Test Problem Class (km^2 × number of sites)								
	15 × 7	15 × 14	15 × 28	30 × 27	30 × 54	30 × 108	45 × 61	45 × 122	45 × 244
0%	0.00017	0.00026	0.00018	0.00009	0.00015	0.00015	0.00010	0.00005	0.00011
5%	0.00034	0.00021	0.00012	0.00011	0.00008	0.00015	0.00010	0.00007	0.00006
10%	0.00036	0.00026	0.00015	0.00006	0.00009	0.00004	0.00008	0.00004	0.00010
15%	0.00032	0.00024	0.00013	0.00004	0.00012	0.00006	0.00007	0.00004	0.00010
20%	0.00031	0.00031	0.00013	0.00004	0.00006	0.00006	0.00008	0.00004	0.00013
25%	0.00032	0.00027	0.00025	0.00004	0.00006	0.00014	0.00007	0.00009	0.00004
30%	0.00019	0.00032	0.00036	0.00006	0.00008	0.00011	0.00009	0.00006	0.00006
35%	0.00021	0.00037	0.00026	0.00010	0.00008	0.00010	0.00012	0.00011	0.00007
40%	0.00021	0.00032	0.00028	0.00009	0.00009	0.00008	0.00011	0.00010	0.00010
45%	0.00023	0.00034	0.00020	0.00007	0.00004	0.00009	0.00012	0.00004	0.00007
50%	0.00028	0.00022	0.00019	0.00012	0.00010	0.00009	0.00011	0.00007	0.00006
55%	0.00027	0.00033	0.00018	0.00010	0.00014	0.00015	0.00005	0.00012	0.00008
60%	0.00027	0.00032	0.00007	0.00010	0.00008	0.00019	0.00006	0.00005	0.00019

Table 5.5: SD figures associated with the ave. values in Table 5.4

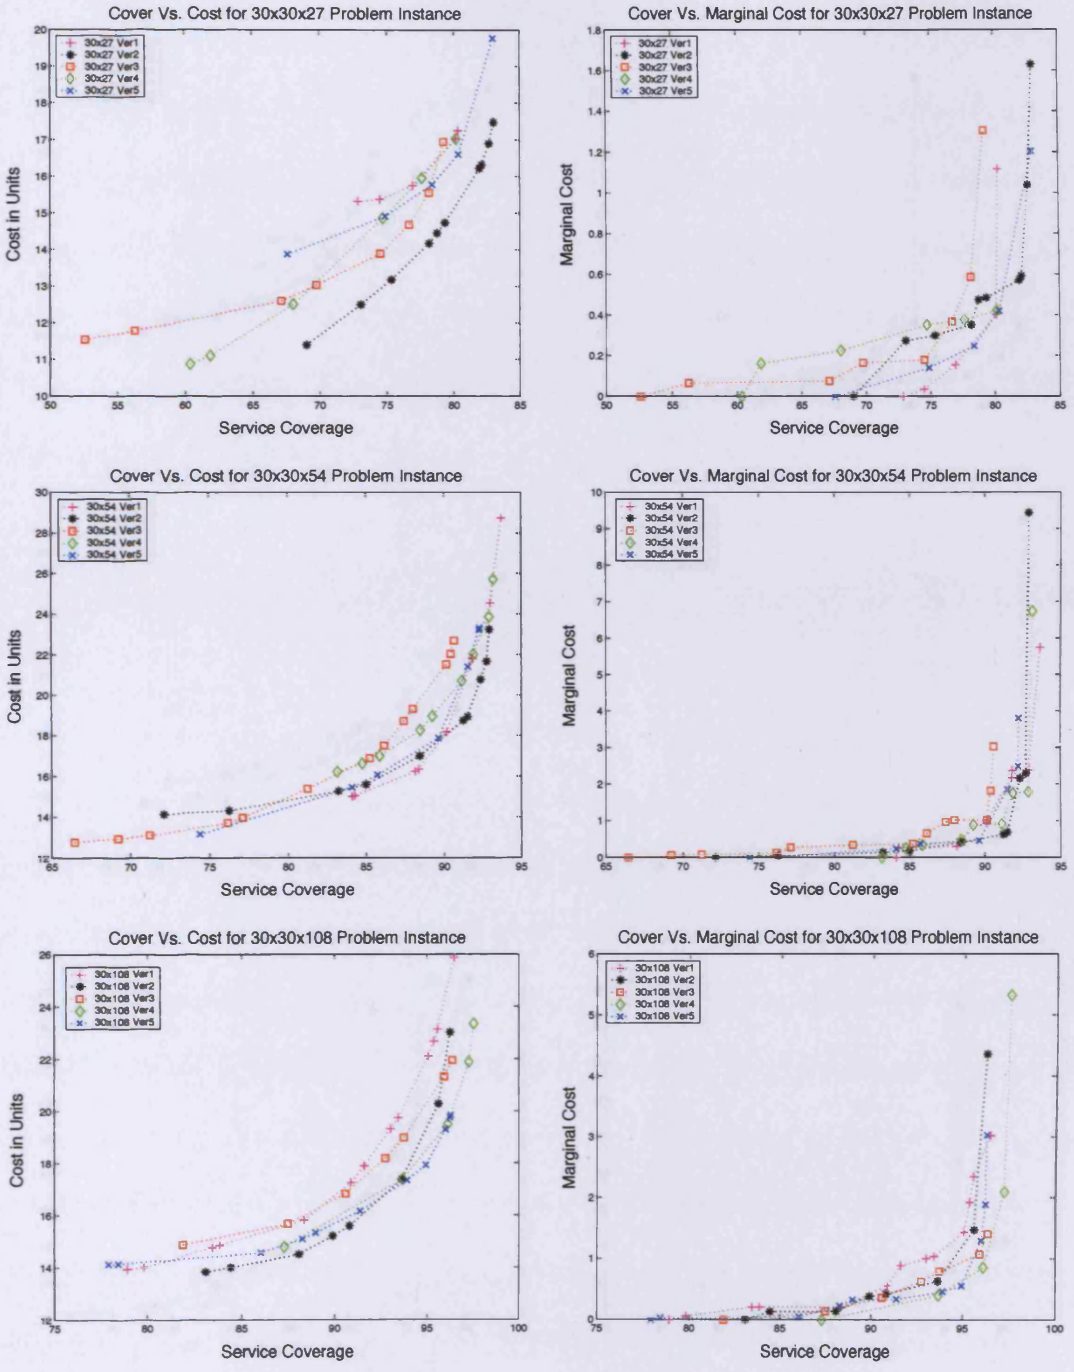

Figure 5.8: Pareto fronts and marginal cost curves for the 30×30 km region with 27, 54, and 108 candidate sites

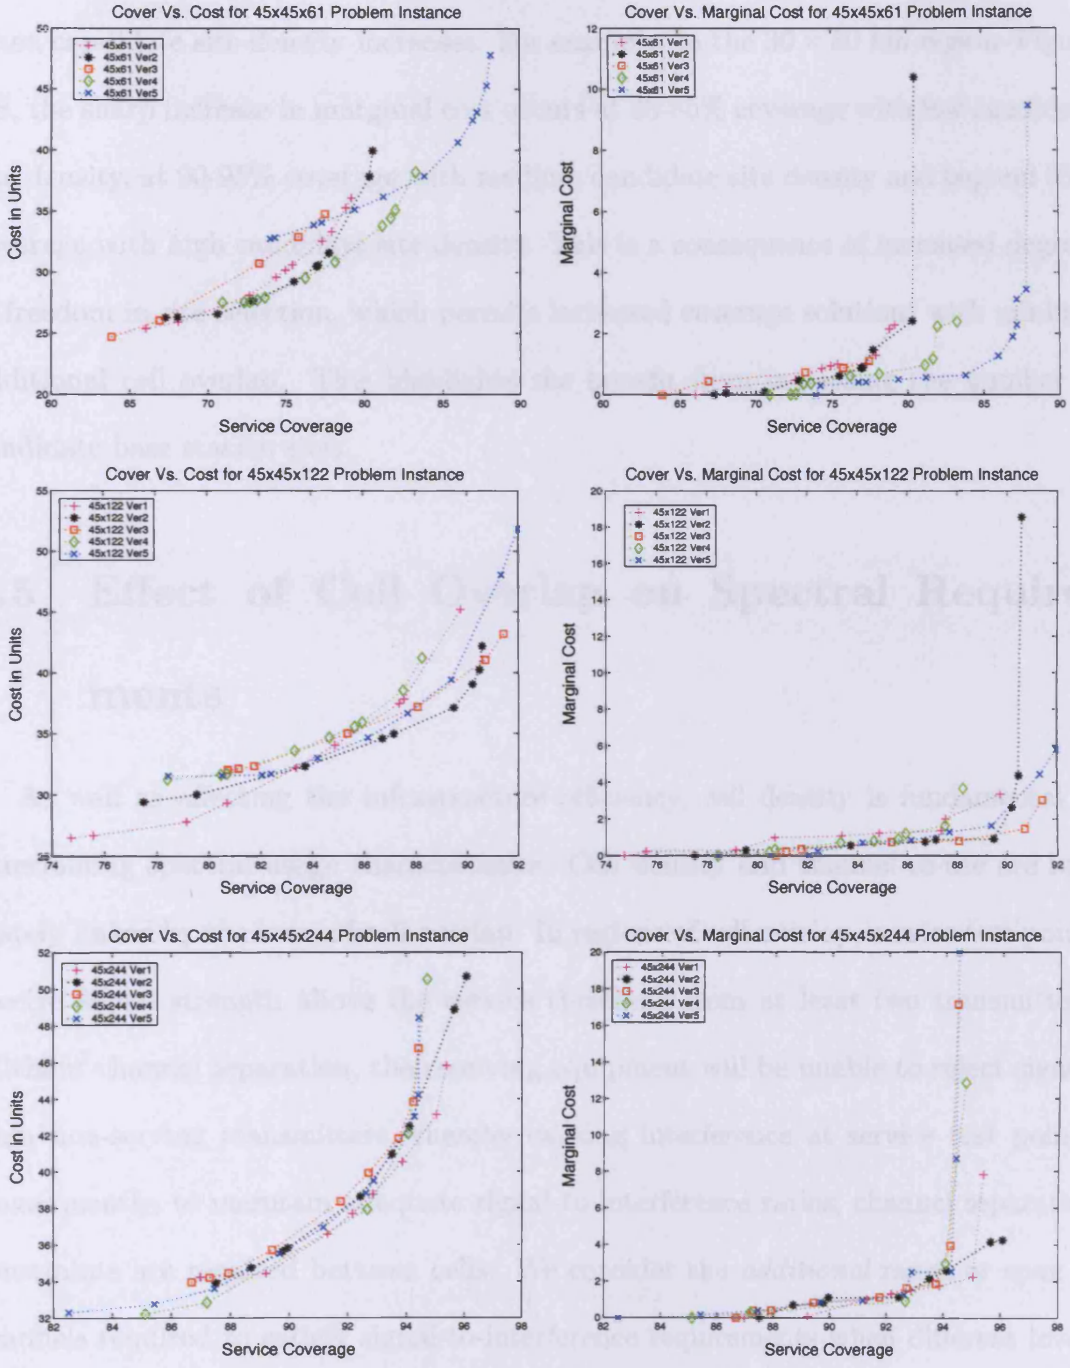

Figure 5.9: Pareto fronts and marginal cost curves for the 45×45 km region with 61, 122, and 244 candidate sites

The sharp increase in marginal cost occurs at progressive higher coverage levels when candidate site density increases. For example, in the 30×30 km region–Figure 5.8, the sharp increase in marginal cost occurs at 80-85% coverage with low candidate site density, at 90-95% coverage with medium candidate site density and beyond 95% coverage with high candidate site density. This is a consequence of increased degrees of freedom in site selection, which permits increased coverage solutions with minimal additional cell overlap. This highlights the benefit from increasing the number of candidate base station sites.

5.5 Effect of Cell Overlap on Spectral Requirements

As well as affecting the infrastructure efficiency, cell density is fundamental in determining spectral usage characteristics. Cell density and channel re-use are intimately linked by the issue of cell overlap. In regions of cell overlap, service test points receive signal strength above the service threshold from at least two transmitters. Without channel separation, the receiving equipment will be unable to reject signals from non-serving transmitters, thereby causing interference at service test points. Consequently, to maintain adequate signal-to-interference ratios, channel separation constraints are required between cells. We consider the *additional* range or span of channels required to satisfy signal-to-interference requirements when different levels of cell overlap (controlled by α) are permitted. We assume the application of FDMA technology and consider spatial service area coverage in an uncapacitated network.

In the following Section, we explain the approximation used.

5.5.1 Channel separation constraints

We consider channel separation between pairs of overlapping cells. Constraints are generated using a 'worse-case' service test point from the overlap region, which is selected using one of two assumptions.

- *Strongest Server*

  Under this assumption, service test points in the overlap region are assumed to be served by the cell providing the greatest received signal strength. Therefore the 'worse-case' service test point is selected as that receiving near-equal signal strength from both cells.

- *Non-Strongest Server*

  Under this assumption, service test points in the overlap region may be served by the cell providing the weakest received signal strength. When all cells are co-channel, the 'worse-case' test point is that with the lowest signal-to-interference ratio.

On selection of the 'worse-case' service test point, channel separation can then be applied between the overlapping cells until signal-to-interference ratio goes beyond the required threshold. For experimental purposes, a required signal-to-interference ratio of 9dB has been applied. In Table 5.6, we tabulate the factor by which interference is reduced due to adjacent channel separation. These are taken to emulate the GSM system. Repeating this process across all pairs of cells, we obtain a set

of necessary constraints which must be satisfied if target signal-to-interference ratios are to be achieved in regions of cell overlap. For a given a set of constraints, the frequency assignment software package FASoft [41] has been used to determine the exact minimum span of channels required to satisfy all constraints. Specifically the exhaustive backtracking approach has been employed.

Channel Separation	1	2	3
Attenuation (dB)	18	50	58

Table 5.6: The effect of adjacent channel separation

5.5.2 Results

We have considered each of the five problem instances from the nine problem classes in Table 5.1, using thirteen settings for α (0 to 60% in 5% steps). For each problem instance and α setting, a non-dominated front of solutions has been obtained. The highest coverage solution from this set has then been considered for purposes of channel assignment.

The average spans have been analysed and are presented in Figures 5.10 and 5.11. The span of channels reported in these figures represents the *additional* spectral requirements which are induced when cells are permitted to overlap, as controlled by α. A number of observations are notable. As to be expected, increasing cell overlap increases the required span. This is most acute when the non-strongest server assumption is applied, as shown in Figure 5.10. However, for low values of α, similar requirements are imposed by the strongest server and non-strongest server assumptions, as evident in Figure 5.11. Furthermore, beyond initial α values, the average

required span increases in a broadly linear fashion. For the settings used in these experiments, the results show that to provide adequate signal-to-interference ratio for the most infrastructure efficient level of cell overlap (approximately 10%), a small additional increase in number of channels is required.

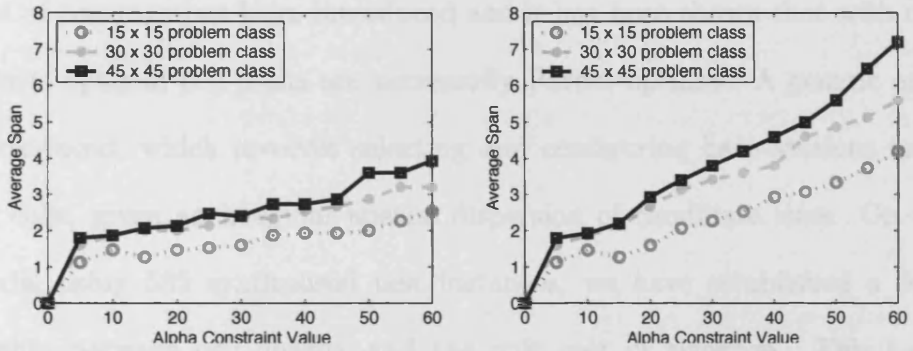

Figure 5.10: Ave. span by α constraint setting given strongest (left) and non-strongest (right) server assumption

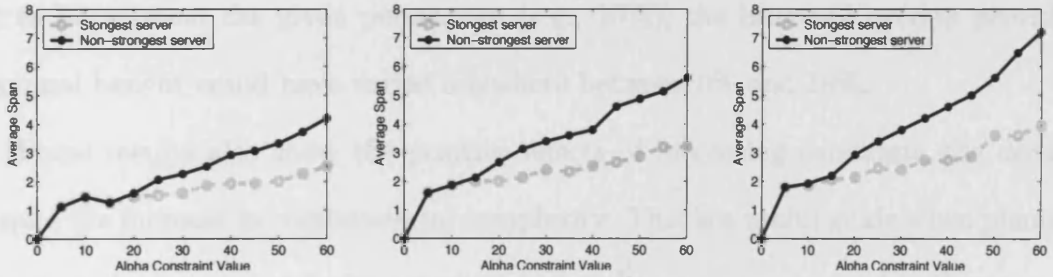

Figure 5.11: Ave. span by α constraint setting for each problem class: 15×15 (left), 30×30 (middle), 45×45 (right)

5.6 Infrastructure Efficiency Conclusions

In this chapter, we have brought together important aspects of radio engineering, economics and combinatorial optimisation for wireless network design. Our focus has been on the efficiency of infrastructure in providing spatial service coverage. The unit cost of coverage has been introduced and it has been shown that with regard to this metric, optimal cell plans are necessarily Pareto optimal. A generic model has been introduced, which involves selecting and configuring base stations to provide wireless cells, given an irregular spatial dispersion of candidate sites. On applying this model using 585 synthesised test instances, we have established a consistent relationship between cell density and the unit cost of coverage. This has shown that there is maximal financial benefit from permitting approximately 10% inter-cell overlap when sites are (uniform) randomly dispersed and sites costs vary (uniform) randomly between 1 and 2. It is worth noting that as the inter-cell overlap is only set to be *at most* the given percentage (e.g., 10%), the inter-cell overlap providing maximal benefit could have varied anywhere between 0% and 10%.

These results also show the positive effects of increasing candidate site density, despite the increase in combinatorial complexity. This is a useful guide when planning operational networks with characteristics related to our general model. However, we note that our modeling assumptions are restricted and cells may well tessellate in an irregular fashion.

A new concept, called marginal cost of service coverage has also been introduced to assess the effects of additional infrastructure expenditure on coverage. The diminishing return in service coverage for additional units of expenditure has been

quantified and assessed. This highlights a progressively rapid decline in return (in terms of service coverage) on investment. This means that for networks where a high level of coverage is required (e.g., greater than 90%) a small reduction in coverage level would most likely lead to a very significant reduction in infrastructure costs. On a practical level, this is useful for judicious operators to keep in mind, as at a critical point, trying to achieve more coverage will result in an overall financial loss. This approach casts the problem closer to economic supply theory and encourages operators to identify a point beyond which additional investment in service coverage cannot be justified economically, i.e., additional investment has little additional impact on service coverage. Finally, for spatial area coverage, the impact of cell overlap on channel requirements has been assessed. This has determined a progressive increase in required span due to increased overlap, but this remains modest at the level required for optimal infrastructure efficiency.

5.7 Analysis of 2D Model

On consideration of Chapters 3, 4, and 5 the main benefit of the 2D model is that it is easy to implement and affords fast computational speed. It provides an ideal testing environment in which to confirm the efficacy of the generic model, algorithms to resolve the CPP, and infrastructure efficiency. Testing with the 2D model and the overall strategy for multi-objective cell planning (i.e., G2PS) proved that this model and approach can produce high quality cell plans in terms of service coverage and cost. However, one has to be conscious of the fact that the evidence in support of the approach is only theoretical due to the following limitations:

- As traffic, terrain, and shadowing are not considered in the 2D model the radiation patterns are idealised and circular, similar to those found in disk graph models. While useful in terms of generalized model and algorithm testing, it is an oversimplification given the shape of even omni-directional antennas is not circular in reality.

- In addition, using only omni-directional antennas is an oversimplification when most operators use large and small directive masts as well. That is, the choice is not simply at what power one omni-directional antenna should be, it is a matter of deciding how many antennas to place (i.e., one to three) at a site location and at what power, tilt, and azimuth.

- In addition, coverage in the 2D model was simply assumed (except during specified testing on spectral cost within this chapter) if the service threshold was met by any antenna at a STP. However, in a realistic scenario, signals 'compete' with one another, with coverage going, for instance, to the antenna providing the strongest signal (i.e., best server model).

- The 2D model allowed measurement of only STP coverage and cost. If traffic test points (TTP) and a best server model were added, traffic capacity, interference, and handover could also be measured. Also, it would be possible to resolve channel assignment more realistically.

Due to these limitations, it would not be prudent to say with confidence that the 2D model can definitely produce high quality cell plans in a real-world environment. This limitation equally applies to statements regarding the efficacy of disk-graph mod-

els and the demand node concept, which are also over-simplifications of a real-world testing environment. Additionally, the computational speed afforded by a simplified model might mean that certain aspects of a given cell planning approach are no longer feasible in a more realistic model. While complexity analysis can be performed to guess at the speed of execution, actual run times given a more realistic simulation are still difficult to foresee with certainty.

In order to allow stronger statements regarding the confidence of the application of any cell planning algorithm to real-world wireless network environments, and to improve the accuracy of the computational cost, more effort is needed to produce a realistic simulation environment. In other words, the efficacy of an approach to the CPP, and the confidence to which it could be applied to a real-world environment, is directly proportional to the realism of the simulation environment. Therefore, if we wish to build efficacious models and have confidence in their applicability to the real-world, the models used must surpass those afforded by the 2D model, disk graph models, and the demand node concept. This will be addressed in the following Chapter.

Chapter 6

3D Model

In this chapter, we engineer a 3-dimensional (3D) GSM physical layer simulation environment which is appropriate and practical for simulating real world problems and then extend the 2D cell planning strategy to the 3D environment. Due to commercial sensitivity, there is very limited public access to data for network planning. As we wish to make this simulation environment freely available to others, no commercially protected data sources are used. In addition to the materials provided here, the interested reader is referred to the Momentum (Models and Simulation for Network Planning and Control of UMTS) Project at http://momentum.zib.de/index.php, which, although designed specifically for UMTS, also provides a number of free data sets for network planning.

To establish a realistic 3D model, the 2D test point simulation environment presented in Chapter 3.1 is extended to meet the following requirements:

- Radiation patterns must not be idealised. They should vary based on the given circumstance, e.g., the surrounding terrain, shadowing effects, traffic levels, and

antenna type and parameter settings.

- Antennas commonly used by GSM operators should be available for use within the model. These include omni-directional, small directive, and large directive antennas. In addition, the model should allow up to one omni-directional or up to three directive antenna at a site (i.e, which is a common operational parameter), and allow all relevant parameters (i.e., power, tilt, and azimuth) to be modifiable.

- Signals at the same service test point from different antenna should be allowed to 'compete' with one another, with coverage going, for instance, to the antenna providing the strongest signal (i.e., best server model).

- Traffic test points (TTP) should be distributed throughout the working area to simulate demand.

To this end, the 2D simulation environment is converted into a 3D one as delineated in Section 6.1. The available antennas used in the 3D model are then discussed in Section 6.2 along with results from a battery of tests in Section 6.2.1. Following this, an introduction to 3D cell planning in Section 7.1 is followed by the details of the strategy in Section 7.2.

6.1 3D Simulation Environment

The technique used here to generate input data for the cell planning strategy is similar to that described in Chapter 3.1. and originally found in [68] and [67], where

network test points and site locations are defined within a working area W and any point within W is defined by it Cartesian coordinates: (x,y). The model is extended here to include height (x,y,z) at every test point, which enables the angle of incidence and propagation loss to be estimated between test points and candidate sites. A major strength of this approach is that the creation of the simulation environment is transparent (i.e., all the details of its creation are known) which makes it useful for in-depth scientific investigation. This model uses simulated down-link received signal strength. This is the standard approach used in frequency division communication systems such as the current GSM mobile telephone system. It is important to realise that we are primarily aiming to create scenarios that have the level of detail and complexity faced in 'real' network planning scenarios, and importantly offer similar challenges. However the scenarios do not map on to actual cities or particular propagation environments, but this does not impede their value for network planning and design.

6.1.1 Components of working area

We first describe the model used within the EU project *ARNO* (*Algorithms for Radio Network Optimization, IT Project 23243*) and reported in [89]. The *ARNO* model involves several components: the network size (km^2), mesh (m), total traffic (Erlangs), S_q (dBm) and number of RTP, STP, and TTP. We closely follow a green-field town scenario, see Table A.1, as given in [89]. These parameters (and simulation environments) have also been used in, e.g., [10, 81]. As the *ARNO* simulation environment and details of its formation are not available to the public, they are not

ideal for scientific investigations. We create five publicly available scenarios. These network are termed SE1, SE2, SE3, SE4, and SE5 for simulation environments 1, 2, 3, 4 and 5 and are available for download (along with the corresponding angle of incidence and propagation loss matrices) at http://www.raisanen.co.uk/datasets.html. These five scenarios were devised to reproduce common cell planning scenarios (e.g., to service a town or road network).

The component parts within W allow us to simulate wireless coverage. These components are:

- A set of reception test points RTP $= \{RTP_1, \ldots, RTP_{n_{RTP}}\}$, at which signal reception quality is measured.
- A set of service test points STP $= \{STP_1, \ldots, STP_{n_{STP}}\}$, where a signal must be received above a minimum specified service threshold, S_q, to ensure a required quality of service.
- A set of traffic test points TTP $= \{TTP_1, \ldots, TTP_{n_{TTP}}\}$, which each carry a traffic load measured in Erlangs.
- A set of candidate base station site locations CBS $= \{CBS_1, \ldots, CBS_{n_{CBS}}\}$, which can be selected for configuring up to 3 antenna,and where each CBS_i has an associated commissioning cost $\$(CBS_i)$.

In this scenario, which is visible in Figure 6.1, the TTP $\subseteq$ STP $\subseteq$ RTP. The RTPs form a rectangular grid of evenly spaced points at a specified mesh increment (e.g., every 200 meters). The STPs and TTPs indicate areas an operator wants to service and could theoretically be any subset of RTP locations. The STPs require a minimum

service threshold S_q to be met, which is the signal level necessary to facilitate wireless communication. In this chapter, S_q is assumed to be -90 dBm.

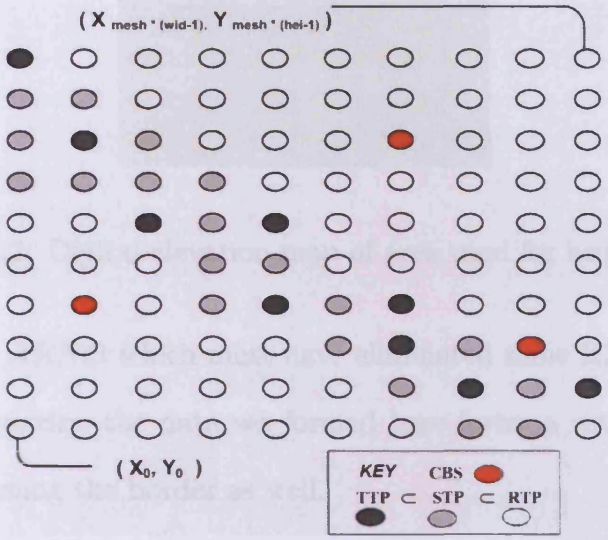

Figure 6.1: Each test point and CBS has an associated (x,y,z) value. The (x,y) values start at (0m,0m) and increase to ((mesh*(width-1)), (mesh*(length-1))), which would be (1800m, 1800m) in the instance portrayed–assuming a 200m mesh increment

SE	#CBS	#RTP	#STP	#TTP	Traffic	Mesh(m)	Working Area(km)
ARNO	568	56,792	17,393	6,656	2,988.12	200	50.0 x 46.0
SE1	568	56,792	17,393	8,590	2,988.27	200	49.4 x 45.6
SE2	568	56,792	17,393	3,985	2,988.75	200	49.4 x 45.6
SE3	568	56,792	17,393	6,602	3,221.84	200	49.4 x 45.6
SE4	568	56,792	56,792	11,953	2,988.25	200	49.4 x 45.6
SE5	568	56,792	56,792	56,792	2,988.12	200	49.4 x 45.6

Table 6.1: Greenfield town network characteristics for ARNO (3_0) and five simulation environments

To form a working area 50,000 x 46,000 m^2 consistent with the *ARNO* data and which has 56,792 RTPs, we form a rectangular grid of size 248 RTPs by 229 RTPs. At a mesh increment of 200m, this forms an area that is 49,400 x 45,600 m^2. This figure

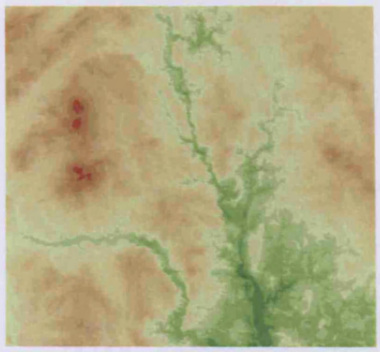

Figure 6.2: Digital elevation map of area used for height data

differs slightly from *ARNO* which must have eliminated some RTP at the edges. To eliminate any irregularity, the data set formed here forms a single closed region of RTP, with RTP forming the border as well.

We extend the *ARNO* model now by explicitly adding a height setting to all components. The height information is taken from a portion of a Digital Elevation Model, see Figure 6.2, of the UK is associated to each RTP [1]. 568 randomly positioned CBS have been added as shown in Figure 6.3. Generally, the layout of traffic in a region involves proximity to a town centre or a road network–Figure 6.4. This has been estimated in five different ways to five different scenarios:

1. SE1 mimics a town centre where traffic distributes as a function of distance–Table 6.2, from a hypothetical town centre (left picture in Figure6.4). The closer to the centre, the more traffic and consequently the greater the difficulty of resolving traffic demand. This scenario was allowed the most variation in

[1]This data is from LANDMAP Project. The Project was funded by The Joint Information Systems Committee (JISC) to provide orthorectified satellite image mosaics of Landsat, SPOT and ERS radar data and a DEM for the whole of the British Isles. Copyright University Of Manchester/University College London Year 2001. Original Landsat 4 & 5 LANDSAT data ©NOAA. Distributed by CHEST under license from Infoterra International.

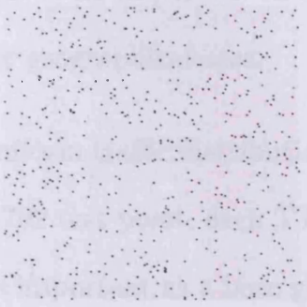

Figure 6.3: Example of each candidate base station location

traffic, with 21 randomly determined values between 0 and 1: 1.00000, 0.99995, 0.63990, 0.59987, 0.39984, 0.37988, 0.35993, 0.33983, 0.33555, 0.31539, 0.30971, 0.26002, 0.23932, 0.21988, 0.18396, 0.15500, 0.13981, 0.12114, 0.11871, 0.07981, 0.03673.

2. SE2 simulates traffic distributed along a hypothetical road network in a town, where every traffic test point (TTP) carries 0.75 Erlangs (middle picture in Figure 6.4). This scenario is expected to be difficult to resolve in terms of traffic capacity, as it has the highest density of traffic per TTP.

3. SE3 combines the approaches of SE1 and SE2, with roughly half the traffic distribution coming from each (right picture in Figure 6.4). This is an important network as it combines two strategies and increases the total amount of traffic by 8%.

4. SE4 follows the general approach of SE1, but the traffic distribution function is modified–Table 6.3, to focus traffic more closely around five hypothetical town centres (left picture in Figure 6.5) instead of one. Each TTP for this network

carries 0.25 Erlangs. This is an important network as it forces resolution of traffic throughout a larger geographical area.

5. Finally, SE5 provides a uniform traffic distribution (right picture in Figure 6.5). Considering there are 56,792 test point, each TTP carries 0.052615157 Erlangs of traffic. This scenario is important as a base guide, and is expected to be the easiest to resolve during planning, and the most likely to achieve an ideal cell tessellation.

Quotient	Probability
0.00	1.00
0.00 to 0.10	0.90
0.10 to 0.20	0.80
0.20 to 0.30	0.70
0.30 to 0.40	0.60
0.40 to 0.50	0.50
0.50 to 0.60	0.40
0.60 to 0.70	0.30
0.70 to 0.80	0.20
0.80 to 0.90	0.10
0.90 to 1.00	0.00

Table 6.2: Probability of TTP placement for SE1 and SE3 where the quotient is formed by dividing the distance to an STP from the town centre by the distance of the furthest STP from the town centre

Two other network characteristics required in each scenario are the angle of incidence matrix (AIM) and propagation loss matrix (PLM). The AIM defines the vertical angles from each CBS to each RTP, and the PLM defines the signal loss from each CBS to each RTP. The AIM is formed by finding the tangent of angle θ between the CBS and the specified RTP as depicted in Figure 6.6. In instances where the straight line of between the CBS and the destination RTP falls between two

Quotient	Probability
0.00	1.00
0.00 to 0.05	0.90
0.05 to 0.10	0.60
0.10 to 0.15	0.50
0.15 to 0.20	0.40
0.20 to 0.25	0.30
0.25 to 0.30	0.20

Table 6.3: Probabilities of TTP placement in SE4, taking the distance to an STP from the town centre divided by the furthest distance from the town centre

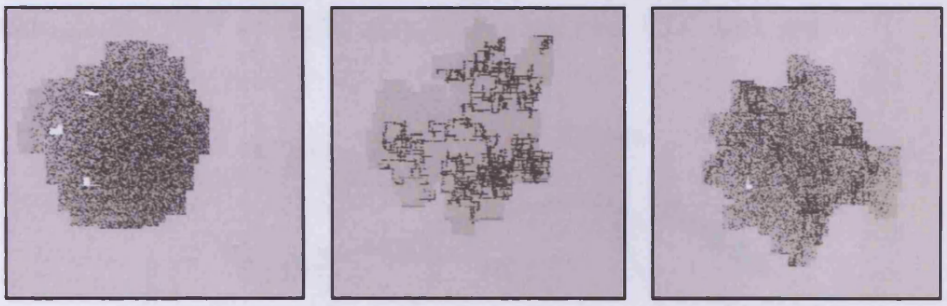

Figure 6.4: SE1 (left) models a town where traffic is distributed as a function of distance from the centre. SE2 (middle) models a town with street-style traffic distribution. SE3 (right) combines SE1 and SE2. RTP are light gray, STP gray, and TTP dark gray

RTP, the height (used for 'opposite' in the tangent formula) is an estimate based on linear interpolation–Figure 6.7. This was considered the most accurate method which does not rely upon exact height measurements taken at points which fall between two points of known height. The PLM is formed by using the distances from each CBS to each RTP in the propagation loss formula–Equation 6.1. In this way, the PLM provides the level of signal loss from every CBS to every RTP in W.

Details regarding how to recover the data (once the files are downloaded from http://www.raisanen.co.uk/datasets.html) is available in Appendix A.

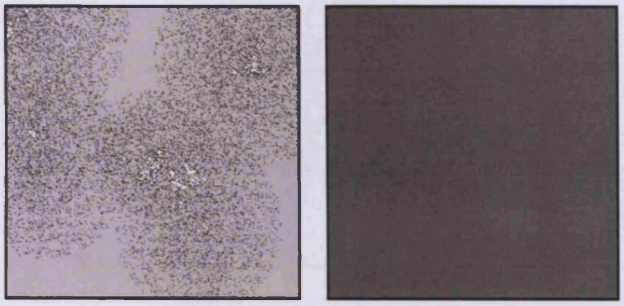

Figure 6.5: SE4 (left) models a town where traffic is distributed as a function of distance from five town centres. SE5 (right) has uniform traffic distribution across the working area. RTP are light gray, STP gray, and TTP dark gray

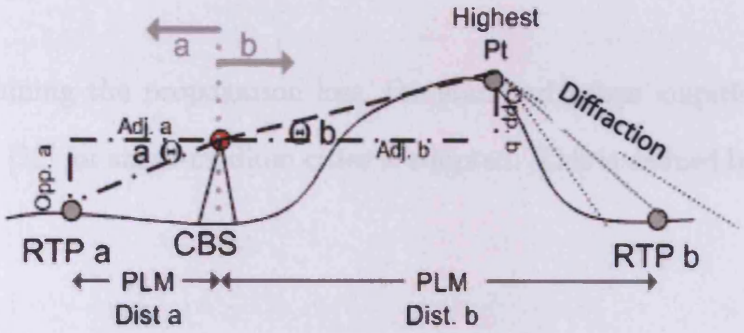

Figure 6.6: *AI* tangents and *PLM* distances in two cases: (a) where RTP_a is the highest point between CBS and RTP_a, and (b) where RTP_b is not the highest point. In case(a), the adjacent and PLM distances are equal, while in case(b) they are not, as the *AI* is determined by the highest point which diffracts the signal to RTP_b

6.1.2 Signal strength losses

There are a number of factors which weaken signal strength. The ones specifically considered here are the propagation loss (PL) which occurs as signal strength attenuates with distance, random shadowing effects, and signal losses (in dB) at antennas as a result of deviation from the signal main axis.

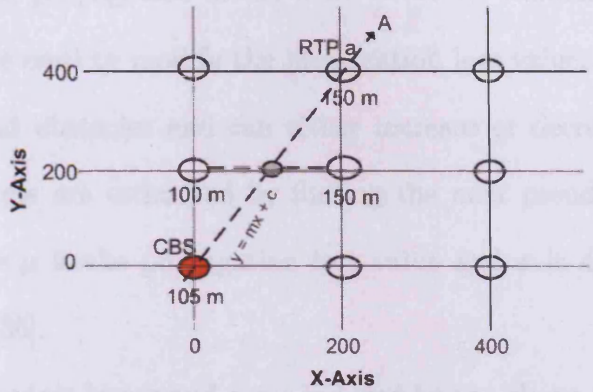

Figure 6.7: Linear height transformation. Given the coordinates of the CBS and RTPa, and the y-axis, computes x using the formula $y = mx + c$. In this instance, x = 125.

For determining the propagation loss, the standard urban empirical model proposed by Hata [32] for small-medium cities is adopted. This is defined by the following formula:

$$PL = 69.55{+}26.16\ \log(f){-}13.82\ \log(h_b){-}a(h_m){+}(44.9{-}6.55\ \log(h_b)){*}\log(R) \quad (6.1)$$

where:

$$a(h_m) = (1.1 *\ \log(f) - 0.7) * h_m - (1.56 *\ \log(f) - 0.8) \quad (6.2)$$

given the variable for frequency f is set to 800MHz, the base station height h_b is the height of the land plus a random number between 30-60m, the mobile receiver height h_m is 1.5m, and the distance R in kilometers from each base station to each RTP are as calculated in Figure 6.6 for PLM. While Hata's model was used, other propagation models and settings could have equally been adopted.

In order to make the propagation model more accurate, random shadowing effects, proposed in [36], are used to modify the propagation loss value. These simulate the effects of clutter and obstacles and can either increase or decrease the strength of reception. The effects are estimated by finding the next pseudo-random Gaussian value (μ, σ)[2], where μ is the propagation loss value and σ is 4 dB, a conservative value suggested in [36].

In addition to random losses and gains induced by the above circumstances, there are signal losses (in dB) at antennas as a result of deviation from the signal main axis. There is one vertical diagram (VDIAG) for omni-directional antenna, and a VDIAG and horizontal diagram (HDIAG) for both directive antennas. An example of these diagrams can be found in [89]. The vertical radiant loss is calculated from two angles: the angle of incidence (AI) measured from base station to RTP (AI° in Figure 6.8, and the vertical tilt which varies from 0,-1...-15 (tilt° in Figure 6.8, where the real incidence angle Θ is measured from the tilt clockwise to AI and is calculated as seen in Figure 6.8: (tilt - AI modulus 360) - (AI - tilt°). The actual signal loss as a result of Θ (denoted Θ.Loss) is then read from the appropriate antenna loss information[3]. For example, if $\Theta = 36$, the loss is 5.99 dB.

The horizontal radiant loss is also calculated from two angles: the antenna azimuth (Az), i.e., the direction the antenna is pointing, and the horizontal angle from 0° to the RTP (RTP°), where the real horizontal loss angle ε is calculated by: (RTP° - Az), see Figure 6.9. The actual signal loss as a result of ε (denoted ε.Loss) is then

[2]This uses the polar method of G. E. P. Box, M. E. Muller, and G. Marsaglia, as described by Donald E. Knuth in The Art of Computer Programming, Volume 2: Seminumerical Algorithms, section 3.4.1, subsection C, algorithm P.

[3]This information must be obtained from a given wireless network operator, as antenna diagrams vary.

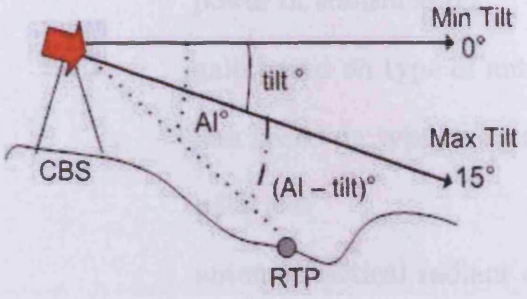

Figure 6.8: Tilt

read from the appropriate antenna loss information. For example, if $\varepsilon = 36$, the loss is 4.44 dB.

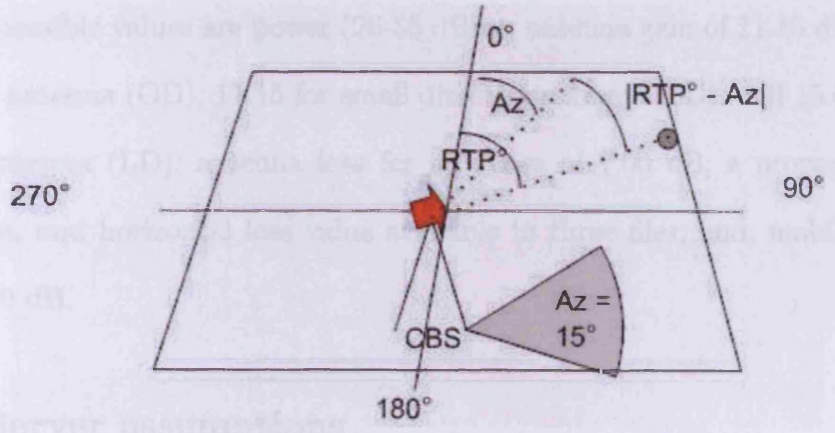

Figure 6.9: Azimuth

The field strength F of a signal from a given antenna antenna is measured at a given RTP. The RTP is said to be covered if F exceeds the service threshold, given here as -90 dBm. F is formally defined as follows:

F = (the final reception strength

Power	power of transmission
+ Ant.Gain	gain based on type of antenna
- Ant.Loss	loss based on type of antenna
- PL	path loss
- Θ.Loss	antenna vertical radiant diagram loss
- ε.Loss	antenna horizontal radiant diagram loss
+ Mob.Gain	gain based on mobile receiver
- Mob.Loss)	loss based on mobile receiver

where the possible values are power (26-55 dBm); antenna gain of 11.15 dB for omni-directional antenna (OD), 17.15 for small directive antenna (SD), and 15.65 for large directive antenna (LD); antenna loss for all types of 7.00 dB; a propagation loss, vertical loss, and horizontal loss value available in three files; and, mobile gain and loss set at 0 dB.

6.1.3 Server assumptions

We consider the server assumption whenever the signal strength at a given STP from two or more antennas is above the service threshold, as we need to decide at this point which antenna is servicing a given RTP for clarifying, for instance, which antenna is carrying the associated traffic load. Generally, there are two approaches to resolving this issue. One is the strongest server assumption, and the other the lowest SIR (i.e., signal-to-interference ratio) server assumption. These are defined as

follows:

- *Strongest Server*

  Under this assumption, service test points where the signal strength is above the threshold from two or more antennas are assumed to be served by the antenna providing the greatest received signal strength.

- *Lowest SIR Server*

  Under this assumption, service test points where the signal strength is above the threshold from two or more antennas may be served by the cell providing the lowest signal-to-interference ratio.

We use the strongest server, or best server, assumption, as it provides a more expedient means for determining which antenna services a given STP. That is, when using the strongest server, one only needs to compute the strength of signals at a point and select the strongest one. With the lowest SIR server, on the other hand, one would need to compute the ratio of the power of the wanted signal to the total residue power of the unwanted signals, which is significantly more intensive, particularly considering the magnitude of times this would need to done over a trial.

6.1.4 Network components and objective measures

Given the working area, where the number of test points in W is denoted n (and therefore, the number of STP, for example, is denoted n_{STP}), and knowledge of the field strength from each CBS to each RTP, the following network components can be re-defined:

- A *cell* corresponds to the set of test points *covered* by one antenna, where the signal received is $\geq S_q$ and stronger than signals from other antennas.

- A *site* may consist of one omni-directional antenna, or up to three directive antennas. Thus, a site may provide up to three cells.

- A *network* CBS' is a subset of all sites within W. CBS' contains those sets with at least one antenna with a non-zero power setting p, where the CBS', or *solution*, has five corresponding network objective values.

The five *network* objective values are service coverage, cost, traffic hold (or capacity), handover, and interference, and are defined as follows:

- *Coverage* is the sum of the service test points (STP) covered in the working area divided by the total number of STP expressed as a percentage.

$$cover_{CBS'} = \frac{\sum_{i=1}^{n_{STP}} 1, if STP_i \ covered}{n_{STP}} \times 100.$$

- *Cost* is the number of sites commissioned with at least one antenna with a non-zero power.

$$cost_{CBS'} = \sum_{CBS_i \in CBS'} \$(CBS_i).$$

- *Traffic hold* is the sum of the traffic at each site location (CBS) in the network; where the traffic at a given site is the sum of the traffic for each cell at the site i (i.e., where the traffic for a given cell is the sum of traffic for all TTP covered

by the cell, up to a maximum of 43 Erlangs) is denoted $CBS_T(i)$; divided by the total network traffic expressed as a percentage.

$$traffic_{CBS'} = \frac{\sum_{i=1}^{n_{CBS'}} CBS_T(i)}{\sum_{i=1}^{n_{TTP}} TTP_i} \times 100.$$

- Call handover is achieved by a cell if there exists STP in its coverage area such that signals $\geq S_q$ from four other cells are within 7 dB. These values were suggested by a radio planner but must ultimately be considered arbitrary. We define the objective *handover* as the number of cells achieving call handover in CBS', denoted $C_H num$, divided by the total number of cells in CBS', denoted $Cnum$, expressed as a percentage.

$$handover_{CBS'} = \frac{C_H num}{Cnum} \times 100.$$

- *Interference* is the sum of signals above -90 dBm at STPs in the network that are not the best server or used in handover, where the sum of the interfering signals for a given STP is denoted I_i. This figure is divided here by the total number of STP for ease of representation.

$$interference_{CBS'} = \frac{\sum_{i=1}^{STP_i} I_i}{n_{STP}} \times 100.$$

6.1.5 Simulation environment conclusions

We have presented a 3D model that captures detail as faced when planning actual cellular networks. Although approximations have been made, the model allows the

measurement of key network design objectives, including service coverage, traffic hold (capacity), interference, handover and cost. Importantly, irregular and conservative path loss has been estimated taking into account aspects of terrain. The model maintains flexibility in that one can change the granularity or resolution of test points in the working area to suit one's need for either computational speed or increased realism.

The limitations or weaknesses of the simulation environment are:

- It would be possible to develop more advanced traffic distribution models to more accurately reflect distribution in the real-world. For example, if traffic estimates were freely available for test points, these could be used instead of generated ones.

- The placement of base station site locations is random, rather than specific to the service area. This means that it would be possible for a site to be in an awkward position for planning. However, this is partially beneficial as cell planning algorithms need to work harder to find optimal sites.

- The computational time is longer due to increased realism.

- Spectral cost and channel assignment are not currently in-built to the simulation model. This is due to the fact that these are also NP-hard problems, and the combined difficulty of resolving them concurrently given the current speed of computers and the chosen cell planning model is untenable.

However, the model has many strengths over the 2D model, demand node concept, and disk graph models. In particular, radiation patterns vary based on the

given circumstance, e.g., the surrounding terrain, shadowing effects, antenna type, and parameter settings. The model allows up to one omni-directional or up to three directive antenna at a site and makes all relevant parameters (i.e., power, tilt, and azimuth) modifiable. Traffic test points are allocated to cells on a best-server model. Finally, traffic test points are not clumped, but are distributed according to characteristics commonly found in small to medium sized towns. In addition, the model is adaptable to a number of situations not covered within this thesis:

- The model can be used in wireless network expansion scenarios, as well as green-field scenarios.
- Information provided by the model could be interpreted into a disk graph or demand node model.
- The channel assignment problem can be considered after a cell plan is developed.
- The model can be adapted to provide either form of base station placement (freedom or limited).

6.2 Single Site Coverage

Now that the simulation environment has been formulated, it is possible to investigate the three antenna types: omnidirectional, small directive, and large directive. This investigation is undertaken in order to clarify the effects of changes to the simulation environment (e.g., shadowing or non-shadowing) and the effects of changes to antenna settings (e.g., azimuths). A specialized simulation environment was created specifically for this purpose (Section 6.2.1).

As previously noted, each antenna can handle a maximum of 43 Erlangs of traffic and each site can have one omnidirectional or up to three directive antennas [61]. Each antenna also has a specified gain and loss–Table 6.4, and an associated antenna diagram. The antenna diagrams, in conjunction with antenna settings, determine the vertical and horizontal radiant losses (Section 6.1.2). This means that the radiation patterns from each are different, and the effects of different combinations of antennas at a site is yet unknown. For this reason, coverage patterns are investigated within a shadowing (with increased realism) and non-shadowing (idealized) environment.

Table 6.4: Antenna gains and losses in dB

Antenna Type	Gain	Loss
Omnidirectional (OD)	11.15	7.00
Large Directive (LD)	15.65	7.00
Small Directive (SD)	17.15	7.00

A typical omnidirectional (OD) coverage pattern is depicted in Figure 6.10 with [#10] and without [#9] shadowing. The OD has only vertical radiant loss, as its horizontal radiation pattern always covers 360°. However, both the large directive (LD) and small directive (SD) antennas have vertical and horizontal radiant loss–as they 'point' in a given direction. Typical radiation patterns for three SD antennas with azimuths set at 0°, 120°, and 240° respectively are given in Figure 6.10 with[#11] and without [#1] shadowing. Similarly, three large directives with [#15] and without [#5] shadowing. Beneath [#1][#5][#11][#15], the radiation patterns when only a single directive antenna is commissioned at 0°, 120°, and 240° degrees respectively. These figures show the typical patterns of coverage for each antenna (by azimuth), and show that shadowing results in less symmetric radiation patterns.

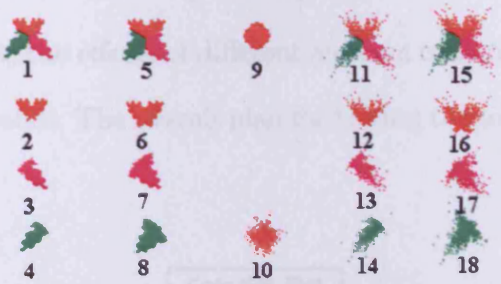

Figure 6.10: Diagrams 1-9 show radiation patterns in a non-shadowing environment. Diagrams 10-18 show patterns in a shadowing environment. All antenna power settings are 55 dBm

6.2.1 Antenna testing

Testing antenna configurations is necessary to determine the best settings to use when attempting to configure a network. It allows insight into which antenna configurations provide the most coverage, which is important when trying to seek lowest cost networks, and which changes to an antenna most effectively reduce or increase coverage, which is important when trying to fine-tune coverage or capacity. For example, if a given cell is overloaded, is it better to reduce the power or the tilt? Also by recording coverage in one testing environment, other testing environments can be compared. For example, all tests here are carried out in both a shadowing and a non-shadowing propagation environment. This information is useful for informing the decoding algorithm used for the SE.

The antenna simulation environment (ASE) used for these tests was designed specifically to test the effects of different antenna configurations. To achieve this, there is only one site (as only one is necessary for measurements) positioned in a central location (in the same size working area as SE1-5). The antenna is set to

a height of 30m and is surrounded by a completely flat surface. By eliminating the variation in land height, the effects of different antenna configurations should be more accurate and generalizable. The overall plan for testing the antenna configurations is given in Figure 6.11.

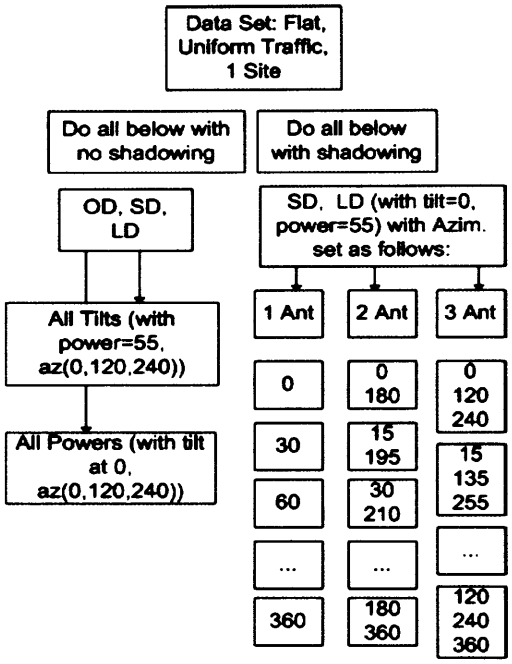

Figure 6.11: Outline of test plan for antenna attributes

Starting at the top of Figure 6.11, all tests are carried out using ASE with and without shadowing. The three main variables to alter at a site (in addition to how many antenna and of which type) are power, tilt, and azimuth. Differences caused by power and tilt can be measured for OD, SD, and LD antennas. Differences caused by azimuth can only be measured for SD and LD antennas.

To simulate this environment, one uses the field strength formula delineated in

Section 6.1.2 to determine the strength of the signal from each base station to each RTP. Then, using the server assumptions in Section 6.1.3, one determines the objective measures delineated in Section 6.1.4 to quantify each cell plan created.

Effects of altering tilt

On the left-hand side of Figure 6.11, the STP coverage is measured for OD (with power set at 55dBm), SD and LD (with power set at 55dBm and three antennas with azimuths of 0°, 120°, 240°) by changing the tilt from 0° to -15° in -1° increments. The results are in Figure 6.12.

There are several notable results: First, antennas in the shadowing environment provide more coverage. Considering the shadowing effects are normally distributed, this means the additive influences are more noticeable than the subtractive ones. However, the *patterns* of antenna coverage in both environments was similar. Second, OD antenna coverage peaks at -2° tilt rather than the expected 0° and steadily decreases on either side. While this may partially be the effect of ASE, it does mean that one cannot assume coverage is necessarily the greatest at 0° tilt. Third, for SD and LD antennas coverage peaks at 0° drops dramatically for -1° and -2° and then rises from -3° to -6° before falling again. These patterns were found to be true at other powers as well, although with slight shifts. For example, while coverage peaks for OD antennas at -2° (from 55-48dBm), it shifts to -3° starting at 47dBm.

For example, one can see that when a SD is tilted (top line of Figure 6.13) the amount of STP coverage does not simply decrease as tilt is increased, nor do the shapes stay the same or necessarily become smaller. Ultimately, this means that tilt-

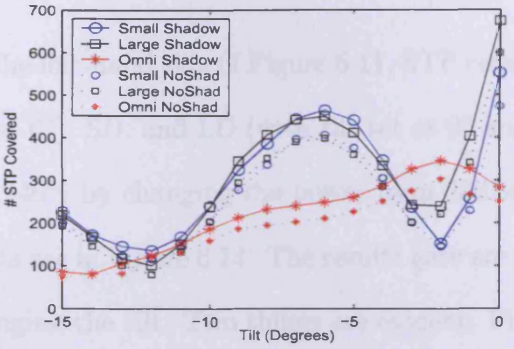

Figure 6.12: Cover by tilt for OD, SD, and LD antennas

ing an antenna to decrease coverage is 'hit-or-miss', as the effects are not monotonic. While, on average, the amount of coverage does go down with tilting, the amount given -1° increments cannot be guaranteed. In addition, as the cell shape can change quite remarkably as a result of tilting, the often finely balanced configurations in a network scenario must be considered at risk during this procedure. Given this scenario uses flat terrain, one could imagine the effects are even more remarkable on hilly terrain.

Figure 6.13: These diagrams show radiation patterns (using ASE, non-shadowing) when reducing SD tilt from 0° to -15° in -1° steps (top row) and reducing power from 55 to 26 in -2dBm steps (bottom row)

Effects of altering power

Continuing down the left-hand side of Figure 6.11, STP coverage is next measured for OD (with tilt set at 0°), SD, and LD (with tilt set at 0° and three antennas with azimuths of 0°, 120°, 240°) by changing the power from 26dBm to 55dBm in 1dBm increments. The results are in Figure 6.14. The results here are more predictable than those afforded by changing the tilt. Two things are evident: First, the earlier finding was confirmed that antennas in the shadowing environment provide more coverage. Second, as power rises, coverage increases, and as power falls, so does coverage.

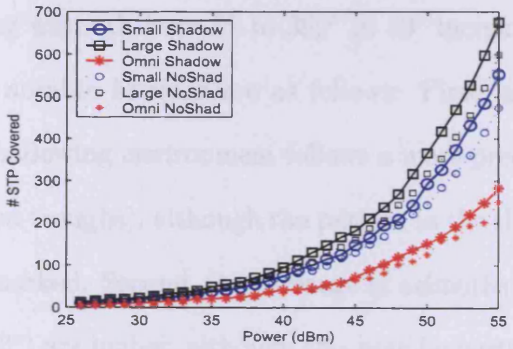

Figure 6.14: Cover by power for OD, SD, and LD antennas

Unlike the 'hit-or-miss' tilting scenario, lowering the power for a SD antenna results in a reliable decrease in STP coverage, as seen in Figure 6.13, while maintaining the original shape of the When considering possible planning implications, it seems that the reliability of lowering power should be preferred to tilting. This is a notable finding, as researchers in the past have not clarified configuration choices. For example, while [89] lowers power to repair overloaded cells, he does not justify why only

random changes to tilt are made. Other researchers (e.g., [40]) simply avoid using tilt during traffic repair. Thus, while lowering power is more straight-forward to use, tilting could be used as systematically if the size of cells based on changes to tilt were measured.

Effects of altering azimuth

Turning now to the right-hand side of Figure 6.11, one, two, and three antenna configurations are tested using both SD and LD antennas. The power is set to 55dBm and the tilt to 0° for all these tests. With one antenna, STP coverage is measured as a result of changing azimuth from 0° to 360° in 30° increments. The results are in Figure 6.15. The notable findings are as follows: First, as would be expected, coverage in the non-shadowing environment follows a more predictable pattern (with more shallow peaks and troughs), although the pattern in the shadowing environment is similar, only more marked. Second, the coverage at azimuths off the main axis (i.e., 0°, 90°, 180°, 270°, 360°) are higher, although this may be partially dependent on the scenario tested. Third, coverage for a LD antenna is considerably larger than that afforded by a SD antenna.

With two antennas commissioned at the site, coverage is measured with azimuth settings spaced evenly apart by 180° starting at 0°-180° and ending at 180°-360° in 15° increments. For example, the second test would be with azimuths at 15°-195°. The results are in Figure 6.16. The findings are less clear than those with one antenna. However, it appears that again coverage in the non-shadowing environment follows a more predictable pattern. Also, coverage for two LD antennas is considerably larger

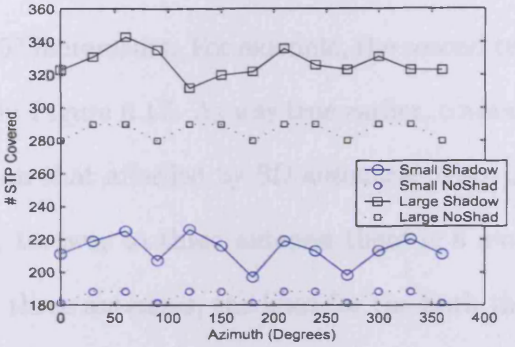

Figure 6.15: Cover by azimuth for one (SD and LD) antenna

than that afforded by two SD antennas. Third, if pressed one might suggest that for two SD antenna, the best settings are 45°-225° and 135°-315° and for two LD 30°-210°, 60°-240°, and 90°-270° as the number of STP covered is highest at these two points.

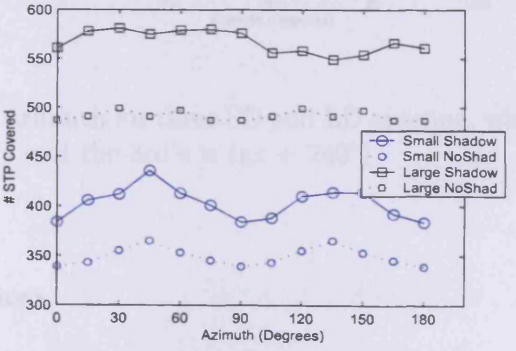

Figure 6.16: Cover by azimuth for two (SD and LD) antennas, where the 2nd antenna's azimuth is (az + 180°)

Finally, coverage using three antennas at the site is measured with azimuth set-

tings spaced evenly apart by 120°. Azimuth settings start at 0°-120°-240° and increase to 120°-240°-360° in 15° increments. For example, the second test would be 15°-135°-255°. The results are in Figure 6.17. As was true earlier, coverage for LD antennas is considerably larger than that afforded by SD antennas. Also, one notices that when progressing from one, to two, to three antenna there is a *smoothing* in the coverage differences. With three antennas, the lines for the both the non-shadowing and shadowing environments are nearly straight, which suggests that any evenly spaced azimuth settings here produce roughly the same coverage.

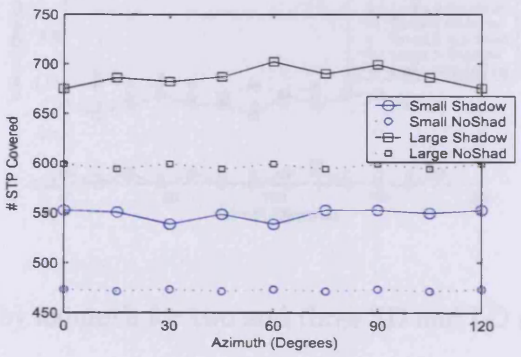

Figure 6.17: Cover by azimuth for three SD and LD antenna, where the 2nd antenna's azimuth is (az + 120°) and the 3rd's is (az + 240°)

Inherent coverage loss

During two and three antenna trials, the STP coverage was measured for each antenna separately and in total. This is possible as antenna at the same site can be trying to cover the same STP. However, in a best-server situation, only one can ultimately cover the STP. If more than one antenna *can* cover an STP, this raises

the interference in the network and makes resolving the channel assignment problem more difficult. This phenomenon is loosely termed 'overlap' here, and simply refers the percentage of STP that can be covered by more than one antenna. In terms of wireless network planning, while some overlap is desirable, as it allows call hand-over, it is preferable to limit it rather than encourage it during initial planning. The results are in Figure 6.18.

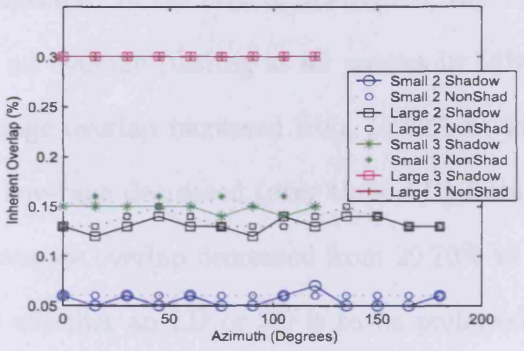

Figure 6.18: Overlap by azimuth for two and three SD and LD antenna combinations

There are two interesting features in Figure 6.18. First, the amount of overlap for LD antenna combinations is twice that of SD combinations. That is, with two antenna SD have roughly 6% overlap and LD 12% overlap, and with three antenna SD have 15% overlap and LD 30% overlap. For example, in one three LD antenna scenario, each antenna individually covered 323, 312, and 326 STP–for a total of 961. However, when taken together they actually covered 675 STP, indicating a coverage loss of 30%. Second, patterns and amounts of overlap are similar in both shadowing and non-shadowing environments.

Effects of antenna combinations

To complete the investigation of antenna configurations, particularly with three antenna at a site, further tests in a non-shadowing environment were run. In each case, the azimuths were set to 0, 120, and 240 and the tilt was set to 0. The purpose of these tests was to see if there was any benefit to mixing the combinations of SD and LD antennas. The two possible combinations were SD-SD-LD and LD-LD-SD. The results were as expected. In the case of SD-SD-LD, the STP coverage increased (over three SD) from an average (testing at all powers in 1dBm steps) of 126.13 to 137.33, while the average overlap increased from 15.60% to 21.03%. In the case of LD-LD-SD, the STP coverage decreased (over three LD) from an average of 159.70 to 152.33, while the average overlap decreased from 29.70% to 24.17%. As it is very difficult to be certain whether an LD or SD is to be preferred in a given situation, combinations of antennas seems appealing.

6.2.2 Antenna conclusions

Results from the antenna trials suggest best practices in terms of what changes to make when trying to increase or decrease STP coverage during antenna configuration. The main findings were that radiation patterns in a shadowing environment appear less symmetric, and therefore more real, than those afforded by a non-shadowing environment. When antenna coverage is measured, results consistently indicate that coverage is higher in a shadowing environment, but that otherwise, they produce similar trends in radiation and coverage patterns. Second, it was found that reducing power should be preferred over reducing tilt during antenna configuration, as reduc-

ing power results in a more reliable decrease in coverage while maintaining a similar radiation pattern. And finally, although LD antenna combinations produce more coverage than SD combinations–and both substantially more than OD, SD combinations result in less overlap, and concomitant interference. Considering site cost tends to greatly exceed antenna cost, placing three random combinations of antennas at each site appears a sensible course of action to maximize coverage and minimize cost. This is due to the fact that commissioning a site is more expensive than adding an antenna to a site. Therefore, as three antennas provides the maximum STP coverage it is preferable to using just one or two. Thus, this configuration is used during initial site configuration in the next chapter.

Chapter 7

3D Cell Planning Strategy

In this chapter, in order to satisfy multi-objective cell planning objectives in this more realistic simulation environment, we describe necessary modifications to the 2D decoder (as described in Chapter 4.1) and 2D cell planning strategy G2PS (as described in Chapter 3.2). Changes to each are necessary to account for the significant increase in simulation complexity and new objective measures (e.g., to meet traffic hold requirements).

7.1 Introduction to 3D Cell Planning

In the 2D model, we introduced a decoder that was able to translate a permutation π, which represented an ordering of CBS, into a cell plan. This approach mimics the way in which the problem might be attempted manually. The decoder was effectively a greedy, sequential algorithm which acted by limiting the permissible amount of overlap, or multi-coverage, that occurs between cells. While effective in the 2D model,

this decoder was not designed to meet cellular traffic load restrictions, and therefore would be ineffective in the 3D simulation. For example, if we used the 2D decoder in the 3D model, a cell could quickly acquire over 43 Erlangs of traffic and thus become *overloaded.* However, we ideally want each cell to handle only up to 43 Erlangs of traffic and no more. Thus, the new decoder needs to have a traffic constraint.

A second difficulty the 3D decoder faces is additional settings at each site. Whereas in the 2D model only one omni-directional antenna existed at a site and only the power of transmission could change, the 3D model needs to account for 1-3 antennas at a site, with settings for power, tilt, and azimuth. The number of potential combination at a site went from 10 in the 2D model:

$\times$ 1 (number of antennas)

$\times$ 10 (possible power settings)

to 162,000 in the 3D model:

$\times$ 3 (number of antennas)

$\times$ 30 (possible power settings)

$\times$ 15 (possible tilt settings)

$\times$ 120 (possible azimuth settings)

Therefore, realism in the 3D simulation means that significantly increased computational complexity needs to be handled efficiently.

Another problem which emerges in the 3D scenario as compared to the 2D scenario is that complete freedom in configuration of a particular site may lead to many cell plan solutions with low performance. This is due to traffic levels within cells. Approximate configurations can be addressed prior to site selection, on a heuristic

basis, to ensure that overloaded cells (i.e., cells covering more than 43 Erlangs of traffic) are avoided. It is also beneficial to commission sites with an estimation of their maximum area coverage under full load to minimize the number of cells deployed and avoid under-capacitated cells. To this end, we introduce a pre-processing stage called TSCA–T_{MAX} site configuration algorithm (Section 7.2.1). With this new pre-processing phase, the original G2PS (Chapter 3.2) is extended to a three phase approach *G3PS*.

7.2 Strategy for 3D Multi-objective Cell Planning

The intent of the strategy for 3D multi-objective cell planning is to find a range of high quality cell plans in terms of all network design objectives. By producing a range of near optimal cell plans the trade-offs between objectives can be assessed, and the need for weighting objectives common in many approaches is avoided. Figure 7.1 is a visual aid to the 3D multi-objective cell planning strategy, termed a generational three phase strategy ($G3PS$) which unlike G2PS (Chapter 3.2) starts with a pre-processing phase that employs a new $TSCA$ algorithm for controlling the permissible traffic per cell (Section 7.2.1), before continuing on to the second phase which determines cell plans, and the third phase which searches for optimal cell plans using a multi-objective algorithm (MOA).

In addition to including the pre-processing stage, we consider including an additional binary representation. This has been included because of the potential in the 3D model for irregular cells, which permit complex and irregular overlap scenarios which may potentially disrupt the overlap-dependent integer representation and

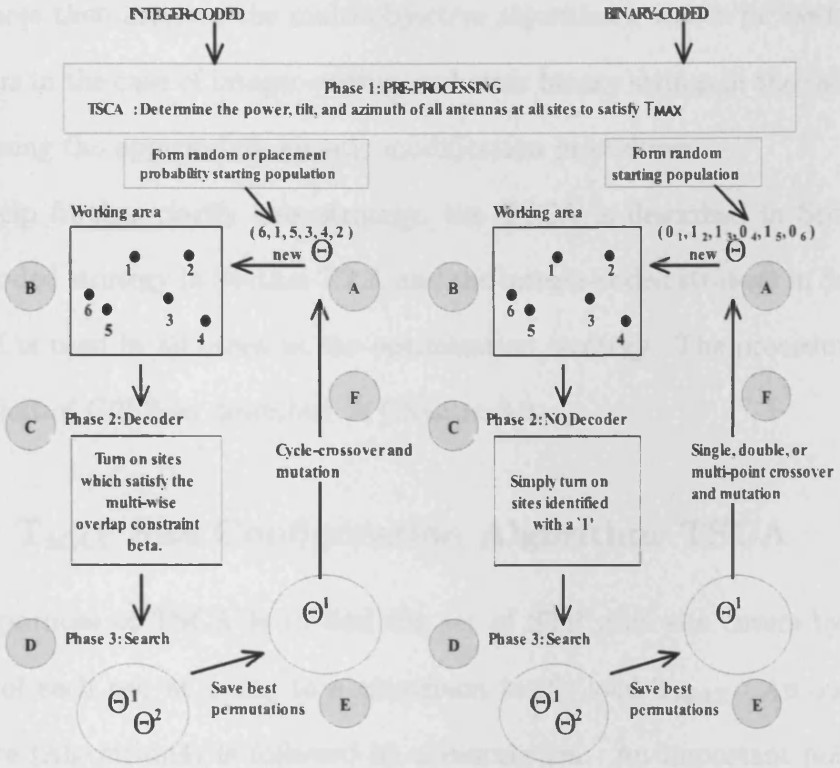

Figure 7.1: 3D multi-objective cell planning strategy

decoder. The potential benefit and difference between the representations will be evaluated in this chapter.

Looking at Figure 7.1, the process starts as either integer or binary-coded. The first pre-processing phase is the same for each and determines the power, tilt, and azimuth of all antennas at all sites to meet the traffic constraint T_{MAX}. A population of permutations is then formed. The integer-coded representation uses the same as that during G2PS, while the binary-coded uses a new procedure (Algorithm 5). In phase two the integer-permutation employs a decoder which limits the permissible

overlap between cells, and the binary-permutation simply turns on sites set to '1'. Phase three then involves the multi-objective algorithm's search for better orderings of integers in the case of integer-coding, or better binary strings in the case of binary-coding using the appropriate genetic modification procedures.

To help further clarify this strategy, the TSCA is described in Sect 7.2.1, the binary-coded strategy in Section 7.2.2, and the integer-coded strategy in Section 7.2.3. NSGA-II is used in all cases as the optimization strategy. The procedure generally follows that of G2PS as described in Chapter 3.2.

7.2.1 T_{MAX} Site Configuration Algorithm–TSCA

The purpose of TSCA is to find the set of STP each site covers by restricting the size of each cell at a site to a maximum traffic load T_{MAX}. An outline of the procedure (Algorithm4) is followed by a description. An important point to make is that this is the *only* time dimensioning is performed. Even more importantly, it is only done *once* per antenna. This differs distinctly from past approaches to this problem.

The process starts by selecting three random antenna (large directive, small directive, or some combination) at a given candidate site. This was chosen as the procedure based on antenna combination findings in Chapter 6.2, in the absence of foreknowledge regarding which type of antenna may be most suitable in a given scenario it is best to use a random combination of small and large directive antennas. The antennas tilts are then set to 0°, 120°, and 240° respectively, as this is a commonly selected starting position and conforms to findings in Chapter 6.2. Each antenna is then indi-

Algorithm 4 T_{MAX} Site Configuration

```
1: Set T_MAX to chosen value
2: for all CBS do
3:     Randomly set each of 3 antennas to LD or SD
4:     Set azimuths to 0°, 120°, and 240° respectively
5:     Set power to maximum (55 dBm)
6:     Set tilt to minimum (0°)
7:     for all Antenna do
8:         Set boolean 'saved configuration' = false
9:         while saved configuration is false do
10:            measure traffic load
11:            if traffic load ≤ T_MAX then
12:                save configuration = true
13:            else if power > 26 then
14:                reduce power by 1dBm
15:            else if power = 26dBm and tilt > -15° then
16:                reduce tilt by 1°
17:            else
18:                save configuration = true
19:            end if
20:        end while
21:    end for
22: end for
```

vidually set to the largest possible power (in the range 26-55 dBm) and largest tilt (in the range 0° to -15°) that does not exceed an alterable hard traffic constraint T_{MAX}. To do this:

- The maximum power is attempted first, followed by progressively lower powers– in 1 dBm steps. If a given power level does not exceed T_{MAX}, then the antenna is commissioned at that power with 0° tilt , and proceed to configure the next antenna.
- If the traffic load is still exceeded at the lowest power (i.e., 26 dBm), begin tilting the antenna starting at -1° and proceeding in -1° steps until the tilt is -15°. If a given tilt setting does not exceed T_{MAX}, then commission at 26 dBm and that tilt, and proceed to configure the next antenna.
- If traffic load is still exceeded, assign this antenna the lowest power (26 dBm) and greatest down-tilt (-15°), and proceed to configure the next antenna.

The set of STP covered by each antenna at each site it then saved–for use in the optimization process.

There are several important points to make about the approach:

1. Traffic hold can be guaranteed by setting T_{MAX} at 43 Erlangs.
2. Given perfect cell tessellation one would choose $T_{MAX} = 43$, and cost would be at a minimum.
3. Given real world tessellation one would choose $T_{MAX} > 43$ to keep cost at a minimum.

4. By guaranteeing traffic hold, other objectives can be used in the optimization process.

Traffic hold is used when determining the sets of STP each site handles, as it is the most difficult objective to satisfy. By ensuring traffic hold is met in this way, the genetic search process can focus on meeting two fundamental objectives: high service coverage and low cost. Service coverage is the focus during searches because if all STP are covered, all TTP are covered by definition, and, given that we can guarantee traffic hold, there are fewer dimensioning problems. We also attempt to keep the cost of the network as low as possible. This is achieved by setting T_{MAX} to where the traffic in each cell is as close as possible to 43 Erlangs.

7.2.2 Using TSCA with binary representation

As noted above, there are two permutation-coded strategies to use with TSCA. The simpler strategy is binary-coded. In this approach, each candidate base location is mirrored by the corresponding permutation index location. For example, candidate site number 42 is located at index 42. In this way, whether a site is on or off is determined by whether a 1 (on) or 0 (off) is at each index location. For example, the representation of a 40 candidate site string may appear as follows:

Binary coded: **1011000000110001010101000000100000010001**

To form a population of binary-permutations, the strategy outlined in Algorithm 5 is employed. It works as follows: After setting the size of the population, the

minimum number of sites to assign as '1', or 'on', in a given binary permutation is set to the total traffic divided by 129 (which is the maximum allowable traffic at a site given three antenna, each carrying 43 Erlangs). This is the minimum number of sites which *could* satisfy 100% traffic hold given perfect tessellation.

Then for each member of the population a chromosome, whose length is equal to the number of candidate sites, is initialized with '0's. For the first third of the population, the total number of sites turned 'on' is between 1 and the minimum number of sites. For two-thirds of the population, the total number of sites turned on is between the minimum and double the minimum. Then, randomly replace the appropriate number of 0's with 1's in the chromosome. This process makes a very diverse starting population highly probable, and should encourage very low coverage and cost networks to those with high coverage and cost.

7.2.3 Using TSCA with integer representation

The integer-permutation strategy is the same as that delineated in Chapter 3.3 and uses a decoder to translate permutation π into a cell plan by limiting the permissible amount of overlap that occurs between cells. The decoder chosen for use during 3D cell planning is the MCO-S decoder delineated in Chapter 4. The reason is that this is the only feasible decoder out of the six potential decoders analyzed in Chapter 4.1. This is due to the fact that TSCA determines all settings (e.g., power, tilt, azimuth, number of antennas), including those which use to be under the control of the decoder. For example, there is no by-power means for selecting the next candidate site, as the power has been pre-determined by TSCA. Thus, the decoder must select

Algorithm 5 Form a population of binary-permutations

1: Set size of population to chosen value
2: Set minAddSize as (total traffic / 129)
3: **for all** population member–i **do**
4: Initialize this chromosome (of length n_{CBS}) to 0s
5: **if** $i <$ size of pop / 3 **then** Set maxToTurnOn = Random.nextInt(minAddSize) + 1
6: **else** Set maxToTurnOn = minAddSize + Random.nextInt(minAddSize)
7: **end if**
8: Set min (maxToTurnOn, n_{CBS})
9: Set j = 0
10: **while** j < maxToTurnOn **do** Choose random index location
11: **if** index == 0 **then** set index = 1; j = j + 1
12: **end if**
13: **end while**
14: **end for**

the next candidate by-site. In addition, it is also significantly more convenient to use the multi-wise cell overlap constraint as these can be checked once 'on-the-fly' as opposed to the pair-wise cell overlap constraints.

Chapter 8

3D Parameter Tuning

In this chapter, the most pertinent parameter settings within the 3D simulation environment, decoder, and genetic algorithm are investigated. This is done in order to find the most suitable testing configuration to use in the final testing in Chapter 9. The two settings tested for in the simulation environment are the number of candidate base station locations (Section 8.2.1) and total amount of traffic (Section 8.2.2). The two settings tested in relation to the new 3D decoder are the maximum amount of permissible traffic per cell (Section8.3.1) and antenna configurations (Section 8.3.2). Following this, five settings for the genetic algorithm are tested. The first compares a placement probability (when forming the initial parent populations) to a random procedure (Section 8.4.1). The second notes differences when altering the population size (Section 8.4.2). The third estimates the ideal number of stagnation generations to allow before termination (Section 8.4.3). The fourth tests changes in the crossover procedure (Section 8.4.4), and the fifth changes to the representation scheme (Section 8.4.5).

An outline of the testing is provided in Figure 8.1. All tests are carried out using SE1, SE2, SE4, and SE5 (Table 6.1) with antenna heights of 30-60 meters. As results here can vary from test to test, all tests are carried out for 10 trials and measure the average and standard deviations for STP coverage, overlap, cost, traffic coverage, number of solutions, number of generations, and time, unless otherwise noted. It is expected that the results from these tests will allow a good estimate of the best parameters and settings to use when determining SE benchmark solutions in Chapter 9.

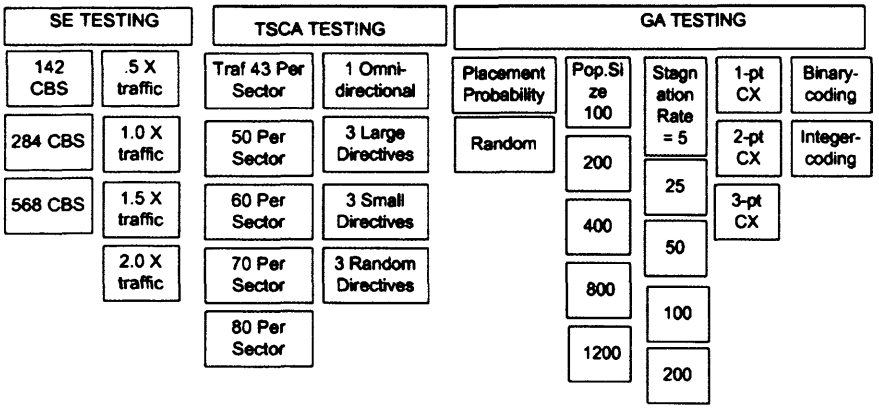

Figure 8.1: Outline of test plan for 3D simulation environment and multi-objective cell planning approach

It is important to note that traffic hold, interference, and handover are not measured here. The main reason is that they require intensive calculations, which would make such extensive parameter testing infeasible. However, the measurement of these objectives will be investigated in Chapter 9 during final comparisons. The objectives traffic hold and interference are therefore approximated here (when reporting final cell plan objective values) as follows:

- The approximation of traffic hold is termed *traffic coverage.* Rather than summing only traffic <= 43 Erlangs in each cell, the sum of all traffic covered (i.e., if a TTP is covered, it is summed) is simply divided by the total amount of traffic, expressed as a percentage.

- The approximation of interference is termed *overlap*, as the two measures are likely to be highly correlated given they involve the same exact STP. The only difference is that overlap is the count of these STP and interference is a measure of the interfering signals at these STP. Overlap is the sum of the number of STP each cell covers, divided by the total number of STP, expressed as a percentage.

Handover, however, cannot be measured without a server assumption and therefore no comparative measure is applied, although handover is later shown to improve along with coverage (Chapter 9). To clarify, during the optimization process, the 2D model used service coverage and cost. We continue to use service coverage and cost in the 3D model as well. However, for purposes of large-scale parameter testing, we approximate values for traffic hold (i.e., traffic coverage) and interference (i.e., overlap). This is done for efficiency for reporting final cell plan objective values. In Chapter 9, true traffic hold, interference, and handover values are reported during less extensive (although more intensive) testing.

By parameter tuning, we are looking for the best trade-off between time and solution quality, and whether a given setting is capable of fully testing the capabilities of the cell planning algorithm employed.

8.1 Standard testing configuration

The standard test configuration uses a population size of 500, run with a stagnation rate of 50 generations, with a T_{MAX} of 60, using multi-point crossover, placement probability, 568 CBS, and binary-coding. These test parameters were chosen based on provisional testing in order to produce reliable tables in a feasible amount of time. These default parameter settings are not optimized and are used, along with binary coding, to initially verify that the configuration of the simulation environment (Section 8.2.1 and 8.2.2) is reasonable and challenging, particularly in terms of number of CBS and traffic.

Perhaps most notable is the choice of binary-coded permutations over integer-coded permutations as used in the 2D model. The main motivation for selecting binary-coded permutations is that there are fewer parameters for configuration and it is considerably faster than the permutation-based approach. When a given variable was tested (e.g., changes made to the population size), all other variables adhered to the standard test configuration.

In order to accurately assess the effect of parameter changes, results (including the average and standard deviation) are measured at two points. The first set of results are based on the objective values of the first solution providing $\geq$ 95% STP coverage, which effectively makes STP coverage constant (at 95%) and allows one to note variations in the other objectives. The second set of results are based on the solution with the highest STP coverage, which allows one to note the ability of the chosen parameters to maximize coverage, and how this affects other objectives. Objectives used during testing and reported in tables throughout use the following

abbreviations: cover (cove.), overlap (over.), cost, traffic (traf.), number of generations (gens.), time, and number of solutions (sols.).

8.2 Simulation Environment Tuning/Testing

The first test on the simulation environment is used to indicate whether 568 candidate sites is necessary, or if using fewer sites is equally effective. This test is performed because it impacts greatly upon computational time and potential solution quality, e.g., in terms of obtainable service coverage. The second test is used to determine if the amount of traffic distributed creates a challenging environment for cell planning, or if more or less traffic is required. This test is performed because the amount of traffic in the service area significantly impacts network configurations. For example, low traffic results in larger cells, and high traffic results in smaller cells.

8.2.1 Changes in number of CBS

The first tests compared allowing 568, 284, and 142 candidate sites to afford a rough assessment of whether the full 568 sites is necessary or beneficial. Given our working area, this equates to 0.25, 0.13, and 0.06 candidate sites per km^2 respectively.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.2, 8.3, 8.4, 8.5) were compared. Interpreting the results (Table 8.1) indicates that using 568 CBS results in a 16.31% cost savings when compared to using 124 CBS, and an 11.00% cost savings versus 284 CBS. There was also a 17.67% savings in terms of overlap versus 124 CBS and 13.22% savings versus 284 CBS. Although this generally came at the expense of roughly 15% more GA

generations and consequently 22% more computational time.

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.7, 8.8, 8.9, 8.10) were compared. Interpreting the results (Table 8.6) indicates that using 568 CBS results in a 7.60% coverage improvement versus 124 and 2.36% versus 284, and similarly for traffic. For example, with 142 candidate sites, the average STP coverage was 85.08%, 87.53%, 92.80%, and 92.67% for SE1, 2, 4, and 5 respectively; with 284 candidate sites, the average STP coverage was 94.25%, 94.14%, 94.17%, and 95.58%; and, with 568 candidate sites, the average STP coverage was 97.05%, 96.59%, 97.48%, and 97.49%. However, this does come at an average cost increase of 19% and overlap increase of 29%.

Overall, the results indicate that using 568 candidate sites is preferable. Most importantly, using 568 sites results in lower cost solutions, with lower overlap (for solutions of equivalent coverage), and a greater number of solutions. Using 568 CBS also results in solutions with the highest STP and traffic coverage. The trade-off is only in terms of a modest 22% increase in run times. Therefore, 568 candidate sites will be preferred in future tests.

#CBS	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
142	0.08	-17.67	-16.31	-0.19	21.89	30.52	24.04
284	0.04	-13.22	-11.00	0.25	8.20	14.04	12.64

Table 8.1: Ave. % change in objective measures for 568 CBS based solutions vs. 142 and 284 across SE (for 1st solutions exceeding 95% STP coverage)

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
142	0.00	0.00	0.00	0.00	187.60	277.64	46.20
	0.00	0.00	0.00	0.00	80.99	96.42	4.42
284	95.13	116.11	52.00	93.93	232.00	354.35	51.90
	0.00	0.00	0.00	0.00	73.79	90.35	3.30
568	95.18	103.80	46.20	94.39	236.70	401.15	61.80
	0.07	9.09	2.23	0.30	110.68	159.51	5.49

Table 8.2: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with number of CBS changes

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
142	0.00	0.00	0.00	0.00	172.00	272.64	46.00
	0.00	0.00	0.00	0.00	109.80	137.13	1.90
284	95.02	155.70	54.67	88.88	211.50	364.51	55.50
	0.01	7.48	1.70	0.58	102.27	144.73	3.41
568	95.10	128.23	47.20	89.90	235.60	412.39	61.60
	0.05	10.85	2.27	0.82	85.10	125.31	7.28

Table 8.3: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with number of CBS changes

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
142	95.04	114.10	60.00	94.80	248.70	791.54	56.10
	0.00	0.00	0.00	0.00	134.66	385.28	5.11
284	95.10	112.63	56.50	94.95	253.50	855.15	62.00
	0.07	10.89	3.40	0.41	114.13	346.49	5.50
568	95.13	98.96	51.20	94.39	332.50	1200.63	73.70
	0.10	10.64	3.22	0.26	168.85	541.09	6.36

Table 8.4: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with number of CBS changes

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
142	95.04	118.95	59.00	95.04	255.50	813.23	57.30
	0.00	0.00	0.00	0.00	111.11	290.77	3.23
284	95.11	104.83	53.80	95.11	308.80	1017.20	66.80
	0.10	11.32	3.06	0.10	215.77	648.34	7.32
568	95.10	99.09	51.11	95.10	299.30	1060.96	73.20
	0.07	9.50	2.64	0.07	161.52	538.47	11.39

Table 8.5: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with number of CBS changes

#CBS	Cover	Overlap	Cost	Traffic
124	7.60	37.99	24.91	9.49
284	2.36	19.29	13.40	3.20

Table 8.6: Ave. % change in objective measures for 568 CBS based solutions vs. 142 and 284 across SE (for highest coverage solutions)

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
142	85.75	98.94	45.30	82.38
	1.43	16.21	4.65	1.67
284	93.64	122.58	51.00	92.14
	1.14	12.05	3.32	1.56
568	97.10	165.82	61.60	96.41
	0.78	21.53	5.90	0.93

Table 8.7: Ave. and SD for highest cover solution in SE1 given 10 trials with number of CBS changes

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
142	87.16	108.16	45.00	79.15
	1.30	7.91	1.90	2.68
284	94.46	159.50	54.80	88.25
	0.71	16.20	3.97	1.21
568	96.51	178.87	61.10	92.51
	0.95	22.67	7.18	1.87

Table 8.8: Ave. and SD for highest cover solution in SE2 given 10 trials with number of CBS changes

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
142	92.84	106.94	55.90	92.55
	1.51	14.29	5.73	1.72
284	95.19	128.40	61.40	94.82
	1.46	17.69	5.61	1.66
568	97.82	170.68	73.30	97.47
	1.07	22.86	6.54	1.25

Table 8.9: Ave. and SD for highest cover solution in SE4 given 10 trials with number of CBS changes

#CBS	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
142	93.81	112.06	56.30	93.81
	0.82	8.96	3.23	0.82
284	96.60	145.17	66.00	96.60
	1.07	22.64	7.27	1.07
568	97.63	171.79	73.20	97.63
	1.55	43.14	12.06	1.55

Table 8.10: Ave. and SD for highest cover solution in SE5 given 10 trials with number of CBS changes

8.2.2 Changes in amount of total traffic

The next variable to consider is the total traffic in the service area. Traffic is an important parameter as it can directly affect cell size. In order to test this parameter, the total amount of traffic was multiplied by 0.50, 1.00, 1.50, and 2.00 to test whether the original amount of traffic was suitable.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.12, 8.13, 8.14, 8.15) were compared. Interpreting the results (Table 8.11) indicates that using 1$\times$ traffic is a reasonable choice, as there is no stark argument for or against using an alternative. For example, in terms of overlap using 1$\times$ traffic results in 16.99% more overlap versus 0.5$\times$ and 4.62% less overlap versus 1.5$\times$ and 9.30% less than 2.0$\times$. This may occur because the number of cells needed is higher in higher traffic environments, which tends to result in more overlap. In terms of cost, using 1$\times$ traffic results in 44.21% higher cost versus 0.5$\times$ and 37.08% lower cost versus 1.5$\times$ and 83.92% less than 2.0$\times$. Again, as more cells are needed in higher traffic areas, the cost rises. Thus, as could be expected, 1$\times$ is a reasonable choice of traffic level. Adding more traffic results in a more difficult scenario, and less, an easier one. It is difficult to say, however, that one choice is definitely better than another. The choice is to some extent arbitrary.

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.17, 8.18, 8.19, 8.20) were compared. Interpreting the results (Table 8.16) indicates that using 1$\times$ traffic, similar to earlier conclusions, appears a reasonable choice, as it provides a 'middle ground' in terms of cover, overlap, cost, and traffic. Overall, the results indicate that there is no evidence that the

amount of traffic should be modified, unless one specifically desires a much harder or easier scenario.

#CBS	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
0.5	-0.13	16.99	44.21	-1.05	6.36	-13.30	9.45
1.5	0.01	-4.62	-37.08	0.31	1.78	13.62	-0.56
2.0	0.04	-9.30	-83.92	0.04	-12.17	10.27	-9.45

Table 8.11: Ave. % change in objective measures for 1× traffic based solutions vs. 0.5×, 1.5×, and 2.0×, across SE (for 1st solutions exceeding 95% STP coverage)

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
0.50	95.30	82.00	22.80	94.78	169.00	369.03	53.50
	0.13	8.07	0.87	0.26	50.64	87.02	2.25
1.00	95.18	103.80	46.20	94.39	236.70	401.15	61.80
	0.07	9.09	2.23	0.30	110.68	159.51	5.49
1.50	0.00	0.00	0.00	0.00	238.00	358.98	64.80
	0.00	0.00	0.00	0.00	75.56	92.12	4.81
2.00	0.00	0.00	0.00	0.00	286.20	385.49	69.30
	0.00	0.00	0.00	0.00	100.20	119.31	5.95

Table 8.12: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with amount of traffic changes

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
0.50	95.27	88.38	22.90	92.91	170.10	362.48	53.40
	0.15	5.46	0.94	0.39	34.12	57.02	3.61
1.00	95.10	128.23	47.20	89.90	235.60	412.39	61.60
	0.05	10.85	2.27	0.82	85.10	125.31	7.28
1.50	0.00	0.00	0.00	0.00	259.40	429.52	62.60
	0.00	0.00	0.00	0.00	100.63	147.21	5.37
2.00	0.00	0.00	0.00	0.00	233.00	366.14	65.00
	0.00	0.00	0.00	0.00	75.81	100.93	6.42

Table 8.13: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with amount of traffic changes

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
0.50	95.20	87.36	31.80	94.67	407.70	1704.42	74.10
	0.12	7.99	2.27	0.31	230.87	875.14	8.80
1.00	95.13	98.96	51.20	94.39	332.50	1200.63	73.70
	0.10	10.64	3.22	0.26	168.85	541.09	6.36
1.50	95.12	106.09	71.00	93.81	302.00	914.54	72.40
	0.06	7.90	2.16	0.59	196.98	525.39	8.92
2.00	0.00	0.00	0.00	0.00	386.20	1081.62	80.30
	0.00	0.00	0.00	0.00	173.60	459.84	12.51

Table 8.14: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with amount of traffic changes

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
0.50	95.23	94.97	32.30	95.23	324.30	1393.60	64.70
	0.12	7.97	2.00	0.12	199.65	812.00	11.56
1.00	95.10	99.09	51.11	95.10	299.30	1060.96	73.20
	0.07	9.50	2.64	0.07	161.52	538.47	11.39
1.50	95.10	101.11	69.25	95.10	273.60	803.30	71.40
	0.08	7.69	2.59	0.08	156.85	405.52	8.67
2.00	95.06	108.31	94.00	95.06	337.40	890.61	81.40
	0.05	0.00	0.00	0.05	140.58	348.55	11.62

Table 8.15: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with amount of traffic changes

Traf. ×	Cover	Overlap	Cost	Traffic
0.5	-2.29	-52.96	9.04	-3.61
1.5	5.32	33.85	0.09	8.02
2.0	9.89	46.70	-9.13	15.06

Table 8.16: Ave. % change in objective measures for 1× traffic based solutions vs. 0.5×, 1.5×, and 2.0×, across SE (for highest coverage solutions)

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
0.50	99.83	283.77	53.50	99.79
	0.07	16.25	2.62	0.11
1.00	97.10	165.82	61.60	96.41
	0.78	21.53	5.90	0.93
1.50	90.63	103.40	64.20	88.49
	1.41	13.28	4.60	1.87
2.00	84.35	73.95	68.50	80.17
	2.90	7.67	5.90	3.51

Table 8.17: Ave. and SD for highest cover solution in SE1 given 10 trials with amount of traffic changes

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
0.50	99.68 0.15	278.68 22.32	53.30 3.87	99.44 0.35
1.00	96.51 0.95	178.87 22.67	61.10 7.18	92.51 1.87
1.50	90.62 1.65	139.81 15.08	61.80 5.46	79.62 3.83
2.00	85.03 1.80	118.07 14.51	64.90 6.88	68.82 3.46

Table 8.18: Ave. and SD for highest cover solution in SE2 given 10 trials with amount of traffic changes

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
0.50	99.46 0.61	257.10 39.40	74.50 10.05	99.47 0.59
1.00	97.82 1.07	170.68 22.86	73.30 6.54	97.47 1.25
1.50	93.32 2.63	107.09 20.05	71.50 9.06	91.58 3.16
2.00	90.07 3.81	90.98 19.93	79.80 12.43	86.55 5.32

Table 8.19: Ave. and SD for highest cover solution in SE4 given 10 trials with amount of traffic changes

Traf. ×	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
0.50	98.97 0.78	230.67 53.27	64.50 12.34	98.97 0.78
1.00	97.63 1.55	171.79 43.14	73.20 12.06	97.63 1.55
1.50	93.82 2.77	105.34 22.20	70.80 8.73	93.82 2.77
2.00	91.19 3.71	84.65 19.87	80.70 11.48	91.19 3.71

Table 8.20: Ave. and SD for highest cover solution in SE5 given 10 trials with amount of traffic changes

8.3 Decoder Tuning and Testing

The first test on the decoder is used to indicate a good T_{MAX} setting. This test is performed because T_{MAX} impacts on the cell density, and therefore the overall network coverage and cost. The second test is used to determine a good antenna combination to use. This test is performed to help clarify whether one combination is generally to be preferred over another, as this also impacts overall coverage and cost.

8.3.1 Changes in max traffic

A variable to consider during TSCA is the allowable maximum traffic (T_{MAX}) per cell. T_{MAX}'s of 43, 50, 60, 70, and 80 were tested in order to determine the settings which provide the highest coverage solutions at the lowest cost. T_{MAX} is important as it determines the allowable traffic per sector during TSCA. Given perfect tessellation one would choose 43 *ipso facto*, as we have defined that as the maximum amount of traffic permitted in a cell. However, given cells do not perfectly tessellate an appropriate setting needs to be found which helps maximize coverage and minimize cost, while providing adequate traffic performance.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.22, 8.23, 8.24, 8.25) were compared. Interpreting the results (Table 8.21) indicate that as T_{MAX} rises and STP coverage is held constant (varying by no more than 0.12%), (1) overlap decreases linearly by 14.84%, (2) cost decreases linearly by 40.34%, and (3) traffic coverage as expected remains relatively constant (varying by no more than 0.89%).

Second, the average (from 10 trials) of solutions providing the highest STP cov-

erage across all SE (Tables 8.27, 8.28, 8.29, 8.30) were compared. Interpreting the results (Table 8.26) indicates that as T_{MAX} rises (1) coverage increases linearly by 5.32%, (2) overlap increases linearly by 61.15%, (3) cost decreases linearly by 9.26%, and (4) traffic coverage increases linearly by 8.56%.

Overall, the results indicate that using higher T_{MAX} settings is preferable in terms of STP coverage, traffic coverage, and cost. However, the trade-off for these benefits is more overlap (and ultimately interference) if the highest coverage solutions are sought. However, given solutions at the same coverage level, raising T_{MAX} results in lower overlap and cost. While there is no empirical test for preferring one over another, one can extrapolate that for every 10% rise in T_{MAX} above 43 there is roughly an additional 15% overlap for a 1.25% increase in coverage, a 2.3% decrease in cost, and 2.14% increase in traffic coverage. However, given solutions at the same coverage level, every 10% rise in T_{MAX} results in an overlap savings of 3.7% and a cost savings of 10%. This is perhaps what would be expected: As T_{MAX} increases the size of cells should also increase, which will have the effect of lowering cost (as you need fewer if covering an equivalent size area) and overlap (as fewer cells is probably positively and directly correlated to lower overlap). Different operators may have different requirements in terms of these objectives.

T_{MAX}	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
50	-0.04	0.83	10.12	-0.20	-4.05	-8.83	0.98
60	-0.07	4.05	21.92	-0.32	7.30	-3.73	3.98
70	-0.10	9.66	32.03	-0.61	8.78	-6.24	7.07
80	-0.12	14.84	40.34	-0.89	11.72	-7.20	8.87

Table 8.21: Ave. % change in objective measures for $T_{MAX} = 43$ based solutions vs. 50, 60, 70 and 80 across SE (for 1st solutions exceeding 95% STP coverage)

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
43	95.00	120.83	66.00	93.86	263.70	399.87	63.90
	0.00	0.00	0.00	0.00	134.31	185.14	9.86
50	95.05	121.56	60.00	93.79	248.30	395.04	62.20
	0.03	5.88	2.55	0.25	64.99	93.15	7.73
60	95.18	103.8	46.2	94.39	236.70	401.15	61.80
	0.07	9.09	2.23	0.30	110.68	159.51	5.49
70	95.20	96.84	39.30	94.39	211.80	378.86	58.40
	0.14	3.93	1.27	0.33	82.18	119.04	4.45
80	95.23	88.20	32.80	94.54	303.90	535.70	60.00
	0.15	2.73	0.75	0.44	106.80	175.77	5.33

Table 8.22: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with T_{MAX} changes

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
43	0.00	0.00	0.00	0.00	253.90	417.59	64.30
	0.00	0.00	0.00	0.00	90.39	126.35	3.85
50	95.05	124.08	54.60	89.83	287.90	497.59	65.10
	0.02	13.64	3.07	0.44	220.72	339.53	8.42
60	95.10	128.23	47.20	89.90	235.60	412.39	61.60
	0.05	10.85	2.27	0.82	85.10	125.31	7.28
70	95.16	112.48	39.40	90.72	203.30	392.91	61.00
	0.09	10.99	1.74	0.86	58.91	89.62	7.38
80	95.14	104.34	34.20	91.48	165.10	346.68	57.60
	0.08	6.96	1.33	0.93	56.31	93.29	4.10

Table 8.23: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with T_{MAX} changes

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
43	95.10 0.05	104.36 8.03	66.71 2.49	93.83 0.48	334.70 153.49	1061.28 461.74	79.00 12.39
50	95.13 0.09	100.00 9.36	59.14 3.23	94.42 0.29	358.90 146.84	1184.00 468.13	74.80 11.93
60	95.13 0.10	98.96 10.64	51.20 3.22	94.39 0.26	332.5 168.85	1200.63 541.09	73.70 6.36
70	95.14 0.07	96.30 9.65	46.10 2.84	94.61 0.34	370.60 193.15	1335.90 685.57	72.40 12.49
80	95.18 0.12	91.72 6.45	41.30 1.68	94.69 0.35	312.70 87.19	1174.66 297.27	70.70 5.06

Table 8.24: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with T_{MAX} changes

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
43	95.08 0.03	99.36 6.65	64.60 2.42	95.08 0.03	336.60 128.25	1033.00 376.55	74.50 7.75
50	95.12 0.08	100.44 5.35	58.20 1.72	95.12 0.08	341.40 111.56	1093.06 325.74	76.60 4.48
60	95.10 0.07	99.09 9.50	51.11 2.64	95.10 0.07	299.30 161.52	1060.96 538.47	73.20 11.39
70	95.12 0.09	97.66 7.97	45.90 2.26	95.12 0.09	315.60 108.21	1138.91 390.63	69.90 9.09
80	95.15 0.11	95.07 6.68	41.60 1.91	95.15 0.11	267.40 112.66	1044.77 410.56	68.20 9.10

Table 8.25: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with T_{MAX} changes

T_{MAX}	Cover	Overlap	Cost	Traffic
50	-2.25	-15.44	1.13	-3.90
60	-3.79	-34.60	4.02	-6.05
70	-4.57	-48.12	7.34	-7.38
80	-5.32	-61.15	9.26	-8.56

Table 8.26: Ave. % change in objective measures for $T_{MAX} = 43$ based solutions vs. 50, 60, 70 and 80 across SE (for highest coverage solutions)

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
43	91.59 2.38	109.15 26.16	63.70 11.04	89.35 3.03
50	94.45 1.33	125.83 23.88	61.50 8.04	93.09 1.50
60	97.10 0.78	165.82 21.53	61.60 5.90	96.41 0.93
70	97.94 0.74	177.77 20.05	57.70 4.45	97.55 0.91
80	99.21 0.29	215.26 33.37	59.60 6.15	99.03 0.40

Table 8.27: Ave. and SD for highest cover solution in SE1 given 10 trials with T_{MAX} changes

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
43	92.58 1.01	143.12 10.54	63.80 3.79	83.85 2.80
50	95.36 1.56	167.67 24.21	65.30 10.08	90.31 3.16
60	96.51 0.95	178.87 22.67	61.10 7.18	92.51 1.87
70	97.71 0.79	206.86 28.69	60.60 7.79	94.83 1.91
80	98.45 0.39	212.63 14.69	56.70 4.05	96.81 0.73

Table 8.28: Ave. and SD for highest cover solution in SE2 given 10 trials with T_{MAX} changes

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
43	95.30 3.20	137.42 30.10	79.00 12.66	93.97 4.21
50	96.22 2.25	146.60 34.77	74.50 12.21	95.65 2.81
60	97.82 1.07	170.68 22.86	73.30 6.54	97.47 1.25
70	98.14 1.26	187.65 42.63	72.00 13.19	98.11 1.31
80	98.70 0.50	195.43 21.55	70.20 5.36	98.57 0.64

Table 8.29: Ave. and SD for highest cover solution in SE4 given 10 trials with T_{MAX} changes

T_{MAX}	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
43	95.49 2.03	125.12 21.66	74.20 8.16	95.49 2.03
50	97.33 0.60	153.45 14.52	75.90 4.57	97.33 0.60
60	97.63 1.55	171.79 43.14	73.20 12.06	97.63 1.55
70	98.20 0.88	185.84 34.50	69.70 9.45	98.20 0.88
80	98.41 0.86	195.94 32.28	68.00 9.39	98.41 0.86

Table 8.30: Ave. and SD for highest cover solution in SE5 given 10 trials with T_{MAX} changes

8.3.2 Changes in antenna configurations

The default antenna configuration during TSCA is three randomly selected directive antenna at each CBS. Three alternative combinations were tested. These were: one omni-directional antenna, three small directive, and three large directive. Results should indicate which antenna configuration to prefer.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.32, 8.33, 8.34, 8.35) were compared. Interpreting the results (Table 8.31) indicates that that as the antenna configuration is changed and coverage is held constant three large directives result in 0.45% increase in overlap and a 6.91% increase in cost for a 15.74% time savings versus three random directives. On the other hand, three small directives result in 2.90% increase in overlap and a 3.86% decrease in cost for a 14.08% time savings versus three random directives. Omnidirectional antenna could not reach 95% cover, and therefore could not be considered a viable option on their own.

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.37, 8.38, 8.39, 8.40) were compared. Interpreting the results (Table 8.36) indicates that when the highest coverage solutions are sought three random, large, and small directives perform similarly. That is, there is little to choose between them. Omni-directional antenna did not perform well here either, resulting in a significant decrease in STP and traffic coverage. In the absence of a strong case for using three large or small directives exclusively, the safest option appears to be three random directive antennas.

Antenna	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
3 Large	-0.02	-0.45	-6.91	0.19	2.08	15.74	0.80
3 Small	-0.04	-2.90	3.86	-0.26	-0.01	14.08	1.87
1 Omni	NA	NA	NA	NA	-1.92	24.52	-6.93

Table 8.31: Ave. % change in objective measures for 3 random directive antenna based solutions vs. 3 large, 3 small, and 1 omni across SE (for 1st solutions exceeding 95% STP coverage)

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Random	95.15	99.12	44.70	94.21	342.60	746.40	69.90
	0.07	7.92	1.62	0.21	113.66	225.17	5.34
Large	95.17	103.33	50.40	94.03	262.90	570.96	65.60
	0.07	10.59	2.62	0.26	114.09	223.20	7.23
Small	95.17	100.25	42.40	94.23	225.50	507.92	61.40
	0.09	7.17	1.43	0.27	65.36	120.92	6.00
Omni	0.00	0.00	0.00	0.00	279.70	487.76	70.50
	0.00	0.00	0.00	0.00	76.94	113.15	5.04

Table 8.32: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with antenna configuration changes

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Random	95.14	116.46	45.50	90.17	291.40	683.54	65.80
	0.05	7.74	1.80	0.46	152.21	320.81	5.91
Large	95.14	118.68	49.20	90.00	322.90	752.49	69.40
	0.08	14.44	2.27	0.69	125.11	265.17	6.86
Small	95.17	114.91	42.70	90.81	277.00	653.17	67.00
	0.07	5.94	1.10	0.58	93.18	189.06	4.07
Omni	0.00	0.00	0.00	0.00	336.40	638.01	71.10
	0.00	0.00	0.00	0.00	126.64	217.55	8.51

Table 8.33: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with antenna configuration changes

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Random	95.13	97.40	51.00	94.39	333.60	1540.92	74.40
	0.07	6.89	1.79	0.30	174.22	748.19	9.41
Large	95.14	95.84	52.89	94.01	288.60	1002.41	73.60
	0.07	8.60	2.38	0.29	160.78	525.07	11.82
Small	95.16	100.67	49.30	94.58	391.50	1373.25	76.90
	0.08	10.39	3.32	0.40	222.48	746.07	10.64
Omni	0.00	0.00	0.00	0.00	306.80	1003.99	82.80
	0.00	0.00	0.00	0.00	200.04	612.84	13.50

Table 8.34: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with antenna configuration changes

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Random	95.06	93.47	49.40	95.06	314.80	1454.62	77.10
	0.05	5.93	1.56	0.05	144.05	606.96	8.79
Large	95.10	90.91	50.90	95.10	370.30	1242.30	76.00
	0.07	8.13	2.34	0.07	194.74	662.21	14.46
Small	95.15	101.37	49.00	95.15	383.50	1323.22	76.70
	0.08	6.36	1.84	0.08	143.11	476.10	11.10
Omni	0.00	0.00	0.00	0.00	373.40	1135.93	82.90
	0.00	0.00	0.00	0.00	192.56	577.61	13.11

Table 8.35: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with antenna configuration changes

Antenna	Cover	Overlap	Cost	Traffic
3 Large	0.64	8.10	1.11	0.90
3 Small	-0.04	-3.18	2.49	-0.29
1 Omni	12.61	60.65	-6.70	17.31

Table 8.36: Ave. % change in objective measures for 3 random directive antenna based solutions vs. 3 large, 3 small, and 1 omni across SE (for highest coverage solutions)

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Random	98.05 0.58	189.40 19.45	69.90 6.07	97.48 0.83
Large	96.66 0.82	151.74 24.10	65.10 7.26	95.83 1.04
Small	97.55 0.82	172.06 21.22	60.50 6.18	96.89 1.05
Omni	82.66 2.16	59.80 8.86	70.10 5.73	78.44 2.89

Table 8.37: Ave. and SD for highest cover solution in SE1 given 10 trials with antenna configuration changes

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Random	97.50 0.54	197.57 21.19	65.40 6.36	94.52 1.27
Large	97.26 0.75	204.48 25.42	68.90 6.96	93.72 1.71
Small	98.09 0.27	210.16 14.26	66.10 3.99	95.83 0.67
Omni	85.13 2.95	101.84 13.94	70.50 8.59	71.25 5.64

Table 8.38: Ave. and SD for highest cover solution in SE2 given 10 trials with antenna configuration changes

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Random	97.77 1.16	171.38 29.70	74.30 10.31	97.34 1.51
Large	97.35 1.46	156.80 34.07	72.70 11.75	96.75 1.84
Small	98.07 1.28	188.97 34.05	76.20 10.39	97.93 1.27
Omni	85.74 5.08	67.36 16.89	82.00 13.70	82.42 6.11

Table 8.39: Ave. and SD for highest cover solution in SE4 given 10 trials with antenna configuration changes

Antenna	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Random	98.32 1.05	185.60 31.88	76.60 9.21	98.32 1.05
Large	97.86 1.49	171.68 47.05	76.00 14.93	97.86 1.49
Small	98.09 1.22	195.35 41.07	76.50 12.04	98.09 1.22
Omni	88.73 4.20	64.91 23.51	83.00 15.03	88.73 4.20

Table 8.40: Ave. and SD for highest cover solution in SE5 given 10 trials with antenna configuration changes

8.4 Genetic Algorithm Tuning and Testing

Five parameter tests were run for the GA. These investigated changes to the starting population, the population size, the stagnation rate, crossover methods, and representation scheme. These variables were considered important in terms of trade-offs between solution quality and run time. For example, while one might prefer to use a population size of 1000 to encourage a high quality solution, this might not be a feasible setting due to extremely long run times. Thus, the tests performed should give an indication of what results in good solution quality in a feasible amount of time. The approach uses binary representation in the GA (due to run time) until we lastly observe the effects of switching to integer representation.

8.4.1 Changes in starting population formation

An important variable to consider during binary-coding is whether a random starting population or one using a placement probability (PP) that uses information regarding the site provides the best trade-off of speed versus solution quality. To compute the PP several additional measures are taken during TSCA, and then applied when forming the starting population for the genetic algorithm. These measures are taken to keep costs low by encouraging the placement of sites (a) that cover more STP than others, (b) that carry a higher traffic load than others, and (c) which minimize overlap.

To calculate the placement probability PP_i for a given site, the following formulae generates a real number $\in [0, 1]$. During the formation of parents, the PP_i is compared to a randomly generated float value. If the float value is less than the PP_i the site is

added, else not. The formulae for determining the PP_i is:

$$PP_i = (estSTP_i * TCF_i);$$

where

$$estSTP_i = (numSTP_i / highest_{STP})/(SCF_i + 1);$$

To encourage sites which cover more traffic, a traffic correction factor (TCF_i) is applied to each site i as follows: 0.0 if traffic at the site is < 10 Erlangs, 0.33 if < 43, 0.66 if < 86, and 0.99 if ≥ 86. Note that in terms of an entire site with three directive antenna, 129 Erlangs would be the optimum and maximum permissible. To encourage sites which cover more STP, $estSTP$ is a calculation which first converts the number of STP covered to a real number between 0 and 1 by dividing the number of STP covered by each site, $numSTP_i$, by the highest number of STP covered by any site, $highest_{STP}$, and then divides it by the number of sites in its signal reception area, SCF_i, $+1$ to avoid division by 0 and to account for itself. The site correction factor (SCF_i) around a given CBS is found by summing the number of CBS which share a common covered STP. In general, sites surrounded by a number of other sites do not need to be commissioned with as high a probability.

As the fewest possible sites needed can be computed as total traffic divided by 129 (which in this instance is $2988.27/129 = 23$, attempt to place a '1' a random number of times between this 1 and this number for 33% of the population and this number and twice this number for 67% of the population in order to obtain a diverse

population which will now naturally vary from very low coverage and cost to high coverage and cost.

Trials were then run to compare the PP procedure to a random one which follows the same procedure as PP except it simply assigns the appropriate number of 1s to randomly determined base stations. The findings were as follows:

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.42, 8.43, 8.44, 8.45) were compared. Interpreting the results (Table 8.41) indicates that when coverage is kept constant using random starting populations raises overlap by 2.85% and cost by 3.86%, with little difference in generations, run time, or number of solutions. Taken on its own, this indicates using PP should be preferred.

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.47, 8.48, 8.49, 8.50) were compared. Interpreting the results (Table 8.46) indicates little difference between the two procedures in terms of final solution quality. That is, there is an expected trade-off that while random has slightly lower coverage and traffic coverage, it also has slightly lower cost and overlap.

Overall, the results indicate that while either PP or random starting populations can be used, there is some justification in choosing PP over random if low overlap and cost is sought at a given level of coverage less than the maximum. However, if only maximum is cover is sought, there is some justification is selecting random base stations. Theoretically, the PP method should encourage use of the best sites, but the random method has the advantage of having more diverse starting populations, as it does not favour the use of one site over another.

Formation	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
random	0.04	-2.85	-3.86	-0.11	1.27	2.16	1.46

Table 8.41: Ave. % change in objective measures for PP based solutions vs. random across SE (for 1st solutions exceeding 95% STP coverage)

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
PP	95.18	103.80	46.20	94.39	236.70	401.15	61.80
	0.07	9.09	2.23	0.30	110.68	159.51	5.49
Random	95.08	112.71	50.38	94.19	227.70	372.64	59.50
	0.05	14.42	4.92	0.17	103.32	149.97	6.41

Table 8.42: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with initial population formation changes

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
PP	95.10	128.23	47.20	89.90	235.60	412.39	61.60
	0.05	10.85	2.27	0.82	85.10	125.31	7.28
Random	95.09	118.92	46.63	90.55	331.60	563.93	65.40
	0.04	10.61	2.34	0.61	217.05	349.39	13.17

Table 8.43: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with initial population formation changes

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
PP	95.13	98.96	51.20	94.39	332.50	1200.63	73.70
	0.10	10.64	3.22	0.26	168.85	541.09	6.36
Random	95.10	107.09	54.60	94.34	182.60	693.81	67.80
	0.05	10.18	2.65	0.31	54.02	163.90	3.71

Table 8.44: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with initial population formation changes

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
PP	95.10	99.09	51.11	95.10	299.30	1060.96	73.20
	0.07	9.50	2.64	0.07	161.52	538.47	11.39
Random	95.09	100.95	51.60	95.09	308.40	1102.82	73.00
	0.07	8.29	2.50	0.07	163.71	546.31	9.17

Table 8.45: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with initial population formation changes

Formation	Cover	Overlap	Cost	Traffic
random	0.72	4.24	1.08	0.81

Table 8.46: Ave. % change in objective measures for PP based solutions vs. random across SE (for highest coverage solutions)

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
PP	97.10	165.82	61.60	96.41
	0.78	21.53	5.90	0.93
Random	95.66	140.85	58.80	94.67
	1.73	23.98	6.68	2.18

Table 8.47: Ave. and SD for highest cover solution in SE1 given 10 trials with initial population formation changes

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
PP	96.51	178.87	61.10	92.51
	0.95	22.67	7.18	1.87
Random	96.47	195.38	66.00	92.75
	1.80	55.31	14.34	3.52

Table 8.48: Ave. and SD for highest cover solution in SE2 given 10 trials with initial population formation changes

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
PP	97.82 1.07	170.68 22.86	73.30 6.54	97.47 1.25
Random	96.59 0.52	150.79 15.84	68.40 4.13	95.92 0.74

Table 8.49: Ave. and SD for highest cover solution in SE4 given 10 trials with initial population formation changes

Method	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
PP	97.63 1.55	171.79 43.14	73.20 12.06	97.63 1.55
Random	97.53 1.25	172.68 32.96	72.40 8.92	97.53 1.25

Table 8.50: Ave. and SD for highest cover solution in SE5 given 10 trials with initial population formation changes

8.4.2 Changes in population size

The next variable to consider is the population size. The population size is an important parameter for the genetic algorithm, as we do not know *a priori* the appropriate number of individuals to employ for searching the objective space. A variety of sizes 100, 200, 400, 800, and 1200 are tested. For each SE 10 trials are run at each population size, resulting in 200 tests.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.52, 8.53, 8.54, 8.55) were compared. Interpreting the results (Table 8.51) indicates that as the population size rises and coverage is held constant overlap decreases by up to 22.97% and cost decreases by up to 18.36%. However, the trade-off in terms of the number of generations and time required is steep. For example, the number of generations run increases by up to 177.40% and time increases up to a staggering 5041.61%. This suggests that larger population sizes lead to better quality solutions (i.e., less cost and overlap), but at the cost of longer run times. Considering the benefits of population size appear to start levelling off after 800 generations, 800 would be the maximum recommended population size, with sizes at and above 200 producing good benefits as well.

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.57, 8.58, 8.59, 8.60) were compared. Interpreting the results (Table 8.56) indicates that when searching for the highest coverage solutions, the normal pattern follows, i.e., increases in coverage come at higher cost and overlap. Similar to above, the benefits do appear to start levelling off around 800, and, again, there are good arguments for preferring even lower settings (e.g., 400), as the

coverage, cost, and overlap trade-offs appear desirable. Overall, the results indicate that a population size between 400-800 provides a respectable trade-off in terms of coverage and cost.

Pop-Size	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
200	-0.05	9.54	9.77	-0.10	-50.87	-203.24	-16.94
400	-0.02	12.08	11.46	0.01	-88.32	-750.59	-32.16
800	-0.02	21.49	16.63	-0.24	-137.30	-2454.26	-50.67
1200	-0.07	22.97	18.36	-0.22	-177.40	-5041.61	-63.36

Table 8.51: Ave. % change in objective measures for population 100 based solutions vs. 200, 400, 800 and 1200 across SE (for 1st solutions exceeding 95% STP coverage)

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
100	95.05	127.68	55.50	93.96	135.80	49.82	49.60
	0.03	6.47	4.50	0.01	65.16	18.36	5.48
200	95.16	112.53	49.14	94.36	218.90	158.94	57.40
	0.05	8.30	2.64	0.32	69.99	42.28	4.36
400	95.18	103.80	46.20	94.39	236.70	401.15	61.80
	0.07	9.09	2.23	0.30	110.68	159.51	5.49
800	95.12	92.31	43.70	94.38	312.90	1280.71	71.20
	0.07	7.47	2.05	0.25	116.87	429.57	5.33
1200	95.16	93.79	43.60	94.37	276.70	2062.26	72.00
	0.11	4.56	0.80	0.17	104.92	701.10	6.03

Table 8.52: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with population size changes

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
100	0.00	0.00	0.00	0.00	129.60	51.89	48.10
	0.00	0.00	0.00	0.00	69.96	20.21	3.51
200	95.14	132.39	48.50	90.40	165.00	136.14	50.80
	0.02	5.13	1.50	0.06	82.13	57.04	5.36
400	95.10	128.23	47.20	89.90	235.60	412.39	61.60
	0.05	10.85	2.27	0.82	85.10	125.31	7.28
800	95.11	105.38	43.50	90.81	329.10	1441.79	69.70
	0.08	11.78	1.80	0.96	112.43	457.72	6.99
1200	95.14	103.43	42.10	90.48	390.40	2967.29	78.30
	0.08	5.65	0.83	0.55	154.97	1113.76	10.79

Table 8.53: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with population size changes

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
100	95.07	119.61	61.00	94.27	173.50	139.28	54.40
	0.05	18.09	6.00	0.39	99.59	64.25	5.43
200	95.15	103.99	54.40	94.20	212.20	358.79	64.30
	0.10	9.23	2.06	0.52	103.72	156.27	11.92
400	95.13	98.96	51.20	94.39	332.50	1200.63	73.70
	0.10	10.64	3.22	0.26	168.85	541.09	6.36
800	95.12	92.83	49.00	94.34	359.10	3008.92	82.70
	0.07	5.14	1.79	0.28	153.29	1275.03	10.60
1200	95.22	87.84	47.60	94.59	418.30	5793.45	87.30
	0.08	7.30	1.50	0.31	166.70	2277.04	11.65

Table 8.54: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with population size changes

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
100	95.16	109.13	57.00	95.16	145.60	112.82	52.10
	0.00	0.00	0.00	0.00	70.47	42.35	4.21
200	95.11	105.11	53.00	95.11	280.50	421.93	66.80
	0.06	8.52	3.23	0.06	139.93	190.29	6.72
400	95.10	99.09	51.11	95.10	299.30	1060.96	73.20
	0.07	9.50	2.64	0.07	161.52	538.47	11.39
800	95.13	92.24	48.30	95.13	375.50	3054.57	84.50
	0.08	4.84	1.79	0.08	167.63	1347.19	9.77
1200	95.17	90.69	47.40	95.17	529.30	7388.61	96.40
	0.11	6.50	2.24	0.11	282.53	3759.45	11.89

Table 8.55: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with population size changes

Pop-Size	Cover	Overlap	Cost	Traffic
200	-3.24	-25.31	-14.72	-3.85
400	-5.26	-51.93	-30.29	-6.89
800	-6.57	-81.12	-48.33	-8.87
1200	-7.05	-103.17	-61.08	-9.62

Table 8.56: Ave. % change in objective measures for population 100 based solutions vs. 200, 400, 800 and 1200 across SE (for highest coverage solutions)

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
100	92.42	110.49	49.80	90.67
	2.33	18.99	6.08	2.78
200	95.84	135.94	56.90	94.96
	1.11	15.46	4.57	1.44
400	97.10	165.82	61.60	96.41
	0.78	21.53	5.90	0.93
800	98.22	193.71	71.20	97.75
	0.78	15.55	5.44	1.05
1200	98.46	195.41	71.60	98.08
	0.51	21.74	6.36	0.57

Table 8.57: Ave. and SD for highest cover solution in SE1 given 10 trials with population size changes

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
100	92.16 1.53	129.46 18.28	48.00 4.49	84.79 2.37
200	94.02 1.28	142.21 24.04	50.00 5.29	87.24 2.57
400	96.51 0.95	178.87 22.67	61.10 7.18	92.51 1.87
800	98.15 0.77	210.92 19.78	68.70 6.99	95.83 1.78
1200	98.73 0.73	251.61 37.85	77.70 10.94	97.09 1.42

Table 8.58: Ave. and SD for highest cover solution in SE2 given 10 trials with population size changes

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
100	92.46 2.72	110.38 18.43	55.50 7.09	91.37 3.26
200	94.95 2.78	135.41 34.93	63.50 11.99	94.31 3.00
400	97.82 1.07	170.68 22.86	73.30 6.54	97.47 1.25
800	98.68 0.83	196.89 32.45	81.70 10.60	98.59 0.81
1200	99.02 0.96	215.40 36.92	86.80 11.61	99.00 0.99

Table 8.59: Ave. and SD for highest cover solution in SE4 given 10 trials with population size changes

Pop. Size	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
100	92.57 2.04	104.20 18.20	53.00 5.69	92.57 2.04
200	96.79 1.20	151.78 21.66	66.80 5.91	96.79 1.20
400	97.63 1.55	171.79 43.14	73.20 12.06	97.63 1.55
800	98.83 0.76	216.61 36.05	84.80 10.27	98.83 0.76
1200	99.45 0.48	256.68 40.87	96.60 12.79	99.45 0.48

Table 8.60: Ave. and SD for highest cover solution in SE5 given 10 trials with population size changes

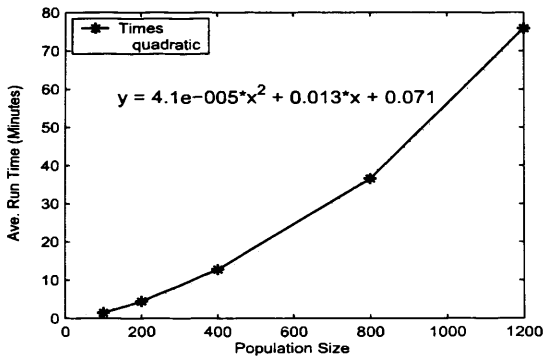

Table 8.61: Ave. run times (for a single trial using SE1) as a result of changing the population size along with a quadratic interpolation

8.4.3 Changes in stagnation rate

The next parameter to consider is the number of allowable generations (when no changes are occurring in the best objective values) before termination. Rates of 5, 25, 50, 100, and 200 generations are tested in order to determine the rate with the best trade-off of speed versus solution quality. One would like to select the lowest rate possible that does not sacrifice significantly on solution quality.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.63, 8.64, 8.65, 8.66) were compared. Interpreting the results (Table 8.62) indicates that as the stagnation rate rises and coverage is held constant overlap decreases linearly starting at 50 by 22.02% and cost decreases linearly starting at 50 by 14.60%. This suggests that large stagnation rates lead to better quality solutions (i.e., less cost and overlap), but at the cost of longer run times. Considering the trade-off in terms of the number of generations run and execution time, a stagnation rate of 50 is recommended. However, if execution time were not a factor, then a setting of 200 or higher should be employed.

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.68, 8.69, 8.70, 8.71) were compared. Interpreting the results (Table 8.67) indicates that as the stagnation rate rises and the highest possible coverage is sought STP and traffic coverage rises linearly (by 9.10% and 11.38% respectively), and overlap and cost increase linearly (by 119.85% and 71.28% respectively).

Overall, when considering all factors, a choice of 50 or 100 seems practical when seeking the highest coverage solution, while trying to minimize cost and overlap. The

general trend is that larger stagnation rates not only find higher maximum STP and traffic coverage, but given a certain level of coverage (e.g., 95%) the cost and overlap is lower than those found with smaller stagnation rates. The only trade-off is run times. However, when seeking the highest coverage solutions, increases in cost and overlap mean a lower stagnation rate (i.e., 50 or 100) is a more judicious choice.

Rate	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
25	NA	NA	NA	NA	-641.90	-603.33	-39.16
50	0.02	13.50	9.39	-0.30	-1461.62	-1436.52	-49.94
100	0.04	17.75	11.91	-0.18	-2878.53	-2931.33	-63.06
200	0.02	22.02	14.60	-0.24	-7982.95	-8659.31	-101.30

Table 8.62: Ave. % change in objective measures for stagnation rate of 5 based solutions vs. 25, 50, 100 and 200 across SE (for 1st solutions exceeding 95% STP coverage)

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
5	0.00	0.00	0.00	0.00	18.80	32.88	43.30
	0.00	0.00	0.00	0.00	10.81	14.28	4.08
25	95.22	114.91	48.75	94.20	85.40	150.49	55.80
	0.15	11.57	3.27	0.29	37.45	52.00	3.06
50	95.20	106.17	46.89	94.38	207.3	354.66	60.10
	0.07	9.410	2.38	0.16	88.20	126.41	7.20
100	95.14	99.01	45.30	94.33	488.70	832.30	64.20
	0.08	8.42	1.79	0.23	221.43	327.69	4.89
200	95.16	94.91	43.60	94.40	873.60	1579.33	72.50
	0.07	6.86	1.62	0.27	476.25	717.60	10.00

Table 8.63: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with convergence rate changes

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
5	95.14	172.04	58.00	89.79	16.00	31.32	43.90
	0.00	0.00	0.00	0.00	8.17	11.63	5.34
25	95.12	129.54	48.11	90.19	133.30	232.88	61.30
	0.07	11.67	2.77	0.67	63.80	95.86	5.33
50	95.12	110.69	44.38	90.82	241.60	431.61	61.00
	0.10	7.34	1.41	0.54	142.17	226.69	9.56
100	95.16	106.82	43.90	90.62	513.40	926.97	68.70
	0.09	4.17	1.14	0.66	235.44	385.70	8.79
200	95.10	102.49	42.60	90.57	1231.80	2330.60	81.00
	0.05	4.25	0.92	0.59	522.54	893.32	7.35

Table 8.64: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with convergence rate changes

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
5	0.00	0.00	0.00	0.00	16.50	58.94	45.80
	0.00	0.00	0.00	0.00	11.06	29.07	5.62
25	95.15	104.90	53.78	94.61	164.80	556.12	65.60
	0.05	6.88	2.20	0.36	69.17	217.20	5.57
50	95.13	99.65	50.90	94.51	390.50	1386.15	74.30
	0.08	7.09	1.92	0.22	265.08	899.58	13.70
100	95.16	95.45	49.78	94.39	527.40	2046.30	76.10
	0.09	9.47	2.78	0.33	420.13	1544.00	15.42
200	95.21	87.97	48.00	94.53	1727.40	6989.43	99.40
	0.10	6.54	1.41	0.30	891.91	3453.76	18.18

Table 8.65: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with convergence rate changes

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
5	0.00	0.00	0.00	0.00	21.00	70.72	45.00
	0.00	0.00	0.00	0.00	16.91	43.94	4.31
25	95.20	103.44	53.20	95.20	143.10	472.80	65.20
	0.11	6.64	2.18	0.11	63.51	179.99	3.92
50	95.18	97.51	50.60	95.18	266.10	945.93	71.90
	0.06	7.92	2.62	0.06	117.86	368.34	11.31
100	95.11	92.84	48.50	95.11	611.00	2236.50	81.60
	0.09	6.97	1.91	0.09	241.57	822.93	9.94
200	95.15	88.83	47.60	95.15	1998.60	7732.60	106.30
	0.10	3.67	1.02	0.10	987.97	3708.49	10.45

Table 8.66: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with convergence rate changes

Rate	Cover	Overlap	Cost	Traffic
25	-6.15	-36.16	-19.31	-7.26
50	-7.00	-49.58	-27.38	-8.53
100	-7.76	-66.04	-37.86	-9.64
200	-9.10	-119.85	-71.28	-11.38

Table 8.67: Ave. % change in objective measures for stagnation rate of 5 based solutions vs. 25, 50, 100 and 200 across SE (for highest coverage solutions)

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
5	90.98	114.48	50.80	89.22
	1.82	13.80	3.09	2.21
25	95.80	140.55	57.00	94.82
	1.10	10.81	2.61	1.41
50	96.60	152.21	59.60	95.78
	0.98	27.02	7.53	1.11
100	97.48	166.23	63.60	96.86
	0.77	18.67	4.82	0.87
200	98.40	196.75	72.00	98.03
	0.96	35.80	10.62	1.10

Table 8.68: Ave. and SD for highest cover solution in SE1 given 10 trials with convergence rate changes

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
5	92.17 1.51	134.74 21.69	51.50 5.43	85.41 2.39
25	96.40 0.99	182.57 25.57	61.10 4.78	92.05 2.21
50	96.77 1.30	176.19 34.90	60.60 9.87	93.48 2.43
100	97.87 0.90	200.22 36.09	68.40 9.05	95.49 1.87
200	98.97 0.38	245.62 32.25	81.70 11.46	97.51 1.00

Table 8.69: Ave. and SD for highest cover solution in SE2 given 10 trials with convergence rate changes

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
5	89.40 2.83	100.55 7.16	53.70 3.29	88.81 2.86
25	96.66 1.23	145.00 19.61	65.70 6.13	96.28 1.37
50	97.68 1.34	173.16 41.71	74.80 14.15	97.49 1.48
100	97.58 2.15	177.52 50.04	75.50 16.01	97.35 2.49
200	99.23 1.21	250.22 57.89	98.70 18.48	99.14 1.37

Table 8.70: Ave. and SD for highest cover solution in SE4 given 10 trials with convergence rate changes

Rate	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
5	90.82 2.51	102.77 11.72	52.80 4.12	90.82 2.51
25	96.81 0.75	146.12 15.12	65.50 4.52	96.81 0.75
50	97.69 1.29	166.87 36.64	71.40 10.97	97.69 1.29
100	98.60 0.95	199.17 30.94	80.70 9.76	98.60 0.95
200	99.77 0.18	284.04 30.38	106.10 10.33	99.77 0.18

Table 8.71: Ave. and SD for highest cover solution in SE5 given 10 trials with convergence rate changes

8.4.4 Changes in crossover

The type of cross-over used can influence the solution quality and run times, as the means by which the genetic algorithm attempts to improve solutions has been changed. For this reason, single, double, and multi-point crossover methods were investigated.

First, the average (from 10 trials) of the first solution providing $\geq$ 95% STP coverage across all SE (Tables 8.73, 8.74, 8.75, 8.76) were compared. Interpreting the results (Table 8.72) indicates that as the crossover method changes from single-point and coverage is held constant double-point crossover leads to marginally higher overlap and cost (i.e., a rise of 3.06% and 2.28% respectively) and longer execution times (i.e, 21.30% longer), and multi-point crossover also leads to slightly higher overlap and cost (i.e., a rise of 11.79% and 9.28% respectively) and marginally longer execution times (i.e, 2.26% longer)

Second, the average (from 10 trials) of solutions providing the highest STP coverage across all SE (Tables 8.78, 8.79, 8.80, 8.81) were compared. Interpreting the results (Table 8.77) indicates that as the crossover method changes from single-point and maximum coverage is sought double-point crossover results in a very marginal increase in coverage for a slight rise in overlap and cost (i.e., 8.63% and 7.90% respectively), and multi-point crossover results in a slight decrease in coverage for a slight rise in overlap and cost (i.e., 3.65% and 4.34% respectively).

Overall, the results indicate that single-point crossover should be preferred over double or multi-point crossover, as it tends to lead to better solutions in terms of coverage, cost, overlap, and execution times.

X-Over	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
double-pt.	-0.03	-3.06	-2.28	0.00	-20.20	-21.30	-7.05
multi-pt.	0.01	-11.79	-9.18	0.02	3.67	2.26	-3.13

Table 8.72: Ave. % change in objective measures for single-point crossover based solutions vs. double and multi-point across SE (for 1st solutions exceeding 95% STP coverage)

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Multi-pt	95.18	103.80	46.20	94.39	236.70	401.15	61.80
	0.07	9.09	2.23	0.30	110.68	159.51	5.49
Double-pt	95.15	111.69	49.20	94.34	285.10	475.72	65.30
	0.10	10.83	3.46	0.19	124.46	180.39	4.69
Single-pt	95.07	125.87	52.22	94.15	238.70	399.39	64.90
	0.07	18.09	3.55	0.33	107.31	159.71	8.65

Table 8.73: Ave. and SD for first solutions in SE1 exceeding 95% STP coverage given 10 trials with crossover method changes

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Multi-pt	95.10	128.23	47.20	89.90	235.60	412.39	61.60
	0.05	10.85	2.27	0.82	85.10	125.31	7.28
Double-pt	95.15	130.38	49.44	89.74	286.40	523.58	66.50
	0.10	13.87	2.41	0.75	153.86	246.86	7.83
Single-pt	95.10	144.57	54.00	90.15	226.20	401.42	61.70
	0.09	19.21	3.93	0.46	105.81	161.23	3.26

Table 8.74: Ave. and SD for first solutions in SE2 exceeding 95% STP coverage given 10 trials with crossover method changes

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Multi-pt	95.13	98.96	51.20	94.39	332.50	1200.63	73.70
	0.10	10.64	3.22	0.26	168.85	541.09	6.36
Double-pt	95.19	99.74	51.60	94.39	341.10	1282.39	77.60
	0.06	8.93	2.42	0.46	183.09	707.72	13.79
Single-pt	95.08	109.54	54.50	94.19	264.00	1016.14	77.50
	0.06	11.81	3.91	0.26	173.81	653.93	16.56

Table 8.75: Ave. and SD for first solutions in SE4 exceeding 95% STP coverage given 10 trials with crossover method changes

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Multi-pt	95.10	99.09	51.11	95.10	299.30	1060.96	73.20
	0.07	9.50	2.64	0.07	161.52	538.47	11.39
Double-pt	95.16	99.88	50.89	95.16	407.60	1409.43	80.00
	0.10	6.61	2.02	0.10	209.73	726.43	16.67
Single-pt	95.21	101.53	52.56	95.21	326.40	1161.08	74.80
	0.10	5.70	2.83	0.10	215.06	806.70	14.16

Table 8.76: Ave. and SD for first solutions in SE5 exceeding 95% STP coverage given 10 trials with crossover method changes

X-Over	Cover	Overlap	Cost	Traffic
double-pt.	-0.14	-8.63	-7.90	-0.05
multi-pt.	0.73	-3.65	-4.34	1.01

Table 8.77: Ave. % change in objective measures for single-point crossover based solutions vs. double and multi-point across SE (for highest coverage solutions)

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Multi-pt	97.10 0.78	165.82 21.53	61.60 5.90	96.41 0.93
Double-pt	97.08 0.63	166.88 17.95	65.60 5.68	96.45 0.73
Single-pt	96.29 1.21	169.99 23.75	65.60 10.39	95.47 1.50

Table 8.78: Ave. and SD for highest cover solution in SE1 given 10 trials with crossover method changes

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Multi-pt	96.51 0.95	178.87 22.67	61.10 7.18	92.51 1.87
Double-pt	96.68 1.35	199.88 29.61	67.10 9.32	92.35 3.14
Single-pt	95.61 0.87	178.53 12.44	60.90 3.33	90.56 2.02

Table 8.79: Ave. and SD for highest cover solution in SE2 given 10 trials with crossover method changes

X-Over	Ave. Cove. SD	Ave. Over. SD	Ave. Cost SD	Ave. Traf SD
Multi-pt	97.82 1.07	170.68 22.86	73.30 6.54	97.47 1.25
Double-pt	97.87 1.25	185.04 41.53	78.00 14.68	97.45 1.63
Single-pt	96.61 2.22	183.81 61.30	78.70 19.82	96.46 2.01

Table 8.80: Ave. and SD for highest cover solution in SE4 given 10 trials with crossover method changes

X-Over	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Multi-pt	97.63 1.55	171.79 43.14	73.20 12.06	97.63 1.55
Double-pt	97.97 1.44	195.39 57.62	79.70 17.10	97.97 1.44
Single-pt	97.72 1.34	179.69 47.94	76.00 14.57	97.72 1.34

Table 8.81: Ave. and SD for highest cover solution in SE5 given 10 trials with crossover method changes

8.4.5 Changes in representation

The type of representation used could potentially have a significant influence over solution quality and run times, as the means by which the genetic algorithm attempts to improve solutions has been changed. For this reason binary and integer-coded representations were investigated. Due to extremely long run times for integer-coding, the test were only run on SE1 and SE2. The overlap constraint for MCO-S was set to 80% based on initial trials to ensure high coverage solutions were found.

First, the first solutions in SE1 and SE2 providing $\geq$ 95% STP (Table 8.83) were compared. Interpreting the results (Table 8.82) indicates that as the choice of genetic representation changes from binary-coded and coverage is held constant that integer-based solutions result in 27.41% less overlap and incur a cost savings of 13.12% for SE1, and 30.90% less overlap and a cost savings of 11.65% for SE2. The trade-off is that run times leap 652.30% for SE1 and 565.68% for SE2.

Second, solutions in SE1 and SE2 providing the highest STP coverage (Table 8.85) were compared. Interpreting the results (Table 8.84) indicates that as the choice of genetic representation changes from binary-coded and maximum coverage is sought STP coverage increases marginally by 0.17% and traffic coverage by 0.44%, while overlap decreases by 28.61%, and cost decreases by 16.55% for SE1; and STP coverage increases marginally by 0.49% and traffic coverage by 1.85%, while overlap decreases by 31.16%, and cost decreases by 6.12% for SE2.

Therefore, if time were not a factor, one would prefer the integer-coded representation and strategy over binary-coded. In short, if the highest quality solutions are being sought, integer-coding should be used. However, if the fastest solutions are

being sought, binary-coding should be used.

Represent.	Cover	Overlap	Cost	Traffic	Gens.	Time	Sols.
Int.(SE1)	0.07	27.41	13.12	-0.13	-30.18	-652.30	63.28
Int.(SE2)	-0.11	30.09	11.65	-1.49	-24.54	-565.68	63.07

Table 8.82: Ave. % change in objective measures for binary-coded solutions vs. integer-coded for SE1 and SE2 (for 1st solutions exceeding 95% STP coverage)

Repres.	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD	Gens. Ave. SD	Time Ave. SD	Sols. Ave. SD
Ran.(SE1)	95.08	71.95	38.83	94.33	446.00	5615.14	25.67
	0.00	0.03	0.37	0.24	100.30	1328.23	3.54
Int.(SE1)	95.15	99.12	44.70	94.21	342.60	746.40	69.90
	0.07	7.92	1.62	0.21	113.66	225.17	5.34
Ran.(SE2)	95.14	116.46	45.50	90.17	291.40	683.54	65.80
	0.05	7.74	1.80	0.46	152.21	320.81	5.91
Int.(SE2)	95.24	81.42	40.20	91.51	362.90	4550.18	24.30
	0.10	4.41	0.60	0.35	111.70	1217.12	2.83

Table 8.83: Ave. and SD for highest cover solution in SE1 and SE2 given 10 trials with genetic representation changes

Represent.	Cover	Overlap	Cost	Traffic
Int.(SE1)	-0.17	28.61	16.55	-0.44
Int.(SE2)	-0.49	31.16	6.12	-1.85

Table 8.84: Ave. % change in objective measures for binary-coded solutions vs. integer-coded for SE1 and SE2 (for highest coverage solutions)

8.5 Conclusions

In this chapter we completed parameter tuning and testing for the simulation environment, decoder, and genetic algorithm. The general trend was that many settings incur a trade-off between solution quality and execution time. Testing was

Repres.	Cove. Ave. SD	Over. Ave. SD	Cost Ave. SD	Traf. Ave. SD
Ran.(SE1)	98.22 0.00	135.22 0.12	58.33 3.54	97.91 0.41
Int.(SE1)	98.05 0.58	189.40 19.45	69.90 6.07	97.48 0.83
Ran.(SE2)	97.50 0.54	197.57 21.19	65.40 6.36	94.52 1.27
Int.(SE2)	97.98 0.16	136.00 10.04	61.40 4.50	96.27 0.33

Table 8.85: Ave. and SD for first solutions in SE1 and SE2 exceeding 95% STP coverage given 10 trials with genetic representation changes

completed in a standardized environment which for the SE used 568 candidate sites and 1× traffic; which for TSCA used $T_{MAX} = 60$ and three random directive antenna; and which for the GA used placement probability, population size = 500, stagnation rate = 50, multi-point crossover, and binary-coding. These settings were based on preliminary testing.

Following testing, the following recommendations can be made:

- SE

  CBS: Use 568 CBS for default and when seeking best solution quality. Use less for faster or more challenging runs.

  Traffic multiplier: Use 1× for general testing. Use larger multiplier for more challenging runs, and small multiplier for faster runs.

- TSCA

  T_{MAX}: Use $T_{MAX} = 60$ for general testing. Use higher T_{MAX} settings to

encourage higher coverage solutions at lower cost, but recognize that this comes at the expense of more overlap and ultimately interference.

Antennas: Use three randomly determined directive antennas.

- GA

  Starting population: Use placement probability for speed and random populations for ultimate solution quality.

  Population size: A default size of 500 is suggested, although any table between 200-800 has benefits in terms of trade-off between speed and solution quality. 200 should be the minimum size used for faster runs, and 800 for ultimate solution quality.

  Stagnation rate: A default rate of 50 is recommended. A smaller rate results in faster runs, while a larger rate results in better solution quality. A maximum size of 200 is suggested.

  Cross-over: Single-point crossover is recommended for binary-coding.

  Representation: Binary-coding is recommended for faster runs and general testing, while integer-coding is recommended for real-world testing and higher solution quality.

In order to test these parameter settings, a final test using all 3D objective measures is suggested. This will be performed in Chapter 9. The relationship between certain objective can be determined with these tests, and by running tests with both GA representations perhaps more detailed advice can be given regarding the preference of one over the other.

Chapter 9

Final 3D Cell Plan Testing

In this chapter, we perform the final 3D cell plan testing. In this testing, we compare results from single-trial tests using SE1, 2, and 3 (Figure 6.1) with antennas at heights of 30-200m using binary-coded default settings, binary-coded optimized settings, and integer-coded optimized settings. This is done to provide evidence as to the most suitable approach to use during realistic 3D cell planning. Second, we investigate the correlational matrix between objective measures, in order to assess the degree of relationships between objectives. And finally, we look at example solutions to each SE and propose a benchmark solution for each.

During testing the maximum degree of coverage by T_{MAX} is as displayed in Table 9.1. As we use the default T_{MAX} of 60, it notable that 100% coverage is not possible in all SE. For this reason, we continue to use the approach followed in Chapter 8 of measuring the average and standard deviation at two points. The first is based on the objective values of the first solution providing $\geq$ 95% STP coverage, which effectively makes STP coverage constant (at 95%) and allows one to note variations

in the other objectives. The second is based on the solution with the highest STP coverage, which allows one to note the ability of the approach to maximize coverage, and how this affects other objectives.

Sim.Env	T_{MAX} value			
	43	60	80	100
SE1	99.15	99.89	100.00	100.00
SE2	99.48	99.86	99.99	100.00
SE3	99.27	99.98	99.99	99.99

Table 9.1: Maximum degree of coverage in simulated environments by T_{MAX} setting

9.1 Beta Tuning for Integer-coded Approach

An additional concern is the affect that β can have on the results. Given cell overlap is restricted to the percentage of overlap β when using the integer-permutation strategy, an appropriate level capable of achieving high levels of coverage was determined. This can be seen in Figure 9.1. This figure shows that a β between 70 to 90 is needed to achieve acceptable levels of coverage. Therefore, for experiments conducted herein, an average value of 80 was used during integer-coded testing.

9.2 Single-trial Comparisons

To compare the effectiveness of each approach, we compare the Pareto fronts of non-dominated solutions each produces. The approaches used are termed *binary-default*, which uses the trial settings as described in Chapter 8.1, *binary-optimized*, and *integer-optimized*. The only differences are that the optimized approaches use a

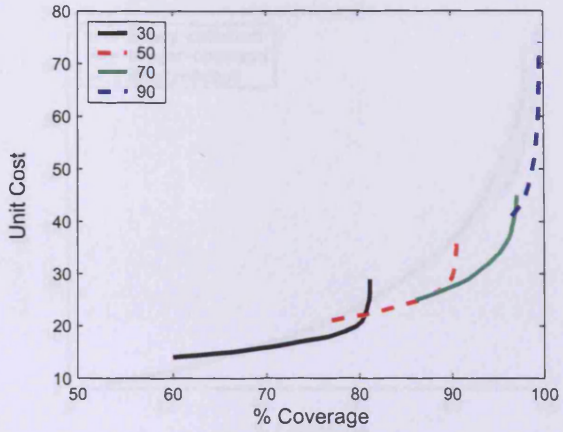

Figure 9.1: General pattern of cost vs. coverage by β value with $T_{MAX} = 80$

population size of 800 instead of 500, a convergence rate of 100 instead of 50, and the binary-optimized uses single-point crossover instead of multi-point.

The results of the comparisons for SE1 appear in Figure 9.2, for SE2 in Figure 9.3, and for SE3 in Figure 9.4. The trends in each figure are similar: First, the integer-optimized approach achieves the best unit cost of coverage, followed by the binary-optimized, and finally by the binary-default. This is seen most clearly in SE1 and SE3, with the binary-optimized set only performing slightly better at very high levels of coverage in SE2. Second, the binary-optimized approach produces the longest Pareto front, and therefore the widest range of solutions. This is seen most clearly in SE1 and SE2. In SE3, the binary-default is equally long, although the solutions are not as optimal. Note that the reason the integer-optimized Pareto front is not as long is that the β setting largely determines the level of coverage, with low settings (e.g., 10%) resulting in low coverage and high settings (e.g., 70%) resulting in high coverage. This is shown to be true later in Figure 9.1. In short, these results show the integer-

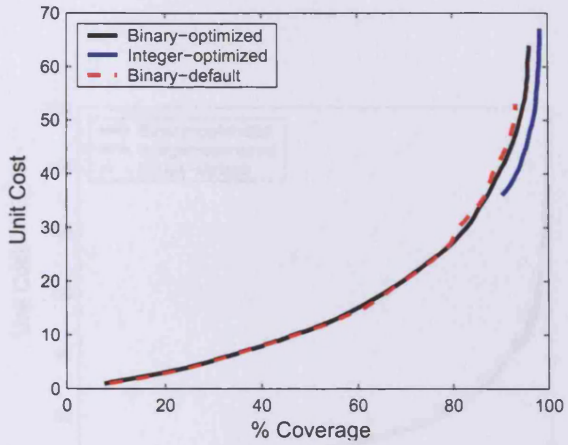

Figure 9.2: STP coverage vs. cost using default binary-coding, optimized binary-coding, and optimized integer-coding SE1

optimized approach creates the highest quality solutions, while the binary-optimized approach creates the widest variety of solutions. For the integer-based approach to create a wide variety of solutions, a wide-ranging number of β settings would need to be used, although this would increase its execution time considerably. For example, given the average run time ends up being 3.5 hours using one β setting, using 10 settings would require roughly 35 hours.

9.3 Relationship Between Objectives

Next, the relationships between objectives for SE1 can be seen visually in Figure 9.5, 9.6, 9.7, 9.8, 9.9, 9.10, and in a Pearson correlation matrix (based on results using binary-optimized) in Table 9.2. The Pearson correlation was chosen as it is the most common measure of the linear relationship between two variables. It ranges from 1 to -1, where a correlation of 1 means that there is a perfect positive linear relationship

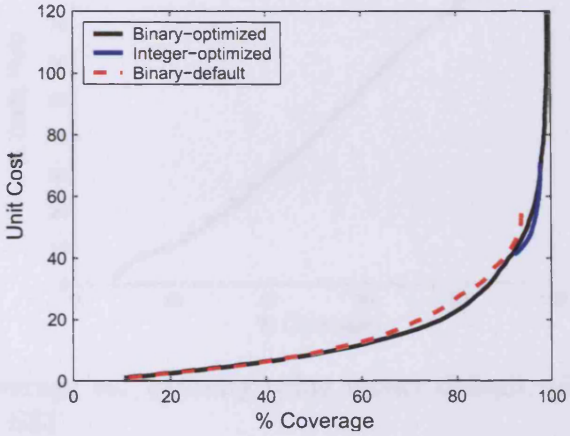

Figure 9.3: STP coverage vs. cost using default binary-coding, optimized binary-coding, and optimized integer-coding SE2

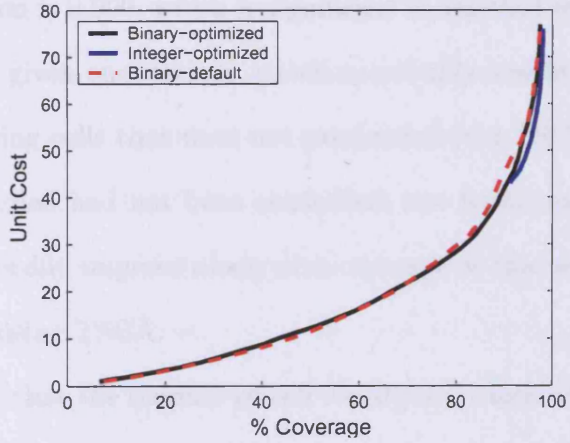

Figure 9.4: STP coverage vs. cost using default binary-coding, optimized binary-coding, and optimized integer-coding SE3

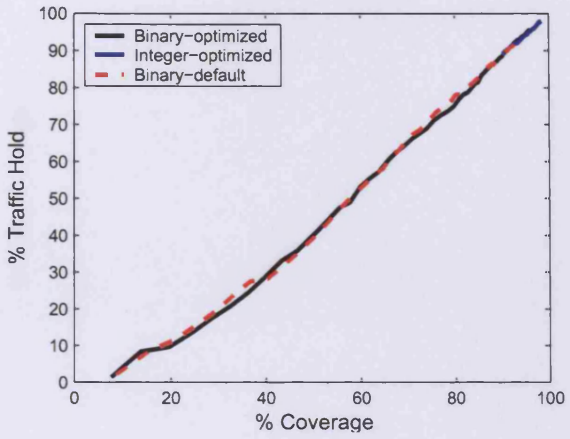

Figure 9.5: STP coverage vs. capacity using binary-default, binary-optimized, and integer-optimized in SE1

between variables. Figure 9.5 shows that there is nearly a perfect positive correlation between traffic hold and coverage. Table 9.2 reveals the correlation is 0.998, which is significant at the 0.01 level (2-tailed). Similarly, Figure 9.6 shows that there is nearly a perfect positive correlation between traffic coverage and STP coverage. Table 9.2 reveals the correlation is 0.999, which is significant at the 0.01 level (2-tailed). These results are expected given our settings, which specifically sought to improve coverage (using NSGA-II) using cells that were not overloaded (via TSCA). If the amount of traffic each cell handled had not been controlled, one would expect that the traffic hold would not so readily improve along with coverage as that seen here. This shows the effectiveness of using TSCA.

Figure 9.7 shows that the number of *cell violations* (which occurs whenever a cell exceeds 43 Erlangs of traffic) climbs linearly until roughly 50% coverage, after which point the number of cell violations linearly decreases. Table 9.2 reveals the correlation is -.603, which is significant at the 0.01 level (2-tailed). Therefore, generally we should

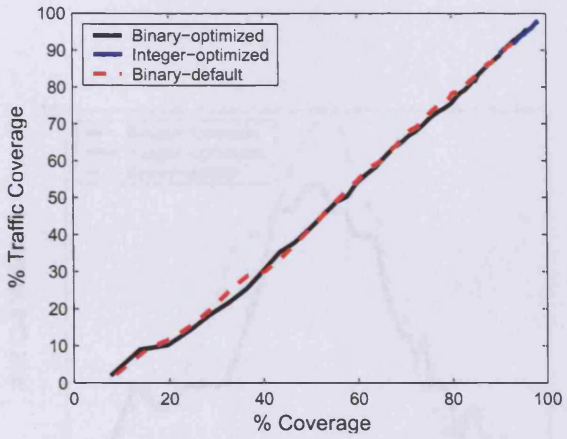

Figure 9.6: STP coverage vs. traffic using binary-default, binary-optimized, and integer-optimized in SE1

expect that as coverage rises, the number of cell violations falls. This pattern occurs because as more cells are added, cells are more tightly packed, which means that the traffic gets 'shared' between them, reducing the overall number of cell violations as traffic is spread amongst more cells.

Figure 9.8 shows that handover and coverage tend to improve along with one another. Table 9.2 reveals the correlation is 0.886, which is significant at the 0.01 level (2-tailed). Therefore, generally we should expect that as coverage rises, handover rises as well. This pattern occurs because as more cells are added, cells are more tightly packed, which means that handover occurs more frequently and easily. If cells were not in contact with one another, coverage would be low and handover could not occur. So, the opposite is true as well: the more contact between cells, the more handover that occurs.

Figure 9.9 and 9.10 show a similar trend. This is that as coverage rises, so do overlap and interference. Indeed, Table 9.2 reveals the correlation between coverage

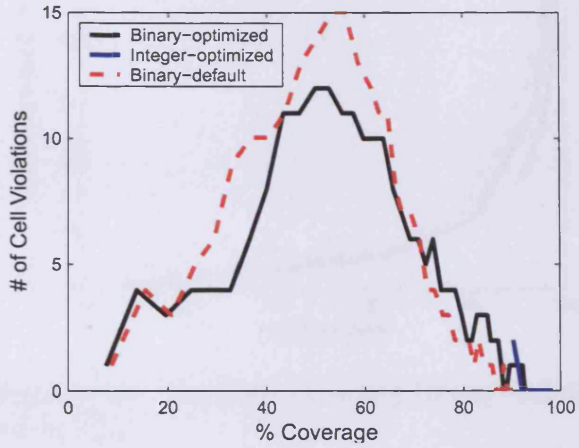

Figure 9.7: STP coverage vs. cell violations using binary-default, binary-optimized, and integer-optimized in SE1

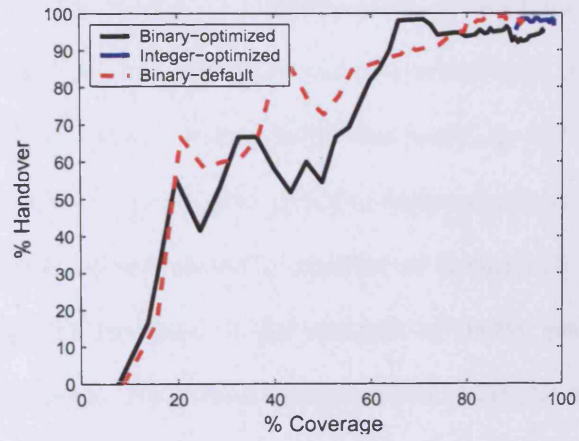

Figure 9.8: STP coverage vs. handover using binary-default, binary-optimized, and integer-optimized in SE1

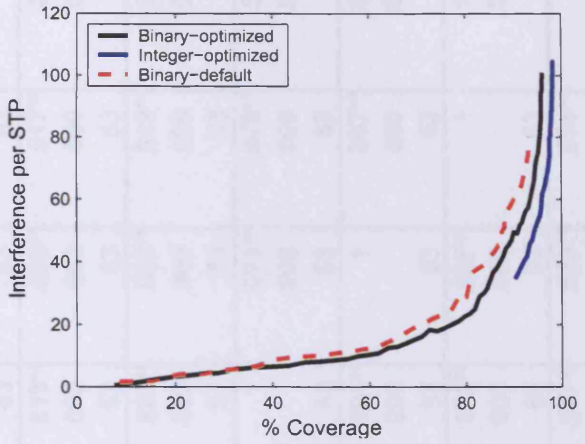

Figure 9.9: STP coverage vs. interference using binary-default, binary-optimized, and integer-optimized in SE1

and interference is 0.806 and between coverage and overlap is 0.875, which in both cases is significant at the 0.01 level (2-tailed). Therefore, generally we should expect that as coverage rises, interference and overlap rise as well. This pattern occurs because as more cells are added and coverage improves, cells are more tightly packed, which results in more cell overlap and concomitant interference. Indeed, it is difficult to imagine *any* realistic cell planning scenario where this relationship would not hold because as has been shown throughout this work, as 100% service coverage is approached cost (and therefore number of cells) increases at an accelerated rate.

The results in these figures shows a number of important trends given our 3D cell planning strategy which controls the amount of traffic per cell. These can be summarized as follows–although it must be understood that the results and summaries are dependent on the algorithms used herein, and that different findings are likely if one employs different cell planning algorithms; in particular, the algorithm used herein directly controls traffic, whereas alternative approaches do not:

		cover1	violations1	traffic1	capacity1	interference1	overlap1	cost1	handover1
cover1	Pearson Correlation	1	-.603**	.999**	.998**	.806**	.875**	.907**	.886**
	Sig. (2-tailed)		.000	.000	.000	.000	.000	.000	.000
	N	63	63	63	63	63	63	63	63
violations1	Pearson Correlation	-.603**	1	-.626**	-.635**	-.749**	-.756**	-.757**	-.369**
	Sig. (2-tailed)	.000		.000	.000	.000	.000	.000	.003
	N	63	63	63	63	63	63	63	63
traffic1	Pearson Correlation	.999**	-.626**	1	1.000**	.819**	.886**	.917**	.870**
	Sig. (2-tailed)	.000	.000		.000	.000	.000	.000	.000
	N	63	63	63	63	63	63	63	63
capacity1	Pearson Correlation	.998**	-.635**	1.000**	1	.821**	.888**	.919**	.868**
	Sig. (2-tailed)	.000	.000	.000		.000	.000	.000	.000
	N	63	63	63	63	63	63	63	63
interference1	Pearson Correlation	.806**	-.749**	.819**	.821**	1	.991**	.979**	.547**
	Sig. (2-tailed)	.000	.000	.000	.000		.000	.000	.000
	N	63	63	63	63	63	63	63	63
overlap1	Pearson Correlation	.875**	-.756**	.886**	.888**	.991**	1	.997**	.633**
	Sig. (2-tailed)	.000	.000	.000	.000	.000		.000	.000
	N	63	63	63	63	63	63	63	63
cost1	Pearson Correlation	.907**	-.757**	.917**	.919**	.979**	.997**	1	.676**
	Sig. (2-tailed)	.000	.000	.000	.000	.000	.000		.000
	N	63	63	63	63	63	63	63	63
handover1	Pearson Correlation	.886**	-.369**	.870**	.868**	.547**	.633**	.676**	1
	Sig. (2-tailed)	.000	.003	.000	.000	.000	.000	.000	
	N	63	63	63	63	63	63	63	63

**. Correlation is significant at the 0.01 level (2-tailed).

Table 9.2: Correlations between objectives measures

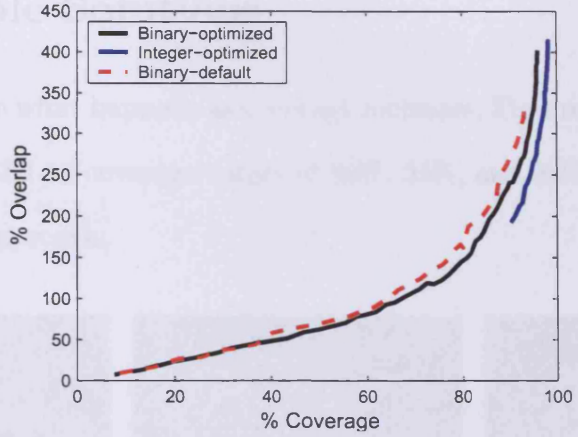

Figure 9.10: STP coverage vs. overlap using binary-default, binary-optimized, and integer-optimized in SE1

- Coverage, traffic coverage, and traffic hold are nearly perfectly positively correlated to one another.
- As coverage rises, interference, overlap, cost, and handover also rise and the number of cell violations drops.

These relationships mean that given *this* 3D cell planning strategy:

- One can safely optimise for cover and cost, given the highly significant correlations between coverage and all other objective measures.
- One would like to minimize the correlation between coverage and overlap and interference further.
- The strategy performs remarkably at maintaining such a close relationship between coverage and traffic hold; one would therefore like to increase T_{MAX} as much as possible while maintaining this correlation (to help reduce cost, overlap, interference, and cell violations).

9.4 Example Solutions

To help visualize what happens as coverage increases, Figures 9.11, 9.12, and 9.13 show solutions on SE1 at coverage values of 90%, 95%, and 98% respectively for the integer-optimized approach.

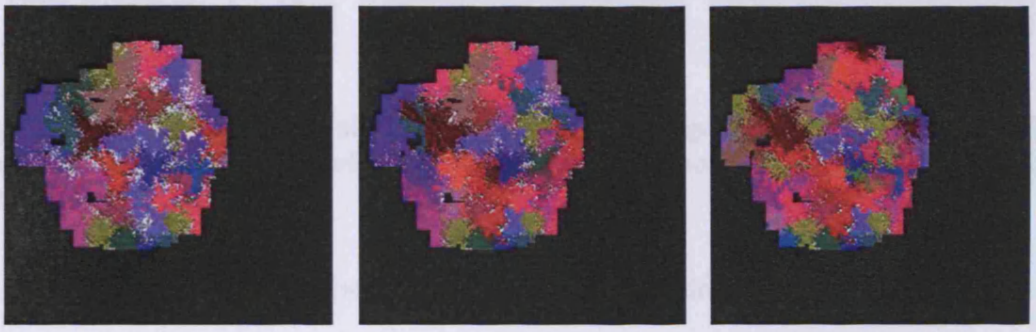

Figure 9.11: Example of solutions to SE1 at coverage values of 90% (left), 95% (middle), and 98% (right), where white indicates an uncovered STP, black RTP, and colored areas covered STP

Figure 9.12: Example of solutions to SE2 at coverage values of 90% (left), 95% (middle), and 98% (right), where white indicates an uncovered STP, black RTP, and colored areas covered STP

The benchmarks objective measures for solutions to SE1, SE2, and SE3 are given

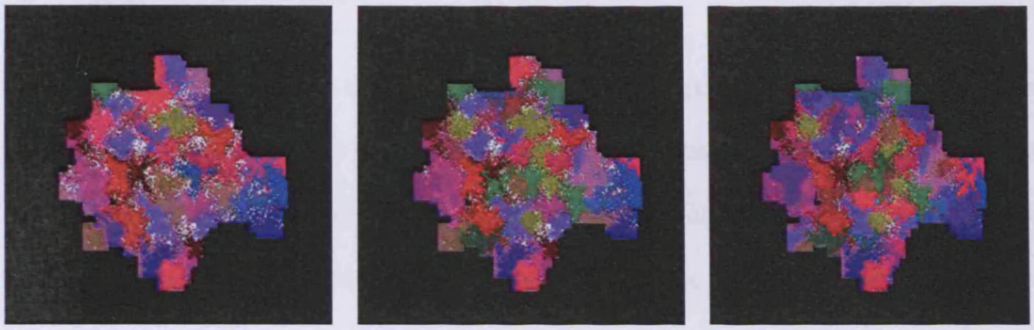

Figure 9.13: Example of solutions to SE3 at coverage values of 90% (left), 95% (middle), and 98% (right), where white indicates an uncovered STP, black RTP, and colored areas covered STP

in Table 9.3. There are two benchmark solutions. The first provides the solution with the lowest cost per unit of coverage, and the second the solution with the highest overall coverage. The former may be preferred by a low-cost operator and the second by a high-quality operator. That is, the first identifies the network with the largest return for investment, and the second the highest quality solution in terms of coverage.

SE	Cover	Cost	Capa.	Viol.	Hand.	Inte.	Over.	Traf.
SE1(UCC)	90.283	36.000	88.989	2.000	95.327	34.540	192.750	89.083
SE1(Cov)	98.074	60.000	97.711	0.000	97.753	88.397	377.140	97.711
SE2(UCC)	92.238	41.000	88.407	3.000	98.347	41.587	216.863	88.833
SE2(Cov)	97.804	67.000	96.194	1.000	97.917	95.042	391.520	96.537
SE3(UCC)	91.180	43.000	88.109	4.000	97.656	43.779	222.831	88.524
SE3(Cov)	98.057	75.000	97.380	0.000	96.279	97.271	396.441	97.380

Table 9.3: Benchmark solutions for SE1, SE2, and SE3 for lowest unit cost of coverage (UCC) and highest coverage (Cov)

9.5 Conclusions

Based on the finding in this chapter, we would suggest that the G3PS approach holds promise. There is clear evidence that it can meet high quality standards in terms of coverage, traffic hold, and handover, which is notable considering each antenna is dimensioned only once. Also, considering the average run time for each trial was roughly 3.5 hours, there is ample room to either add complexity to this model, (e.g., it would be easy to reduce interference by using the field strength formula to reset all antennas to the smallest transmission power necessary to service the set of STP covered by a given cell), or to use this strategy'as is' if fast network configuration were necessary, as might be the case in some military or emergency scenarios, or if dynamic dimensioning were needed. Some final considerations and future work will next be delineated in Chapter 10.

Chapter 10

Conclusions

The goal of this thesis was to scientifically investigate approaches to automatic mobile network design to improve the understanding of fundamental conflicts, tensions, and relationships between design objectives and to help operators to meet the demands of their subscribers more efficiently. This goal can be subdivided as follows:

- Engineer appropriate 2D and 3D wireless network planning simulation environments and make these publicly available.
- Apply novel and appropriate multi-objective cell planning strategies and tune their performance.
- Investigate the fundamental relationships and tensions between design objectives and suggest how this impacts cell planning.

The various chapters in this thesis have gone towards meeting these goals. The contributions are summarized in Section 10.1, and avenues of promising future research in Section 10.2.

10.1 Summary of Contributions

The main contributions following the introduction (i.e., Chapter 1) and review of the literature (i.e., Chapter 2) and before the conclusions (i.e., Chapter 10) are the following:

- In Chapter 3: **A 2D simulation model and multi-objective cell planning strategy was introduced and tested.** The 2D cell planning model is a simplified model appropriate for theoretical investigations. Within Chapter 3, it was used to investigate the suitability of four multi-objective algorithms to find an optimal ordering of candidate sites, where the order of sites was used by a greedy decoder to determine a cell plan. It was found that NSGA-II performed the best, and therefore it was used as the main MOA throughout this work.

- In Chapter 4: **Six decoders used to determine cell plans were designed and tested.** The decoders, which differed by how they controlled cell overlap and how they selected the next candidate site, were tested on their ability to optimise coverage and cost by controlling cell density. Results indicated that the most complex decoder, PMCO-P, was the most effective, as it used more criteria for determining cell density. However, all decoders were capable of finding good network designs and PCO and MCO decoders were also considerably faster. Thus, each decoder was useful in specific circumstances.

- In Chapter 5: **The cell density at which infrastructure efficiency is maximized was investigated.** The rate at which infrastructure cost changes is fundamental in determining the amount of coverage, and related capital invest-

ment, which the operator should employ to maximize profit. Results indicated that as 100% coverage is approached there is progressively rapid decline in return on investment. This clarifies the relationship between coverage and cost and suggests why investigators and cost-conscious operators should not simply seek to maximize coverage.

- In Chapter 6: **A 3D simulation model was engineered, design objectives defined and single site coverage tested.** An existing model was extended to create a realistic wireless simulation environment. Network components capable of allowing the measurement of all pertinent design objectives were also defined. This 3D model has been made freely and publicly available. This assures results presented herein can be challenged and extended in future scientific works.

- In Chapter 7: **A 3D cell planning strategy was extended from the 2D model.** Due to the increased simulation complexity (in particular, traffic load and increased site complexity) a pre-processing phase (termed TSCA) was introduced to the cell planning strategy to configure cells before optimizing on coverage and cost. TSCA helps guarantee that the most difficult aspect of cellular network design (i.e., meeting traffic hold or capacity requirements) is met. This frees the network designer to concentrate on ascertaining high service coverage at low cost. An additional binary-representation model was also added to compare and contrast to the standard integer-representation used during 2D planning.

- In Chapter 8: **Pertinent 3D variables and parameters were tuned and**

tested. Results generally indicated a trade-off for many parameter settings (e.g., GA population size) between solution quality and speed of execution, where different settings are more appropriate depending upon the application. From this work, an optimized parameter configuration was suggested for both the integer and binary based approaches.

- In Chapter 9: **Tests to compare 3D cell planning strategies and objective relations were completed.** Results indicated that the integer-based approach should be preferred if the highest quality network solutions are sought, and the binary-based approach preferred if execution time is paramount, or if the largest Pareto front of non-dominated alternative networks is desired. It was found that STP coverage, traffic coverage, and traffic hold were nearly perfectly positively correlated when using *this* 3D cell planning strategy and optimized settings, and that using this strategy as coverage rises overlap, interference, cost, and handover also rise. This indicates that when using the optimized settings, optimizing for coverage and cost sufficiently accounts for meeting all pertinent network design objectives while using the G3PS approach.

10.2 Future Work

The research conducted herein has given rise to several promising avenues of future research on wireless network modeling, multi-objective cell planning, and the impact of design objective relations. These are as follows:

- The 3D wireless network model can be improved once access to more GPS data

(e.g., traffic load or population density at given Cartesian coordinates) is made publicly available. In this way, the theoretical traffic modeling could be eliminated in favour of real data. The potential of joining DEM data, GPS data, and processing to calculate PLM and AIM via a GRID infrastructure is also exciting. This would allow instant access to extremely realistic network simulation environments to participating partners, and hopefully the wider academic community via updates to an accessible database.

- The multi-objective cell planning model designed herein could be extended in future work. For example, it may be possible to make the model more dynamic when processing speeds improve. This would make it possible to increase the T_{MAX} setting to the furthest limit before the significant correlations between design objectives deteriorates. This should aid in discovering more low-cost networks. Also, it would be interesting to encourage more low-cost operators to consider the evidence provided. For example, the research shows that reasonably high network coverage can be obtained more cheaply by optimizing the infrastructure efficiency. These cost improvements could translate into more value for money to their customers.

- The work on the tensions and relationships between network design objectives can also be worked on further. Two important considerations already manifest is the diminishing return on investment as coverage rises, and that due to the correlations between objectives, it is possible to optimise based on two objectives. Future work could investigate the validity of this by testing more combinations of objectives in increasingly complex simulation environments. Also, pairs of

different objectives could be used in the optimization process to see if the same correlations are found.

Thus, although this thesis has gone a long way to improving the development of our understanding of simulation environments, multi-objective GSM wireless network planning, and the relations and tensions between design objectives, there is further work to be done to help meet the growing needs of people who are on the move, and who want reliable access to telecommunication services wherever they are.

Bibliography

[1] E.H.L. Aarts and J.H.M. Korst. *Simulated Annealing and Boltzmann Machines.* John Wiley and Sons, 1989.

[2] R.G Akl, M.V. Hedge, M. Naraghi-Pour, and P.S. Min. Multicell cdma network design. *IEEE Transactions on Vehicular Technology*, 50:711–722, 2001.

[3] S.M. Allen, S. Hurley, R.K. Taplin, and R.M. Whitaker. Automatic cell planning of broadband fixed wireless networks. In *Proceedings of the IEEE VTC Conference (Spring)*, pages 2808–2812, Rhodes, Greece, May 2001.

[4] E. Amaldi, A. Capone, and F. Malucelli. Improved models and algorithms for UMTS radio planning. In *Proceedings 54th IEEE Conference on Vehicular Technology*, volume 2, pages 920–924, 2001.

[5] E. Amaldi, A. Capone, and F. Malucelli. Optimizing base station siting in UMTS networks. In *Proceedings 53rd IEEE Conference on Vehicular Technology*, volume 4, pages 2828–2832, 2001.

[6] H.R. Anderson and J.P. McGeehan. Optimizing microcell base station locations using simulated annealing techniques. In *Proceedings 44th IEEE Conference on Vehicular Technology*, pages 858–862, 1994.

[7] J. D. Bagley. *The Behavior of Adaptive Systems which Employ Genetic and Correlation Algorithms.* PhD thesis, University of Michigan, 1967.

[8] R. Battiti, A. Bertossi, and D. Cavallaro. A randomized saturation degree heuristic for channel assignment incellular radio networks. *IEEE Transactions on Vehicular Technology*, 50(2):364–374, 2001.

[9] Roberto Battiti and Alan Albert Bertossi. Greedy, prohibition, and reactive heuristics for graph partitioning. *IEEE Trans. Comput.*, 48(4):361–385, 1999.

[10] A. Berny and O. Sarzeud. Optimization de resaux de radiotelephonie mobile par recherche locale et selection. *JNCP*, pages 953–966, 2000.

[11] S. Bleuler, M. Brack, L. Thiele, and E. Zitzler. Multiobjective genetic programming: Reducing bloat using spea2. In *Proceedings of the Congress on Evolutionary Computation 2001*, volume 1, pages 536–543, 2001.

[12] Heinrich Braun. On solving travelling salesman problems by genetic algorithms. In *PPSN I: Proceedings of the 1st Workshop on Parallel Problem Solving from Nature*, pages 129–133, London, UK, 1991. Springer-Verlag.

[13] P. Calegari, F. Guidec, P. Kuonen, and D. Wagner. Genetic approach to radio network optimizations for mobile systems. In *Proceedings 47th IEEE Conference on Vehicular Technology*, volume 2, pages 755–759, 1997.

[14] B. Chamaret, S. Josselin, P. Kuonen, M. Pizarroso, B. Salas-Manzanedo, S. Ubeda, and D. Wagner. Radio network optimization with maximum independent set search. In *Proceedings of the IEEE VTC'97 Conference*, pages 770–774, Phoenix, AZ, May 1997.

[15] D.W. Corne, J.D. Knowles, and M.J. Oates. The pareto envelope-based selection algorithm for multiobjective optimization. In *Proceedings of the Sixth International Conference on Parallel Problem Solving from Nature*, pages 839–848, 2000.

[16] Charles Darwin. *On The Origin of Species by Means of Natural Selection, or the Preservation of Favoured Races in the Struggle for Life.* John Murray, London, 1859.

[17] K. Deb. *Multi-objective optimization using evolutionary algorithms.* Wiley, Chichester, England, 2001.

[18] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *Lecture Notes in Computer Science*, volume 1917, pages 849–858, 2000.

[19] K. Deb and S.Jain. Running performance metrics for evolutionary multi-objective optimization. *KanGAL Report No.2002004*, pages 1–18, 2002.

[20] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable test problems for evolutinary multi-objective optimization. In *Kangal Report No. 2001001*, pages 1–27, 2001.

[21] E. Ekici and C. Ersoy. Multi-tier cellular network dimensioning. *Wireless Networks*, 7:401–411, 2001.

[22] C.M. Fonseca and P.J. Fleming. Genetic algorithms for multiobjective optimization: Formulation, discussion, and generalization. In *Proceeding of the Fifth International Conference on Genetic Algorithms*, pages 416–423, 1993.

[23] C.M. Fonseca and P.J. Fleming. On the performance assessment and comparison of stochastic multiobjective optimizers. In *Fourth International Conference on Parallel Problem Solving from Nature*, pages 584–593, 1996.

[24] M. Galota, C. Glaßer, S. Reith, and H. Vollmer. A polynomial-time approximation scheme for base station positioning in UMTS networks. In *5th Discrete Algorithms and Methods for Mobile Computing and Communications Conference*, pages 52–59, 2000.

[25] A. Ganz, C.M. Krishna, D. Tang, and Z.J. Haas. On optimal design of multiter wireless cellular systems. *IEEE Communications Magazine*, pages 88–93, February 1997.

[26] B. Gavish and S. Sridhar. Economic aspects of configuring cellular networks. *Wireless Networks*, 1:115–128, 1995.

[27] B. Gavish and S. Sridhar. The impact of mobility on cellular network configuration. *Wireless Networks*, 7:173–185, 2001.

[28] W.K. Hale. Frequency assignment: Theory and applications. In *Proceedings of the IEEE*, volume 68, pages 1497–1514, 1980.

[29] J.K. Han, B.S. Park, Y.S. Choi, and H.K. Park. Genetic approach with a new representation base station placement in mobile communications. In *Proceedings 54th IEEE Conference on Vehicular Technology*, volume 4, pages 2703–2707, 2001.

[30] M. P. Hansen and A. Jaszkiewicz. Evaluating the quality of approximations to the non-dominated d set. *Technical Report IMM-REP-1998-7, Technical University of Denmark*, pages 1–6, 1998.

[31] Q. Hao, B. Soong, E. Gunawan, J. Ong, C.B. Soh, and Z. Li. A low-cost cellular mobile communication system: A hierarchical optimization network resource planning approach. *IEEE Journal on Selected Areas in Communications*, 15(7):1315–1326, 1997.

[32] M. Hata. Empirical formula for propogation loss in land mobile radio services. *IEEE Transactions on Vehicular Technology*, 29(3):317–325, August 1980.

[33] J.H. Holland, editor. *Adaptation in natural and artificial systems.* The University of Michigan Press, Ann Arbor, 1975.

[34] J. Horn, N. Nafpliotis, and D.E. Goldberg. A niched pareto genetic algorithm for multiobjective optimization. In *Proceedings of the First IEEE Conference on Evolutionary Computation*, pages 82–87, 1994.

[35] I. Howitt and S-Y. Ham. Base station location optimization. In *Proceedings of the IEEE VTC'99 Conference*, volume 4, pages 2067–2071, 1999.

[36] X. Huang. *Automatic cell planning for mobile network design: Optimization models and algorithms.* PhD thesis, Universitt Karlsruhe, 2001.

[37] X. Huang, U. Behr, and W. Wiesbeck. Automatic base station placement and dimensioning for mobile network planning. In *IEEE Vehicular Technology Conference*, volume 4, pages 1544–1549, 2000.

[38] X. Huang, U. Behr, and W. Wiesbeck. Automatic cell planning for a low-cost and spectrum efficient wireless network. In *Proceedings of Global Telecommunications Conference (GLOBECOM)*, volume 1, pages 276–282, 2000.

[39] X. Huang, U. Behr, and W. Wiesbeck. A new approach to automatic base station placement in mobile networks. In *International Zurich Seminar on Broadband Communications*, pages 301–306, 2000.

[40] S. Hurley. Planning effective cellular mobile radio networks. In *IEEE Transactions on Vehicular Technology*, volume 51, pages 243–253, March 2002.

[41] S Hurley, D.H. Smith, and S Thiel. FASoft: A system for discrete channel frequency assignment. *Radio Science*, 32(5):1921–1939, 1997.

[42] L.J. Ibbetson and L.B. Lopes. An automatic base station placement algorithm. In *Proceedings of the IEEE VTC'97 Conference*, pages 770–774, Phoenix, AZ, May 1997.

[43] B.A. Julstrom and G.R. Raidl. A permutation-coded evolutionary algorithm for the bounded-diameter minimum spanning tree problem. In *2003 Genetic and Evolutionary Computation Conference's Workshops Proceedings, Workshop on Analysis and Design of Representations*, pages 2–7, 2003.

[44] N. Gerlich K. Tutschku and P. Tran-Gia. An integrated approach to cellular network planning. In *presented at 7th Internaltional Telecommunication Network Planning Symposium*, 1996.

[45] R.H. Katz. Adaptation and mobility in wireless information systems. *IEEE Personal Communications Magazine*, 1:6–17, 1994.

[46] V. Khare, X. Yao, and K. Deb. Performance scaling of multi-objective evolutionary algorithms. *KanGAL Report No.2002009*, pages 1–15, 2002.

[47] J. D. Knowles and D. W. Corne. Local Search, Multiobjective Optimization and the Pareto Archived Evolution Strategy. In B. et al. McKay, editor, *Proceedings of Third Australia-Japan Joint Workshop on Intelligent and Evolutionary Systems*, pages 209–216, Ashikaga, Japan, 1999. Ashikaga Institute of Technology.

[48] J.D. Knowles and D. Corne. On metrics for computing non-dominated sets. In *World Congress on Computational Intelligence*, pages 711–716, 2002.

[49] Joshua D. Knowles. *Local-search and hybrid evolutionary algorithms for pareto optimization.* PhD thesis, The University of Reading, 2002.

[50] I. Laki, L. Farkas, and L. Nagy. Cell planning in mobile communication systems using sga optimization. In *Proceedings of International Conference on Trends in Communications*, volume 1, pages 124–127, 2001.

[51] N. Laumanns, M. Laumanns, and H. Kitterer. Evolutionary multiobjective integer programming for the design of adaptive cruise control systems. In *Proceedings of the Fifteenth International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, pages 1–10, 2002.

[52] C.Y. Lee and H.G. Kang. Cell planning with capacity expansion in mobile communications: A tabu search approach. *IEEE Transactions on Vehicular Technology*, 49(5):1678–1691, March 2000.

[53] Nissan Lev-Tov and David Peleg. Exact algorithms and approximation schemes for base station placement problems. In *SWAT '02: Proceedings of the 8th Scandinavian Workshop on Algorithm Theory*, pages 90–99, London, UK, 2002. Springer-Verlag.

[54] R.M. Mathar and T. Niessen. Optimum positioning of base stations for cellular radio networks. *Wireless Networks*, 6:421–428, 2000.

[55] R.M. Mathar and M. Schmeink. Optimal base station positioning and channel assignment for 3G mobile networks by integer programming. *Annals of Operations Research*, 107:225–236, 2001.

[56] R.M. Mathar and M. Schmeink. Integrated optimal cell site selection and frequency allocation for cellular radio networks. *Telecommunication Systems*, 21:339–347, 2002.

[57] H. Meunier, E. Talbi, and P. Reininger. A multiobjective genetic algorithm for radio network optimization. In *Proceedings of the 2000 Congress on Evolutionary Computation*, volume 1, pages 317–324, 2000.

[58] A. Molina, G.E. Athanasiadou, and A.R. Nix. The automatic location of base-stations for optimised cellular coverage: A new combinatorial approach. In *Proceedings of the IEEE VTC'99 Conference*, pages 606–610, 1999.

[59] A. Molina, G.E. Athanasiadou, and A.R. Nix. Optimised base-station location algorithm for next generation microcellular networks. *Electronics Letters*, 36(7):668–669, 2000.

[60] A. Molina and G.E. Nix, A.R. Athanasiadou. The effects of delay spread for cellular network planning using the combination algorithm for total optimisation. In *Proceedings of the 1st International Conference on 3G Mobile Communication Technologies*, pages 171–175, 2000.

[61] M. Mouly, M-B. Pautet, and T. Haug. *The GSM System for Mobile Communications.* Telecom Publishing, 1992.

[62] L. Raisanen and R.M. Whitaker. Multi-objective optimization in area coverage problems for cellular communication networks: Evaluation of an elitist evolutionary strategy. In *Proceedings of the 2003 ACM Symposium on Applied Computing*, pages 714–720, 2003.

[63] L. Raisanen, R.M. Whitaker, and S. Hurley. A comparison of randomized and evolutionary approaches for optimizing base station site selection. In *Proceedings of the 2004 ACM Symposium on Applied Computing*, pages 1159–1165, 2003.

[64] Larry Raisanen and Roger M. Whitaker. Comparison and evaluation of multiple objective genetic algorithms for the antenna placement problem. *MONET*, 10(1-2):79–88, 2005.

[65] D.P. Reed. The cost structure of personal communication services. *IEEE Personal Communications Magazine*, pages 102–108, 1993.

[66] Colin R. Reeves. Hybrid genetic algorithms for bin–packing and related problems. 1993.

[67] P. Reininger and A. Caminada. Connectivity management on mobile network design. In *presented at the 10th Conf. Eur. Consortium for Mathematics in Industry*, Goteborg, Sweden.

[68] P. Reininger and A. Caminada. Model for gsm radio network optimization. In *presented at the 2nd ACM Int. Conf. Discrete Algorithms and Methods for Mobility*, Dallas, Texas.

[69] P. Reininger, S. Iksal, A. Caminada, and J.J. Korczak. Multi-stage optimization for mobile radio network planning. In *Proceedings of the IEEE VTC'99 Conference*, volume 3, pages 2034–2038, 1999.

[70] J. Sarnecki, C. Vinodrai, A. Javed, P. O'Kelly, and K. Dick. Microcell design principles. *IEEE Personal Communications Magazine*, pages 76–82, 1993.

[71] J.D. Schaffer. *Some experiments in machine learning using vector evaluated genetic algorithms.* PhD thesis, Nashville, TN: Vanderbilt University, 1984.

[72] J. Schott. *Fault tolerant design using simple multicriteria genetic algorithms.* PhD thesis, MS Thesis, Department of Aeronautics and Astronautics, MIT, Cambridge, Massachusetts, 1995.

[73] H.D. Sherali, C.H. Pendyala, and T.S. Rappaport. Optimal location of transmitters for micro-cellular radio communication system design. *IEEE Journal on Selected Areas in Communications*, 14(4):662–673, May 1996.

[74] N. Srinivas and K. Deb. Multi-objective function optimization using non-dominated sorting genetic algorithms. *Evolutionary Computation*, 2:221–248, 1995.

[75] T. Starkweather, S. McDaniel, K. Mathias, D. Whitley, and C. Whitley. A comparison of genetic sequencing operators. In Rick Belew and Lashon Booker, editors, *Proceedings of the Fourth International Conference on Genetic Algorithms*, pages 69–76, San Mateo, CA, 1991. Morgan Kaufman.

[76] D-W. Tcha and Y-S. Myung. Base station location in a cellular cdma system. *Telecommunication Systems*, 14:163–173, 2000.

[77] K. Tutschku. Interference minimization using automatic design of cellular communication networks. In *IEEE Vehicular Technology Conference*, pages 634–638, 1998.

[78] C.L. Valenzuela. A simple evolutionary algorithm for multi-objective optimisation(SEAMO). In *IEEE Congress on Evolutionary Computation*, pages 717–722, 2002.

[79] David A. Van Veldhuizen and Gary B. Lamont. Evolutionary computation and convergence to a pareto front. In J.R. Koza, editor, *Late Breaking Papers at the Genetic Programming 1998 Conference*, pages 221–228, Stanford,CA, 1998. Stanford University Bookstore.

[80] David A. Van Veldhuizen and Gary B. Lamont. On measuring multiobjective evolutionary algorithm performance. In *Proc. of the 2000 Congress on Evolutionary Computation*, pages 204–211, Piscataway, NJ, 2000. IEEE Service Center.

[81] M. Vasquez and J-K. Hao. A heuristic approach for antenna positioning in cellular networks. *Journal of Heuristics*, 7:443–472, 2001.

[82] D.A. Veldhuizen and G.B. Lamont. On measuring multiobjective evolutionary algorithm performance. In *Congress on Evolutionary Computation*, pages 204–211, 2000.

[83] N. Weicker, G. Szabo, K. Weicker, and P. Widmayer. Evolutionary multiobjective optimization for base station transmitter placement with frequency assignment. *IEEE Transactions on Evolutionary Computation*, 7(2):189–203, 2003.

[84] R.M. Whitaker and S. Hurley. Omni-directional cell planning. In G. Anandalingam and S. Raghavan, editors, *Telecommunications Network Design and Management*, chapter 2, pages 25–41. Kluwer Academic Publishers, Netherlands, 2002.

[85] R.M Whitaker and S. Hurley. Evolution of planning for wireless communication systems. In *36th Annual Hawaii International Conference on System Sciences*, pages 1–10, January 2003.

[86] J.K.L. Wong, M.J. Neve, and K.W. Sowerby. Optimisation strategy for wireless communications system planning using linear programming. *IEE Electronics Letters*, 37(17):1086–1087, 2001.

[87] M.H. Wright. Optimization methods for base station placement in wireless systems. In *Proceedings of the IEEE VTC'98 Conference*, pages 387–391, 1998.

[88] T. Kawano Y. Okumura, E. Ohmori and K. Fukuda. Field strength and its variability in vhf and uhf land mobile service. *Rev. Elec. Comm. Lab*, 16:825–873, September-October 1968.

[89] J. Zimmermann, R. Hons, and H. Muhlenbein. ENCON: evolutionary algorithm for the antenna placement problem. *Computers and Industrial Engineering*, 44:209–226, 2003.

[90] E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation*, 8(2):173–195, 2000.

[91] E. Zitzler, M. Laumanns, and L. Thiele. Spea2:improving the strength pareto evolutionary algorithm. *Technical Report 103, Computer Engineering and Networks Laboratory(TIK), ETH Zurich, Switzerland*, 2001.

[92] E. Zitzler, M. Laumanns, L. Thiele, C. Fonseca, and G. da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *Technical Report 139, Computer Engineering and Networks Laboratory(TIK), ETH Zurich, Switzerland*, 2002.

[93] E. Zitzler and L. Thiele. Multiobjective optimization using evolutionary algorithms–a comparative case study. In *Parallel Problem Solving from Nature*, pages 292–301, 1998.

[94] Eckart Zitzler. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications.* PhD thesis, Swiss Federal Institute of Technology Zurich, Switzerland. TIK-Schriftenreihe Nr. 30, 1999.

Appendix A

Guide to Data Recovery

A.1 Simulation Environment

There are five simulation environments, termed SE1, SE2 ... SE5. Each simulation environment file is in ASCII format and provides the complete set of reception test points (RTP), service test points (STP), traffic test points (TTP), and candidate base stations (CBS), as provided in Figure A.1.

Figure A.1: Greenfield town network characteristics for five simulation environments

SE	#CBS	#RTP	#STP	#TTP	Traffic	Mesh(m)	Working Area(km)
SE1	568	56,792	17,393	8,590	2,988.27	200	49.4 x 45.6
SE2	568	56,792	17,393	3,985	2,988.75	200	49.4 x 45.6
SE3	568	56,792	17,393	6,602	3,221.84	200	49.4 x 45.6
SE4	568	56,792	56,792	11,953	2,988.25	200	49.4 x 45.6
SE5	568	56,792	56,792	56,792	2,988.12	200	49.4 x 45.6

The general layout of the file is:

Each plan begins with a header. The header gives the size of the working area (in number of test points); the distance between each test point (in meters); the

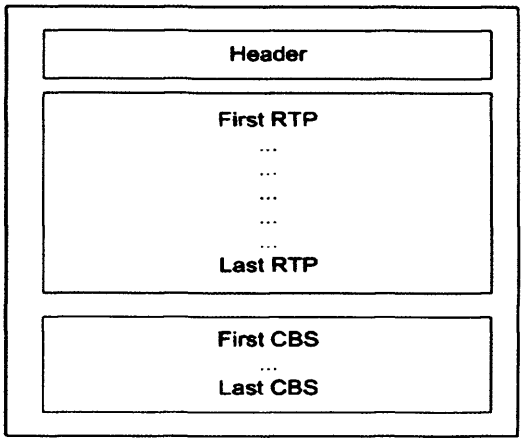

Figure A.2: Simulation environment file layout

number of RTP, STP, TTP, and CBS; and the total amount of traffic in the working area (in erlangs). For example:

Long seed: 1419697713325834287

Size of working area: 248 x 229

Number of RTP: 56792

Number of STP: 17393

Number of TTP: 8590

Number of CBS: 568

Total traffic: 2988.2662442999754

Following this, there is one line for each RTP after the guidance header:

******RTP#,X,Y,Z,STP,TTP ******

Each line, therefore, starts by providing the number of the given RTP (in the range

0 to (NumberOfRTP - 1)); its X, Y, and Z location (in meters); a 1 if this test point is also an STP (otherwise a 0); and then an amount of traffic–which also indicates this test point is a TTP (otherwise a 0). For example,

0 ...

14678,9200,11800,239,1,0.335545

14679,9400,11800,218,1,0.0

14680,9600,11800,233,1,0.0

... (NumberOfRTP - 1)

the second line indicates RTP number 14678 is at an x-axis location of 9200 and a y-axis location of 11800 at a height (z-axis) of 239 meters. This test point is an STP and TTP carrying 0.335545 erlangs of traffic. After this, there is a gap of one line, and then one line for each CBS after the guidance header:

****** CBS#,X,Y,Z ******

Each line, therefore, starts by providing the number of the given CBS (in the range 0 to (NumberOfCBS - 1)), followed by its X, Y, and Z location (in meters). For example,

0 ...

8,3200,8600,256

9,9600,36000,216

10,39800,24200,174

... (NumberOfCBS - 1)

the second line indicates CBS number 8 is at an x-axis location of 3200 and a y-axis location of 8600 at a height (z-axis) of 256 meters.

A.2 Propagation Loss Matrix

The propagation loss information is stored in ASCII format. The general layout of the file is:

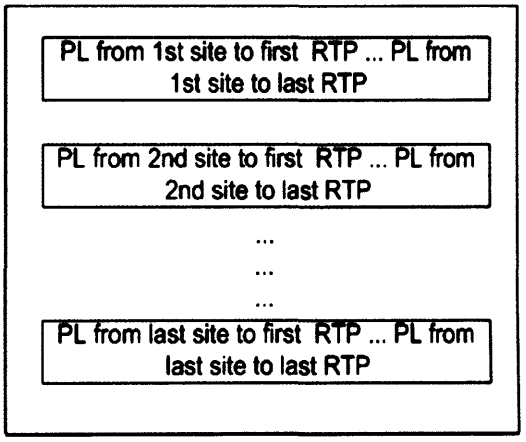

Figure A.3: Propagation loss file layout

For example, the propagation loss for the first three CBS could look as follows:

14187 14125 13896 14910 ... 14944 14523 13961

16493 16003 16849 15844 ... 16484 16466 16330

15193 16103 16589 16508 ... 16052 16178 16113

To read an entry, simply retrieve the value and divide it by 100. For example, the first entry of 14187 would be a propagation loss of 141.87.

A.3 Angle of Incidence Matrix

The angle of incidence (AI) information is stored in ASCII format. The general layout of the file is:

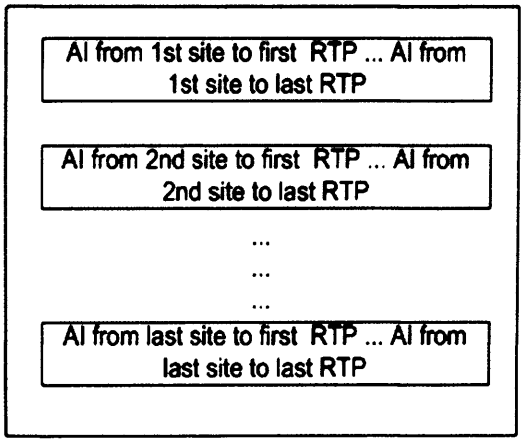

Figure A.4: Angle of incidence file layout

For example, the AI for three CBS could look as follows:

1246 1241 1235 1226 ... 1220 1211 1205

782 785 787 791 793 ... 797 799 803 807

1180 1181 1152 1152 ... 1153 1154 1154

To read an entry, simply retrieve the value and divide it by 100. For example, the first entry of 1246 would be an AI of 12.46.

A.4 Horizontal and Vertical Loss Diagrams

Horizontal diagrams are available for small and large directive antennas, and vertical diagrams are available for small, large, and omni-directional antennas. The data is stored in order from 0 to 359 degrees for each situation.

To read an entry, simply retrieve the value and divide it by 100. For example, an entry of 448 would be a loss of 4.48.