



BINDING SERVICES
Tel +44 (0)29 2087 4949
Fax +44 (0)29 20371921
e-mail bindery@cardiff.ac.uk

Optimising BFWA Networks

Thesis submitted to

School of Computer Science,
Cardiff University

for the degree of

Doctor of Philosophy

By

A.A. Wade

2005

UMI Number: U584742

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U584742

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



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

Acknowledgements

Many thanks and much appreciation to my supervisors, Stuart Allen, Ken Craig and Steve Hurley for their help and encouragement over the course of this work. Also to Richard Taplin, Isa Usman, Bill Miners, Silvia Cirstea, Nick Thomas and Mike Willis for all the questions answered and advice given at various times along the way. Finally many thanks to my parents and to Floyd, Rosie and the late, lamented Trodgie, who have provided very gratefully received financial and emotional support.

I gratefully acknowledge the Radio Communications Research Unit at the Rutherford Appleton Laboratory for financially sponsoring this work.

This thesis is dedicated to Tamsin, for putting up with my wittering for the last 3 years.

Abstract

Broadband Fixed Wireless Access (BFWA) networks are an attractive alternative to cable-based technologies, in offering low-cost, high-speed data services, telephony and video-on-demand to residential and business users. However, in order to compete successfully with available alternative telecommunications solutions, the planning and design of efficient networks is crucial.

This thesis presents two tools that enable the planning and evaluation of BFWA networks. AgentOpt is a network design and optimisation tool. A detailed account of the novel scheme, using the principles of emergent, self-organising systems, which AgentOpt employs for finding profit-optimal networks is given. The use of two distinct types of agent entity allows the multi-objective profit/coverage nature of the network planning problem to be satisfied. AgentOpt networks are compared with designs produced by other methods to establish to what extent this decentralised agent approach can optimise BFWA networks.

The Network Validation Tool (NVT) analyses the network designs produced by AgentOpt and other automatic cell planning tools (ACPs). This is achieved through simulating the subscription take-up of the potential users in the network. By repetition of this process, statistical data about the various design configurations of the network is produced. This allows a planning engineer to compare and contrast network solutions that may differ in design but perform similarly in terms of expected profit. In this work the NVT is used to formulate some general guidelines about the best-practice use of ACPs.

Key-words: BFWA, the network planning problem, automatic cell planning, heuristic optimisation, emergent intelligence, wireless network optimisation, wireless network evaluation.

Contents

Abstract	i
Contents	ii
List of Figures	vii
List of Tables	xii
Glossary of Terms	xiii

Chapter 1. Introduction to BFWA	page 1
1.1 The Last-Mile Problem.	1
1.1.1 Cable-based Systems	2
1.1.2Wireless-based Systems	4
1.2 Why Wireless?	9
1.3 The EMBRACE Project	10
1.4 Thesis Outline	12

Chapter 2. Cell Planning Techniques	page 14
2.1 The Cellular Network Planning Problem	14
2.1.1 BSPP	15
2.1.2 BSCP	16
2.1.3 FAP	16
2.2 Cell Planning Techniques	17
2.2.1 Manual Planning	17
2.2.2 Simple Cell Planning Strategies and Methods	20
2.2.3 Integrated Planning Tools	23
2.2.4Other Approaches	28

2.3 The Frequency Assignment problem	30
2.4 Propagation and Building Data	31
2.4.1 Building the 3-D model	32
2.4.2 Propagation Models	35
2.5 The ECHO tool	38
2.5.1 Metrics	38
2.5.2 Constraints	38
2.5.3 User Lists	39
2.5.4 The Network design algorithms	39
2.5.5 Multi-objective considerations	43
2.5.6 Pareto Watchers	44

Chapter 3. Heuristic Techniques for combinatorial optimisation

problems	page 45
3.1 Combinatorial Problems	45
3.2 NP-Completeness	46
3.3 Heuristics	48
3.3.1 Local Neighbourhood Search	48
3.3.2 Simulated Annealing	49
3.3.3 Genetic Algorithms	51
3.3.4 Tabu Search	52
3.4 Multi-Objective Optimisation	54
3.4.1 Weighting Method	55
3.4.2 Constraint Method	56
3.4.3 The need for better algorithms	57
3.5 Emergent Systems	57
3.5.1 The ACO heuristic	59
3.5.2 Applications of ACO	61
3.5.3 Particle swarm Optimisation	62

3.5.4 Agent Technology	62
3.6 Conclusions	64
Chapter 4. Modelling BFWA	page 65
4.1 The network model	65
4.2 The problem size	72
Chapter 5. AgentOpt	page 73
5.1 Adapting emergent intelligence techniques to cell planning	73
5.2 The AgentOpt scheme structure	75
5.3 User-Agents	77
5.3.1 Neighbours	79
5.3.2 Distraction	80
5.3.3 Resentment	81
5.3.4 Subscription charge ratio	82
5.3.5 Changing sites	83
5.3.6 Dormancy	84
5.4 Site-Agents	86
5.4.1 The Site-Phase	87
5.4.2 Is a new sector required?	88
5.4.3 Mergeing same site sectors	89
5.4.4 Improving sectors	91
5.4.5 Reducing the cost of sectors	92
5.4.6 Increasing coverage	93
5.4.7 Improving the channel assignment	94
5.4.8 Increasing capacity	95
5.4.9 Destroying sectors	95
5.4.10 Merging sectors from different sites	96
5.4.11 Reassessment of each Site-Agent	97

5.5 Evaluation of the effect of system parameters on AgentOpt	98
5.5.1 Starting the optimisation	98
5.5.2 The neighbourhood radius	101
5.5.3 The petition threshold	104
5.5.4 The coverage threshold	107
5.5.5 Randomness	109
5.6 Conclusions	110

Chapter 6. Evaluating BFWA Networks **page 111**

6.1 The problem with probability	111
6.1.1 The AgentOpt approach	111
6.1.2 The planning process	113
6.2 The Network Validation Tool	114
6.2.1 The validation approach	114
6.2.2 The up-link architecture	115
6.2.3 NVT design	117
6.2.4 Uplink interference checking	120
6.2.5 Post processing interference analysis	122
6.3 NVT validation	123
6.3.1 NVT output	124
6.3.2 An example network	126
6.3.3 Randomness	127
6.3.4 Length of validation	128
6.3.5 Power equalisation	129
6.3.6 Path-loss errors	130
6.4 Conclusions	131

Chapter 7. Evaluating AgentOpt networks **page 132**

7.1 The AgentOpt networks	132
---------------------------	-----

7.2 Comparisons with manual plans	135
7.3 Improving the AgentOpt networks	140
7.4 Comparisons with ECHO generated networks	143
7.5 Conclusions	146
Chapter 8. Investigations into BFWA networks	page 148
8.1 Subscription probabilities and take-up rates	148
8.2 Availability and quality of service	159
8.3 Maximum channel numbers	165
8.4 Capacity requirements	168
8.5 Return periods	172
8.6 Diversity	179
8.7 Known-subscription networks	183
8.8 Conclusions	190
Chapter 9. Conclusions and further work	page 192
9.1 AgentOpt	192
9.2 NVT	193
Chapter 10. Bibliography	page 196
Appendix I.	page 204
Appendix II.	page 206
Appendix III.	page 208
Appendix IV.	page 209

List of Figures

Figure 1-1. A BFWA network example.....	page 7
Figure 2-1. A BS is placed on the highest ground.....	18
Figure 2-2. Then at the second highest.....	18
Figure 2-3. ...and the third.....	19
Figure 2-4. The Best-Effort plan.....	19
Figure 2-5. A simple illustration of the hexagonal cell system.	
Each different colour represents a different frequency.....	21
Figure 2-6. A 3-D model of the town of Malvern (UK).....	32
Figure 2-7. A section of the layered Malvern dataset.....	33
Figure 2-8. Each vertical wall has an outline in its respective layer.....	33
Figure 2-9. The roof ridges are stored in a separate layer.....	34
Figure 2-10. A visual representation of path-loss levels.....	35
Figure 2-11. A path-loss profile.....	36
Figure 2-12. The 3-D model.....	37
Figure 2-13. The filter input.....	37
Figure 2-14. The ECHO Tool.....	38
Figure 2-15. The ECHO Tabu-Search Algorithm.....	42
Figure 2-16. A Pareto front as found by ECHO.....	44
Figure 3-1. A simple Local Neighbourhood Search algorithm.....	49
Figure 3-2. An example of the Simulated Annealing algorithm.....	50
Figure 3-3. A simple Genetic algorithm.....	51
Figure 3-4. The Tabu-Search algorithm.....	53
Figure 3-4. Ants may have a choice of paths from the nest to the food source.....	60
Figure 5-1. The flow of local information through the AgentOpt system.....	76
Figure 5-2. A pseudo-code representation of the User-Phase.....	79
Figure 5-3. An example of a) a user with no alternative,	
and b) a user surrounded by sites (though not sectors).....	81
Figure 5-4. The Site-Phase algorithm.....	88
Figure 5-5. The segment chosen by a Site-Agent.....	88

Figure 5-6. A pseudo-code representation of the sector improvement method.....	91
Figure 5-7. A comparison of six networks produced using three different starting strategies.....	99
Figure 5-8. A random start network cell plan.....	100
Figure 5-9. A best path-loss network cell plan.....	100
Figure 5-10. The performance of each starting method over time.....	100
Figure 5-11. The neighbourhood radius.....	101
Figure 5-12. A comparison of networks generated using different radius values.....	102
Figure 5-13. Petition Threshold comparisons.....	104
Figure 5-14. Petition Threshold - Comparison of two optimisation runs over time.....	105
Figure 5-15. The PT=100 network.....	105
Figure 5-16. The PT=750 network.....	106
Figure 5-17. After 400 iterations.....	107
Figure 5-18. User-Agents petitioning Site-Agents after 400 iterations.....	107
Figure 5-19. A comparison of different Coverage Threshold levels.....	108
Figure 5-20. A comparison of different levels of randomness within the optimisation....	109
Figure 6-1. The Network Validation Tool.....	118
Figure 6-2. The main validation algorithm.....	119
Figure 6-3. The C2I checking algorithm.....	120
Figure 6-4. A Malvern network.....	126
Figure 6-5. Differences in random seeds.....	127
Figure 6-6. Levels of Expected Profit.....	128
Figure 6-7. The Number of valid users as a percentage of the total.....	128
Figure 6-8. Measures of income and outlay, with and without Power Equalisation.....	129
Figure 6-9. As the error increases the profit decreases.....	130
Figure 6-10. Invalid users mean less profit.....	131
Figure 7-1. Cell plan for network AO-1.....	133
Figure 7-2. Cell plan for network AO-2.....	133
Figure 7-3. Cell plan for network AO-3.....	133
Figure 7-4. A comparison of the profit over time throughout the three runs.....	134
Figure 7-5. Cell plan for network M-1.....	136
Figure 7-6. Cell plan for network M-2.....	137
Figure 7-8. A comparison of the expected profit levels.....	137
Figure 7-9. A comparison of the service levels.....	138
Figure 7-10. A comparison of expected income, outlay and profit.....	139

Figure 7-11. A comparison of the percentage of capacity supplied that is utilised.....	139
Figure 7-12. Cell plan for network AO-1-E.....	141
Figure 7-13. Cell plan for network AO-2-E.....	141
Figure 7-14. Cell plan for network AO-3-E.....	141
Figure 7-15. Cell plan for network E-1.....	144
Figure 7-16. Cell plan for network E-2.....	144
Figure 7-17. Cell plan for network E-3.....	144
Figure 7-18. A comparison of the expected profit levels of the improved AO and ECHO generated networks.....	146
Figure 7-19. A comparison of the percentage of capacity utilised in each network.....	146
Figure 8-1. Network 10%-1.....	150
Figure 8-2. Network 10%-2.....	150
Figure 8-3. Network 10%-3.....	150
Figure 8-4. Comparison of expected profit at different take-up levels.....	150
Figure 8-5. Comparison of users served at different take-up levels.....	151
Figure 8-6. Comparison of capacity used at different take-up levels.....	151
Figure 8-7. Network 20%-1.....	152
Figure 8-8. Network 20%-2.....	152
Figure 8-9. Network 20%-3.....	153
Figure 8-10. Expected profit over different take-up levels.....	154
Figure 8-11. Number of invalid users due to lack of capacity.....	155
Figure 8-12. Percentage of valid users in the network.....	155
Figure 8-13. Comparison of expected profit levels at 20% subscription.....	156
Figure 8-14. Comparison of outlay at 20% subscription.....	157
Figure 8-15. Comparison of valid users at 20% subscription.....	157
Figure 8-16. Comparison of post-processing C2I errors at 20% subscription.....	158
Figure 8-17. A comparison of expected profit at 10% subscription.....	159
Figure 8-18. A comparison of expected profit at 30% subscription.....	159
Figure 8-19. A comparison of expected profit over differing Av levels.....	161
Figure 8-20. A comparison of service percentage over differing Av levels.....	161
Figure 8-21. A comparison of expected profit after NVT analysis.....	163
Figure 8-22. Expected profit levels for each network validated at different Av levels.....	163
Figure 8-23. The number of users removed due to up-link difficulties as a percentage of the total removed, at different Av levels.....	164

Figure 8-24. Number of valid users as a percentage of those chosen, at different Av levels.....	165
Figure 8-25. A comparison of expected profit levels for networks designed with differing numbers of channels.....	167
Figure 8-26. A comparison of the NVT predicted profit levels for each network.....	167
Figure 8-27. A comparison of the profit expected from each network at varying capacity requirements.....	171
Figure 8-28. A comparison of the percentage of supplied capacity that is used by each network at varying capacity requirements.....	171
Figure 8-29. A comparison of the profit expected and the % of users served when the networks are analysed with a take-up rate of 30%.....	172
Figure 8-30. A comparison of the number of users removed and that have PPI errors when the networks are analysed with a take-up rate of 30%.....	173
Figure 8-31. Network SRP-1.....	174
Figure 8-32. Network SRP-3.....	174
Figure 8-33. Network SRP-5.....	175
Figure 8-34. The optimisation of network SRP-1.....	176
Figure 8-35. The optimisation of network SRP-3.....	176
Figure 8-36. The optimisation of network SRP-5.....	176
Figure 8-37. Network MRP-1.....	177
Figure 8-38. Network MRP-3.....	178
Figure 8-39. Network MRP-5.....	178
Figure 8-40. A comparison of expected profit for each Sevenoaks scenario network at the different return periods.....	179
Figure 8-41. A comparison of expected profit for each Malvern scenario network at the different return periods.....	180
Figure 8-42. A comparison of expected profit and diversity percentage for each network.....	183
Figure 8-43. A comparison of profit predicted for each network at three different subscription take-up levels.....	183
Figure 8-44. Network CN-1.....	185
Figure 8-45. Network CN-2.....	186
Figure 8-46. Network CN-3.....	186
Figure 8-47. A comparison of expected profit values.....	187
Figure 8-48. A comparison of invalid user numbers.....	187

Figure 8-49. Network PN-1.....	188
Figure 8-50. Network PN-2.....	188
Figure 8-51. Network PN-3.....	189
Figure 8-52. A comparison of the expected profit across the different groups of users.....	189
Figure 8-53. A comparison of the percentage of possible revenue collected by each network across the different user groups.....	190
Figure 8-54. A comparison of the percentage of possible subscribers served by each network across the different user groups.....	190
Figure AI-1. A map of the Malvern scenario. Potential base sites are shown as white dots.....	204
Figure AII-1. A map of the Sevenoaks scenario. Potential base sites are shown as white dots.....	206
Figure AIV-1. Network 30%-1.....	210
Figure AIV-2. Network 30%-2.....	210
Figure AIV-3. Network 30%-3.....	210
Figure AIV-4. Expected Profit at different levels of subscription.....	210
Figure AIV-5. A measure of invalid users at different subscription levels.....	211
Figure AIV-6. The percentage of valid users at different subscription levels.....	211
Figure AIV-7. Network 40%-1.....	212
Figure AIV-8. Network 40%-2.....	212
Figure AIV-9. Network 40%-3.....	213
Figure AIV-10. Expected profit at different subscription levels.....	213
Figure AIV-11. Percentage of valid users at different subscription levels.....	214
Figure AIV-12. Initial and final errors at a 40% take-up rate.....	214
Figure AIV-13. Network 50%-1.....	215
Figure AIV-14. Network 50%-2.....	215
Figure AIV-15. Network 50%-3.....	216
Figure AIV-16. A comparison of profit expected by NVT at different subscription levels.....	216
Figure AIV-17. A comparison of valid user percentage over different subscription levels.....	217
Figure AIV-18. The number of users made invalid through various means.....	217

List of Tables

Table 5-1. Average user numbers.....	page 103
Table 6-1. ECHO predicted network metrics.....	126
Table 7-1. Three AgentOpt generated networks.....	132
Table 7-2. The NVT analysis of the three AgentOpt networks.....	135
Table 7-3. The manually constructed networks.....	136
Table 7-4. The NVT analysis of the manually planned networks.....	138
Table 7-5. The fine-tuned AO networks.....	140
Table 7-6. The NVT analysis of the improved AO networks.....	142
Table 7-7. The ECHO generated networks.....	143
Table 7-8. The NVT Analysis of the ECHO generated networks.....	145
Table 8-1. Networks generated at a 10% take-up rate.....	149
Table 8-2. Networks designed for 20% take-up.....	152
Table 8-3. A survey of removed users at 20% actual take-up.....	154
Table 8-4. A comparison of capacity issues.....	168
Table 8-5. Networks planned for varying traffic demands.....	170
Table 8-6. The Sevenoaks RP networks.....	175
Table 8-7. The Malvern RP networks.....	178
Table 8-8. The low diversity networks.....	182
Table 8-9. The higher diversity networks.....	182
Table 8-10. The randomly generated user groups.....	185
Table 8-11. The <i>certain</i> networks.....	185
Table AI-1. The Malvern scenario in numbers.....	205
Table AII-1. The Sevenoaks scenario in numbers.....	207
Table AIII-1. The AgentOpt initialisation parameters.....	208
Table AIV-1. Networks designed for 30% take-up.....	209
Table AIV-2. Networks designed for 40% take-up.....	212
Table AIV-3. Networks designed for 50% take-up.....	215

Glossary of Terms

ACP	Automatic Cell Planning
ADSL	Asymmetric Digital Subscriber Line
AO	AgentOpt
BFWA	Broadband Fixed Wireless Access
BPL	Broadband over Power Lines
BRAN	Broadband Radio Access Networks
BS	A <i>Base Station</i> (BS) is a collection of infrastructure, such as a generator and a mast, at a particular site. A BS houses antennae which can transmit a signal to a limited portion of the network.
BSCP	Base Station Configuration Problem
BSPP	Base Station Placement Problem
Cell	See Sector
CNPP	Cellular Network Planning Problem
Coverage	A covered user must receive a signal from a transmitting antenna, such that the received signal strength exceeds a given threshold.
DSL	Digital Subscriber Line
ECHO	Efficient Cell-planning Heuristic Optimiser
EMBRACE	Efficient Millimetre Broadband Radio Access for Convergence and Evolution
ETSI	European Telecommunications Standards Institute

FAP	Frequency Assignment Problem
FDD	Frequency Division Duplex
FWA	Fixed Wireless Access
GA	Genetic Algorithm
HIPERACCESS	High Performance Radio Access
IEEE	Institute of Electrical and Electronics Engineers
ISDN	Integrated Services Digital Networks
IST	Information Society Technologies (EU)
LMDS	Local Multi-point Distribution Service
LNS	Local Neighbourhood Search
MOP	Multi Objective Problem
NVT	Network Validation Tool
PLT	Power Line Telecommunications
RAL	Rutherford Appleton Laboratory
RCRU	Radio Communications Research Unit
RPD	Rapid Pipeline Development
RTP	Receiver Test Point
SA	Simulated Annealing
SDSL	Symmetric Digital Subscriber Line
Service	A <i>served</i> user must be covered, must also have been allocated a block of capacity within the network and must have a signal-to-interference ratio that exceeds a given threshold.
Sector	A <i>sector</i> , or a <i>cell</i> , can be defined as a combination of a BS, a particular antenna housed at that BS and a set of users that can be associated with (are either covered by or served by) that antenna.
SHDSL	Symmetric High-speed Digital Subscriber Line
SOP	Single Objective Problem

TS	Tabu-Search
TDD	Time Division Duplex
Wi-Fi	wireless networks using IEEE 802.11 standard
WLAN	Wireless Local Area Networks
VDSL	Very high-speed Digital Subscriber Line

Chapter 1

Introduction to BFWA

This thesis develops a system to automatically plan Broadband Fixed Wireless Access (BFWA) networks. Through the development of a software tool for the evaluation of automatically planned networks, general rules about the ways these planned networks behave are formed and utilised in order to improve the final networks selected for implementation.

This thesis also presents a novel optimisation scheme, utilising the principles of emergent intelligence, as an alternative to existing automatic cell planning tools. Alternative schemes are compared and contrasted and the networks produced by the new scheme are subjected to investigations.

1.1 The Last-Mile Problem

The *Last-Mile* problem can be simply defined as “How can we move data from the high speed fibre optic cables (known as the *core* or *backhaul* network) that run the length and breadth of the country, to the homes and workplaces of the people who require high-speed data access?” The last twenty years has seen networked personal and business computing become ubiquitous, and the continued expansion of a data transfer economy requires ever increasing broadband data transfer rates. This thesis informally defines a broadband connection as an always on connection delivering speeds of over 56kbps. Both a down-link (for data from backhaul network to customer) and an up-link (for data from customer to network) must be provided, usually in the same manner. The backhaul network is commonly considered to be operating at less than it’s capacity, because of the data bottle-neck formed around the last

miles. Therefore, the demand for higher and higher data rates is likely to continue to expand to fill whatever is available. This problem of under-utilisation or wasted capacity continues to be an annoyance to telecommunications companies and governments alike. There are at this point in time two general categories of technology that are employed to solve this problem – *cable* and *wireless*.

1.1.1 Cable-based Systems

The vast majority of existing *last-mile* broadband access (or high bandwidth data connection) in the UK is through some form of physical connection or cable. This section gives a brief survey of the technologies in use and some of the advantages and disadvantages associated with them.

XDSL

One of the most common, certainly amongst residential users, is xDSL (*Digital Subscriber Line*). This was originally designed for delivery of video content over the existing analogue, twisted copper telephone lines. It has now been adapted to deliver data at high speed in an 'always on' broadband connection, primarily to the residential and small business user. There are a number of different types of DSL, the most widely implemented being Asymmetric DSL. In the UK, ADSL provides data rates of up to 8Mbps on the down-link and 512kbps on the up-link¹. This service now reaches 90% of UK households², with plans for 99% penetration in 2005. At the moment ADSL can be supplied as long as the existing phone line is less than 6 km in total length from an enabled (i.e. modern fibre-optic cable connected) telephone exchange, although BT is planning to relax this requirement.

¹ The commercial packages available usually restrict down-link speeds to <3Mbps

² <http://www.bt.com/broadband/>

Obviously, the main benefit of xDSL is that in using existing twisted pair telephone cable, the network is already installed. This gives minimal install and disruption times (e.g. digging up the road) in connection to individual users, which is an important consideration. Indeed the majority of users are able to install the technology themselves.

xDSL comes in a number of different versions. SDSL (Symmetric) is being developed to cope with users who require large amounts of data to be regularly both sent and received, such as in a business environment. There are also rarer forms such as SHDSL (Symmetric High-speed) and VDSL (Very high-speed), which relies on fibre-optic cabling.

ISDN

ISDN (Integrated Services Digital Networks) is essentially a dedicated digital data line, allowing simultaneous transmission of voice and data. The digital line replaces the old analogue line providing up to 128kbps. ISDN has been in circulation since the mid 1980's in one form or another and is now generally seen as having been superseded by xDSL and Cable.

CABLE

Fibre-Optic cabling forms the backbone of the existing telecommunications networks and also has a place in solving the last mile problem. As well as the connections between telephone exchanges, fibre-optic cable is already installed throughout a lot of UK towns and cities, generally to a single box (or node) in most streets and then on into individual buildings, providing subscription television services and data. The two main cable television companies in the UK offer a variety of broadband data services ranging from a basic 1Mbps service up to 3Mbps. Well over one million residential users are connected to some form of cable broadband service in the UK, with 40% of the

population having access to this technology³. For businesses close to data exchanges there is always the choice of having a dedicated cable line, although this is likely to be very expensive and potentially disruptive depending upon the location.

PLT

PLT (Power Line Telecommunications), sometimes also called BPL (Broadband over Power Lines) offers a more novel approach to solving the last-mile problem. The system uses existing power lines to transmit and receive data at broadband speeds. The technology is still being developed and a number of problems need to be overcome, but the obvious advantage with this approach is the ubiquity of the infrastructure. The approach is somewhat controversial however, especially among users of the frequency bands set aside for amateur radio enthusiasts. The majority of power lines are un-insulated, making interference a large problem and complaints about the technology have been made from a number of quarters, from radio clubs to national broadcasters. It remains to be seen whether the technology can be developed into a viable and marketable product.

1.1.2 Wireless- based systems

Given the superior market-penetration of the cable-based technologies at the moment, and their generally higher level of technological maturity, how do wireless data provision technologies, still largely in their infancy, fit into the Last-Mile picture? How can they expect to compete within the data-provision market, already well populated with established infrastructure and tried and tested technology? There are a number of different wireless technologies, at various stages in their development, and each with their own advantages and disadvantages.

³ <http://news.bbc.co.uk/1/hi/technology/3309543.stm>

SATELLITES

Satellite data provision systems are, at present, generally hybrids. Although the growth of satellite television has furthered the receive-only technology and also reduced prices, affordable systems can usually only provide broadband access on the down-link. All up-link traffic must travel through another technology, probably cable-based. The physical distances involved must also be considered. A request for certain data must travel firstly to the service provider, then up to the satellite and back down again, leading to slight ($\approx 1/4$ of a second) time delays. This, in turn, can lead to receive-confirmation problems, unless special satellite-adapted transmission-reception schemes are implemented. Also the cable up-link connection can add significant latency. These delays are likely to frustrate users of real-time interaction systems, such as on-line gaming or video conferencing. However, two big positives for satellites are the ease of installation at the user premises and the very wide range of coverage - both particularly useful in remote areas.

WI – FI

Wi-Fi, or wireless local area networks (WLANs) are widely in use today. The IEEE 802.11⁴ standard was originally designed to replace cable based Ethernet and other LANs within a single office block or factory. The Wi-Fi frequency bands (in Europe and America the 900 MHz, 2.4 GHz and 5.8 GHz bands) are unlicensed in most countries. This means that Wi-Fi particularly lends itself to operation by small companies and individuals. In addition to this Wi-Fi *hotspots* can be set up in areas such as airports, hotels or cafes, allowing nomadic users high speed data access. Depending on the particular implementation, each Wi-Fi base station operates at around 10Mbps providing comfortable internet access for around 50 people.

⁴ There are a number of versions of the standard, with 802.11.b being the most widely used.

The success of this idea has led to rural communities (and companies contained within dispersed offices) providing their own broadband networks through the use of the 802.11 technology. This may be achieved through the use of a number of Wi-Fi base stations, connected by point to point radio links, and finally connected to the backhaul network through either cable or another microwave link [1]. This is obviously a very useful aspect of the technology especially for small rural communities who are unlikely to be within reach of any cable-network.

There are, of course, some negatives to be considered with the implementation of the 802.11 standard. As the Wi-Fi technology allows unlicensed use, in dense urban or commercial areas there can be interference problems with neighbouring Wi-Fi and other wireless networks (such as Bluetooth). There can be problems with spectral inefficiency, and security issues must be closely attended to. Also, because of the frequencies and transmission power levels that the technology uses, the range of the Wi-Fi base stations are limited to a few hundred meters at maximum power, making it unsuitable (due to high infrastructure costs and general inefficiency) for networks that are geographically large but contain very widely dispersed users with relatively low data demands, as well as networks with densely packed high data demand users.

OPTICAL WIRELESS

Optical wireless is starting to be seen as another feasible alternative to provide a last-mile solution. This technology can provide a dedicated link, using light to transmit data between a user access location and the local exchange or backhaul network access point. The link length is limited to around 1 km and of course must have line of sight (LOS), but for some self-contained office blocks or buildings, such as an out-of-town business park, this need not represent a problem. Another advantage to the optical wireless

solution is its seamless integration with optical fibre, because of the fundamental similarity between the two technologies.

BFWA

Broadband Fixed Wireless Access systems (also called LMDS, or Local Multipoint Distribution Service) are usefully summed up by their name. They aim to provide a high bandwidth wireless service using fixed-position antennas. In the context of the Last-Mile problem, these networks are point-to-multipoint systems, operating in the microwave spectrum band, although point-to-point microwave links are in use as backhaul carriers and may also be considered as BFWA systems. Some BFWA systems may include repeater stations in order to get signals into difficult areas.




Figure 1-1. A BFWA network example

There are a number of standards in development across the microwave FWA field. Wi-Max is the name of a new BFWA IEEE infrastructure and practices standard - 802.16⁵. The Wi-Max standard covers the 2-66 GHz frequency range, although the hardware adhering to the standard will be restricted to a

⁵ <http://grouper.ieee.org/groups/802/16/>

certain band within that range. IEEE has also created a new working group, 802.20⁶, whose aim is "to enable worldwide deployment of affordable, ubiquitous, always-on and interoperable multi-vendor mobile broadband wireless access networks that meet the needs of business and residential end user markets." This system will operate in bands lower than 3.5GHz, and whilst the standard is being developed for mobile end-users, it may well be adapted for fixed purposes in a similar way to 802.11 (Wi-Fi).

The ETSI (the European Telecommunications Standards Institute) project, named BRAN⁷ (Broadband Radio Access Networks) has also been creating BFWA standards for infrastructure and software. HIPERACCESS (High Performance Radio Access) covers frequencies above 11 GHz, and is "intended for point-to-multipoint, high speed access (25 Mbps typical data rate) by residential and small business users to a wide variety of networks including the UMTS core networks, ATM networks and IP based networks"⁸, although as it was approved after 802.16, it may already have lost ground to the IEEE standard. HIPERLAN1 & 2 are other ETSI standards, this time for frequencies around the 5 GHz range, and aimed primarily at smaller WLANs.

MESH NETWORKS

Mesh networks are a variant of BFWA using a similar technology base but enabling a user to receive service even if the link to the base station is not line-of-sight. The technology, developed by Radiant Networks⁹ in the UK and Caly Networks in the USA, relies on the equipment at each user's premises to provide service to every other. This means a very robust network, with the ability to guarantee very high availability levels, as any interference or disruptions on one route from transmitter to user can be avoided by a dynamic alteration of the path. High costs and under-utilisation are the

⁶ <http://grouper.ieee.org/groups/802/20/>

⁷ http://portal.etsi.org/portal_common/home.asp?tbkey1=BRAN

⁸ <http://portal.etsi.org/bran/kta/Hiperlan/hiperlan1.asp>

⁹ <http://www.radiantnetworks.com>

obvious disadvantages of using the high specification and adaptable equipment.

1.2 Why Wireless?

All these Wireless-based systems; BFWA, Hybrid-Satellite (to a lesser extent), WLANs, Mesh networks – have a number of advantages over cable-based systems. The ability for a simple and quick roll-out of the required network infrastructure leads to a faster realisation of revenue for the service provider. The infrastructure is also generally cheaper than miles of cable, certainly when the cost of the installation is included. In as yet un-cabled areas, such as rural areas, underdeveloped countries, etc, this is obviously of great benefit. However wireless systems have further advantages. Satellite and BFWA can also be used for broadcast in the same way as cable, but the networks (if the equipment at the user site remains the property of the provider) are much more resistant to costs incurred through turnover of customers, known as churn. There is also an inherent flexibility in the rollout of the network. Portions of the network do not have to be in place until there is sufficient demand.

It may be that Broadband Fixed Wireless Access systems can provide a cheaper, more adaptable and more flexible solution to the 'last mile' problem than cable, Wi-Fi, satellite or other wireless networking technologies. That will depend a great deal on the equipment developed and the standards implemented across the field. However it is not within the remit of this thesis to exhaustively compare and contrast these competing technologies. Rather, this thesis aims to show how, through the use of software tools, these BFWA networks can be planned to enhance their efficiency, adaptability and ultimately their profitability and competitiveness.

1.3 The EMBRACE project

EMBRACE¹⁰, Efficient Millimetre Broadband Radio Access for Convergence and Evolution, was an Information Society Technologies (IST) project funded within the European Commission Fifth Framework, running from Jan 2000 to Dec 2002. The project aimed to build on the ACTS project, specifically AC215 CRABS¹¹, which had developed and demonstrated an LMDS BFWA system specification, produced a propagation prediction tool and performed a number of trials including user surveys, technical tests and propagation measurements. There were many partners involved in the EMBRACE project, based all over Europe, including Telenor (Norway), Joanneum Research, The University of Salzburg (Austria), Cardiff University, the Rutherford Appleton Laboratory and Telewest Ltd (UK). A full list of contributors and background information on the project can be found on the project website.

The central objective of EMBRACE was to develop a low-cost and efficient BFWA solution to the last mile problem suitable for the mass-market. In order to successfully compete with cable-based technologies, the EMBRACE BFWA network takes advantage of the particular wireless network properties that have been detailed in the last two sections. The 'convergence' included within the EMBRACE acronym indicates that the system created within the project would need to provide users with digital television, video on demand, high speed internet access and voice communications - in short all data requirements within one system. Obviously this is an advantage over the xDSL and ISDN competitors, at least for potential residential subscribers. To fulfil this specification, the EMBRACE BFWA system was designed as an asymmetric interactive broadcasting network with point to multi-point distribution (or down-link) and point to point return channel (up-link)

¹⁰ <http://www.telenor.no/fou/prosjekter/embrace/>

¹¹ <http://www.telenor.no/fou/prosjekter/crabs>

offering an efficient and sophisticated method to merge the individual data demands. A full specification of the technical aspects of the EMBRACE network can be found in [2].

The proposed EMBRACE system operates specifically within the 40.5 to 43.5 GHz frequency band, but is also designed to be applicable to other millimetre bands, above 20 GHz, which may be allocated to this type of service. There were two main reasons for specifying the 40GHz band. Firstly, at the start of the project, this band had recently been allocated to BFWA usage in Europe. Secondly the technology surrounding the 40GHz millimetre band had matured to a point where the infrastructure required would be in mass production and prices would start to fall.

A mass-market fixed-wireless solution to the last mile problem, of the type that EMBRACE prescribed, indicated a network where both business and residential users requirements could be accommodated and every building within the network coverage area, be it a house, an office block or a factory, could be offered a service if required. If every different type of user must be able to be served, then the data capacity rates that are provided by the network must have enough flexibility to attract the whole range of users. This means, in urban areas, and to a lesser extent suburban areas, a high capacity, geographically dense network. In rural cases it would most probably imply quite the opposite. It must also be considered that if the network is to be run successfully by a provider it must have some ability to evolve, to grow as capacity demands and user numbers increase.

In order to achieve such a network whilst retaining profitability for the service provider, complex cellular coverage patterns and innovative spectrally efficient frequency allocation schemes must be carefully planned, bespoke, for each particular implementation. To this end, some software tools were

developed for EMBRACE – a radio wave propagation prediction tool, and a cellular network planning tool.

1.4 Thesis Outline

The thesis starts by giving an account of the problems to be considered in automatic cell planning and details previous work in this area. This includes brief discussions of manual planning methods, of the early analytical and geometric approaches that led to modern propagation prediction tools, of the frequency assignment problem and of integrated tools that attempt to automatically generate a feasible, or indeed optimal network. The chapter then briefly details the two software tools produced under the auspices of the EMBRACE project, that have formed the starting point for the work in this thesis. RPD (Rapid Pipeline Development) is a 3D ray-tracing propagation modelling tool, produced at the Rutherford Appleton Laboratory, whilst ECHO (Efficient Cell-planning Heuristic Optimiser), developed at Cardiff University, is a network planning tool that aims to produce network designs within operator defined constraints.

There is an over-view of combinatorial problems, into which category the Cellular Network Planning Problem falls. Following this there is a discussion of the heuristic techniques employed in the literature to tackle the problem, and a view of other methods with particular relevance to this thesis.

The component parts of the BFWA system are formally defined with particular reference to the problem of planning of such a network. This allows the Cellular Network Planning Problem to be formally defined.

AgentOpt is an automatic network planning tool, developed for this thesis, which incorporates a novel Emergent Intelligence optimisation scheme. The

optimisation method developed is discussed at length and experimental results are displayed to further the explanations.

The Network Validation Tool (NVT) is introduced. This is a software tool which has been developed in this thesis in order to evaluate the networks produced by the AgentOpt and ECHO tools. As both network planning tools rely upon various statistical probabilities in the planning procedure, the NVT is needed to establish the operational performance of the generated networks.

Using analysis from the NVT, networks generated by the AgentOpt tool are compared and contrasted with networks generated by other methods. Ways to improve the networks are discussed and some measure of the value of the AgentOpt generated networks is formulated.

A range of different BFWA scenarios are investigated, in which otherwise similar near-optimal networks are shown, by the use of NVT, to have very different characteristics. This leads to the development of a number of strategies that enable planning engineers to make the most of tools such as AgentOpt and ECHO.

Finally, an overview of the results discussed earlier in the thesis and the conclusions drawn from them is presented. Ideas for the improvement of the AgentOpt optimisation scheme are proposed and the evolution of the NVT is discussed along with more general comments on the furtherance of automatic cellular network planning.

Chapter 2

Cell Planning Techniques

This chapter starts with an informal definition of the Cellular Network Planning Problem and its constituent parts. This is followed by a review of the previous work completed within the area of cellular network planning. This is followed by a brief outline of the ECHO automatic network planning tool and the RPD radio-wave propagation prediction tool.

2.1 The Cellular Network Planning Problem

The advent of mobile or cellular telephone technologies in the 1980s initiated a new area of research – that of cellular network planning. Cellular networks have the following defining characteristics.

- A *Base Station* (BS) is a collection of infrastructure, such as a generator and a mast, at a particular site. A BS houses antennae which can transmit a signal to a limited portion of the network.
- In order to be considered *covered*, a user (or some subset of the network area) must receive a signal from a transmitting antenna, such that the received signal strength exceeds a given threshold.
- A *served* user must be covered, must also have been allocated a block of capacity within the network and must have a signal-to-interference ratio that exceeds a given threshold.

- A *sector*, or a *cell*, can then be defined as a combination of a BS, a particular antenna housed at that BS and a set of users that can be associated with (are either covered by or served by) that antenna.

The Cellular Network Planning Problem (CNPP) can be formally defined in a number of ways, dependant upon the method used to produce solutions. There is a formal definition of the problem, as it applies to this thesis, in section 4.1. Informally, the CNPP can be defined as minimizing the cost (either strictly financial or in terms of infrastructure i.e. the number of base stations or sector antennas) required to provide the desired level of coverage (i.e. a signal strength above a given threshold) to the area in question. The coverage may be measured as a percentage of the whole area, or as a percentage of some pre-defined discrete points, across the proposed network. Some definitions of the problem place more emphasis upon the maximization of the coverage levels instead of minimizing the cost.

With regard to this informal definition, the CNPP can be simplified further and seen as a combination of the following:

1. The Base Station Placement Problem (BSPP)
2. The Base Station Configuration Problem (BSCP)
3. The Frequency (or channel) Assignment Problem (FAP)

2.1.1 BSPP.

A simple description of the BSPP starts with the geographical area of the network and a set of potential, or candidate, BS sites. Usually a BS cannot be located anywhere within the network area due to terrain and planning considerations, so some subset of the area consisting of particular co-ordinates is provided. Each of these candidate sites has associated costs and

some notion of the potential area of coverage. The problem is to select a subset of the candidate sites within the network area while maximising some measure of network performance, for example, to ensure a desired coverage percentage is met, whilst also minimizing the overall cost. In real world situations BS costs are not uniform so it does not automatically follow that a viable network with the fewest sites will also be the cheapest.

2.1.2 BSCP.

The BSCP assumes that either the BSPP has been solved or that it is being solved concurrently. In either case the BSCP starts with a list of BS sites and a list of potential hardware (sector antennae) to be placed at each site. Here the problem is to place infrastructure and configure it at the sites, again so that cost is minimized whilst attaining the coverage constraints. Each solution to the BSCP must assign a value to every relevant operational parameter for each BS.

The inclusion of capacity considerations into either of these problems, adds another level of complexity. Some way must be found to represent the traffic demands and the capacity offered across the network. Both the BSPP and BSCP problems must take account of these levels when determining coverage.

2.1.3 FAP.

The FAP assumes that both previous problems have been solved or are being solved concurrently, and a network of fully configured BS have been placed at sites throughout the area. The problem also assumes the allocated frequency band for which the network has been planned, has been divided into discrete portions or channels. The problem is the further assignment of these channels, to the transmitting antennas, so as to maximise the quality of service by minimising the Carrier-To-Interference (C2I) ratio across the

network, thus making the planned for coverage and supplied capacity achievable.

It is obvious that in the context of the CNNP each of these sub-problems are to a certain extent inter-dependant. By solving the BSPP using a certain pre-defined coverage area, the solutions for the BSCP are limited in regard to keeping the solution globally optimal. Equally by optimally configuring the sites in terms of cost, the coverage of the network as whole may be affected. The FAP will obviously be affected by the solutions to the other problems and it is possible to envisage a network that permits no good solutions to the FAP. This is unlikely however, and, because there would generally be no monetary difference between alternative assignments, the FAP is generally considered separately as a post-processing phase.

An automatic cell planning tool (ACP) is defined as a software tool that produces complete network designs without any interaction with a planning engineer. Such a tool should be able to solve each one of these sub-problems, not so that the individual solutions are optimal, but in such a way that the combined solution – the plan for the final operational network – produces a good sub-optimal global solution.

2.2 Cell Planning Techniques

2.2.1 Manual Planning

Early cellular networks were constructed in a fairly ad-hoc manner that followed certain rule-of-thumb planning rules. BS would be placed at the top of hills to ensure maximum coverage, then gaps plugged with extra antennas as and when they were necessary. In figures 2-1 to 2-4, a simplified example of the manual planning approach is demonstrated using the Malvern data set

defined in Appendix I. The figures show the network plan at each stage during the process. The small grey crosses each represent a potential network user in the scenario, with the candidate BS sites shown as a white circle outlined in black.

In this scenario the aim is to produce a network that will provide a data link service to at least 75% of the potential customers in Malvern.




Figure 2-1. A BS is placed on the highest ground




Figure 2-2. Then at the second highest...

Firstly, a base station is placed on the highest available point within the network area and populated with 4 90° beam-width antennae configured in a clover-leaf pattern, as shown in Figure 2-1. Each antenna is represented in the figure as a different coloured arrow. The potential users that can then receive service from these antennae are shown as small blocks of the same colour. At this stage the network wide coverage level (not represented within the figure) is a respectable 70%, but the service level (the percentage of potential users that have been coloured), which takes into account sector capacity and interference concerns, is only 27%.

Adding in a second base station (Figure 2-2) brings the coverage level up to 82%, and a third (Figure 2-3) pushes it to 97%. Again service lags behind with 62% and 78% respectively.

It is obvious that some of the sectors in these plans, for example the blue and yellow, south-west and north-west sectors, are contributing very little benefit to the overall plan in terms of coverage or service levels. A Best-Effort plan, as shown in figure 4, removes these ineffective and possibly interfering sectors. This has no discernable effect upon the coverage or service levels, which remain at 97% and 78%, but it does reduce the infrastructure costs by more than 50%.




Figure 2-3. ...and the third




Figure 2-4. The Best-Effort plan

This 'clover- leaf' manual network planning approach is compared against results from an automatic planning tool in [3].

Whilst the process certainly produces usable network plans, there are obvious issues with this approach. For instance, the coverage levels may well be acceptable but the example shown above utilises a modern planning tool to obtain the information presented. Methods and tools were needed to analyse the proposed networks and predict coverage levels. There can be no indication that this network is a 'good' design, which immediately leads to

the question of how to quantify this notion of 'goodness'. As mobile and wireless network subscribers flocked to use the new technologies, the service providers realised that their networks must become more efficient and cost effective. Also, the size of the networks to be planned expanded, making the use of the manual planning method infeasible. New methods needed to be found that would lead towards better solutions.

2.2.2 Simple Cell Planning Strategies and Methods.

In [4], Gamst assesses the body of work in 1987 and lays out some ground rules in cellular network design. Three areas of data, crucial to the process of planning efficient radio networks are defined:

- Geographic Data – All data characterizing the physical service area.
- Radio System Data – E.g. frequency band, channel bandwidth, etc.
- Quality Requirements – E.g. The levels of coverage, service, interference, etc.

These broad categories are vital to finding optimal operational solutions to the cell planning problem.

In the same paper Gamst also divides contemporaneous approaches into two categories – Geometric (described as “powerful but imprecise”) and Analytical (“more exact but non-constructive”).

GEOMETRIC

The Geometric approach is typified by the hexagonal cell system (see figure 2-5), which is primarily concerned with providing a solution to the BSPP. Omni-directional antennae are placed in the centre of regular hexagons which form a regular tiling across the entire network. Cells are assigned frequencies such

that there is adequate separation between neighbouring cells. Directional antennas, that have higher signal gain but are limited in coverage area can replace the omni-directional antennas in order to introduce sectorisation of the cells. Splitting the cell into, for example, four 90° sectors allows the system to adapt to varying traffic densities, or to cover a larger area. Without this option, the assumption of a fixed cell coverage radius would severely limit the value of this approach for real world networks.




Figure 2-5. A simple illustration of the hexagonal cell system.
Each different colour represents a different frequency.

In the systems most simple form an inadequate isotropic coverage prediction model and a frequency reuse scheme that fails to account for changes in traffic density also posed problems. A further flaw with this approach is the arbitrarily chosen BS sites and the problem of their real world availability.

ANALYTICAL

GRAND (Generalised Radio Network Design), one of the first radio network planning tools, depends on a generalization of the notion of a cell, as a fuzzy set, which is introduced in [4,5,6]. GRAND integrates what Gamst describes as the three basic design tasks:

- Generation of Base Station Configurations (BSCP)
- Radio-Wave Analysis (propagation, coverage, interference)
- Frequency assignment (FAP)

As Gamst notes, the tool does not and can not propose any efficient planning strategies for the solving of the BSPP. This is left to the engineer manipulating the tool. It is seen more as an environment where planned networks can be tested, rather than planning itself.

There are now many commercially available tools which play a similar role. Planet DMS, from Marconi¹², Hexagon's PARTNER¹³ software suite, Agilent's 6474A optimisation tool¹⁴, and ICS Telecom (ATDI)¹⁵, as well as many other similar tools, offer various features such as technology evaluation, business modelling, spectrum optimisation and network proving.

However, this type of tool still relies to a large extent on the skills, intelligence and hard-work of network planning engineers and so cannot be called an automatic cell planning tool. The engineer will have to make a lot of choices during the building of the network, most importantly the BS positioning and so, in the context of the BSPP at least, these tools offer no improvement upon the manual planning algorithm described earlier in this chapter. These tools

¹² <http://wnp.marconi.com/planet/index.shtml>

¹³ http://www.hexagonltd.com/partner_suite.html

¹⁴ <http://we.home.agilent.com/USeng/nav/-536890803.0/pc.html>

¹⁵ http://www.atdi.com/ics_tel_rf_network_planning_software.htm

are most certainly of use however as they do allow for in depth simulation and proving of networks before implementation, and are useful for building a business case. However, as concluded in [6], if it can be shown that cost-optimal networks are not necessarily intuitive, then the planning engineer should be removed from the process.

The analytical planning approach, typified by the tools mentioned above, may also be said to be too focused on the radio design aspects of network planning whilst ignoring other interdependent aspects of the cellular network planning problem. This can lead to deficiencies in the planned networks in terms of base-site placement and configuration, and end-user considerations such as customer behaviour or traffic demands.

2.2.3 Integrated Planning Tools.

So called Integrated Planning tools may be said to be driven not only by coverage considerations, but also by some form of customer demand. They may be seen as an integration of the analytic with the geometric model – not solely concerned with the radio system design and blanket coverage, but also with ensuring service. However, some examples of these tools suffer from a number of the same real-world implementation problems as the basic analytical and geometric models.

In [7], Martinez-Dalmau et al. present an approach leading to a GSM network configured in terms of traffic density which relies on a rigorous study of the proposed network area before the planning starts. In this approach the target network area is divided into discrete zones of homogeneous traffic density. The proposed algorithm calculates the required number of base stations, and the corresponding number of channels per station, starting with the densest traffic zones and continuing across the network area. A simple cost function is

used by the planner to compare different base station configurations. The cost function is as follows:

$$Cost = An_i + Bn_A + Cn_C$$

where A is the cost per BS site, n_i is the number of BS sites, B is the cost per antenna, n_A is the number of antennas, C is the cost per user and n_C is the number of users.

The authors suggest the use of optimisation techniques to improve the network, before analytical methods are used for final validation.

This approach realises the importance of the network infrastructure planning (e.g. BS locations) that is not considered within purely analytic tools. Also there is recognition of the (financial) cost factor in the planning process – ‘good’ networks should not be seen simply as coverage-optimal networks. However, the integration of these two aspects of the radio planning problem can be seen to be incomplete. A truly optimal network must take account of both aspects at the same time – poor locations in terms of monetary cost may provide superior coverage. This approach does not strike a balance, as there is no real attempt at an optimisation of whole solutions. The need is not only to be able to automatically plan a network – there must be some form of justification that the chosen plan is in some way one of the best possible.

A recognition of this need for engineers to have tools that can independently optimise a number of different parameters at one time (instead of glorified propagation calculators such as described in the previous section) led to the development of planning tools such as PLATON [8], WISE [9], STORMS [10] and ICEPT [11]

PLATON [8] was the first real automatic cellular network planning tool (GSM). Taking simple inputs, such as a terrain database and generic

representations of cell types, the tool produces two outputs: A good radio configuration with frequency assignment, and an evaluation of the solution with respect to infrastructure cost. In PLATON the BS placement is not dependant on radio coverage but traffic density considerations. The final network configuration is optimised based upon parameters such as quality of service, efficient use of spectrum and the overall cost of the solution

Focussing on indoor and campus sized wireless systems (what would now be called Wi-Fi networks) Bell Labs devised WISE (Wireless System Engineering) [9]. This tool, designed to optimally locate BSs, includes a physical database of the network area, a ray-tracing propagation prediction model and an optimisation component – based upon the Nelder-Mead simplex method [12]. It is noted that this approach produces a large number of similarly well performing networks. The authors contend that it is not necessarily vital to find the true global optimal network because, due to approximation in propagation methods and other unconsidered constraints and assumptions, the networks produced are themselves approximations. A far more realistic aim is to find a selection of *good* base station placements.

Another tool for the automatic planning of mobile radio networks, named STORMS (Software Tools for the Optimization of Resources in Mobile Systems) [10] aimed to perform the radio planning task in two steps. Taking in a candidate set of BS sites, and assuming a uniform coverage area and cost per site, the problem is tackled by attempting to choose a subset of BSs such that the required coverage level is satisfied and cost kept low. The problem is defined in this way in order to apply two graph theoretical approaches to the choice of subset.

Firstly there is the Maximum Independent Set (MIS) model, defined in [10] and [13]. This approach employs a greedy algorithm to select the subset of BS. Within this model an individual candidate site is connected to a neighbour if

their respective predicted coverage areas are deemed to overlap to excess. Then with each candidate site as a vertex, and connections responding to edges, an independent set will consist of sites whose coverage areas do not encroach too greatly upon any other within the set. The MIS will, as a consequence, maximize the coverage area.

The second approach uses the concept of the Dominating Set which is explored in [10] and [14]. This model introduces a set of 'potentially covered locations' and employs a genetic algorithm. The genetic algorithm approach is explored further in the section on meta-heuristics, 2.4.

The Integrated Cellular Network Planning Tool, or ICEPT [11,15] was designed specifically to obtain good cellular mobile radio networks without any interaction from an engineer. As noted by the authors it was envisaged as an adaptable tool, able to plan a network from scratch or optimally redesign or extend existing configurations. The paper reinforces the four major design considerations for cellular networks:

- Radio Transmission – The physical propagation behaviour of the radio waves.
- Mobile Subscriber – In this context, the behaviour of the end-user.
- Resource allocation – Frequency and channel allocations.
- System Architecture – The underlying system hardware.

It can easily be seen that these areas are linked. For example in any wireless network the signal propagation determines the interference at some point which in turn determines the optimal frequency allocation. Again the geography of the network area is discretised and represents the spatial distribution of traffic demand. These discrete points, called *demand nodes*, then form the basis of the optimisation. The problem becomes a *Maximum Covering*

Location Problem (MCLP), that is, locate the base stations to maximise coverage of the demand nodes. The Automatic Base-station Positioning Algorithm (ABPA), introduced in [16], is employed to solve the MCLP. The ICEPT tool also includes a dynamic ray-tracing propagation prediction method, [17], which allows the tool to produce propagation information for any position in the network region. This allows the optimisation a great deal of freedom in generating placement solutions. Although ray-tracing methods can be computationally expensive this approach provides a trade-off option, allowing the planner to decide between generating an accurate solution over a long duration or a (relatively) fast but inaccurate solution.

The ARNO (Algorithms for Radio Network Optimisation) project supported the approach detailed in [18]. This method for antenna placement within GSM networks uses a three phase approach:

- A constraint based phase which eliminates poor solutions.
- An optimisation phase using a Tabu-Search algorithm.
- A fine tuning phase which fixes the antenna configurations.

As in other methods the network area is discretised. A set of Reception Test Points (RTP) covers the entire network area. The signal strength across the network can be measured at these points. Two sub-sets of the RTPs are also defined – Signal Test Points (STP) and Traffic Test Points (TTP). Each STP has an associated signal-threshold value, whilst each TTP represents a certain level of traffic. This allows the network constraints to be transferred to the two sets. For example, it is then possible to demand that a certain percentage of STP must receive a signal higher than the defined threshold. More information on the Tabu-Search heuristic may be found in chapter 3.

Another approach, in [19], utilises the Simulated Annealing heuristic to produce designs for BFWA networks. This work goes further than others mentioned here in that the tool provides solutions to both the CNPP and the further problem of back-haul optimisation. The network model is similar to those already discussed and a brief discussion on the SA algorithm used resides in chapter 3.

2.2.4 Other Approaches.

There is a sizable body of work concerned with the optimisation of one or more of the component problems rather than with the production of an all-inclusive automatic planning tool. These works are all broadly similar in the network model that they utilise and emphasise the need for good information. The authors of [20, 21] and [22] offer a very business oriented view of the issues and general methodologies involved in producing good solutions to the network planning problem, whilst [23] and [24] describe the planning of rural and campus size networks respectively, using expert systems and a branch and bound method. A three-stage hierarchical optimisation method is proposed in [25] which also utilises the heuristic technique of simulated annealing in its approach to planning low-cost mobile networks. Very analytical approaches, based upon field strength measurements and calculations, are used in both [26] and [27].

The BSPP has been investigated separately by many people in the past, such as in [28], which demonstrates a method utilising a number of nonlinear optimisation algorithms, whilst in [29] the optimisation of indoor wireless BS positions is carried out using an integer programming formulation. The concept of a traffic driven network design process was investigated in [30]. A number of different algorithms for automatic base site placement were tested. This approach appreciates the need for human-independent solutions and

offers novel methods for obtaining final network plans based upon generic, well-behaved sectors. In [31, 32, 33] and [34], a variety of heuristic techniques are used, including Tabu-Search, Genetic Algorithms and Simulated Annealing. These algorithmic techniques are discussed in depth in chapter 3.

The importance of all these approaches in their recognition of the multifaceted problem of automatic planning should not be underestimated. The discretisation of the geographical network space – based on traffic demand points – is a very important step in modelling the CNPP, as is the idea of pre-defining candidate sites for antenna placement. Both these concepts reduce the complexity of finding good solutions to the problem, by either disallowing infeasible network solutions or ignoring unimportant or easily assumed calculations.

The majority of the approaches discussed have been concerned with networks providing mobile-telephony services. In general the main concern of a GSM- or UMTS-network oriented CNPP is one of coverage levels (with reasonable capacity allocation) and thus service provision. However once a near-complete coverage network has been found the cost of infrastructure is minimized. This is particularly important to the mobile-service provision technologies because of the returns (in subscription revenue and goodwill) gained by complete coverage. Whilst coverage is obviously very important within all wireless networks, fixed-wireless service providers may find that the returns expected from hard-to-cover areas are relatively less. This may be especially true coming into an already well provided for market. The fact that near-100% coverage across BFWA networks may not be necessary or desirable is a point worth noting.

2.3 The Frequency Assignment Problem.

In many technologies the Frequency Assignment Problem (FAP) must be solved across a network in order to provide service of a particular quality for the users. Each transmitter must have a sufficient allocation of channels to provide sufficient capacity to its users, whilst not interfering with nearby sectors or being interfered with itself. In general, this can be easily achieved provided there is a sufficient range channels to choose from. In this case the problem is trivial. However, in heavily loaded networks, for example, those with strictly regulated channel bands, the assignment may not be so straightforward. For this reason a large body of work has been carried out in relation to the subject. With regards to this thesis the BFWA-FAP has largely been assumed to be relatively simple and so not of great import. This is because the networks considered function at a frequency around 40GHz. At this level the adjacent channel attenuation is 25dB, rather a high value which means that adjacent channel interference is not a major consideration. Also, the scenarios considered in this thesis have been empirically found to have easy allocations. For these reasons this section is very brief and considers only the most important results. For more detailed treatments of the problem see [35] and [36]

The problem itself can be approached in two different ways. Firstly there is the Minimum Span problem (MS-FAP). In this formulation the problem is to determine the minimum number of channels required, whilst achieving a network-wide minimum interference level.

The second formulation of the problem is known as the Fixed Spectrum Problem (FS-FAP), and is in general more pertinent to the researcher interested in practical ACP. In this case the level of interference across the network is minimized, under the constraint of a fixed allocation of frequency channels. This means not only an interference free network, but also that

transmission power levels across the network can be reduced – an important consideration for service providers and their lawyers.

The two FAP formulations may be represented as graph colouring problems and numerous authors have taken advantage of this fact. Work, such as that described in [37,38,39] details graph theoretical optimisation techniques to produce minimum span bounds for particular problem instances. In [40] a set of functions and methods is presented for a generalised object-oriented approach to programming for the FAP. Heuristic techniques are very successful in attempting to find solutions. Genetic Algorithms are used in [41,42,43], Simulated Annealing in [43,44] an Adaptive local search algorithm in [45], Tabu-Search in [46] and variations on the Ant Colony Optimisation Heuristic [47,48]. Other tools such as FASoft [49], developed at Cardiff University, offer a large variety of techniques in order to assess the relative performances of different schemes.

2.4 Propagation and Building Data

RPD (Rapid Pipeline Development) is a tool that has been developed by the Radio Communications Research Unit (RCRU) at Rutherford Appleton Laboratory (RAL)¹⁶ to model the propagation of radio waves at microwave frequencies. The tool is used to produce the datasets and propagation information used in this thesis. RPD uses a 3-D model of the area the proposed network will cover, and uses ray-tracing techniques to construct link profiles from which path loss values are calculated. Figure 2-6 shows the town of Malvern as modelled in 3-D by the RPD.

¹⁶ Principle authors: John Biddescombe and William Miners (RCRU-RAL)



Figure 2-6. A 3-D model of the town of Malvern (UK)

2.4.1 Building the 3-D model

The data input utilised by RPD to create the 3-D models are ASCII AutoCad DXF (Drawing Interchange Format¹⁷) files. The information contained therein conforms to the Technical Working Party on Mobile and Terrestrial Propagation's 3-D Dataset standard[50]. When no such up-to-date data is available it must be constructed. Various private companies provide this service. To ensure accuracy of information, aerial photographs of the proposed network area are digitised by hand, recording the corners and heights of all features, including buildings or other erections, and all vegetation, along with other information such as roof ridge lines and roads. This information is then saved in separate layers in the DXF file. A separate file contains terrain information in Digital Terrain Map (DTM) format. A small portion of the information in the Malvern DXF file is shown, in graphical form in figure 2-7. The different coloured layers represent different aspects of the scenery, with a single building possibly being composed of a

¹⁷ Copyright of AutoDesk, Inc.

number of layers. In figure 2-7 the red lines represent the 1st or base building layer, the yellow lines represent the second building layer while the dark and light blue lines represent roof ridges and vegetation outlines respectively.




Figure 2-7. A section of the layered Malvern dataset

Figures 2-8 and 2-9 show the way that buildings are divided into the separate layers during the digitisation process.




Figure 2-8. Each vertical wall has an outline in its respective layer




Figure 2-9. The roof ridges are stored in a separate layer

The data can then be piped into RPD, keeping the distinct layers separate. The 3-D graphical engine within the RPD integrates the Visual Tool-Kit (VTK) software package. This is an open source, freely available software system for 3D computer graphics, image processing, and visualization¹⁸.

The conversion from the two dimensional DXF layers into three dimensional structures proceeds in various stages, which are briefly outlined here:

- First build the terrain
- For each layer of data, drop the points onto the terrain
- Construct vertical 'walls' down from each of the layers
- Finally construct the horizontal cells e.g. roof coverings

A more detailed explanation can be found in [51].

Once the model has been constructed, a number of decisions must be made. When producing propagation data, potential BS sites are chosen with reference to local planning issues and site availability. A choice must be made on the position of the Reception Test Points (RTPs) around the proposed network. In figure 2-10, below, each building has had RTPs placed around the roof edges at 1 meter intervals. In this figure the RTPs are coloured depending

¹⁸ <http://public.kitware.com/VTK/>

on the value of the path loss from the site to the RTP. The transmitter is shown as a grey cone. The next section gives details of how these path-loss values are obtained.




Figure 2-10. A visual representation of path-loss levels

2.4.2 Propagation Model

The path-loss values generated by the RPD are calculated at each RTP from each candidate base station, taking into account interference and blocking from buildings and vegetation and signal-loss due to distance (Figure 2-10).

RPD gives a choice of three propagation models

1. Line Of Sight (LOS) – Simple yes-no model
2. Cost210 [52]
3. ITU-R [53], [54].

The object-oriented design of the RPD ensures that other models can be easily inserted into the tool. All the work in this thesis is based upon path-loss calculations utilising the ITU-R propagation models.

When an accurate model of the area has been created and a location chosen for a potential transmitter (figure 2-10), path-loss values from transmitter to each RTP are determined using elements from ITU-R P.452-10 [54] and ITU-R P.526-7 [53]. Although these recommendations are only intended for use with frequencies of up to 30 GHz, the model is suitable for extension up to frequencies of 42 GHz.




Figure 2-11. A path-loss profile

RPD creates 2-D profile slices from the transmitter to a RTP. The profile shown in figure 2-11 is in the vertical plane. Rays are fired to the test point diffracting over or around any obstacles on the path – terrain, buildings or vegetation. The path loss is then a combination of free space loss along the path plus any losses from diffraction. There is no true multi-path predictions included in the model at present although paths are also calculated in the horizontal plane and integrated into the final value (A total of three paths will be considered – diffracting over the obstacle, diffracting to the left and diffraction to the right- if there is no obstacle the paths will obviously be identical). Work is continuing to include a larger collection of diffraction paths and reflections.

The RPD module (called a filter) that performs the path-loss calculations takes in three inputs - transmitter information, set of reception test points and scenery information. The scenery information includes the building roofs, the terrain and vegetation (see figures 2-12 and 2-13).



Figure 2-12. The 3-D model.




Figure 2-13. The filter input

This means that vegetation is treated as solid objects just the same as buildings. At a frequency of 42 GHz, the ability of a signal to pass through any vegetation is negligible. However, trees and shrubs will allow some signal to pass through at lower frequencies (by scattering), the amount dependant upon the size and type of foliage, the time of year and also the depth of vegetation. Work is ongoing to produce a vegetation model for the RPD to increase the accuracy of the path-loss data across a wider range of frequencies.

2.5 The ECHO tool

ECHO (Efficient Cell-Planning Heuristic Optimiser) is a software tool produced within the EMBRACE project by the Centre for Intelligent Network Design at Cardiff University [3,55,56]. It is a tool which creates designs for BFWA networks to the specifications laid out by the EMBRACE project.




Figure 2-14. The ECHO Tool

2.5.1 Metrics

Within ECHO the term *profit* relates to the cost function for the optimisation. Put simply, it is the revenue generated over the time period minus the initial cost of all infrastructure and the associated annual maintenance costs over this period. Other metrics used with the tool are *service*, *coverage* and *diversity*.

2.5.2 Constraints

The ECHO tool allows a planner to define minimum network levels of three of the most important network metrics. These metrics – coverage, service and diversity – are formally defined in section 4.1. The ability to set the minimum levels of these metrics can be very useful for service providers. With increasing regulation of the radio spectrum the ability to guarantee certain

levels of, for instance, coverage across an area, will become an important part of any tendering process. There are also optimisation considerations that are discussed in the next section.

2.5.3 User lists

User lists are sets of potential users (or *subscribers*) within the network, which the planner can manipulate as a group. The default user list is all users in the network. When a constraint is set over the network as a whole, it is set on the list of all potential users. The planner is able to define any number of user lists for use in the optimisation. For example, with prior knowledge that a certain organization, with numerous buildings across the network, has pledged to take up a service subscription, all the buildings (classified as separate users) are grouped together with a service constraint of 100%. Note that these users must also be given a subscription probability of 100% in order for the required service to be planned for. Another example would be that of extending an existing network. In this case there are already subscribing users which must continue to be served in any new network. ECHO offers two ways to achieve this. The first method is the same as above, but this doesn't guarantee that the existing network infrastructure remains unchanged. In order to allow this ECHO offers the planner the ability to 'lock' sectors and users, thus leaving the existing network unaltered whilst the extension is placed around it. This also allows a service provider to investigate the financial implications of extending a network as compared to a complete re-planning.

2.5.4 The Network Design Algorithms

ECHO's network design algorithm is based upon a number of optimisation operators that change various aspects of the network. There are 37 optimisation operators used in the latest version of ECHO. As well as

operators that can alter the bearing and tilt of a sector, there are operators that add or remove a sector to a (random) site, that alter the power used to transmit, that choose the best antennae for each user and that can increase or decrease the capacity of a sector.

The whole network design algorithm proceeds in three phases:

1. **Initialisation.** A number of sectors, s , are added to generate an initial serving network.
2. **Repair.** A number, o , of specialised operators are enacted to improve the design to satisfy any defined constraint percentages.
3. **Optimisation.** Improves the design using the customised Tabu-Search (TS) optimisation algorithm.

The initialisation file, which can be altered by the planner, contains the values for s and o as well flags for the production of network statistics and values particular to the TS scheme.

The optimisation phase uses an algorithm based upon the Tabu-Search heuristic (see section 3.3.4). In this scheme, a neighbourhood is defined as a list of m operators that would have some effect on the network, if enacted from the current solution. Operations that would have no financial effect on the network, or ones that produce a configuration that does not meet the network constraints, are not included in the neighbourhood. The value of m , set in the initialisation file, is determined by the planner before the optimisation (the default value is set at 10). Each operator is enacted in turn and the effect it has on the network is examined. If the effect is positive it is added to the neighbourhood before being undone. This list is then looped over to find the best allowed move and the best tabu move. Allowed moves

are ones in which the constraints set upon the network continue to be satisfied. This check has the effect of dramatically reducing the size of the feasible search space. It may also mean that the final optimal network is decided by the random operators enacted within the repair stage. It is possible to conceive of discrete pockets of this 'constrained space' across the solution space as a whole. As the algorithm is forced to stay within a single constrained space by this check, the effect may be similar to that of a hill-climbing algorithm becoming marooned in a local optima.

The Recency (or Tabu) list holds the last j operators that have been enacted, where j is again defined by the planner in the initialisation file (this default is set at 10 also). These operators are considered tabu and each operator in the neighbourhood is checked against the list, to determine its status. Tabu moves are only chosen for enaction if they are appreciably better than the best allowed move.

The algorithm for the ECHO TS optimisation phase is as shown in Figure 2-15, over the page.

```

Repeat
  Create Neighbourhood
  Loop over Neighbourhood
    Get this operator's effect on profit
    Is the operator Tabu?
      Yes:  Compare effect against best Tabu move
           If better, save
      No:   Compare against best allowed move
           If better, save
    End Neighbourhood loop

  Is BestTabuProfitChange + CurrentProfit > BestProfit  and
    BestTabuProfitChange > BestAllowedProfitChange?

    Yes:  Enact Tabu operator
    No:   Enact allowed operator

  Is CurrentConstrainedMetricLevel < MinimumAllowedLevel?

    Yes:  Undo operator
    No:   Place operator on Tabu list

  Is CurrentProfit > BestNetworkProfit?

    Yes: Save current network as best

Until termination criteria are met

```

Figure 2-15. The ECHO Tabu-Search Algorithm

There is no explicit stopping condition included within the network elements of the scheme, so the length of the optimisation (the number of iterations through the search space) is set by the planner in the initialisation file.

As each of the operators make use of a uniformly distributed random value to dictate their behaviour there is no specific inverse of a move to forbid. Most operations, such as *Change Sector Bearing*, if used twice in sequence, would be very unlikely to generate the exact same random value that would be required to take the network back to a previous solution. Using the last 10 moves as tabu makes the chance even more remote, as the move would have to improve on the current network- the one it moved to from this state previously. Other operators which have more obvious inverses, for example,

Add Sector and *Remove Sector*, would not be able to oscillate more than once for the same reasons.

2.5.5 Multi-Objective considerations

The multi-objective nature of the CNPP is addressed in two ways by the ECHO tool. Firstly, the introduction of the service metric, and of subscription probabilities for the potential users, allows the two main objectives to be combined. As discussed in chapter two most approaches to solve the CNPP go for maximum coverage levels as the more dominant objective, and then secondarily consider the minimization of the infrastructure costs. By using a cost function that includes revenue as well as cost, and by tying that revenue directly to the levels of coverage, the multi-objective nature of the problem is diminished.

Secondly there is a constraint based approach to ensure a network is planned, and the optimisation scheme is pushed in the right direction. This means that the planner has the ability to place minimum requirements on the levels of any or all of the three main network metrics; coverage, service and diversity. These constraints may be placed upon the entire network or across a user list, and the network planned so as to adhere to them. From a multi-objective optimisation point of view this enables trade offs to be made between, for example, the overall coverage level of the network and the global profit levels.

It is conceivable that a potential network could exist that would have no solutions that are profitable. In this case ECHO would produce the obvious answer and not place any infrastructure in the area. However, the network may have to be built (e.g. to fulfil some contractual obligations) and so some plan would be necessary. Another reason to plan an unprofitable network over an area would be to attempt to find the reasons that it was unprofitable. In either of these cases the constraint levels can be set to force the optimisation

scheme to ignore the no-infrastructure solution and find the best network available.

2.5.6 Pareto Watchers

ECHO also has an option to estimate the Pareto front (see section 3.4) for any pair of network metrics. As the ECHO scheme uses a financially based cost function to drive the optimisation and hard constraints to ensure the multi-objective nature of the problem is considered, the addition of 'Pareto watchers' allows for a more balanced multiobjective approach. Any pair of network metrics may be chosen, such as coverage fraction and profit, or network income and network spend.




Figure 2-16. A Pareto front as found by ECHO

Throughout the duration of the optimisation, the levels of the particular 'watched' metrics are recorded. If the current network solution gives a metric level that exceeds the best so far, then that network instance is saved. The line produced, an example of which is shown in figure 2-16, cannot necessarily be assumed to be a true pareto front, however it does represent the front of the set of explored networks. By increasing the number of network iterations the edge of the explored network set will approach the Pareto optimal set.

Chapter 3

Heuristic techniques for combinatorial optimisation problems.

This section details a particular group of techniques used to produce near-optimal solutions to combinatorial problems, such as the CNPP.

3.1 Combinatorial Problems

Definition 3-1. *A combinatorial optimisation problem*

For a problem, X : Maximize $f(\bar{s})$

where S , is defined as the set of all possible solutions to X ,

$\bar{s} \in S$ is a single solution vector,

$f : S \rightarrow \mathbb{R}$ is the cost or objective function

Problems, such as those defined at the start of chapter 2, where the solution, \bar{s} , is an collection of discrete objects (such as a list of site co-ordinates in the BSPP), are called combinatorial optimisation problems. The value of the solution is decided by a cost function, $f(\bar{s})$, allowing solutions to be compared for desirability. The definition above is of a problem where the most desirable solution gives the maximum value of $f(\bar{s})$, but the problem may equally desire the minimum value of $f(\bar{s})$. One recognised feature of this type of problem is that of global and local optima. The solution space, S , of such a problem can be envisioned as a 3-d grid-map of peaks and troughs, similar to a landscape. If the problem is one of maximization the global-

optimal is then the highest point upon the highest peak across the entire space. Local-optima are the highest points upon all the other peaks across the space.

3.2 NP-Completeness

In any mathematically defined problem, the ability to analyze the entire solution space means no optimisation methods are required, as the global-optima can be easily calculated. In problems where the solution space is too large for complete enumeration, regardless of how much computing power is at the disposal of the researcher, the next step would be to employ an exact algorithm, such as Nead-Melder Simplex [12]. However, the computing time necessary to solve a problem varies with its size and type and, in the case of some problems, exact algorithms will be as impractical as a systematic full solution space search.

Within the set of all decision problems (that is problems which can be formulated so that there is a strict *yes* or *no* answer) there exist some problems which have been proved to be impossible to write an algorithm for, such as Turing's Halt prediction problem. These problems are classed as *undecidable*. All the rest of the decidable problems (ones that are algorithmically solvable) are classed depending upon the effort that must be used. Some problems can be solved with algorithms whose total number of steps (or time to a solution) can be represented by a low-order polynomial function of the size of the problem. Such problems are classed as **P** (Polynomial).

Definition 3-2. *A decision problem, x , is classed as in P iff there exists an efficient algorithm whose running time is $O(P(n))$, where $P(n)$ is a polynomial in the size of the input required, n*

Other problems can be proved to be 'difficult' and the length of any algorithm employed to solve them will require a time defined by an exponential function of the problem size.

The class of NP (Non-Deterministic Polynomial) problems are characterised by the fact that a positive answer to the problem can be easily verified using a (possibly nondeterministic) polynomial-time algorithm. Problems in NP (but not P) are ones that appear to require exponential-time algorithms to obtain solutions, but this cannot be conclusively proved

Definition 3-3. NP-Hard. *A problem x is NP-Hard if every problem in NP is polynomially reducible to x .*

Definition 3-4. NP-Complete. *A problem x is NP-Complete iff x belongs to NP and x is NP-Hard.*

NP-Complete problems form another subset of NP. Any particular NP-complete problem has the property that all other NP problems can be reduced to it using a polynomial function. It then follows that if a polynomial-time algorithm can be found to solve one NP problem, then a polynomial-time algorithm exists to solve all NP problems and thus $P = NP$. It also follows that if any NP problem can be shown to be exponentially difficult then algorithms to solve the NP-Complete problem must also require exponential-time. This makes the NP-Complete set of problems the most difficult across the NP set.

One such difficult problem is the graph edge-colouring problem. As mentioned earlier in this chapter the FAP may be represented as a graph colouring problem which can be shown to be NP-Complete [57]. Problems such as the FAP are called NP-Hard, because whilst they may not actually have been shown to belong to the set NP, the fact that they can be reduced to an NP-Complete problem (and thus all problems in NP) show them to be of equivalent complexity. As the FAP is also a subset of the larger CNPP it then

follows that the CNPP is also NP-Hard. [58], takes a closer look at formal classification of cellular planning problems. For more detailed investigations into the concepts of NP-Completeness and NP-Hardness see [57],[59] and [60].

3.3 Heuristics

In the (presumed) absence of reasonably efficient algorithms for NP-Hard problems, techniques known under the cover-all name of *Heuristics* have been developed. As defined in [61], “A *Heuristic* is a technique which seeks good (i.e. near-optimal) solutions at a reasonable computational cost without being able to guarantee either feasibility or optimality, or even in many cases to state how close to optimality a particular feasible solution is.”

3.3.1 Local Neighbourhood Search

Definition 3-5. *Local Neighbourhood Search concepts.*

$N(\bar{s}) \subseteq S$, is defined as a neighbourhood of solution \bar{s} .

\bar{s}_{opt} , is defined as a near-optimal solution.

Local Neighbourhood Search (LNS) is one such heuristic method. This introduces the concept of a Neighbourhood. A solution, \bar{s} , is said to have the neighbourhood $N(\bar{s})$, defined as the set of alternative solutions that are accessible by performing a simple operation upon \bar{s} . This operation may be a simple addition or subtraction, or any other basic change to some part of solution s , and a set of acceptable operations must be defined in advance. To extend the earlier analogy, a neighbouring solution equates to one that can be found moving to a neighbouring co-ordinate upon the solution grid. The basic local search optimisation method starts at an arbitrary solution point \bar{s}_0

within the space, It proceeds with a systematic stepping between the current solution and one of its objectively better neighbours.

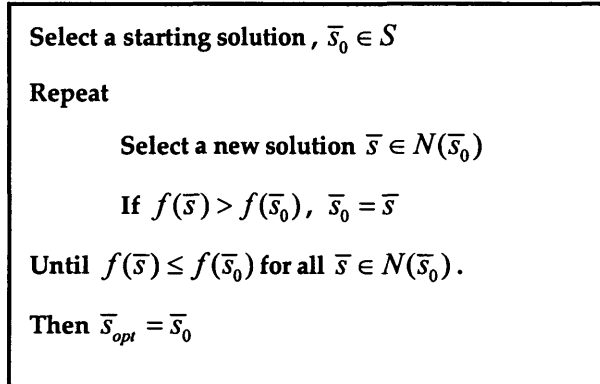


Figure 3-1. A simple Local Neighbourhood Search algorithm

This improved neighbouring solution can be found either randomly (Monte Carlo Method) or from an exhaustive search of the neighbourhood (Descent Method). Once the method finds a solution, \bar{s}_{opt} , where $N(\bar{s}_{opt})$ contains no superior solutions, or $\nexists \bar{s} \in N(\bar{s}_{opt})$ s.t. $f(\bar{s}) > f(\bar{s}_{opt})$ then \bar{s}_{opt} is declared near optimal. This final optimal solution depends entirely on the starting solution. This leads to an obvious problem - how can it be known that \bar{s}_{opt} is globally optimal? Is it even representative of a good solution?

3.3.2 Simulated Annealing

Simulated Annealing (SA) is a heuristic optimisation algorithm that is based upon the physical process of annealing (heating then cooling a metal in order to promote toughness). It can be seen as an extension of the LNS. For more detailed descriptions see [62,63].

Definition 3-6. *Simulated Annealing concepts.*

t , is defined as the temperature of the optimisation,

$\phi(t)$, is defined as the cooling schedule of the optimisation.

Simulated Annealing attempts to get around the limitations of the basic LNS method by allowing solutions to be accepted that are not an improvement on the current solution. This allows the optimisation to work its way out of local minima.

In the classic SA algorithm (shown in figure 3-2) a starting temperature, t_0 , is defined at initialisation, which will control the optimisation. As the optimisation progresses the temperature is reduced by some function, $\phi(t)$, in a process known as 'cooling'. The function governing the speed of this cooling is known as the 'cooling schedule'. A starting solution is chosen and from there a neighbourhood of solutions is constructed.

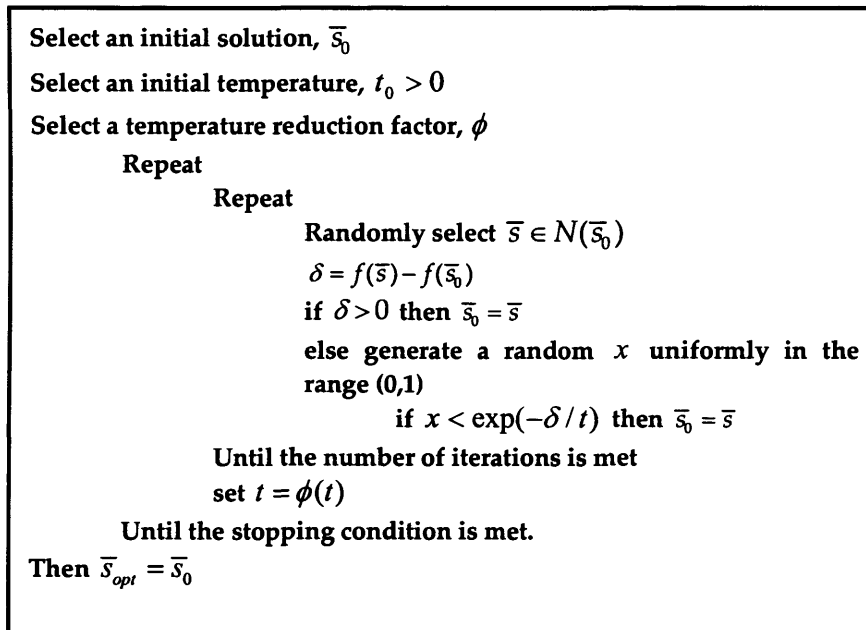


Figure 3.2. An example of a maximising Simulated Annealing algorithm

In this algorithm, the smaller the difference between two solutions the more likely it is that an inferior solution will be accepted. Also, the 'cooling' of the optimisation leads to a stabilizing effect, making the acceptance of inferior solutions less likely and forcing a final calmer hill-climbing search.

Simulated Annealing has been used by various teams within the area of cellular network planning. The Adaptive Base-station Positioning Algorithm [16,17] employs SA techniques in order to achieve “quasi-optimal base station locations”, as does [33]. The authors use the technique to minimize the error between the objective function, where the desired conditions are simultaneously satisfied (for example no cost but total coverage (the ideal)), and the current system configuration. Louta et al, go a step further in [19], using the scheme to optimise both position and configuration of antennae in a BFWA network.

3.3.3 Genetic and Evolutionary Algorithms

The Genetic Algorithm (GA) concept is based upon the concept of survival of the fittest. Starting with a population of randomly generated solutions through breeding, mutation and death the overall fitness of each solution is raised over successive generations. For more information see [64] .

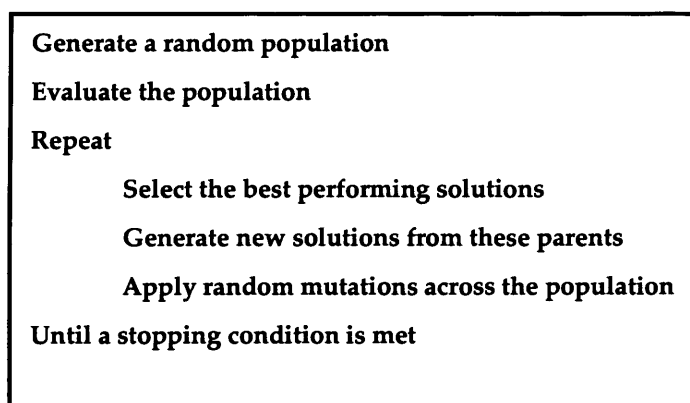


Figure 3-3. A simple Genetic Algorithm

From the initial random population weighted random pairs are chosen. The solutions have a probability of being chosen based upon their fitness determined by a cost function. At a random point within the vector the two solutions are merged and added to a new population. This generation should

then be fitter as a whole than the last generation. There is also the small possibility of a mutation (a random change) within all of the solutions. As the optimisation progresses the overall fitness of the population increases. The optimisation is then stopped either after a certain number of generations or once some other metric has been satisfied.

The STORMS project [14,65] utilised GAs in its work on cellular network planning. The concept was extended by introducing subsets of the population, called *islands*, which evolve independently under separate GAs. Some migration of solutions to neighbouring island was allowed under strict regulation. This approach also allowed the *parallelisation* of the algorithm, with each population running on separate processors to reduce computation time. Other approaches that concentrate on using GAs to obtain solutions for cellular planning problems include [66] and [32], where a GA is shown to outperform a greedy LNS algorithm.

3.3.4 Tabu-Search

The Tabu-Search (TS) heuristic optimisation algorithm, developed during the 1980's (see [67,68] and [69]) attempts to utilise intelligent decision making, through the use of a *memory*. It is designed to allow the downhill steps (if you are maximizing) taken by other strategies such as Simulated Annealing, but also to control their use more effectively. The strategy tries to achieve this through the use of a virtual memory. It both compiles and refers to a *Recency List* that holds a number of the recent moves made by the optimisation.

Definition 3.7 *Tabu-Search concepts*

T , is the *Recency or Tabu List*,

$N_C(\bar{s}_0)$, is the candidate set of solution s_0 , s.t. $\forall \bar{s} \in N_C(s_0)$, $\bar{s} \in N(\bar{s}_0)$ and $\bar{s} \notin T$.

$N_T(\bar{s}_0)$, is the tabu set of solution s_0 , s.t. $\forall \bar{s} \in N_T(\bar{s}_0)$, $\bar{s} \in N(\bar{s}_0)$ and $\bar{s} \in T$,

k is the aspiration criteria.

The Recency List allows the optimisation to partition the moves available in a neighbourhood into two sets: the set of Tabu moves and the set of non-Tabu moves. As their name suggests, the Tabu moves can be thought of as ones that are best avoided. An aspiration condition is placed on the use of Tabu moves - they have to give a considerable advantage (i.e. greater than k in a maximising algorithm) over the allowed moves to be accepted. A move can be labelled Tabu by many criteria but the most common are moves that are the inverse of, or the same as, a recently accepted move. These criteria will stop the optimisation oscillating between two solutions and will have the effect of steering the algorithm down a variety of paths in the solution space.

```

Select a starting solution,  $\bar{s}_0 \in S$ 
Record the best,  $\bar{s}_{best} = \bar{s}_0$ 
And the best cost,  $c_{best} = f(\bar{s}_0)$ 
Clear the recency list,  $T$ 
Repeat
    Determine the candidate and tabu sets  $N_C(\bar{s}_0)$ ,  $N_T(\bar{s}_0)$ 
    Select  $\bar{s}_{Cmax} \in N_C(\bar{s}_0)$  such that  $\bar{s}_{Cmax} = \max f(N_C(\bar{s}_0))$ 
    Select  $\bar{s}_{Tmax} \in N_T(\bar{s}_0)$  such that  $\bar{s}_{Tmax} = \max f(N_T(\bar{s}_0))$ 
    If  $f(\bar{s}_{Tmax}) - f(\bar{s}_{Cmax}) > k$  then  $\bar{s}_0 = \bar{s}_{Tmax}$ 
    Else  $\bar{s}_0 = \bar{s}_{Cmax}$ 
    If  $f(\bar{s}_0) > c_{best}$  then  $c_{best} = f(\bar{s}_0)$  and  $\bar{s}_{best} = \bar{s}_0$ 
    Reconstruct the recency list
Until the iteration limit is reached
Then  $\bar{s}_{opt} = \bar{s}_{best}$ 

```

Figure 3-4. The Tabu-Search algorithm

The Tabu-Search heuristic was applied to the antenna placement problem by Vasquez and Hao in [18] and has also been found to be effective when applied to the FAP (see [35,49]). Work such as [31,70], have compared Tabu-Search with SA and GAs in the area of the antenna location problem. Although all three methods were shown to be very good, the Tabu-Search algorithm tended to give the most consistent results. Details of how the Tabu-Search scheme applies to ECHO can be found in Chapter 2.

3.4 Multi-Objective Optimisation

The techniques described above have all assumed Single-Objective Problems (SOP). That is, they have a solution that is dependant on the optimisation of only one factor, or objective. A problem involving the simultaneous optimisation of several objectives (such as our combined CNPP) is described as a Multi-Objective Optimisation Problem (MOP). In such problems the individual objectives are inter-dependant and must be considered simultaneously in order to find a near-optimal global solution.

In a SOP each solution can be ordered according to the objective function, f , and so an optimal solution is usually clearly defined. E.g. for any two solutions x_1, x_2 , it will be the case that either $f(x_1) \geq f(x_2)$ or $f(x_2) \geq f(x_1)$ and so the solution that gives the maximum value of f is the optimal solution.

When solving a MOP there will not necessarily exist a unique optimal solution, one that gives a global maximum with respect to all the objectives. This is because the objectives, being interdependent, may conflict. For example, in the cell planning case, adding a sector to take the coverage level to 100% would satisfy the coverage level objective. From a cost objective point of view however, the last ten percent of coverage may increase the infrastructure outlay ten-fold.

Instead, with MOPs, a set of optimal solutions, which represent the compromises necessary between competing objectives, is more desirable. This set of solutions can be defined as optimal if there exist no other solutions within the search space that are superior, when all objectives are considered individually. These are called Pareto¹⁹ optimal solutions.

For formal definitions of Pareto optimal sets and fronts see [71]. Informally, a Pareto optimal front consists of a set of solutions which are globally optimal for one objective at some or all possible values for each other objective in the problem. What is required is an optimisation method that can find solutions along the Pareto optimal front. These solutions will represent the alternative trade-offs between the various objectives, allowing a suitable compromise solution can be chosen.

Traditional approaches to this problem attempt to turn the MOP into a more manageable form. The two most popular methods are weighting and constraint.

3.4.1 Weighting Method

The weighting method converts the original MOP into an SOP by forming a linear combination of the problems objectives, as for example:

$$\text{Maximise } y = f(x) = w_1 \cdot f_1(x) + w_2 \cdot f_2(x) + \dots + w_m \cdot f_m(x),$$

Where w_i , $i \in [1, m]$, is the weight for the i -th objective f_i .

Usually, $0 < w_i < 1$, and $\sum_{i=1}^m w_i = 1$.

The weight w_i indicates the relative importance of the objective f_i and it must be specified for each of the m objectives beforehand. Solving the above SOP

¹⁹ Named for Vilfredo Pareto (1848-1923) – a noted economist and sociologist.

for a certain number of different weight combinations will give a set of solutions which approximates the Pareto optimal set. However, in [71] Deb details some problems with the weighting method, the main disadvantages being that determining the weights can be difficult, and the method itself can not find Pareto optimal solutions in MOPs with a non-convex Pareto optimal front.

3.4.2 Constraint Method

Another technique for solving a MOP is to optimize one objective (perhaps the most important one, if that can be determined) and to treat the remaining $m-1$ objectives as constraints. Thus the MOP is converted into, and can be subsequently solved as an SOP as follows:

$$\text{maximise } y = f(x) = f_j(x)$$

$$\text{subject to } f_i(x) \geq \theta_i, \quad i = 1, \dots, m, \text{ and } i \neq j.$$

The lower bounds, θ_i , should be chosen by the optimiser for each of the $m-1$ objectives. The Pareto optimal front can then be found by solving the above constraint SOP with different values of θ_i . However, there also exists a difficulty with this method. If the lower bounds are not carefully chosen it may be the case that no feasible solution can be obtained for the corresponding SOP. The only way to avoid this situation is to know a range of suitable values for θ_i before starting.

3.4.3 The need for better algorithms

Although there exists a few other methods such as goal programming and the min-max approach, all of these methods have certain difficulties with MOPs.

Based upon the discussions laid out above, concerning the weighting and constraint methods, it can be shown that

- Some kind of knowledge about the problem is required, which is not always available.
- Certain techniques, such as the weighting method are particularly sensitive to the shape of the Pareto optimal front.

Despite these problems, these methods are still widely used, generally because they can be implemented within well-studied SOP algorithms, such as those described earlier in the chapter.

Genetic algorithms have long shown potential for providing solutions to MOPs. This is because they deal simultaneously with a large set of possible solutions, namely the population. This means that a set of pareto optimal solutions may be found in a single optimisation run. In addition GAs appear to be less susceptible to the shape or continuity of the pareto front. GA algorithms specifically tailored to MOPs, such as [72,73], have been shown to be particularly effective in converging and finding diversified solutions upon the pareto-optimal frontier.

When considering the constraint- and weight- based approaches, and also the GA approach to a lesser extent, the main problem appears to be that the optimisation is still essentially concentrating upon a single objective. One relatively new area of optimisation research suggests an approach that seems more able to deal with multiple objectives separately but simultaneously. This concept utilises the decentralised, self-organisation properties inherent in emergent intelligence systems.

3.5 Emergent Systems

Colonies of Harvester ants constantly adjust the number of workers actively foraging for food. No individual ant has the cognitive capacity to assess the colony's food needs, and each has only a meagre chemical vocabulary. However, as a colony, ants can find the shortest distance to a food source and then prioritise that source, depending on distance and ease of access. This is an example of swarm logic or emergent behaviour- the ability of complex systems to emerge from numerous simple components. The organisation of the colony is carried out from the bottom up, as opposed to the top down organisation our society is more used to. As seen in countries, companies and football clubs alike, Heads of State, coaches and CEO's control from the top giving orders to people below them who in turn give orders to people below them. Global information and individual actions determines local actions. In ant colonies local information and collective action produces global behaviour.

Ants communicate through the secretion and detection of pheromones, and this simple method for dissemination of information plays the most important part in the colony intellect. The secretion of different chemicals can represent a number of things from task recognition - i.e. the fact that this ant is hunting for food; and trail attraction - this trail leads to food; to alarm behaviour. These basic indicators are extended by the ability of ants to detect gradients and frequency in the pheromone trails. This means that ants can determine the direction of a food source and tell the difference between ten ants travelling the trail in the last hour and 100. This may mean that ants conduct a sort of statistical study based on random encounters with other ants. Seeing the colony as a whole is a perceptual and conceptual impossibility for any individual ant. However when an ant meets a lot more colleagues than it is

expecting it may follow a certain rule- find something else to do, return to the nest, etc. This is an example of local feedback.

In his book, *Emergence* [74], Johnson gives five fundamental principles that should lead to a system where “macro-intelligence and adaptability derive from local knowledge”.

1. **More is different.** The statistical nature of ant (or agent) interaction demands a sufficient population to explore all options and make an intelligent assessment of the global state.
2. **Ignorance is useful.** Emergent systems grow unwieldy when individuals become too complicated. A densely connected system of simple elements will allow more sophisticated behaviour to ‘trickle up’.
3. **Encourage random encounters.** Arbitrary individual encounters between elements become indicators of the global state of the system, when considered all together.
4. **Look for patterns in the signs.** Elements need the ability to communicate – pattern recognition. Meta-information can filter through a population from patterns within those patterns. Knowing one ant walked in a particular direction means nothing, knowing 50 did in the last hour means a lot.
5. **Pay attention to your neighbours.** Local information can lead to global wisdom. This follows from 3 and 4, without neighbours sharing information you simply have a collection of stupid individuals.

One example of artificial emergent behaviour used for optimisation is Ant Colony Optimisation (ACO), described by Dorigo in [69,75,76] and others. This is a computational optimisation paradigm that is based upon the

emergent behaviour of ant colonies. Using the pheromone communication as its base the ACO can be used to find solutions for graph theoretical problems.

3.5.1 The ACO

Experimental observation has shown that a population of ants will find the quickest route to a food source after a short transitional phase. In the simple diagram shown below, the ants nest is connected to a food source by two paths of unequal length. Ants must choose between the paths to get the food. After a transitory phase, the majority of ants use the shortest path. It is also observed that the probability of selecting the shortest path increases with the difference in path length.




Figure 3-4. Ants may have a choice of paths from the nest to the food source

The individual ant chooses either path, influenced to some extent by the amount of pheromones on the path. At the start there is no difference so ants choose either path at random. Ants choosing the shortest path will obviously get to the food first and then make a decision about which path to return on. The short path has a trail of pheromones increasing the probability of that

path being chosen again. This will start a feedback loop, ants will choose the path with more pheromone trails, and by choosing this path they increase the probability of other ants to choose the path, and so on. This idea can then be applied to finding shortest paths through graphs, and other graph theoretical problems, using artificial ant agents.

3.5.2 Applications of the ACO

The simple Ant Colony Optimisation heuristic can be used to find the shortest path through a graph, $G = (V, E)$, where V is the set of vertices in the graph and E is the set of edges between them.

Definition 3.13 *Ant Colony Optimisation*

k is defined as an ant at vertex i ,

$N(i)$ is defined as the set of vertices adjacent to i ,

$\tau_{ij}(t)$ is defined the intensity of the pheromone trail on edge (i, j) at time t ,

$$p_{ij}^k = \begin{cases} \tau_{ij} & \text{if } j \in N(i) \\ 0 & \text{if } j \notin N(i) \end{cases}$$
 is the probability that k will move to vertex j .

A population of ants is set loose within the graph. The aim of each ant is to reach a given destination vertex and then return to the start point as quickly as possible. Each ant in the system makes decisions at every vertex, depositing pheromone information as they go. This pheromone then evaporates at a steady rate. The effect of the pheromone is the same as in the real world example described above, leading to a convergence toward the shorter paths. The evaporation of the pheromone promotes exploration, making it less likely that the system will converge upon a single sub-optimal path.

The ACO, and variations upon it, have been employed in a various areas, such as Data Mining [77,78], routing in telecommunications networks [79], and the FAP [47,48], as well as classic optimisation problems such as the travelling salesman [80] and bin packing [81] problems.

A lot of work has gone into developing the family of ACO-type algorithms, and it has been shown that many of the best performing implementations improve on the solutions found by the ants with simple local search algorithms [82].

3.5.3 Particle Swarm Optimisation

Particle swarm optimisation [83, 84, 85] is another optimisation technique that utilises emergent intelligence. First proposed in 1995, a PSO system contains a population of candidate solutions (known as *particles*) to the problem in hand. These particles co-exist and collaborate, adjusting their position within the search space (and thus their suitability as a solution) according to a simple memory, called *experience*. Furthermore, these experiences can be shared amongst other particles that may be encountered within the system.

3.5.4 Agent Technology

These examples of optimisation schemes may be included under the banner of what is being called Agent Technology. The concept of an Agent as an autonomous entity, assessing and reacting to information but motivated by its own purposes is a relatively new one in Artificial Intelligence, and covers a wide variety of Research. A broad-definition of Multi-Agent systems as a population of interacting agents, utilising the power of emergent intelligence to work towards a system-wide goal, will obviously cover both the ACO and Particle Swarm optimisation paradigms. However, the majority of the work classed as Agent Technology concerns Agents of varying descriptions, bound

by a common language, trying to solve (often dynamically) real-world problems. Work such as [86] concerns increasing the effectiveness of traffic control systems, in terms of reacting to traffic congestion, the time of week or day and the proximity of roadworks, an accident or some other hindrance to the motorist. Roozmond proposes separate classes of agent performing different, very specific roles. These agents include:

- Traffic Signalling Agents (TSA) – Responsible for a single intersection.
- Authority Agents – Responsible for a small group of TSAs.
- Road Segment Agents – responsible for a certain stretch of road.

These agents may all share a common goal – such as maximising traffic flow through their domain- but must work with many other agents, of all types, in order to achieve it. If each TSA achieves its goal, it is reasonable to assume the system is working optimally. The real-world applications for these multi-agent systems are diverse, including railway traffic-scheduling [87], train coupling and sharing [88], dynamic packet routing in mobile networks [89], rapid prototyping and simulation systems [90], the job-shop scheduling problem [91], and foreign exchange markets [92].

This latter work concentrates upon the estimation of the parameters required for the system to be effective. The system allows each agent to change its behaviour over time, either by a random mutation or by being directly persuaded by a fellow system agent, in a similar manner to the worker ant previously described. Methods are described for estimating the parameters required for optimal performance. A similar approach is undertaken to get solutions to one aspect of the CNPP. Lissajoux et al describe the use of a genetic algorithm to develop a multi agent system for the BSCP in mobile cellular networks [93]. Certain aspects of GA generated networks are

quantified in order to guide the development of a limited but relatively successful multi-agent system.

Socha and Kisiel-Dorohinicki utilise an evolutionary multi-agent system to find solutions to multi-objective optimisation problems in [94]. In this system, agents may reproduce and may die, and each agent is ruled by an amount of non-renewable resource the authors have designated *life energy*. In this scheme, each agent represents a solution to the MOP in question, knowing its own quality to each particular objective. The optimisation proceeds in a manner broadly similar to a GA with the population of agents. The life energy is evenly distributed at the start of the optimisation but with each action the agents are rewarded or penalised by boosting or reducing the level of life energy. As the energy is conserved within the system, for every agent rewarded, another is equally penalised. The higher the life energy (i.e. the more the agent has been rewarded) the higher the likelihood of reproduction. Agents with low energy levels have relatively bad solutions and so are more likely to die.

3.6 Conclusions

As stated, this multi-agent system approaches the problem of MOP optimisation from an evolutionary perspective. How could the MOP problem be approached using emergent intelligence? An adaptation of the ACO can be envisaged where there are distinct *species* of ant-agent, one for each objective in the problem. A set of pheromone levels may be placed upon the edges of the graph, with each species reacting more strongly to their own kind. Set loose within the same graph, but with differing primary objectives, could the competing colonies provide optimal or good sub-optimal solutions?

Chapter 4

Modelling BFWA networks

This chapter formally defines the BFWA network model and the cellular network planning problem as used within this thesis.

4.1 The network model

Definition 4-1. *The BFWA network.*

A potential network consists of the following:

- A set of potential base station sites, $\mathbf{B} = \{B_1, \dots, B_{N_{BS}}\}$
- A set of reception test points (RTPs), $\mathbf{R} = \{r_1, \dots, r_{N_{RTP}}\}$
- A set of potential users, $\mathbf{U} = \{u_1, \dots, u_{N_{user}}\}$

The network planning procedure creates and configures a set of sectors, $\mathbf{A} = \{A_1, \dots, A_{N_{sectors}}\}$ such that each sector is assigned to a particular base station site, and provides a service to a specified set of users.

Definition 4-2. *Base Station Sites.*

A Base Station Site, $B_i = \{x, y, z, F_I, F_A, \delta_{on}\}$, is an exact geographical position where sector antennae may be placed to form a network.

It consists of: x, y, z denoting the geographical location of the site.

F_I, F_A denoting the initial and annual costs associated with housing a sector antenna at the site.

$$\delta_{\text{on}}(B_i) = \begin{cases} 1 & \text{if } B_i \text{ is in use} \\ 0 & \text{otherwise} \end{cases}$$

Definition 4-3. Reception Test Points.

RTPs, $r_i = \{x, y, z, F_I, F_A\}$, are exact positions where user-antennae may be placed to enable a user to receive service from the network.

They consist of : x, y, z denoting the geographical location of the RTP

F_I, F_A denoting the initial and annual costs associated with
using the RTP in the network

Definition 4-4. Potential Users

A potential User,

$$u_i = \{\delta_{\text{covered}}, \delta_{\text{served}}, \delta_{\text{diverse}}, p, R, S, r, av, b^U, b^D, \text{Ant}_U, P_U, \beta_U, \gamma_U, A, C2I_w, F_I, F_A\},$$

is a potential subscription service receiver and has a number of associated parameters.

These are: $\delta_{\text{covered}}(u_i) = \begin{cases} 1 & \text{iff } u_i \in U_{\text{covered}}, \text{ a coverage flag (see Defn. 4.8),} \\ 0 & \text{otherwise} \end{cases}$

$$\delta_{\text{served}}(u_i) = \begin{cases} 1 & \text{iff } u_i \in U_{\text{served}}, \text{ a service flag (see Defn. 4.9),} \\ 0 & \text{otherwise} \end{cases}$$

$$\delta_{\text{diverse}}(u_i) = \begin{cases} 1 & \text{iff } u_i \in U_{\text{diverse}}, \text{ a diversity flag (see Defn. 4.10),} \\ 0 & \text{otherwise} \end{cases}$$

$p \in [0,1]$ denoting the probability that this user will subscribe,

R denoting the predicted annual subscription revenue if
this user is served,

$S = \{r_{i_1}, \dots, r_{i_n}\}$ denoting a set of RTPs that define the user's geographical position,

r denotes a single RTP chosen for the user in this network configuration,

av denotes the availability, or quality of service, requirement of the user,

b^U, b^D denote the uplink and downlink minimum bit rate requirements of the user,

Ant_U denotes the antenna to be used by the user,

P_U denotes the transmission power for the user,

β_U, γ_U denote the tilt and azimuth respectively for the user's antenna

A denotes the active sector for the user,

$C2I_w$ denotes the worst case Carrier To Interference ratio in the present network configuration,

$F_I(u_i) = F_I(r_{u_i}) + F_I(Ant_U),$
 $F_A(u_i) = F_A(r_{u_i}) + F_A(Ant_U)$ denoting the initial and annual costs for serving the user.

Definition 4-5. User Antenna Type

Any number of different user antenna types, $Ant_{u_i} = \{G, L, F_I, F_A, RP\}$, can be defined within the network model.

These consist of: G, L denoting the reception gain and loss respectively for the antenna type

F_I, F_A denoting the initial and annual cost associated with the antenna type

RP denoting the radiation pattern for the antenna type

Definition 4-6. Sectors

A sector, $A_i = \{Ant_s, B_s, P_s, \beta_s, \gamma_s, ch^D, ch^U, U_{cap} F_I, F_A\}$, is a transmission antenna along with a number of associated parameters.

They consist of : Ant_s denoting the antenna type of the sector,

B_s denoting the base site at which the sector is positioned,

P_s denoting the power at which the sector transmits,

β_s, γ_s denoting the tilt and azimuth respectively,

ch^D, ch^U denoting the channels available to the sector on the downlink and uplink respectively,

U_{cap} denoting the set of users that the sector has allocated capacity for in this network configuration.

$F_I = F_I(Ant_s)$
 $F_A = F_A(Ant_s)$ denoting the initial and annual costs of the sector.

Definition 4-7. Sector Antenna Type

Any number of different sector antenna types, $Ant_s = \{G, L, F_I, F_A, RP\}$, may be defined.

These consist of: G, L denoting the reception gain and loss respectively for the antenna type

F_I, F_A denoting the initial and annual cost associated with the antenna type

RP denoting the radiation pattern for the antenna type

Definition 4-8. The set of covered users .

This set, $U_{\text{covered}}(A_j) \subseteq U$, contains all the users that are considered covered by the sector A_j within the network.

$u_i \in U_{\text{covered}}(A_j)$ iff $\exists r_k \in S(u_i)$ such that

$$A_j = A(u_i)$$

$$r_k = r(u_i)$$

$$P_{\text{recieved}}(A_j, r_k, u_i) - L_{\text{rain}}(A_j, r_k, av(u_i)) > P_{\text{threshold}}$$

where $P_{\text{threshold}}$ denotes the required signal threshold throughout the network,

$$\begin{aligned} P_{\text{recieved}}(A_j, r_k, u_i) = & P_s(A_j) + G(\text{Ant}_s(A_j)) - L(\text{Ant}_s(A_j)) \\ & - Q(B(A_j), r_k) + D(A_j, r_k) \\ & + G(\text{Ant}_u(u_i)) - L(\text{Ant}_u(u_i)) \end{aligned}$$

$L_{\text{rain}}(A_j, r_k, av(u_i))$ denotes the rain attenuation that must be included to ensure that $av(u_i)$ is satisfied. The value is calculated using the formulae given in [2, 96],

$Q(B_s(A_j), r_k)$ denotes the path-loss data, provided by RPD (see section 3.2) for the path between site $B_s(A_j)$ and RTP r_k ,

$D(A_j, r_k)$ denotes the directional signal loss calculated with reference to the location, equipment and configurations at A_j and r_k

The set U_{covered}^N can then be defined as the set of covered users across the network as a whole:

$$u_i \in U_{\text{covered}}^N \text{ iff } \exists A_j \text{ such that} \\ u_i \in U_{\text{covered}}(A_j)$$

Definition 4-9. *The set of served users.*

This set, $U_{\text{served}}(A_j) \subseteq U$, contains all the users considered served by the sector A_j within the network.

$$u_i \in U_{\text{served}}(A_j) \text{ iff } u_i \in U_{\text{covered}}(A_j) \\ \text{and } u_i \in U_{\text{Cap}}(A_j) \\ \text{and } C2I_W(u_i) > C2I_{\text{threshold}}$$

where $C2I_{\text{threshold}}$ is the minimum required C2I threshold, chosen by the planner before optimisation.

The set U_{served}^N can then be defined as the set of served users across the network as a whole:

$$u_i \in U_{\text{service}}^N \text{ iff } \exists A_j \text{ such that} \\ u_i \in U_{\text{service}}(A_j)$$

Definition 4-10. *The set of diverse users.*

This set, $U_{\text{diverse}} \subseteq U$, contains all the users that can receive an adequate signal (i.e. one that is above threshold) from more than one sector across the network.

$$u_i \in U_{\text{diverse}} \text{ iff } \delta_{\text{covered}}(u_i) \\ \text{and } \exists A_j \text{ such that,} \\ A_j \neq A(u_i) \\ \text{and } u_i \in (U_{\text{covered}}(A_j))$$

The set U_{diverse}^N can then be defined as the set of served users across the network as a whole:

$$u_i \in U_{\text{diverse}}^N \text{ iff } \exists A_j \text{ such that} \\ u_i \in U_{\text{diverse}}(A_j)$$

Definition 4-11. *The Cellular Network Planning Problem.*

The CNPP can then be formally stated, with reference to the previous definitions, as follows:

$$\begin{aligned} &\text{maximize} && f_{\text{COST}} \\ &\text{where} && f_{\text{COST}} = \sum_u p(u) \delta_{\text{served}}(u) R(u) T_{\text{RP}} \\ & && - \sum_B \delta_{\text{on}}(B) (F_I(B) + F_A(B) T_{\text{RP}}) \\ & && - \sum_A (F_I(A) + F_A(A) T_{\text{RP}}) \\ & && - \sum_u p(u) \delta_{\text{served}}(u) (F_I(u) + F_A(u) T_{\text{RP}}) \\ &\text{subject to} && \frac{|U_{\text{covered}}^N|}{|U|} \geq \rho_{\text{covered}} \\ & && \frac{|U_{\text{served}}^N|}{|U|} \geq \rho_{\text{served}} \\ & && \frac{|U_{\text{diverse}}^N|}{|U|} \geq \rho_{\text{diverse}} \end{aligned}$$

where: T_{RP} , is the return period for the optimisation,

$\rho_{\text{covered}}, \rho_{\text{served}}, \rho_{\text{diverse}} \in [0,1]$ denote constraints defined by the planner at the start of the optimisation.

4.2 The Problem Size

As discussed in chapter two, the CNPP is classified as NP-Hard. The following is included to give some idea of the complexity of the problem defined above.

Assuming a simplified network (based upon the Malvern scenario described in appendix I) the following would apply for the base stations:

- 20 candidate base station sites.
- 4 different antenna types.
- 360 possible azimuth pointings
- 30 possible tilt pointings.

The problem would then have 2^{864000} possible network configurations (solutions). If the user information is also included-

- 7000 potential subscribers.
- 10 RTPs for each user.
- 3 different antenna types.
- 360 possible azimuth pointings.
- 30 possible tilt pointings,

-then the number of configurations grows to $2^{864000 \times 2268000000}$ or $2^{1.9 \times 10^{15}}$. This definition allows for only one antenna at each sector, and does not take into account the potential power levels of the transmitters.

Chapter 5

The AgentOpt paradigm

This chapter presents a novel scheme to find profit-optimised BFWA networks by using the principles of emergent intelligence.

5.1 Adapting emergent intelligence techniques to cell planning

A simple way of introducing the ideas behind the scheme is to reference the rules quoted in Chapter 3, from the book *Emergence* [74]. The five fundamental principles that, as laid out in the work, should lead to a system where “macro-intelligence and adaptability derive from local knowledge” are applied to this problem as follows:

1. **More is different.** The obvious place to start when attempting to devise an emergent system for the CNPP is with the most populous entity – the potential users of the network. Giving each potential user autonomy and simple reactive rule-based behaviour will drive the optimisation scheme. We shall call these entities *User-Agents*. However, a network must have infrastructure to provide service to the users so we introduce a second tier of autonomous agents based upon the potential base sites. These are called *Site-Agents*.
2. **Ignorance is useful.** The *User-Agents* have no direct input into the global optimisation. Their only consideration is themselves and their actions are limited to a small number of choices. The *Site-Agents* are a slightly different matter. The *Site-Agents* are far less populous than the

User-Agents and are not expected to display swarm characteristics. This means that they require more intelligence if the scheme is to be successful. However, each individual Site-Agent is still only concerned with their own profitability and has no concept of the global solution. This general ignorance of the overall situation enables choices to be made by the agents that may be considered globally bad, in a similar manner to the ability of SA and Tabu-Search to take backward steps.

3. **Encourage random encounters.** The majority of decisions taken within the optimisation have a random element, as explained further in the chapter.
4. **Look for patterns in the signs.** Both User- and Site-Agents must have certain pattern recognition abilities in order for the optimisation to progress. But the question is what sort of patterns should they recognise? Whilst the User-Agents must look for signs that will lead them to service, the Site-Agents will employ more sophisticated data analysis to identify patterns that inform them when and how to configure network infrastructure.
5. **Pay attention to your neighbours.** Communication between the individual User-Agents is vital in order to disseminate local information and create the feedback loops so important in emergent systems.

The use of two types of entities in one system is an attempt to address the multi-objective nature of the CNPP, as discussed in chapter 2. An individual User-Agent in the network does not care if the network is profitable. The Site-Agents are interested in service levels only so far as the revenue they raise. Having the two different species of autonomous agents, motivated by different factors toward different aims allows, as a direct consequence, the dual-objective nature of the problem to be considered, without resort to

constraint or weighting methods. There is no aspect of the scheme that aims to optimise the global network profit - the global optimisation comes from the co-operation between all the elements.

5.2 The AgentOpt scheme structure

The optimisation scheme implemented within the AgentOpt (AO) tool is iterative in nature with three distinct phases: User, Site and Control. During the User-Phase each User-Agent in the system is activated. Once activated the agent assesses its current situation and makes an analysis of any local information it is party to. The User-Agent then decides on and takes action, its behaviour dictated by the rules described in section 5.3. The Site-Phase involves a similar process for each Site-Agent, and is described at length in section 5.4. The Control-Phase acts separately from the emergent system, assessing each network design generated and saving the best found so far as the current best solution. Figure 5-1 gives a simple overview of the way that local information flows through the system.




Figure 5-1. The flow of local information through the AgentOpt system

The behaviour of the optimisation is governed by a number of parameters that are set in the initialisation file before the tool is run. A number of the more important of these are described in the following outline of the tool structures. A full list of the parameters may be found in Appendix III.

The optimisation starts with each User-Agent petitioning the Site-Agent from which they receive their strongest signal (this signal is purely the path loss value generated by the RPD and passed into the tool). Other options for the start of the scheme include a random starting pattern or the geographically closest site. Investigations into these options are discussed in detail in section 5.5.1.

5.3 User-Agents

The potential users within the network are a constant quantity. That is to say that the number of potential users or buildings within the network scenario will be the same regardless of the final network planned. These potential users will require service in order to become a part of the network. Therefore potential users can be envisaged as having a some sort of desire to be served. These are two good reasons for choosing the potential user as an intelligent entity in any agent based scheme, as well of course as the large population level which is very important in emergent systems. The focus on service in particular, should continue to give the optimisation the impetus to search widely for networks with high levels of coverage and service. In order to obtain service the User-Agents will *petition* the Site-Agents, demanding coverage. This petition is a formal declaration by the User-Agent to a Site-Agent that the User-Agent wishes to be served and furthermore wishes to be served by that Site-Agent.

The network structure defined in section 4.1, forms the basis of the User-Agent entity in AgentOpt. The transformation of a potential network user into a User-Agent involves the addition of parameters such as are needed to allow each user in the network to behave independently. A formal definition of the User-Agent is given over the page.

Each User-Agent exists in one of the four states set out in definition 5.2, also shown over the page. The state of each User-Agent provides the system with basic information about the network being planned.

Definition 5.1 User-Agents

U_A	is defined as the Set of User-Agents
$u_{Ai} \in U_A$	is an individual User-Agent defined as follows:
$u_{Ai} = \{u_i, d, r, s, e, c, Nu_i\}$	
	where:
$u_i \in U$	represents the user parameters from the network structure,
$d \in [0,100]$	denotes the User-Agents Distraction level,
$r \geq 0$	denotes the User-Agents Resentment level,
$s \in \{0,1,2,3\}$	denotes the User-Agents current State,
$e \in \{0,1,2\}$	denotes the User-Agents reason for its current State
$c \in [0,100]$	denotes the relative charge ratio,
$Nu_A = \{u_{Aj}, \dots, u_{Ak}\}$	denotes the set of Neighbours of u_{Ai}

Definition 5.2 User-Agent States

- $s = 1$ - The User-Agent is petitioning site x for service.
- $s = 2$ - The User-Agent is covered by site x and petitioning for service.
- $s = 3$ - The User-Agent is served by site x .
- $s = 0$ - The User-Agent is dormant.

The User-Phase proceeds as shown in the pseudo-code description provided in Figure 5-2 below.

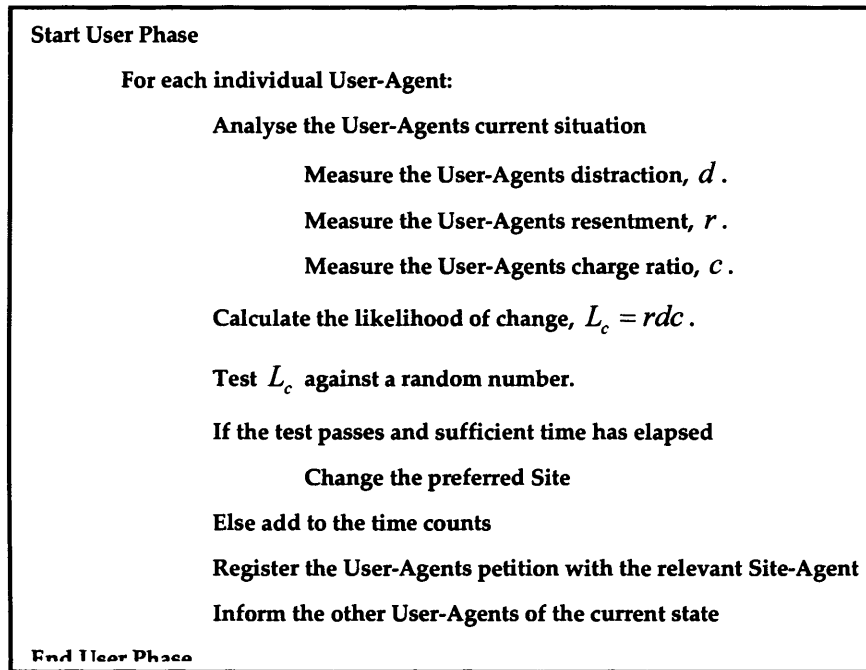


Figure 5-2. A pseudo-code representation of the User-Phase

5.3.1 Neighbours

One important aspect of swarm logic and emergent systems is the random encounters between individuals in the system. The User-Agents within this scheme are obviously geographically static and so, in the absence of physical movement which would enable these chance meetings, each User-Agent has a set of neighbours with whom they can communicate. This set, Nu_A , is defined as all the User-Agents that are positioned within the area bounded by a circle of some radius, R , centred on user u_{Ai} . The size of R will obviously affect the number of neighbours within Nu_A , and so also the efficiency and

effectiveness of the optimisation scheme. This is explored further in section 5.5.

Neighbours are able to inform each other of their current intentions, thus providing the system with local information. User-Agents are able to communicate their current state, along with their preferred site, to their neighbours during each User-Phase. A User-Agent makes an assessment of its own current situation, and of the situations of its neighbours. Local feed-back loops are then created as the information gathered by each User-Agent from their neighbours informs any changes that are made to the individuals current situation.

5.3.2 Distraction

Each User-Agent keeps track of which Site-Agents their neighbours are preferring. The percentage of a User-Agent's neighbours that are not considering that User-Agent's preferred Site-Agent is called the level of Distraction, d . A User-Agent can be said to be highly distracted (e.g. $d \geq 50$) if a large proportion of its neighbours are directing their attention elsewhere. This is further refined by a consideration of the neighbours state. A neighbour is only worth listening to if they are in the same or a better state. This means that a served User-Agent will take no notice of User-Agents who are only covered. These inferior neighbours become irrelevant- regardless of how many there are. There is no comradely solidarity between the agents.

This parameter can be seen as an indicator to the 'goodness' of the current solution. In figure 5-3 a) the highlighted User-Agent has no real choice. Where only one site can serve a set of users then the distraction level will be zero. In figure 5-3 b) however, we note that a number of sites could serve this particular user and more importantly, the neighbours. Therefore, we may expect stages where the distraction level would be high, with various users

petitioning or receiving service from various sites. This would not be a good idea in terms of the global optimisation. Over this area we would expect a single site to be sufficient, or perhaps a combination of sites that concentrate on the surrounding areas. So the User-Agents are encouraged to migrate to the most popular sites. This change in configuration will lower the Distraction, stabilise the optimisation and increase the 'goodness' of the solution.

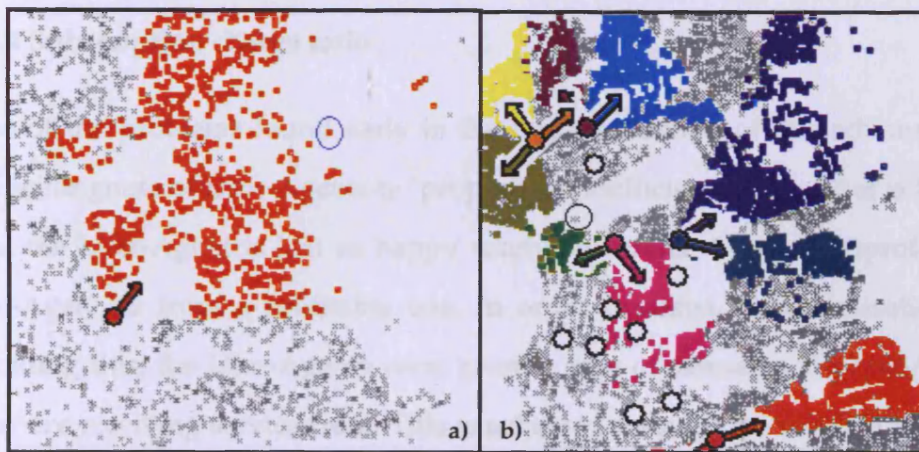


Figure 5-3. An example of a) a user with no alternative, and b) a user surrounded by sites (though not sectors).

5.3.3 Resentment

User-Agents respond to the activities of their neighbours, but also to the activities or lack thereof of their chosen Site-Agent. A User-Agent is justifiably unhappy with a Site-Agent that is paying it no attention, and if this continues over time then the level of unhappiness with the current network solution (from their own selfish point of view) grows to significant levels. This unhappiness is called Resentment, r , and is measured in the following way:

Definition 5.3 Resentment

- In State 1(petitioning) $r = t$
- In State2 (covered) $r = \frac{t}{10}$
- In State3 (served) $r = \frac{t}{100}$

where t is the time in the current state.

5.3.4 Subscription charge ratio

One of the problems found early in the implementation of this scheme was the willingness of User-Agents to ‘prop up’ an inefficient sector. That is to say that the User-Agent is just as happy receiving service from an unprofitable Site-Agent as from a profitable one. In order to assist the optimisation in avoiding this, the User-Agents were given a way of assessing the Site-Agent they are receiving service from. This is achieved through the use of a relative subscription charge ratio, c . In the standard network model each potential user is offering a certain amount of revenue, c_{offer} , to the network in exchange for service. In the AgentOpt scheme each User-Agent keeps this value but also calculates a subscription charge that reflects the number of users served by the particular Site-Agent. The Site-Agent assesses its total costs and divides this number by the total number of User-Agents served, to give an income per User-Agent necessary for the Site-Agent to be profit neutral, c_{demand} . The User Agent receives this charge and compares it with the agreed revenue charge.

Definition 5.4 *Subscription Charge Ratio*

$$c_1 = (c_{demand} / c_{offer})$$

$$c = \begin{cases} 1 & \text{if } s \neq 3 \\ c_1 & \text{if } c_1 < 2 \\ 2 & \text{otherwise} \end{cases}$$

So the Charge Ratio is defined as the fraction by which the charge exceeds the agreed subscription rate, providing that ratio does not exceed 2 and the User-Agent is being served. This value is then used to either increase or reduce the overall probability of a User-Agent changing it's preferred site. It is worth being clear that this value is only used by the User-Agent in order to determine its next action. The Site-Agent and indeed the global network income levels only consider the User-Agent's offered subscription rate.

The User-Agent then takes this value into consideration when deciding to change it's preferred site. This can assist in the migration of users between sectors as explored in the Site-Agent section.

5.3.5 Changing Sites

The likelihood of a user making a change of preferred site, L_c , is determined by the User-Agent's level of displeasure and the amount of neighbours behaving differently as in definition 5.5.

Definition 5.5 *Likelihood of a change of preferred site*

$$L_c = rdc$$

A random number, n , such that $0 \leq n < 100$, can then be generated and compared with L_c so that:

If $n < L_c$, then the User-Agent determines to change it's preferred Site-Agent.

The new Site-Agent to whom the User-Agent will petition for service is chosen from the two most popular in the User-Agent's neighbourhood. If the current preference is the most popular, another random decision, weighted by the percentage popularity of the Site-Agent in question, is taken to determine whether this Site-Agent is worth persisting with.

As can be seen, the User-Agents levels of Distraction and Resentment drive the optimisation into investigating new solutions. These are the feedback loops so important in any emergent system. So, User-Agents in Site-Agent dense areas, with high levels of distraction, are more unstable. If a User-Agent has no real choice about the Site-Agent then d is low – perhaps zero – and there is little likelihood of a change. The same applies to the resentment levels. User-Agents should usually only have high levels of resentment when ignored for a long time and receiving no response from the preferred Site-Agent. However, resentment does grow over time even in the User-Agents that are served. This destabilises the network and enables the tool to try a different solution. Because the new solution is driven by the information gathered during the optimisation of the previous network, the tool is automatically steered away from unfeasible solutions.

5.3.6 Dormancy

This aspect of the scheme can create difficulties however and so there is a need for a cooling effect within the optimisation scheme. As previously discussed the scheme is iterative in nature and the number of iterations is determined by the user of the tool. However the scheme was also envisaged to run in cycles, albeit random and inconsistent cycles of the schemes own making, as described above. The ability of a User-Agent to become *dormant*

has been included in the scheme in an attempt to remove from the potential network a certain level of instability, in order for a more focused period of optimisation, similar to the concept of cooling in Simulated Annealing.

In some network geographies, or under certain constraints or network demands, there may well be a set of User-Agents, perhaps around the periphery, that are unlikely to ever be served. These User-Agents will however still be petitioning the Site-Agents and thus add instability to the network without any positive result. The dormant state is included to recognise these users and temporarily remove them from the optimisation. When a User-Agent stays in the same petitioning state, with no recognition, for an amount of time they will become dormant. This amount of time should be such that if the User-Agent has a number of options as to its preferred site it will be able to try them all out. Whilst in this state the User-Agent does not petition Site-Agents. A random decision is made as to when the user will come out of its dormant state – this is based upon the time in state and the happenings in the neighbourhood.

When in the petitioning or covered state the preferred Site-Agent will inform the User-Agent of the reason for the lack of service as shown in definition 5.6.

Definition 5.6 *Reasons for lack of service*

- $e = 0$ - No reason given.
- $e = 1$ - Not enough capacity at the sector.
- $e = 2$ - A problem with the C2I.

Although the User-Agent does keep track of the current reason, no benefit to the User-Agent can be gained by acting upon it. This information is only used by the Site-Agents, as set out in the next section.

5.4 Site-Agents

Another constant in the network definition of chapter 4 is the list of potential sites. As with the potential users this list will remain constant regardless of where sectors are actually placed. Unlike the User-Agents however, Site-Agents cannot be allowed to focus solely on fulfilling their potential for provision of service. Instead, Site-Agents are driven by profit. In order to realise this profit a Site-Agent can create, modify and destroy sectors at their own geographic site, and so create revenue for themselves by serving users. Global profit optimisation comes from the competition amongst the individuals within the system. As opposed to User-Agents, Site-Agents are more active and will attempt to make a change to their own configuration each iteration. However, as in the TS optimisation scheme, this change will not always be accepted.

Definition 5.2 Site-Agents

BS_A	is the set of Site-Agents,
$bs_{Ai} \in BS_A$	is an individual Site-Agent and is defined as follows:
$bs_{Ai} = \{bs_i, Lbs_i, Sbs_i, T_{petition}\}$	
	where:
$Lbs_{Ai} = \{l_0, \dots, l_{359}\}$	denotes the set of petition lists with $l_j = \{p_m, \dots, p_n\}$ as a single petition list for degree j and p_i is the petition containing information from u_{Ai} ,
$Sbs_{Ai} = \{s_0, \dots, s_n\}$	denotes the list of active sectors at the site with s_k being the k^{th} sector at the site,
$T_{petition} > 0$	denotes the Petition Threshold.

Broadly speaking, the Site-Agents act in a similar manner to the controlling operations of the ECHO Tabu-Search scheme, but with local influence only. Based upon local information gathered and processed the Site-Agent is able to:

- Create a new sector at its Site
- Modify an existing sector at its Site
- Destroy a sector at its Site

The local information comes in the form of requests from the User-Agents. Each Site-Agent receives a *petition* for service from a number of User-Agents in the potential network. When the Site-Agents receive petitions from User-Agents these petition details, p_{ui} , are stored, indexed by the bearing from the potential site to the potential user. This allows the site to build up a geographic picture of its petitioning users and enables a quicker assessment of the information offered as a whole.

5.4.1 The Site-Phase

In a manner similar to the User-Phase, the Site-Phase involves iterating through each Site-Agent, in a random order, analysing the information that each individual has received and acting accordingly. A pseudo-code overview of the Site-Phase is shown in figure 5-4, and each aspect is then described in more detail.

Start the Site Phase

Randomise the list of sites and loop through it.

For each individual Site-Agent:

Determine if new sector is required (Section 5.4.2).

Try to merge existing sectors (5.4.3).

Try to improve existing sectors (5.4.4).

Try to destroy an existing sector (5.4.9).

If a new sector is required then create a new sector

Re-randomise the list of Site-Agents and loop through again.

For each individual Site-Agent:

Try to merge sectors at different sites (5.4.10).

Re-randomise the list of Site-Agents and loop through again.

For each individual Site-Agent:

Reassess each sector.

End the Site Phase.

Figure 5-4. The Site-Phase algorithm

5.4.2 Is a New Sector Required?

In order for a network to exist, some sectors must be installed at the potential base-sites. In order to determine whether it is necessary for a Site-Agent to




Figure 5-5. The segment chosen by a Site-Agent.

place a new sector at its site, a number of calculations must be made. The first consideration is the area in which the Site-Agent should concentrate. As mentioned previously each petition, once received, is stored by the Site-Agent with the index corresponding to the direction from Site-Agent to User-Agent.

The Site-Agent will find the segment of its surroundings that contains the most petitioners, as shown in Figure 5-5. In all

the investigations in this thesis a segment of a maximum of 90° is used, which corresponds with the widest antenna beam-width used in the scenarios. In other investigations this value may be altered within the initialisation file, depending upon the beam-width of the available antennae. If there are existing sectors at the site, the chosen segment will not intrude upon their domains. In this way the segment chosen may be less than the maximum defined.

Once the segment has been found the Site-Agent makes a comparison. If the total number of petitions within the chosen segment exceeds the petition threshold, t_p , the Site-Agent decides that a new sector is required. Although this assessment of the need for a new sector is performed at the start of the phase, the new sector is not configured until after the existing sectors have been assessed. This is in order to both increase the likelihood of a new sector being needed whilst also giving those existing (and so more developed) sectors a chance to pick up all the new petitioners that they can.

If the Site-Agent decides to create a new sector the configuration depends upon a combination of local information and random chance. The azimuth and tilt are determined by taking an average of the vectors to the petitioners within the chosen segment. The antenna is determined by the width of the segment. All other aspects of the configuration, including transmission power and the number of capacity channels, are randomly selected.

5.4.3 Merging Same-Site Sectors

The AgentOpt optimisation scheme can be thought of as evolutionary in its process and so it is important for the individual Site-Agent to have the ability to merge together its own sectors. This allows the Site-Agent to keep hold of subscribing users whilst at the same time reducing the overall number of

sectors at the site. The distinction between merging and destroying a sector must be made as the actions of the Site-Agent are considered by the User-Agents. If a sector was merely removed, and thus the service to users halted, each affected User-Agent would have elevated levels of resentment towards the Site-Agent and take its petition elsewhere. The other consideration is whether a merger between two sectors located at different sites would be useful. This is not considered until later in the phase, however.

The method of finding a successful merger is very simple. Each sector at the site is compared with every other, with the Site-Agent looking for a pair of sectors such that:

- Each sector has been in the network for the minimum period,
- Each sector is not of the maximum beam-width,

and either

- The domain of one sector encroaches on the domain of another,

or

- The antennas domains separately cover an area that could be covered by a larger antenna.

If a pair of sectors is found such that either, or both, of these conditions are met, both sectors are tested to determine their suitability to provide service to the other sector's users. This suitability is measured in terms of the percentage of new users that a sector can provide a signal above threshold to. Once the best sector is determined all the users are transferred over to it and the losing sector is removed from the network.

5.4.4 Improving Sectors

Site-Agents must also be able to configure every aspect of the sectors that they place. The decision making process behind the sector configuration is again driven by the local knowledge the Site-Agent has stored about its petitioners. For example, the Site-Agent will be able to determine that of the petitioners in a certain area, where there exists a certain sector, a larger percentage of these petitioners are covered rather than served by this sector. This is achieved in a number of steps as laid out in figure 5-6.

```
Start the sector improvement methods
  For each sector at the site
    Get the list of petitioners within the sectors domain and determine the size,  $N_p$ .
    If  $N_p \neq 0$ , i.e. the list is populated,
      Determine the number of petitioners that are covered but not served,  $N_{cp}$ .
      Determine the 'Covered' percentage,  $P_c = (N_{cp} / N_p)100$ .
      Determine the number of petitioners that have capacity problems,  $N_{ip}$ 
      Determine the 'Traffic' percent,  $P_t = (N_{ip} / N_p)100$ .
      Determine the number of petitioners that have interference problems,  $N_{ip}$ 
      Determine the 'C2I' percent,  $P_i = (N_{ip} / N_p)100$ .

      If the sector has merged in this iteration, try to increase the capacity (5.4.8),
      else if  $N_p < m_p$  or  $t_{idle} > T_{time}$ , try to reduce the sectors costs (5.4.5),
      else if  $P_c < T_c$  try to increase the coverage (5.4.6),
      else if  $P_t > P_i$  try to improve the channel assignment (5.4.7),
      else try to increase the capacity (5.4.8).

End the sector improvement method
```

Figure 5-6. A pseudo-code representation of the sector improvement method.

5.4.5 Reducing the cost of sectors

A Site-Agent may attempt to reduce the costs associated with a successful sector for one of two reasons. As described in figure 5-5, once the Site-Agent has collected and analysed the local information the first test made is to determine whether the number of petitioners within the chosen area exceeds a minimum threshold, m_p , set by the planner in the initialisation file (see Appendix III). If the minimum threshold is not reached the decision is made to attempt to consolidate the existing configuration.

The second part of the decision process concerns the amount of time the sector has existed in its current state without a change. Whilst the sector is being regularly altered the costs can be somewhat ignored. However, once the sector has reached a level of optimality, and no changes are being made, a timer is set to keep watch on how long this state lasts. Once this timer reaches a pre-defined threshold, the Site-Agent can attempt to reduce the costs of the sector.

It is worth noting that the costs the Site-Agent is attempting to reduce are not limited strictly to the financial, but are mainly the 'cost' to the network as a whole, including inefficient use of transmission power and capacity channels.

The method proceeds by first assessing the capacity made available by the sector and determining its use. If there is found to be excessive capacity provision a further survey of the subscribing User-Agents determines the channel that is most likely to be expendable. This channel is then removed, the new sector configuration analysed and a determination of the changes worth is made. This worth is determined by the rise and fall of two parameters- the overall number of subscribing User-Agents and the cost of the reception equipment used by these subscribers- along with a small random element to encourage the investigation of the search space. If the

change is thought to be worthwhile it will stand, if not then the old configuration will be restored.

The method then proceeds to attempt to reduce the transmission power used by the sector in a similar manner, all the while attempting to choose cheaper infrastructure options whilst maintaining the integrity of the original configuration.

5.4.6 Increasing Coverage

When attempting to improve a certain sector configuration, the Site-Agent assumes that theoretically all the petitioners in the associated segment could receive a signal above threshold. That is, every User-Agent within the segment *could* be covered by this sector. Therefore, a useful measure of the 'goodness' of the current sector configuration is the percentage of petitioners that actually are covered, P_c . The coverage threshold, T_c , is used to determine whether the Site-Agent needs to increase the sector's coverage through some means or would be better served by concentrating upon improving service levels. If $P_c < T_c$, where T_c has been set by the planner in the initialisation file, then the Site-Agent will concentrate upon the coverage.

The Site-Agent has four options to attempt to increase P_c :

- Adjust the antenna azimuth.
- Adjust the antenna tilt
- Adjust the transmission power
- Change the sector antenna

The decision as to which option to concentrate upon is made randomly, with each different aspect following the same basic pattern.

For example, if the Site-Agent has chosen to concentrate upon the sector azimuth, the first task is to survey the petitioners and obtain an average azimuth. The existing bearing is then nudged a small but random amount toward the average. The new configuration is then tested and will be accepted dependant upon bringing new users into the sector. There is also a small random chance that a bad solution may be accepted.

In order to enhance the opportunity for fine-tuning, the Site-Agent will repeat the process a number of times, stopping if an improvement is found or a certain number of iterations is reached.

5.4.7 Improving the channel assignment

As outlined in figure 5-6 the Site-Agent is aware of the percentage of unserved User-Agents that have problems with interference, P_i , and the percentage that are traffic limited, P_t . When $P_i > P_t$ the Site-Agent attempts to improve the channel assignment. In order to provide service to a user, a sector must have a channel assignment such that for each channel in the assignment the user will have a C2I ratio above threshold. So in order to improve the channel assignment and serve more users, the Site-Agent can either change the assignment or change the polarisation of the channels. The network scenarios used in this thesis allow each antenna to have either a vertical or horizontal polarisation, thus effectively doubling the number of channels available. This option is scenario specific but can be very useful if available.

The likelihood of changing the polarisation is arbitrarily set at 20%. This level is chosen because although a change in polarisation can be very effective it is also rather drastic and therefore it is preferable to force the optimisation into investigating different channel assignments. As with all other changes to a

sector's configuration the change of polarisation is enacted, then analysed to consider whether to keep the alteration.

The alteration of the channel assignment is carried out in a semi-random fashion. Each subscribing User-Agent is surveyed for the channel that inflicts the worst C2I ratio. This channel is then swapped with another randomly chosen one. Again, an analysis is made of the suitability of the change and the decision made to either keep or repeal the change.

5.4.8 Increasing capacity

If the number of capacity limited users is higher than those suffering interference then the Site-Agent must attempt to add more capacity to the sector. Firstly, a survey of the shortfall in required capacity is taken. Then new channels, chosen at random, are added to the sector, up to the level of the shortfall. Again an analysis is made of the impact of each new channel added, with the option to remove any new channels that adversely affect the sector.

5.4.9 Destroying sectors

Site-Agents are able to remove a sector from the network, if that sector is deemed unfeasible or unprofitable, and this is the next consideration of the Site-Agent within each Site-Phase. A sector should be removed if all its users have deserted it, or, over a period of time, it is unprofitable and not attracting new users. As the User-Agents swarm from one site to another, over the course of the optimisation, many sectors are created and destroyed. However, Site-Agents are encouraged to be patient with the sectors they create. As with the User-Agent's level of resentment, Site-Agents initially make only smaller, more focused changes.



When looking for redundant sectors the Site-Agent cycles through each one surveying the profit that the sector brings in and also the number of User-Agents that are served. The worst performing sector – either one with no users or the one with the lowest profit - is then chosen. Further tests are made before this sector is destroyed. The sector must have been in the network for a minimum number of iterations. This number is set by the planner in the initialisation file and is designed to stop the optimisation creating and destroying the same sector repeatedly. This will also encourage the destruction of sectors that have been modified but continuously fail to show any improvement. The final test that will condemn this type of sector to death is a measure of profit level. The sector will only be destroyed if the profit level falls below a given threshold (the default value is 0). The removal of the sector must also involve informing its subscribing User-Agents (if there are any), and thus increasing their resentment against the Site-Agent.

5.4.10 Merging sectors from different sites

Once each Site-Agent has gone through the steps outlined above the option to merge together sectors from different sites is explored. This aspect of the Site-Agents' behaviour can be thought of as an extension of the co-operation, or competition, that exists between themselves. In a situation where two Site-Agents have sectors attempting to serve the same set of users the network will be fundamentally unstable. This comes from high levels of distraction amongst the User-Agents. As each Site-Agent attempts to satisfy more of the users a 'critical mass' point is reached and one of the competing sectors will become infeasible. At this point it will be removed as described above. However, it may occur that two barely viable sectors can co-exist. In this case the User-Agents are aware of the numbers of their neighbours being served from elsewhere and this information is passed to the Site-Agents. Inter-Site-

Agent communication takes place, in order to agree to a merging of these sectors.

Each Site-Agent makes an analysis of its associated User-Agents – both subscribers and petitioners – to find the competing Site-Agent that is causing the most distraction in the flock. If two Site-Agents find that they are each competing with the other, negotiations are opened into the possibility of merging two sectors. Both Site-Agents put forward a sector belonging to it's competitor that is believed to be the most promising candidate for a merger. This is determined by the subscribers that each Site-Agent serves. As the Site-Agent loops through each of its own sectors, it interrogates each subscriber for the main source of the subscribers distraction. In this way the Site-Agent builds a picture of it's competitor and the associated sector that is causing most offence. The two sectors are then tested to determine which of the two sectors are more able to cope with the other's subscribers. This test is limited to the ability to provide signal above threshold to the new subscribers as other aspects – such as interference problems and capacity provision – can be resolved after the decision has been taken.

5.4.11 Reassessment of each Site-Agent

At the end of each Site-Phase a final analysis is carried by out each agent to determine how their sectors and subscribers have been affected by the other Site-Agents in the network. This is a reasonably simple retesting of each Subscriber and petitioner in order to build up accurate information for the start of the next Site-Agent phase.

5.5 Evaluation of the effect of system parameters on AgentOpt

This section looks into how various changes in some system parameters affect the network solutions produced by AgentOpt. All of the investigations make use of the Malvern data scenario that can be found in appendix I.

5.5.1 Starting the optimisation

The AgentOpt scheme depends upon local information to evolve the network designs, so some thought must be given to the start point for this evolution. Starting with a completely blank slate will not work – the individual User-Agents must have some information to disseminate. During the design of the AgentOpt tool, two alternative options to start the optimisation process were considered.

The first option is to start the optimisation with each individual User-Agent petitioning to a random Site-Agent. This option is easily achieved and also encourages the diversity of solutions that is desirable in meta-heuristic schemes. However, it must also be considered that any local information is essentially worthless in the early stages of a network design.

The second option is to start the optimisation with each individual User-Agent petitioning a Site-Agent that can be considered the best option by some metric. This approach has the advantage of giving meaningful information from the start of the run. Two metrics were chosen for testing. Either the User-Agent would petition it's closest site or they would concentrate upon the site with the best path loss.

Figure 5-7 over the page shows the expected network profit of six AgentOpt planned networks – two networks planned with each of the three starting options. All six networks were produced using the basic Malvern scenario as

described in Appendix I and the basic AgentOpt system parameters as shown in Appendix III. As would be expected from a system dependant upon local information, the four networks planned using the second approach – that of finding a meaningful start point – performed best. The method which chooses the site with the best path-loss value, performed the best on average.




Figure 5-7. A comparison of six networks produced using three different starting strategies

Figures 5-8 and 5-9 show examples of the cell plans for two of the approaches. In each of these figures the candidate BS sites are large white dots whilst the potential users are small black crosses. Sector antennae are represented by coloured arrows, with users that can receive service from that sector shown in the same colour. The random-start cell plan is visually very untidy and does not inspire confidence. The best path-loss start network plan appears to be an improvement. Although there are holes in the coverage, the cells generally look well formed. It is also interesting to compare the graph of expected profit over time for each approach. Figure 5-10 shows the progress made by the tool during the planning of two networks – one starting the User-Agents with randomly chosen sites and the other starting with the best path-loss site – over 1000 iterations.




Figure 5-8. A random-start network cell plan



Figure 5-9. A best path-loss network cell plan




Figure 5-10. The performance of each starting method over time

As can be seen, the random method gives a far more volatile optimisation. As time goes on the expected profit cycles up and down fairly regularly for three-quarters of the run. After this point the performance drops off slightly but the run is still lively especially when compared to the path-loss method. This later approach starts well and quickly finds a peak, but afterward there is a long slow descent apart from the occasional blip. What this indicates is the trade off between the two start point philosophies. The random start keeps the optimisation unstable for a longer period of time as the self-organisation occurs. The path-loss approach offers a level of organisation from the start, which the tool quickly improves upon. After this however the profit drops off

as small amounts of User-Agents become unhappy and drop out of the network. The organisation is such that they cannot easily be readmitted – at least not unless it is at the expense of another. It is worth noting that the path-loss performance is in general better than that of the random method throughout the whole time period and the differences between the best profit for the two respective methods is large. For this reason the rest of the investigations in this thesis will use the path-loss start method unless otherwise stated. However, the ideal situation would be to find some combination of system parameters that forces the AO optimisation to behave as it does with the random-start approach, whilst reaching the profit levels attained by the path-loss start method.

5.5.2 The neighbourhood radius

We now look at the system parameter that exerts the most control over the information that is distributed throughout the system.




Figure 5-11. The neighbourhood radius

The neighbourhood radius governs the number of neighbours in each individual User-Agent's neighbourhood. In figure 5-11, the User-Agent represented by the red dot obviously has a much larger

number of neighbours than the User-Agent represented by the green dot. During the design of the scheme it was decided that by choosing the neighbourhood by locality (e.g. every User-Agent within a 1000 meter radius) rather than forcing some pre-defined quantity (e.g. the 100 closest User-Agents), the information generated by each individual would be more pertinent to the actual situation. For example, if the green User-Agent above

was required to have the same amount of neighbours as the red User-Agent, green's information, and thus the choices that it makes, could be affected by neighbours so physically far away that the two situations are incomparable. This would render the information useless. Also, the User-Agent analyses it's information in terms of percentages of it's neighbours. The radius of the neighbourhood is of most interest in the way that it affects the optimisation.

Figure 5-12 compares pairs of networks generated using different values for the neighbourhood radius.




Figure 5-12. A comparison of networks generated using different radius values

As the size of the radius increases so does the expected network profit. As can be seen in appendix I the Malvern scenario contains 7085 potential subscribers, distributed fairly evenly across an area of 25km². This gives an average of 283.4 subscribers per km². Table 5-1 gives values for the average number of neighbours in a User-Agents neighbourhood for various radius values.

Table 5-1. Average user numbers

radius (m)	area (m ²)	Average users- Malvern
10	314	< 1
50	7854	2
100	31416	9
250	196350	56
500	785398	223
750	1767146	501
1000	3141593	890
1250	4908739	1391
1500	7068583	2003
1750	9621127	2727
2000	12566371	3561

Chapter 7 details a trio of Malvern scenario networks produced using a different ACP. These networks have very high levels of service – close to 100% - and the number of sectors employed ranges from 8 to 13. This gives a range of 545 to 885 for the average number of users per sector. It is interesting that the AO optimisation is at its most effective when the average neighbourhood size exceeds the average number of users in a sector. It is possible that there may be some connection between the two numbers. It is most likely however that this size of neighbourhood – on average 10-20% of the overall population - is simply large enough to enable the individual User-Agent to make a valid statistical analysis of it's situation, but small enough not to encourage unfeasible options.

5.5.3 The Petition Threshold.

The petition threshold governs how often a Site-Agent may place a new sector. Once the number of petitions in a particular segment of the Site-Agents locality reaches the threshold, a new sector may be placed. Figure 5-13 compares pairs of networks generated using different petition thresholds.




Figure 5-13. Petition Threshold comparisons

Figure 5-13 implies that the AO optimisation scheme is generally more successful when the threshold is set at 250 petitioners or less. Intuitively this would seem to be correct. A lower threshold would give a Site-Agent more chances to place a sector which then can be fine-tuned to improve the overall network. With a higher threshold the number of chances to place a sector would be reduced and search space much more limited. Additionally without many competing sectors the optimisation is somewhat forced to make do with what it has got. This is born out in figure 5-14, over the page. This compares the performance of AgentOpt over time for two networks, one with a threshold of 100 and one with 750.




Figure 5-14. Petition Threshold - Comparison of two optimisation runs over time




Figure 5-15. The PT=100 network

When using the lower threshold value the optimisation is far more active over its run. On at least eight separate occasions this run produces a network plan that produces more profit than the best network found using the higher threshold. The network it produces (see figure 5-15) contains 18 sectors, some of which do appear to be underutilised if not

completely superfluous. The PT=100 run also suffers from the drop off over time discussed in 5.3.1. The extreme petition requirement in the PT=750 network paralyses the optimisation from the start, allowing only minor adjustments to the three sectors that are placed at commencement. Figure 5-16 shows the final best network whilst figures 5-17 and 5-18 are of a sub-optimal network found after 400 iterations. Figure 5-17 shows that the network has hardly changed at all and 5-18 shows which Site-Agents the User-Agents are

petitioning. As mentioned in the last section, previously generated network solutions give an average sector size of 545 to 885 subscribers for the Malvern scenario. At first glance it seems strange that a Site-Agent has a problem collecting enough petitioners to place a sector that may still be under-subscribed, but this can be explained by the behaviour of the User-Agents. A sector may eventually grow to serve many more than the 750 subscribers required. However, to get that many impatient, resentful User-Agents together at one time – close to ten percent of the population – will be difficult. The problem will be further exacerbated by the fact that the three Site-Agents that can place a sector early on in the run will be providing constant distraction. Figure 5-18 gives an indication of the chaotic behaviour of resentful User-Agents.




Figure 5-16. The PT=750 network




Figure 5-17. After 400 iterations




Figure 5-18. User-Agents petitioning Site-Agents after 400 iterations

5.5.4 The Coverage Threshold

The coverage threshold will affect the actions of a Site-Agent during sector reconfiguration. If the percentage of covered petitioners is less than the threshold then the Site-Agent will concentrate on increasing coverage. If not then the Site-Agent will concentrate on providing service. On some occasions the lack of coverage will be able to be remedied for the majority of petitioners. On other occasions it will not be possible. The threshold should help the Site-

Agent to concentrate upon the most lucrative area in any given situation. Figure 5-19 compares the expected profits from ten AgentOpt generated networks.




Figure 5-19. A comparison of different Coverage Threshold levels

The general trend in figure 5-19 points to the threshold being most effective at a level of 50% or above. This is likely to be because service is not purely dependant upon available capacity at a sector but also the individual User-Agent's carrier-to-interference ratio. Although this ratio may well be improved by the alteration of the channel assignment, it will of course also be affected by a change in carrier signal. The Site-Agent will concentrate more of its efforts upon improving the reception of the transmission signal for User-Agents that are not considered covered if the coverage threshold is set at a higher level. In this case those User-Agents that are covered will likely also see an improvement in their received signal strength, thus also likely improving the C2I ratio and making it easier to receive full service.

5.5.5 Randomness

Although the AgentOpt scheme depends heavily upon deterministic methods there are random elements built into the scheme in an attempt to build search space diversification into the solutions that are investigated. Figure 5-20 gives a comparison of the profit levels for ten networks generated using different levels of the most prominent random system parameter, called Random Choice Percent (RCP, see appendix III). This parameter governs the acceptance of changes in the configuration of sectors. The change will automatically be accepted if an analysis shows that it is beneficial. However a change may also be accepted if a randomly generated percentage exceeds the RCP threshold.




Figure 5-20. A comparison of different levels of randomness within the optimisation

The implication of the results displayed above would appear to be that the scheme is at its best when the random element is completely removed. The diversity of solutions encouraged by the random element seem to hold the optimisation back to a certain extent. In some ways this is understandable as the changes that are made when there is no random element are then more

intelligently guided by the local system information. It must also be remembered that there are other random elements at play within the scheme, some of which, such as the User-Agents semi-random decision to change state, will feed back into the system information.

5.6 Conclusions

The AgentOpt tool achieves its primary aim of automatically producing a BFWA network plan to fit a given scenario. The effectiveness of these network solutions is most affected by the parameter that controls the amount of information available to the individual User-Agents in the system. This indicates that the optimisation is indeed being driven by the statistical analysis of local information and feedback loops that define an emergent intelligence system. The majority of the other system parameters are concerned with the decision making facilities of the Site-Agents. Whilst generally having some effect on the chosen solution, it seems clear that the Site-Agents, being more intelligent and less numerous, are less affected by the system parameters than they are by the activities of the User-Agents.

Chapter 7 goes on to compare the networks produced by the AgentOpt tool with networks generated by other means. This gives some indication of the worth of not only the network solutions but of the AgentOpt tool itself.

Chapter 6

Evaluating BFWA networks

This chapter details the design of the Network Validation Tool (NVT), which has been produced for this thesis, to assess BFWA network designs produced by ACP tools such as AgentOpt. The design of the NVT is explained in detail and results are shown which highlight various aspects of the tool.

6.1 The problem with probability

As outlined in chapter two, the preponderance of existing cell planning tools are concerned only with maximising or achieving specific coverage levels whilst minimising infrastructure costs. This is because the majority of cell planning tools, such as [8] and [11], have concentrated on GSM or UMTS networks. These networks have generally in the past only been planned with respect to the allocation of coverage across an area and perhaps some expected traffic demands. Channel assignments can be produced after the coverage network is planned, either as a final phase or using a separate tool.

6.1.1 The AgentOpt approach

In a mobile network the planner does not know where the user will physically be at any point in time. The planner also does not know how many users will be utilising the network at any time. The information used to model the networks is based on statistical probabilities relating to the network region's physical properties. Urban networks are planned with higher user density levels (and thus higher coverage requirements and expected traffic or capacity

levels) than those used in suburban or rural planning. Coverage levels of 100% are desirable to avoid users being trapped without service and becoming disgruntled with their provider, although in reality these levels are rarely achievable. In other words an automatic tool usually tries to maximize coverage whilst minimising infrastructure costs.

The approach used in AgentOpt (and also the ECHO tool) differs from this planning process in a number of ways, with user *service* as the driving factor. The aim of the optimisation is to get the largest profit from the network by providing a service to as many users as possible whilst minimizing the infrastructure costs. This makes two of the most important parameters within the optimisation the frequency assignment and the potential users. Without a suitable frequency assignment there can be no guarantee of service, and without served users there can be no income. As opposed to a GSM network model, the network model defined in chapter 4 *does* know exactly where its potential users are. This implies that the model also has a hard limit on the number of potential users across the network, i.e. service could be seen as a far more important metric than coverage, because service means that network profit can be reasonably calculated.

In planning real world BFWA networks coverage is only one requirement. A user will not pay to be covered – they will pay to be served. As already detailed (in definition 4.10 on page 70), the network profit is derived from the subscription fees of served users. Therefore the frequency assignments have to be solved as part of the optimisation process.

However, there is a statistical problem to deal with. A tool may know exactly where all of its *potential* users are, but how can the tool know which of these potential users will decide to subscribe? If the planners do know – for example if an existing network were to be replaced - then the tool can act on

that information. Otherwise we are back to statistical probability. For the EMBRACE project a basic user subscription probability of 20% was chosen, to be a reasonable level of desired market penetration across an area. This value will obviously depend on many factors; the user's income level, the availability of competing technology, the cost of subscription, and the user's current data service status are some examples. So a small business user with no access to anything except ADSL or ISDN being offered a special business rate would be far more likely to subscribe than a suburban user who has happily subscribed to cable for 10 years.

6.1.2 The planning process

During network planning AgentOpt concentrates on the down-link, assigning capacity from the base stations to users in proportion to their subscription probability. Covered users are then said to be served when apportioned capacity and passing interference requirements.

As an example, this means that 4 users with a data requirement of 128kbps and subscription probability of 25% would each need to be assigned 32kbps to be considered served. So although the network only serves one in four of the users it says it does, it also takes only a quarter of the subscription income from each user.

All this leads to the problem that because the service is based on statistical probabilities of subscription the planned network may not be the optimal implemented network.

It is reasonable to assume that this statistical approach should produce good network solutions if the users wishing to subscribe are distributed geographically evenly. But what happens if users are not evenly distributed? Would two differently configured networks of similar profit levels handle the

situation differently? If so, could one network be said to be superior to the other? And how would this 'superiority' be defined?

Then there are other questions that may be asked. What if the information used (propagation information for example or customer details) is inaccurate? Can a tool handle this? Should a tool over-engineer its solutions in order to make them more robust for implementation in the real world? How can we be sure that this approach can adequately plan the networks?

The following sections describe the design and results from a tool that was produced to answer some of these questions, and to validate the designs produced by AgentOpt.

6.2 The Network Validation Tool

The Network Validation Tool (NVT) has been produced to accurately evaluate BFWA networks, and attempt to answer the questions raised in the previous section. It does this by simulating a real world implementation of the network – i.e. users subscribing - and validating the network user-to-sector up-link.

6.2.1 The validation approach

The tool takes a network design and randomly decides which users wish to subscribe. From this point an assessment of the network's up-link signals can be made. Users are then assigned a frequency channel and interference calculated. This process is described in more detail in the next subsection. The end result is a subset of the potential users that receive an acceptable signal from a particular sector, have no up-link interference errors and have their minimum data requirements satisfied.

NVT produces statistics about networks as they *may* be implemented. This will show differences between seemingly similar designs and hopefully answer some of the questions posed earlier in this chapter. As well as giving averaged figures for outlay, income and profit, there is also information generated on the number of Carrier-To-Interference (C2I) or capacity errors that would stop a user from being provided a service.

Other options the NVT provides include the ability to set the subscription probability of the potential users or decide whether to utilise power equalisation strategies, vary the number of up-link channels or incorporate a random error into the path loss values. These properties are explained in the next section and investigated later in the chapter.

6.2.2 Up-link architecture

When considering up-link implementation we need to investigate different duplexing methods for the up-link architecture. The two to be considered are the Frequency Division Duplex (FDD) and the Time Division Duplex (TDD).

FDD is a technique where down-link and up-link transmissions are sent within separate, widely spaced frequency blocks. This is the existing, proven technology, considered to be more robust and cheaper choice from a hardware point of view. There are obvious financial and regulatory problems regarding the allocation of the extra frequency needed to achieve adequate spacing between the up-link and down-link blocks. Also the up-link to down-link bandwidth ratio is difficult to change, being a hardware constraint, meaning that as the traffic on the network changes over time (minutes and hours or months and years) there will be varying levels of spectrum use inefficiency.

TDD is a newer technique where down and up links are separated in time instead of frequency. Within this definition there are a number of degrees of freedom and two distinct types of synchronisation. Frame Synchronization means that the time frames defined for the network are aligned but the proportion of the frame dedicated to up and down-links can be different. When the time frames are aligned and also up and down-link intervals are identical the synchronisation is known as Signal Direction Synchronisation. Degrees of synchronisation can also be defined i.e. across the whole network, across a base station, within a sector, or none at all.

Conclusions drawn from [95] state:

use FDD if:

- Traffic is essentially broadcasting or constant up-link/down-link ratio
- Spectrum efficiency is less important than cost
- Predictable traffic patterns prior to development

use TDD if:

- Traffic is bursty (e.g. data traffic) up-link/down-link ratio is varying and/or broadcasting constitutes a minor part of the traffic
- Spectrum efficiency is more important than cost
- Uncertainties in traffic pattern mean flexibility is desirable

The NVT was designed to accommodate both the FDD model laid out in [95] and TDD with Whole Network Signal Direction Synchronization (WNSDS) across the entire network. This is because at any given time, using both these duplexing strategies, either the up-link or the down-link is active but never both. Therefore AgentOpt guarantees that the down-link channelisation is

sound and NVT guarantees the up-link is sound. It is worth noting that using TDD-WNSDS degrades a number of the benefits of using TDD in the first place, since it's flexibility is curtailed. Future work could perhaps include software changes that the NVT requires to check both up-link and down-link when calculating C2I and thus handle all types of TTD synchronisation.

6.2.3 NVT design

The Network Validation Tool (NVT) (see Figure 6-1) utilises the same BFWA network model defined in chapter 4. NVT takes the network designs and attempts to simulate actual users subscribing to the network infrastructure. It extends the designed network by considering the hitherto ignored up-link (signal from user to sector), attempting to ensure all potential users who may wish to subscribe will be free of interference on both links and have sufficient sector capacity reserved.

The tool attempts to model the physical implementation of the AgentOpt planned network. We call this a validation. The length of the validation – i.e. the number of simulations or individual networks to be produced – can be controlled by the user of the tool. This enables a large number of simulations to be produced easily in order for that user to build a statistical picture of the various properties of the network in question. A simplified pseudo-code description of the individual network validation used by the tool can be found in Figure 6-2 over the page. In this description, a 'good signal' is defined as a received signal strength being greater than the required threshold. Similarly, a 'good C2I level' exceeds the defined target.

Network Validation

Load Data Load Network Validate

No. of Validation Iterations: 1000 % Probability of User Subscribing: 20 Propagation Error (+dB): 5

Produce Network Files? ☐ Save Best Networks? ☒ Use All Channels? ☐

Power Equalisation? ☒

1136940.00 Average Outlay over a return period of 3.0 years

5470200.00 Average Income over a return period of 3.0 years

4333260.00 Average Profit or a return of 381.1% over a period of 3.0 years

1432.70 or 20.2% Av initially chosen users in network, with subscription probability of 20 %

549.30 or 38.3%, Av initially invalid user errors in the Network

387805.80 Kb Av of Initial capacity errors in the Network

911.70 or 12.9% Valid Users in the network from a potential total of 7085 (7085.0)

644295.60 Kb (or 62.0% of 1040000 Kb) Av of Served capacity in the network

Figure 6-1. The Network Validation Tool

The validation process is concerned with iterating through the list of potential users. For each user a check is made that a sector is assigned to serve, and the down-link signal and carrier to interference ratio (C2I) are above threshold. A random number is generated and checked against the user's subscription probability. If it passes the user is thought to have requested to receive service from the network. A propagation error value (in dB) can also be generated within user defined limits and would be added to the signal loss from the sector-to-user path. This error would affect all signals both to and from the user. Some experiments using this option are shown in later sections.

```

Loop over the list of sectors and reset capacity values to maximum.
Randomly order the list of users.
For each user in the list:
    Generate a random path-loss error and insert into the signal calculation.
    Compare a random number against the user's subscription probability.
    If the user chooses to subscribe:
        Check the users down-link signal against threshold.
        If it fails attempt to find a good signal from another sector.
        Check the users down-link C2I ratio against the threshold .
        If it fails attempt to find a good signal and good C2I level from
another sector.
        If we have a good signal and a good down-link C2I:
            Attempt to remove the desired capacity and validate the
up-link C2I (section 6.2.4) from a randomly chosen channel
at the user's designated sector.
            If the channel has insufficient capacity remaining, or the user fails
the up-link C2I check then try the other channels.
            If the user's capacity can not be supplied:
                Add the user to invalid counts.
                Attempt to add the user's capacity to another sector.
                If the user is still invalid, add to remove counts and remove from
network
            Else add the user to the invalid counts and remove from the network.
Check for further C2I errors in valid users.
End Method.

```

Figure 6-2. The main validation algorithm

Users requiring a service are then interrogated for the amount of reserved capacity they need. The investigations in this thesis assume a symmetric bandwidth requirement from all users. NVT then attempts to add the user to their assigned sector. The sector has a list of channels inherited from AgentOpt down-link assignments. One of this sector's available channels is randomly chosen and the sector interrogated for capacity and interference problems- does the channel have sufficient remaining unallocated capacity?

Are there any up-link C2I problems? If either of these questions are answered in the negative, the NVT cycles through all available other channels and repeating the process. Another option the tool includes is to use all the available channels for the up-link, in an attempt to make the frequency assignment easier. This is only possible in a FDD network as a TDD based one would be restricted to only down-link channels.

6.2.4 Uplink interference checking

Loop over each sector in the list:

If the sector is not the users serving sector:

Add the strongest signals (co-channel and adjacent channels combined) from the interfering sector to the serving sector to the total.

If the sector is the users serving sector:

Add the strongest (combined adjacent channels) signal to the total.

Compare the (carrier signal)/(worst case total interfering signal) with threshold.

End Method.

Figure 6-3. The C2I checking algorithm

The C2I check (the pseudo-code algorithm is shown in figure 6-3) is on the up-link path (the down-link is known to be acceptable by the time this check is made) against a pre-defined threshold (this may be, but does not have to be, set at the same level as the down-link threshold). It is assumed that the network infrastructure will be such that only one user in any sector may be transmitting on each channel, on the up-link, at any time. The strongest interfering co-channel and adjacent-channel signals from users served by the other sectors, are calculated. These signals are then added together to create an absolute worst case scenario – each of the strongest interferers in the

network are transmitting at the same time. As previously stated, interference from down-link paths in the network is ignored as either:-

- there is assumed to be sufficient separation between up and down link channels (FDD).
- the directional signals are separated in time (TDD-WNSDS)

The signal from the chosen user is then calculated and set against the 'worst-case' signal. The C2I can then be checked against the threshold. If the user has a satisfactory C2I ratio the NVT determines whether the power equalising option is turned on before proceeding. If it is, then the user's transmission power is reduced so that the C2I ratio is 1 dB above threshold. This will help other users in the network to pass the C2I check.

The signal calculated by the NVT is the sum of the path-loss from a user to a sector, the user's transmission power, the gain from the user's antenna and the gain from the sector antenna. The gain from the sector antenna is determined by calculating the angle between the signal received from the user and the bearing of the antenna and comparing this value with the relevant radiation pattern.

6.2.5 Post processing interference analysis.

The NVT does not attempt to simulate an operating network as such. Rather the aim is to simulate the take up of the potential users, in order to assess the suitability of the planned network, and of the subscription probability approach. The up-link frequency assignment is performed in a quasi-random fashion. Where possible the allocation of the channels imitate that of the down-link. However, the AgentOpt tool does not precisely assign an individual channel to each user because of the uncertainties of user

subscription. NVT does so randomly, in the order in which the users choose to subscribe. As shown in the previous section, there is a worst-case up-link C2I check for each user at their point of subscription. A user choosing to subscribe will be more likely to easily pass this check at the start of the validation, when the number of other users in the network is low. As more users are added to the active network, not only will their C2I check become more difficult to pass but if they do pass it is more likely that their signal will affect other users.

One way to handle this would be to choose all the users first and then produce an optimal frequency assignment. This approach does not take into account the real-world user subscription process. Users will want immediate service. If they have to wait until the optimal number of subscribers have subscribed, they will go elsewhere. Another approach would be to re-check each of the established users C2I every time a new user wishes to subscribe. This would probably be the approach taken by the service provider. The NVT takes a third approach of allowing the user to enter the network (as long as it passes its own check) and then reassessing the network in its entirety, using the same checking algorithm (see figure 6-3), at the end of the validation. The reason for this is that it is not within the remit of the NVT to solve the problems within the network, but merely to make the planner aware of them. A large number of errors in the post-processing C2I check highlights the network as having a difficult frequency assignment requirement. Another aspect to consider is that all the C2I checks performed in the tool consider the worst-case. It may be that users will have very little interference for the majority of the time, giving them an acceptable level of service. It may be that a problem affecting a number of users could be solved by altering one user's transmitter. Or perhaps the duplexing software can dispose of the problem. In any case NVT uses the check to produce another assessment of the networks

value. It should be noted that any Post-Processing Interference (PPI) errors are considered as *softer* errors than those encountered during the C2I uplink checks, in the main validation algorithm. PPI errors imply that although there are interference problems a certain level of user subscription can consistently be found across the numerous random validations. This implies that the network has some kind of inherent flexibility. A network that consistently finds high C2I up-link errors within validation does not offer such flexibility in its configurations. For this reason C2I up-link errors are considered in the expected profit, whilst PPI errors are not. They should however be considered alongside the profit level.

6.3 NVT Validation

This section presents a number of experiments carried out using the NVT. Investigations are made into various properties of the NVT tool, in particular the randomness, the length of the validation and the effects of some special options.

6.3.1 NVT output

The output from the tool depends upon certain choices made by the planner. The basic output is a file containing a line of metrics for each network tested, and a summary of the average value for each metric across the validation. The main metrics are outlined below:

- **Outlay:** The total sum of both the initial and annual infrastructure costs over the defined time period
- **Income:** The total income taken from the valid subscribing users in the network.

- **Profit:** $\text{Income} - \text{Outlay}$
- **Return:** The percentage return on the infrastructure outlay over the defined time period.
- **Chosen Users:** The total number of users that are randomly chosen to take up subscription. A percentage of the total number of users is also included.
- **Subscription Probability:** Defined individually for each user in the project file or defined across the entire network by the planner.
- **Initial Invalid Users:** The number of chosen users that initially cannot take up their subscription. A percentage of the number of invalid users is also included.
- **Invalid Signal:** The number of invalid users due to insufficient signal strength.
- **Invalid C2I_d:** The number of invalid users due to insufficient downlink C2I ratio.
- **Invalid Capacity:** The number of invalid users due to insufficient capacity at their serving sector.
- **Invalid C2I_u:** The number of invalid users due to insufficient uplink C2I ratio at their serving sector.
- **Removed Users:** The number of chosen but invalid users removed from the network due to insurmountable problems.
- **Removed Signal:** The number of removed users due to insufficient signal strength.

- **Removed C2I_D:** The number of removed users due to insufficient down-link C2I ratio.
- **Removed Capacity:** The number of removed users due to insufficient capacity at their chosen sector.
- **Removed C2I_U:** The number of removed users due to insufficient up-link C2I ratio.
- **Post-Processing C2I_U (PPI) Errors:** The number of valid users included in the network, subsequently found to have up-link C2I problems.
- **Valid Users:** The number of subscribing users in the network.
- **Total %:** The percentage of valid users across the network
- **Chosen%:** The percentage of valid users across the set of chosen users.
- **Served Capacity:** The amount of capacity allocated to subscribers (in Kb).
- **Total Available:** The total amount of capacity available in the network (in Kb).
- **Capacity %:** The percentage of capacity utilised by the subscribers.

In addition to these two outputs, the planner has the option to store the networks generated for further inspection. The planner can decide to store either all the networks generated or a selection that represent extreme values of certain metrics. These include the network with the highest profit value, the network with the lowest outlay value, the network with the most number of valid users and the network with the lowest number of invalid users.

6.3.2 An example network

Appendix I describes the Malvern network scenario which is used as a basis for this section. Figure 6-4 and table 6-1 detail a Malvern network design (based on a coverage constraint of 80% and a return period of three years).




Figure 6-4. An example Malvern network

Table 6-1. The predicted network metrics

Number of sites in use	5
Number of sectors deployed	12
Expected Income	16562400
Required Spend	1645604
Expected Profit	14916796
Coverage %	99.7
Diversity %	71.8
Number of channels in use	37
Bandwidth supplied	1480000

6.3.3 Randomness

One important aspect of the NVT tool is the choice of users subscribing. A randomly generated value is compared against a probability of a particular user subscribing. If the user chooses to subscribe the network will attempt to provide service. As this randomness is such an integral part of the tool, it is important to show that any results obtained will not be unduly influenced by the random number generator (RNG) used. Figure 6-5 shows results from 100 different full validations of a network, similar to the one detailed in 6.3.2, using different seeds for the RNG. The variation in the number of subscribers admitted to a validated network are all within 0.3% of the average value.




Figure 6-5. Differences in random seeds.

6.3.4 Length of validation

As previous described, the validation is in essence the random construction of a number of valid networks to provide statistics about the BFWA network being evaluated. So the question in this subsection is how many networks should we construct for a reasonably accurate picture of the network plan? The simple answer – as in all statistics is – the more the better, but the graphs below present some typical results from runs of 10, 50, 100 and 1000 iterations

(i.e. 10, 50, 100 and 1000 networks generated to evaluate the single network design).




Figure 6-6. Levels of Expected Profit




Figure 6-7. The Number of valid users as a percentage of the total

As can be seen in Figure 6-6, all of the runs perform similarly, when looked at on a large scale. The difference is slightly more exaggerated in Figure 6-7. The number of iterations required to obtain sufficiently representative data is of course arguable as the tool is designed to be random in nature. Unless

otherwise stated all experiments carried out for this thesis have used 500 iterations.

6.3.5 Power Equalisation

The NVT attempts to model the implementation of a network in a real-world way. That is to say that the users are integrated into the network in a random order and not all at the same time. An NVT valid network grows over time in a similar way, although once every user has been interrogated the process halts. The Power Equalisation (PE) option in the tool attempts to reduce a served user's interference effects on users yet to subscribe. The users transmission power is reduced as far as possible whilst ensuring that the transmitted signal still meets the required coverage and C2I threshold. This is a standard technique and should generally be turned on when using the tool. The following results show the differences between validations with and without the power validation.




Figure 6-8. Measures of income and outlay, with and without Power Equalisation

These results are typical of what is obtained and show the worth of implementing PE - an approximate 12% increase in the number of valid users

leads to the higher Income level. The infrastructure outlay is also slightly increased because of the need for extra user premises equipment (due to more users being able to subscribe). When implementing a network the providers will use a similar approach. For this reason PE is used within all of the following experiments.

6.3.6 Path-loss errors

Across the Malvern data set (see appendix I), the average path-loss is approximately 142 dB on each user RTP-to-base site link. So a 7dB error introduced into the link would represent a 5% deviation from the average. The NVT includes a feature that introduces an error into the link, up to a value specified by the tool user. Note that the errors are created randomly and will not necessarily be of the magnitude specified. For example, in the case of a 5%, or 7dB, error specification, the actual addition to each link, e , will be in the range $-7\text{dB} \leq e \leq 7\text{dB}$.




Figure 6-9. As the error increases the profit decreases

This result is obviously to be expected. Figure 6-9 shows the level of expected profit falling as the errors increase, whilst Figure 6-10, shows the

reason. As the magnitude of the error increases the number of invalid users in the network also increases. The signal threshold may no longer be attained or they may be suffering from interference.




Figure 6-10. Invalid users mean less profit.

Of course, large errors within the propagation information defeat the point of accurate and detailed network planning. However, this feature can be useful if the propagation information is thought to be unreliable or simply in terms of considering a worst-case scenario.

6.4 Conclusions

The NVT is able to evaluate networks produced by ACP tools such as AgentOpt and ECHO. This evaluation enables the planning engineer to investigate the properties of individual network designs and so make comparisons between otherwise similar networks. Chapters 7 and 8 investigate some particular aspects of BFWA network planning with the aid of analysis performed by the NVT.

Chapter 7

Evaluating AgentOpt networks

This chapter investigates the fitness of AgentOpt generated networks. These networks are evaluated by the NVT, improved upon and compared with networks generated by other methods.

7.1 The AgentOpt Networks

The three networks detailed in table 7-1 and figures 7-1, 7-2 and 7-3, are among the best networks that AgentOpt has produced, using the Malvern scenario, in terms of expected profit.

Table 7-1. Three AgentOpt generated networks.

Network ID	AO-1	AO-2	AO-3
Number of sites in use	5	5	5
Number of sectors deployed	9	9	9
Expected Income	6781200	6776400	6933600
Required Spend	1690740	1689180	1731320
Expected Profit	5090460	5087220	5202280
Coverage %	99.7	99.7	85.6
Service %	79.8	79.7	81.6
Diversity %	82.9	83.6	83.1
Number of channels in use	22	23	26
Bandwidth supplied	880000	920000	1040000




Figure 7-1. Cell plan for network AO-1.




Figure 7-2. Cell plan for network AO-2.



Figure 7-3. Cell plan for network AO-3.

The three networks are obviously very alike, however they were not chosen for their similarity of configuration or for their similarity of expected profit. The similarities in both profit and configuration come from the use of the start method discussed in Chapter 5. Although there is a random element in the way the optimisation proceeds all three runs would have started from the same point. A comparison of the expected profit over time for each network, figure 7-4, confirms this. Each optimisation run follows a broadly similar path – peaking early and then slowly bottoming out. It seems therefore that AgentOpt produces very similar looking networks very quickly.




Figure 7-4. A comparison of the profit over time throughout the three runs.

Table 7-2 contains data produced by an NVT analysis of the three networks at a 20% subscription take-up level.

7.2 Comparisons with manual plans

As described in Chapter 2, a simple way to construct a network plan manually is to add sectors in a clover-leaf formation to the highest sites in the

Table 7-2. The NVT analysis of the three AgentOpt networks.

Network ID	AO-1	AO-2	AO-3
Income	6889440	6953520	7077000
Outlay	1725330	1745890	1804660
Profit	5164110	5207630	5272340
% Return	299.31	298.28	292.15
Chosen Users	1424.08	1421.40	1426.66
% Chosen	20.10	20.06	20.14
Initial Invalid Users	366.32	336.44	283.68
Removed Users	190.48	153.08	75.78
Valid Users	1148.24	1158.92	1179.50
% of Total	16.21	16.36	16.65
% of Chosen	80.63	81.53	82.68
Kb Av of Served capacity in the network	870922	895434	953721
Total Available Capacity	880000	920000	1040000
% of capacity used	98.97	97.33	91.70

In each of the three networks, the NVT predicts a greater profit than that expected by the AgentOpt tool. This is a consequence of the agent-based system. A potential subscriber may be able to receive service from a particular site within the optimised network. However, if the User-Agent representing that subscriber was concentrating upon another site at the time that the best network was found, then the subscriber would not be included as served. The NVT is then able to find the site that can serve the user and so the expected profit increases. The revenue generated is linked to the number of subscribers, and the NVT measure of the percentage of total users chosen is broadly comparable to the service percentage. Therefore this increase in profit is equal to the increase in service – around 2% in this case.

7.2 Comparisons with manual plans

As described in Chapter 2, a simple way to construct a network plan manually is to add sectors in a clover-leaf formation to the highest sites in the

geographical network area. Table 7-3 and figures 7-5 and 7-6 describe two such plans – the first follows the basic rule laid out above, whilst the second improves on this plan by removing some unnecessary sectors. The data and figures have been produced by reading the manually constructed plans into AgentOpt for analysis.

Table 7-3. The manually constructed networks.

NetworkID	M-1	M-2
Number of sites in use	3	3
Number of sectors deployed	12	9
Expected Income	6612000	6621600
Required Spend	3051400	1354920
Expected Profit	3560600	5266680
Coverage %	97.85	97.04
Service %	77.8	77.9
Diversity %	74.04	68.91
Number of channels in use	48	36
Bandwidth supplied	1920000	1440000



Figure 7-5. Cell plan for network M-1.




Figure 7-6. Cell plan for network M-2.

One basic aim for any automatic planning tool would be to produce solutions that outperform a manually constructed plan. The three AgentOpt networks can be said to do that. Figure 7-8 compares the expected profit from all five networks as predicted by the AgentOpt tool. The three AO networks outperform the M-1 network by more than 40 percent. When compared to the M-2 network however the AO networks do not come out quite so well, all three show a slightly lower expected profit. Figure 7-9 compares the networks on the overall service level. In this comparison the AO networks all show a superior level of service than those of the manually planned networks.




Figure 7-8. A comparison of the expected profit levels.




Figure 7-9. A comparison of the service levels.

The AO networks do reach a larger number of the potential subscribers across the scenario. The NVT provides a similar picture of the five different networks. Table 7-4 contains the data produced by the NVT on the two manually constructed networks.

Table 7-4. The NVT analysis of the manually planned networks.

Network ID	M-1	M-2
Income	6807600	6663720
Outlay	3123120	1371870
Profit	3684480	5291850
% Return	117.97	385.74
Chosen Users	1419.40	1428.24
% Chosen	20.03	20.16
Initial Invalid Users	350.90	360.64
% of Chosen	24.72	25.25
Removed Users	280.50	307.02
Valid Users	1134.60	1110.62
% of Total	16.01	15.68
% of Chosen	79.94	77.76
Kb Av of Served capacity in the network	804063	791581
Total Available Capacity	1920000	1440000
% of capacity used	41.88	54.97

Figure 7-10 compares the three financial aspects of each networks as predicted by the NVT. Again, the AO networks expect to receive slightly less profit when compared to M-2 but the reason for this is clear. Despite the fact that the AO networks receive more income than the M-2 network and serve a greater percentage of the total users, the lower profit is as a result of the higher outlay cost.




Figure 7-10. A comparison of expected income, outlay and profit.




Figure 7-11. A comparison of the percentage of capacity supplied that is utilised.

The NVT also identifies another area in which the AO networks may be considered to outperform the manual networks. As shown in figure 7-11 the AO networks make far more efficient use of the spectrum resources that are supplied. This is of course a direct result of planning the network with knowledge of the potential subscribers as opposed to following a few simple rules.

7.3 Improving the AgentOpt networks

Intuitively, a simple way to improve upon the network plans shown in figures 7-1, 7-2 and 7-3 would be to manually merge the two sectors at the centre of each plan together at the same site. However, this is not desirable in a thesis based around automatic cell planning tools. As described in chapter 2, the ECHO tool is an ACP, utilising the Tabu-Search meta-heuristic to produce network solutions. The networks described over the page in table 7-5 and figures 7-12, 7-13 and 7-14, have been produced by using the ECHO tool to fine-tune each of the three AO networks. In each case the ECHO tool performed 100 iterations, a much smaller number than the usual ECHO run.

Table 7-5. The fine-tuned AO networks.

Network ID	AO-1-E	AO-2-E	AO-3-E
Number of sites in use	5	4	4
Number of sectors deployed	10	9	9
Expected Income	8419200	8347200	8462400
Required Spend	1935440	1782740	1775780
Expected Profit	6483760	6564460	6686620
Coverage %	0.9968	0.9969	0.9982
Service %	0.99	0.982	0.995
Diversity %	0.6957	0.6663	0.6193
Number of channels in use	33	31	32
Bandwidth supplied	1320000	1240000	1280000




Figure 7-12. Cell plan for network AO-1-E.




Figure 7-13. Cell plan for network AO-2-E.




Figure 7-14. Cell plan for network AO-3-E.

The ECHO fine-tuning has a very positive effect upon the AO networks, offering more than a 20% increase in the profit expected for each one. This is due in most part to the better use of the infrastructure. The amount that each network requires in outlay is not greatly affected but the percentage of potential subscribers that can receive service jumps greatly from around the 80% mark to virtually 100%. The cell plans themselves remain very similar, not only to each other but to the networks that they improve upon. In two of the three cases the intuitive improvement mentioned at the start of this section is performed. Table 7-6 shows the NVT analysis of the new networks, which confirms the similarities between the three networks. It is also worth noting that the new networks outperform the manually planned networks on every front.

Table 7-6. The NVT analysis of the improved AO networks.

Network ID	AO-1-E	AO-2-E	AO-3-E
Income	8245500	7941000	8103300
Outlay	1915750	1711580	1734640
Profit	6329750	6229420	6368670
% Return	330.41	363.96	367.15
Chosen Users	1426.90	1428.20	1425.95
% Chosen	20.14	20.16	20.13
Initial Invalid Users	95.05	124.75	93.25
Removed Users	45.65	88.10	73.50
Valid Users	1374.25	1324.50	1350.55
% of Total	19.40	18.69	19.06
% of Chosen	96.31	92.74	94.71
Kb of Served Capacity	975163	945931	954830
Total Available Capacity	1320000	1240000	1280000
% of capacity used	73.88	76.28	74.60

7.4 Comparisons with ECHO generated networks

The original AgentOpt networks can easily be improved using the ECHO tool, up to a level far in excess of the manual plans. However, if the ECHO tool produces networks from scratch that are far superior when compared to the improved AgentOpt networks this would count for nothing. Tables 7-7 and 7-8 and figures 7-15, 7-16 and 7-17, all to be found over the next two pages, give details of three networks generated by ECHO. These networks have been generated at random using the exact same scenario as was used for the AgentOpt networks.

Table 7-7. The ECHO generated networks.

Network ID	E-1	E-2	E-3
Number of sites in use	5	4	6
Number of sectors deployed	8	10	13
Expected Income	8295600	8210400	8383200
Required Spend	1761820	1773580	2210940
Expected Profit	6533780	6436820	6172260
Coverage	99.52	98.74	99.65
Service	97.6	96.6	98.6
Diversity fraction	60.76	53.75	59.97
Number of channels in use	29	41	39
Bandwidth supplied	1160000	1640000	1560000




Figure 7-15. Cell plan for network E-1.



Figure 7-16. Cell plan for network E-2.




Figure 7-17. Cell plan for network E-3.

Table 7-8. The NVT Analysis of the ECHO generated networks.

Network ID	E-1	E-2	E-3
Income	7522870	8209320	8443200
Outlay	1671360	1781050	2219940
Profit	5851520	6428270	6223260
% Return	350.11	360.93	280.335
Chosen Users	1420.48	1418.98	1421.2
% Chosen	20.05	20.03	20.0593
Initial Invalid Users	187.00	68.74	20.3
Removed Users	151.77	41.32	8.9
Valid Users	1253.81	1368.22	1407.2
% of Total	17.70	19.31	19.8617
% of Chosen	88.27	96.42	99.0149
Kb Served Capacity	895713	972631	997084
Total Available Capacity	1160000	1640000	1.56E+06
% of capacity used	77.22	59.31	63.9156

These ECHO generated networks vary quite considerably in almost all regards. The cell plans are visually very different, and although two of the potential sites are used by all three networks they house sectors that are configured in an assortment of ways. The number of sites and sectors in use is also unique to each network. The expected profit also varies quite considerably, and this is amplified by the NVT analysis. E-1 and E-2 show a difference in profit of almost 10%.

Figures 7-18 and 7-19 compare the improved AgentOpt networks with the three ECHO generated networks. Figure 7-18 concentrates upon the expected profit. Network E-2 comes out on top in this comparison, but the three AO networks are all within 3%. Perhaps more importantly the AgentOpt tool appears to produce far more consistent networks, and more quickly, than ECHO. This consistency is further confirmed in figure 5-19 which looks at the percentage of supplied capacity that is used.




Figure 7-18. A comparison of the expected profit levels of the improved AO and ECHO generated networks.




Figure 7-19. A comparison of the percentage of capacity utilised in each network.

7.5 Conclusions

The present incarnation of the AgentOpt tool produces networks that are generally inferior when compared to an equivalent ECHO generated network. However, the using the ECHO tool to fine-tune the AgentOpt networks produces networks that are comparable in terms of predicted profit. The last section demonstrated the fact that the networks produced using a

combination of the two tools are reasonably consistent in configuration. On the one hand this could be somewhat of a problem. By essentially restricting the search space to a much greater degree than is necessary, there will always be the possibility that the AgentOpt tool is producing sub-optimal networks. However, this can also be seen as a benefit to the planning engineer. There would be little temptation to continue generating different network solutions searching for the best possible. It is also possible that having a number of very similar solutions will allow the planner to be somewhat flexible about the final network to be implemented.

Perhaps the most convincing argument for the AgentOpt/ECHO combination is the duration of the optimisation. Using a machine with a 1.33GHz Celeron processor and 256 Mb of RAM, a 500 iteration network generation using the Malvern scenario, which was the standard used for the networks described above, will take AgentOpt approximately 2 hours. The improvement phase using the ECHO tool over 100 iterations will add another 50 minutes. The networks described above that were generated using the ECHO tool only, were produced over 1000 iterations. On the same machine this run would last for a little over 8 hours.

Chapter 8

Investigations into BFWA networks

This chapter details some investigations into various aspects of BFWA networks, and promotes the use of the NVT in successful network planning. The chapter shows how results obtained from NVT suggest some basic strategies to enable network planning engineers to use BFWA ACP tools such as AgentOpt and ECHO to full effect. The investigations in this chapter make use of both the Malvern and Sevenoaks scenarios found in appendices I and II respectively. The networks chosen for discussion have been picked as representative of the networks that AgentOpt and ECHO may produce under those particular conditions.

The first section in this chapter looks at the influence of subscription probabilities on the planned network, whilst section 8.2 investigates quality of service levels. Section 8.3 considers the impact of limiting the number of frequency channels available, whereas sections 8.4, 8.5 and 8.6 evaluate the effects of varying capacity requirements, return periods and diversity levels respectively. Finally section 8.7 scrutinises the statistical approach that forms the basis of the AgentOpt and ECHO tools.

8.1 Subscription probabilities and take-up rates.

As described at the start of chapter 6, tools such as AgentOpt and ECHO rely upon an analysis of the likelihood of each individual user's subscription. In this thesis we assume a scenario wide level of subscription take-up, based upon a service provider's expectations of market penetration. The

subscription probability value should obviously have a large impact upon the final AgentOpt generated network. Any tools that use this approach must assume that the subscription probability level will be accurate, and optimise the network accordingly. However, as with the exact physical position of users which will subscribe, the planner should be aware that the total number of subscribing users is variable. If possible a provider should try to implement a network that can both minimize the loss of profit from under-subscription and still be able to take advantage of unexpected interest. This entire chapter is in essence devoted to determining how the effect of an unknown take-up level can be mitigated, without excessive over-engineering or effect on the final operational network profitability. This specific section looks at the impact of the subscription probability alone. The experiments shown below involve 15 networks generated using AgentOpt and ECHO, three at each of 10%, 20%, 30%, 40% and 50% scenario-wide subscription probability levels.

The three networks generated at a 10% probability level are detailed in Table 8-1 and figures 8-1, 8-2 and 8-3.

Table 8-1. Networks generated at a 10% take-up rate

Network ID	10%-1	10%-2	10%-3
Number of sites in use	2	2	4
Number of sectors deployed	7	4	8
Expected Profit	3217320	3221580	2927130
Expected Outlay	985080	877020	1249470
Expected Income	4202400	4098600	4176600
% of Maximum Income	98.86	96.41	98.23
Coverage %	99.4	97.4	98.2
Service %	99.2	96.1	99.1
Diversity %	27.3	28.3	47.8
Number of channels in use	16	14	24
Bandwidth supplied	640000	560000	960000
Maximum Bandwidth Required	500201	500201	500201

Network 10%-3 was produced by using an AgentOpt network generated at 20% expected subscription and then optimised by ECHO for the 10% rate. This explains the network's better performance across the range of subscription values, that is shown in figures 8-4, 8-5 and 8-6. At the 10% level this network does not perform quite as well as the two rival solutions but is broadly comparable.




Figure 8-5. Comparison of users served at different take-up levels.




Figure 8-6. Comparison of capacity used at different take-up levels.

The other two 10% networks are very limited when it comes to higher take-up levels – as would be expected. Network 10%-1 has a highest take up level of approximately 12% whilst network 10%-2 cannot even serve the 10% it was designed for. Despite the fact that the main error encountered by the NVT is a lack of capacity, there is still at least 10% of the supplied capacity remaining unallocated in each network at each take-up level.

The networks generated at a 20% probability level are described in Table 8-2 and figures 8-7, 8-8 and 8-9.

Table 8-2. Networks designed for 20% take-up

Network ID	20%-1	20%-2	20%-3
Number of sites in use	5	4	5
Number of sectors deployed	8	10	9
Expected Profit	6533780	6436820	6256440
Expected Outlay	1761820	1773580	1854360
Expected Income	8295600	8210400	8110800
% of Maximum Income	97.57	96.57	95.3
Coverage	1.00	0.99	0.9989
Service	0.98	0.97	0.954
Diversity fraction	0.61	0.54	0.6217
Number of channels in use	29	41	30
Bandwidth supplied	1160000	1640000	1200000
Maximum Bandwidth Required	1000402	1000402	1000402



Figure 8-7. Network 20%-1




Figure 8-8. Network 20%-2




Figure 8-9. Network 20%-3

The results shown below highlight the differences between the three networks. With the 10% take-up networks there is nothing much to choose between the first two networks. At 20% subscription, however, 20%-1 is shown to be more inflexible. Whilst 20%-2 shows an expected profit very similar to the value predicted by the ACP tool, 20%-1 loses 800000 monetary units. This is because of relatively high levels of problems, both with the uplink C2I and the available capacity, that do not occur in 20%-2. As shown in table 8-3, the difference between the number of invalid users due to capacity problems is almost entirely responsible for the difference in the number of valid users in the network. Also, note that although 20%-2 has more valid users, it also has a lot more post-processing C2I errors – approximately 360% more than 20%-1. Network 20%-3 has been generated by the AgentOpt/ECHO combination and gives improvements on a number of NVT metrics. The most interesting of these is the complete lack of problems with the uplink C2I ratio.

Table 8-3. A survey of removed users at 20% actual take-up

Network ID	20%-1	20%-2	20%-3
Number of removed users	151.77	41.32	30
as a % of the number of users choosing to subscribe	10.68	2.91	2.12
Removed due to Signal	7.17	17.14	5.4
Removed due to C2I-D	1.02	6.58	19.4
Removed due to C2I-U	25.052	16.86	0
Removed due to Capacity	118.528	0.74	5.2
Post Processing C2I errors	153.69	564.12	441.2
Valid Users	1253.81	1368.22	1349.3

Figure 8-11. Funding of invalid users due to lack of capacity.

As the take-up level is increased the disparities between the three networks becomes even more pronounced (figure 8-10). At its peak 20%-2 can serve nearly 25% of users across the network, whereas 20%-1 does not quite reach the expected 20% service level (figure 8-12). This can be attributed for the most part to the extra capacity that 20%-2 provides (see figure 8-11). To a lesser extent the cell plan is also better suited to the up-link C2I problem than that of 20%-1, although the difference is much smaller than in the capacity case.




Figure 8-10. Expected profit over different take-up levels




Figure 8-11. Number of invalid users due to lack of capacity.




Figure 8-12. Percentage of valid users in the network

Similar analysis and discussions on the networks designed at 30%, 40% and 50% can be found in appendix IV.

Using a 20% standard subscription level the following figures compare the relative performances of pairs of the networks designed for 20%, 30% 40% and 50%. The 10% networks are excluded, because the results obtained above indicate that they are unsuitable for a 20% take-up level.

Figures 8-13 and 8-14 show unexpected results. As previously mentioned the networks designed for each subscription level were chosen as representative of a larger set of generated networks. The (comparatively) large difference between the two 20%-designed networks could be unfortunate, in that it skews the contrast with the other network sets. However this is another consideration a planner must make when choosing a final design. The randomness included in tools such as AgentOpt and ECHO will have an influence upon the final network, even to the extent that the networks designed for 20% take-up (almost) bookend the range of results shown below. Whilst this may make the results look untidy, a search for a network similar to the better performing one would not highlight this issue for the planner. It also strengthens the case put forward in chapter 7 for the more consistent network designs that are produced by AgentOpt.




Figure 8-13. Comparison of expected profit levels at 20% subscription




Figure 8-14. Comparison of outlay at 20% subscription

The most interesting thing about figures 8-13 and 8-14 is the fact that, taking the results at face value, the networks designed for 30% take-up are the best performing at this level, both on average and at the individual level. The 40%-designed networks also narrowly out-perform the 20%-designed ones when considering the average profit levels. The 30% networks have a lower average infrastructure cost when compared with the 20% networks, and this is the reason why they are more profitable.




Figure 8-15. Comparison of valid users at 20% subscription




Figure 8-16. Comparison of post-processing C2I errors at 20% subscription.

As can be seen in figure 8-15 the best-profit 20% network is also the network with the highest average number of valid users. However, as noted earlier in this section, 20%-2 does suffer from high levels of post-processing C2I errors. Figure 8-16 shows a different story for the 30% networks. Their lower PP-C2I levels combined with the higher projected profits points to the use of these networks in place of either of the 20% networks.

These investigations assume the 20% percent take-up rate will be accurate. Figures 8-17 and 8-18 show the comparative performances of the same networks when the actual subscription take-up rate varies. Again however it is the 30% networks that perform best overall.

For this Malvern scenario with a 20% expected subscription rate across the town, the preceding results indicate that a network planned for a 30% take-up is likely to prove to be the best option. At a perfect 10% or 20% subscription rate, the difference between the 30% and 50% networks is relatively small. 30%-2 has a better profit but 30%-1 can match slightly more users. At the higher 30% take-up rate the networks planned for that level are the obvious winners. It must be remembered that the ability required for any network




Figure 8-17. A comparison of expected profit at 10% subscription.




Figure 8-18. A comparison of expected profit at 30% subscription.

For this Malvern scenario with a 20% expected subscription rate across the town, the preceding results indicate that a network planned for a 30% take-up is likely to prove to be the best option. At a perfect 10% or 20% subscription rate, the difference between the 20% and 30% networks is relatively small. 30%-2 has a better profit but 20%-2 can reach slightly more users. At the higher 30% take-up rate the networks planned for that level are the obvious winners. It must be remembered that the outlay required for any network

depends upon the amount of users that do actually subscribe. A sizable proportion of the costs will be that of user equipment. In the examples discussed above the number of sites in use and the number of sectors deployed are almost identical for the 20% and 30% networks. So in this case a 50% increase in the number of users served can be achieved with the increase in outlay being only the extra user premises equipment. Of course this example may well be scenario specific and any planning decision is subject to a number of other considerations, but in a situation where the rate is thought to be unreliable the benefits of using a supposedly over-engineered network could be considerable.

8.2 Availability and Quality of Service

One problem with wireless systems in general, but especially at millimetre-wave level, is that the link between user and provider can be affected by the weather, or more specifically, by heavy rainfall events. For example, a system with an availability (A_v) level of 99% means that the user can expect to lose service for 3.65 days every year, due to intense rain cells passing through the link from transmitter to user. These drop-outs would occur throughout the year, for varying lengths of time, and at varying times of day. For a casual residential internet user this would be unlikely to be more than a slight nuisance. For a business user, dependant upon always-on communications channels, this level would be insufficient. An A_v level of 99.999 would produce an average signal downtime of only 5.2 minutes every year. This level would be comparable with cable systems, which although not affected by weather-born interference, can fall foul of other problems. The AgentOpt and ECHO tools utilise the formula presented in [96] in order to add sufficient attenuation to each individual signal to ensure a certain minimum A_v level.

Investigations into networks designed for a variety of different Av levels throw up a number of interesting points.




Figure 8-19. A comparison of expected profit over differing Av levels.




Figure 8-20. A comparison of service percentage over differing Av levels.

Figures 8-19 and 8-20 show results from networks, generated by the ECHO tool using the Sevenoaks scenario, for each of four different user Av requirements: 99%, 99.9%, 99.99% and 99.999%. These levels equate approximately to service down-times of 4 days, 9 hours, 53 minutes and 5

minutes respectively, over the period of a year. The networks were generated with a service requirement of 80%, except for the 99.999% networks which used a 50% constraint.

From these results it can be seen that the provision of a 99.9% user availability level makes very little difference to the profit and service levels found in the 99% network. In this case it is easy to see that a service provider with an initial aim of 99% availability could increase this level to provide a more robust service to users with very little change to the business case. Indeed the provider may decide to increase the network availability levels higher, with the 99.99% networks attaining on average 95% of the best profit, and the 99.999% networks attaining 72%. The planning engineer would need to consider whether the gain in goodwill and contentment with the service that the increase in Av would provide would be sufficient compensation for operating the network at a lower profit level.

As shown in Figure 8-21 below, analysis by the NVT reduces the expected profit levels marginally, but the trends are preserved. Little difference between the average expected profit for the 99 % and 99.9% Av networks, with less than 5% and 25% drops in the average expected profit for the 99.99% and 99.999% Av networks respectively.




Figure 8-21. A comparison of expected profit after NVT analysis.

Next we take one network generated for each Av value and test the flexibility of the design. The following results (shown in figures 8-22, 8-23 and 8-24) are produced by validating each network using the range of Av values.




Figure 8-22. Expected profit levels for each network validated at different Av levels.

Figure 8-20 shows that the 99.999% network is limited at all Av levels, i.e. that the search for the high levels of availability has introduced some flaw into the design. This can be traced in part to the fact that using this high Av level does

mean that the planning tool will not always be able to return a solution. In this particular case, the use of an 80% service constraint, as with the other Av level networks, meant that the ECHO tool could not progress beyond its second *repair* stage. Although the final networks do have service levels of over 80%, it was found that in this case the tool would not proceed to the third *optimisation* stage at service constraint levels over 50%. This implies that for the Sevenoaks scenario the aim of an optimal network offering an availability level 99.999% to every user is unrealistic. That said, the eventual service levels did exceed 80%, but this must be compared with the levels of over 90% and in most cases over 95% achieved with lower Av requirements (see figure 8-18). This relatively lower service level, combined with higher infrastructure costs due to the need for higher gain user antennas, accounts for the lower profits. In the 99.999% networks the invalid users are excluded because of down-link signal and C2I problems even at the lower Av levels, indicating that the plan is simply not serving enough users.




Figure 8-23. The number of users removed due to up-link difficulties as a percentage of the total removed, at different Av levels.




Figure 8-24. Number of valid users as a percentage of those chosen, at different Av levels

Figures 8-23 and 8-24 display mostly expected results. The 99.999% network performs consistently poorly because it is limited in overall service levels but obviously not by signal-strength or C2I issues, for those that are served.

The majority of invalid users in the other sets of networks are excluded because of up-link issues at or below the Av level they were designed for, and then signal strength deficiencies thereafter, as shown in figure 8-23. Each network performs marginally worse at the 99% level because lowering the attenuation built into the transmission signal increases the interference problems. It is interesting to note that with these individual networks, whilst it is the 99.9% network that overall performs the best, the three lowest Av networks are broadly comparable. This indicates a level of over-engineering in the network solutions. At an availability level of 99.99%, the network designed for a down-time 100 times greater than that performs admirably. It must be that the actual Av levels are much greater than planned for – i.e. the down-link signal exceeds the minimum threshold by a fair margin. This is discussed further in section 8.6. From a planning engineers point of view this may be seen as bonus, as it is likely that a higher availability requested by an

individual user may be accommodated. Of course this level may be provided for explicitly if the request is known and considered before any network design or implementation.

8.3 Maximum channel numbers.

With ever increasing numbers of wireless technologies soon to be vying for spectrum, the ability to increase 'spectral efficiency' within a network is desired. Without any specific definition, any measure of spectral efficiency would be increased by reducing the amount of spectrum used. The AgentOpt and ECHO tools allow the maximum number of channels used within the network to be set. These channels will be contiguous, and thus decreasing the number of channels increases the likelihood that any two channels in use are liable to interfere with one and another (either the same or adjacent within the frequency block).

Figure 8-25 shows the profit levels of 5 networks each designed with a different maximum channel constraint. In these particular network solutions the heavily constrained 2-channel network still attains an expected profit 89% of that predicted for the 10-channel network.




Figure 8-25. A comparison of expected profit levels for networks designed with differing numbers of channels

The difference is more pronounced after the networks have been passed through the NVT. The new expected profit levels for each network are shown in figure 8-26. Not only is the NVT predicted profit for the 2-channel network much reduced (only 68% of the previous level), but the difference between the two extremes is extended. The 2-channel network now only achieves 70% of the 10-channel networks expected profit.




Figure 8-26. A comparison of the NVT predicted profit levels for each network

Another interesting point is the way that the 6- and 8-channel networks outperform the 10-channel example. This is explained with reference to table 8-4 below. Both the 6- and 8-channel networks provide more capacity and so have fewer users invalidated due to capacity.

Table 8-4. A comparison of capacity issues.

Maximum Channels	6	8	10
Total Removed Users	142.54	109.06	139.69
Removed due to Capacity	57.44	81.55	122.74
Number of channels in use	34	32	29
Bandwidth supplied	1360000	1280000	1160000

The 8-Channel network has the fewest invalid users overall which leads to it also having the highest expected income level and in turn the highest projected profit.

These results come from individual randomly chosen networks, and as such make no guarantee that this will always be the case. The point is that the number of available channels should be another consideration for a planner. It has been shown that networks planned with reduced channel numbers *can* be competitive when compared with those networks that are using more. This will of course be scenario dependant, and in some cases may produce vastly inferior networks. However, forcing the optimisation method to work harder by allowing a smaller range of spectrum could have real benefits. As well as the spectrum efficiency argument, a network planned with fewer channels than are available is somewhat future-proofed. An amount of extra capacity could be kept in reserve, ready to be added into the network at any time without any change to the existing configuration.

8.4 Capacity requirements

Broadband data-link provision companies, such as BT, offer services with a maximum data rate, advertising “up to ten times faster than 56k dial-up connections”. This is because the technologies use a contention-ratio system. A basic residential user would most probably be connected to a 512 kbps line, with a standard contention ratio of 50-1. But this means that the line is shared with 49 other users, and so the *maximum* data rate that could be guaranteed would be only 10kbps, although it is highly unlikely that all the subscribers on the loop would be utilising the service at the same time. The network definition in chapter 3 approaches the allocation of capacity in a similar way, but from a different start point. Each user is assigned a *minimum* data rate guarantee, and this value is subtracted from the serving sectors available capacity. The contention ratio at each sector (or each channel even) varies depending upon which subscribers are assigned to it and their stated minimum bandwidth requirement.

Required data-rates will keep growing over the next few decades to accommodate increases in demand for services such as on-line gaming, video-on-demand, software-on-demand, online software updates and so on. For this reason any discussion of specific capacity requirements will become irrelevant in a short period of time. Therefore this section makes note of techniques that can be employed by a network planner independent of the actual capacity required. As can be seen in appendices I and II, each scenario has different capacity requirements. The Sevenoaks scenario makes a minimum guarantee of 64kbps to the basic residential user – higher than today’s ADSL basic broadband package, when we consider the contention ratio. Higher values are assigned to small and medium sized businesses in the area. Looking further

forward to the future, the Malvern scenario guarantees each user a minimum 706kbps rate at all times.

The Sevenoaks scenario offers a reasonable data-rate to potential users if the networks were to be implemented now. However, with the expected growth in bandwidth demands a sensible planner would be advised to look ahead and ensure the planned network can evolve to cope with its users increasing demands. (This assumes that a network should be planned to exist in its original state for as long as possible. The trade-off between re-planning and re-implementing the network every few years and planning an initially over-engineered network to accommodate growth over a longer time-frame would make a good case study but is not discussed in this thesis.)

The five networks briefly described in table 8-5 have been planned using the same subscription fee for each user but varying their capacity demands by a certain percentage. NVT analysis of the networks provides the data used in figures 8-27 and 8-28.

Table 8-5. Networks planned for varying traffic demands.

Network ID	CSN-1	CSN-2	CSN-3	CSN-4	CSN-5
Capacity Requirement	standard	+ 25%	+ 50%	+ 75%	+ 100%
Expected Profit	4888910	4889710	4824200	4659680	4665630
Service %	93.8	93.4	94.2	96.0	95.7
Number of channels in use	11	10	14	20	16
Bandwidth supplied	440000	400000	560000	800000	640000

As the capacity demands rise the expected profit drops off. As there is no equivalent rise in subscription fees this is to be expected. However, a 50% rise in the capacity requirement translates into a less than 2% drop in profit, whilst a 100% increase means less than a 5% cut. The NVT results displayed in figures 8-27 and 8-28 have been generated by analysing each network at each

of the different capacity requirements. NVT indicates that at the standard capacity level using network CSN-3 will actually show an increase of around 2% when compared with CSN-1. As would be expected, the higher capacity networks perform better when the capacity requirement is raised.




Figure 8-27. A comparison of the profit expected from each network at varying capacity requirements




Figure 8-28. A comparison of the percentage of supplied capacity that is used by each network at varying capacity requirements

The percentage of capacity used in each network is shown in figure 8-26. Even for network CSN-1, attempting to serve double the capacity that it was designed for, there is still well over ten percent of the supplied capacity not being used.

The planning engineer should also remember that increasing the users capacity demand will have a similar effect to increasing the subscription rate. Figure 8-29 displays both the expected profit and the percentage of total users served when three of the CSN networks are analysed by the NVT at the increased subscription rate of 30%. Figure 8-30 shows a comparison of the number of users that are removed or that have PPI errors under the same analysis.

When taking all these results into account network CSN-3 seems to be by far the best choice for implementation – at any of the five capacity levels. This network solution offers good performance under a range of different conditions and, as emphasised in figure 8-30, has very low interference levels.




Figure 8-29. A comparison of the profit expected and the % of users served when the networks are analysed with a take-up rate of 30%




Figure 8-30. A comparison of the number of users removed and that have PPI errors when the networks are analysed with a take-up rate of 30%

The planning engineer should always try planning a network at a higher capacity level than is required. This will not always produce satisfactory results, but will have the effect of forcing the ACP to work harder. In certain network scenarios the capacity can be greatly increased without any adverse effect on the expected profit levels.

8.5 Return periods

The AgentOpt and ECHO tools utilise a planner-set parameter called a *return period* (RP) for the network they are designing. This value, in units of years, determines how much subscription revenue is generated by the network plan. The user's subscription charge is based upon an annual rate and so, for example, if a return period of three years is defined then the network will be created assuming each user will be charged three times that value. In scenarios such as Sevenoaks, where different users are charged different rates, brief return periods should force the tools to give greater priority to higher paying users. Since the optimisation is based upon the revenue generated by the network, the return period can have a dramatic effect upon the final

network produced. This section looks at the effect of varying the RP in both the Malvern and Sevenoaks scenarios.

The three cell plans shown in figures 8-31, 8-32 and 8-33 represent typical networks generated by the ECHO tool for the Sevenoaks scenario at three different return periods – 1 year, 3 years and 5 years. Some simple information about each is shown in table 8-6.




Figure 8-31. Network SRP-1




Figure 8-32. Network SRP-3



Figure 8-33. Network SRP-5

Table 8-6. The Sevenoaks RP networks

Network ID	SRP-1	SRP-3	SRP-5
Return Period	1	3	5
Expected Income	1793100	7073850	12346000
Required Spend	1667700	2218240	2567340
Expected Profit	125400	4855610	9778660
Coverage %	84.2	94.1	99.6
Service %	75.3	95.4	99.3
Diversity fraction %	15.2	54.3	82.7
Number of channels in use	5	12	13
Bandwidth supplied	200000	480000	520000

It can be seen that the return period has an enormous difference on the design created. A brief glance at the cell plans shows the viewer that SRP-1 is far less complex and seems underdeveloped. The financial values cannot be compared as of yet but the differences in the coverage, service and diversity metrics is obvious. With a higher return period the optimisation appears to work a lot harder. Figures 8-34, 8-35 and 8-36 show the best profit and current profit values for each iteration within the optimisation run for each network.




Figure 8-34. The optimisation of network SRP-1.




Figure 8-35. The optimisation of network SRP-3.




Figure 8-36. The optimisation of network SRP-5.

The effect of the lower return period is to reduce the comparative value of each user to the network. For a user to be served a user antenna must be placed at the premises and the associated cost accounted for within the optimisation. Most users within the scenario expect to pay 750 units annually and the user antennas range from 500 to 2000 units in cost. The lower return period translates into a lower worth (in cases where the higher-gain antennas are required, serving a user would make a loss, even before the cost of the sector is considered) for each individual user to the network as a whole, and vice-versa. SRP-1 just exceeds its constraint of coverage for 80% of its users, whilst SRP-5, which places a far higher value upon a single user, drives far harder to achieve complete service.

Similar experiments were also performed with the Malvern scenario. Figures 8-37, 8-38 and 8-39 show cell plans for networks optimised with return periods of 1 year, 3 years and 5 years. Table 8-7 presents some basic information about each network.




Figure 8-37. Network MRP-1.



Figure 8-38. Network MRP-3.




Figure 8-39. Network MRP-5.

Table 8-7. The Malvern RP networks.

Network ID	MRP-1	MRP-3	MRP-5
Return Period	1	3	5
Expected Income	2686000	8344800	13858000
Required Spend	1465100	1813360	2017960
Expected Profit	1220900	6531440	11840040
Coverage %	99.1	99.8	99.9
Service %	94.0	98.4	98.7
Diversity %	27.2	58.4	72.2

The results show similar trends to the Sevenoaks scenario results, although not as pronounced. Visually, the MRP-1 network is again less complex than the other two, and there are some obvious areas that could be improved. However, it performs only slightly worse with regard to coverage and service levels. Unlike in the Sevenoaks scenario, a Malvern user can never be served at a loss – the worst case would be profit-neutral – before the sector costs are considered. This is further enhanced by the geographical aspects of the scenario – the highest gain antennae are rarely required.

How do these networks compare against each other on a level playing field? Starting with the Sevenoaks scenario, the results obtained when each network is analysed by the NVT. Each network is validated using the three different RP values to give fair comparisons, as shown in figure 8-40.




Figure 8-40. A comparison of expected profit for each Sevenoaks scenario network at the different return periods.

SRP-1 is confirmed to be an inferior network design, even at the level it was planned for.

The Malvern scenario results are shown in figure 8-41.




Figure 8-41. A comparison of expected profit for each Malvern scenario network at the different return periods.

MRP-5 is shown to be a slightly inferior network in this case, mainly because of increased infrastructure costs. However, there is an argument that says the differences in profit values can be put down to the random nature of optimisation scheme. The differences in these networks are so slight that the RP level can not be definitively blamed.

The planner must be aware of the affect the return period could have upon the optimisation. As has been shown however, the magnitude of the effect will be highly scenario dependant.

8.6 Diversity

As defined in section 4.1, the concept of *diversity* refers to users who can receive service above threshold from more than one BS. This concept received particular interest as a means of increasing availability (Av) or quality-of-service, in BFWA. Experiments following the method laid out in [97] however, showed that Av levels within some ACP generated networks could not be significantly improved by the provision of diverse signal reception [98]. Networks with naturally high levels of user Av show that any

(proportionally small) gains made by employing diversity are outweighed by the increased interference and cost to the network. Section 4.1 designates a user as *diverse* if said user receives a signal, above threshold, from a secondary transmitter. This diversity constraint was shown to be insufficient to the problem.

Usman showed that this constraint in no way guarantees an increase in A_v . A new constraint was then tested, in which a user designated diverse would then see an increase in availability. This new constraint takes into account the ratio between the path-link lengths of the primary and secondary transmitters to the users, and also the angular separation between these links. The constraint also required that a user must have availability of less than 99.9% before diversity was taken into account, in order that the improvement was worthwhile. When implemented in an earlier version of the ECHO tool, this constraint was seen to be to exceptionally hard – even at the very modest level of 5% - for the tool to incorporate into the network[98].

Could the concept of diversity be shown to be useful in another area, that of mitigating the problems involved with using a subscription probability network planning method. Intuitively, the more BSs that can provide a signal to an individual user the likelier it is that that user can receive service from one of them.

The networks described in the two tables below have been generated using the standard Malvern scenario, over a return period of three years. Table 8-8 gives details of three low diversity networks, whilst table 8-9 details three higher diversity networks. Figure 8-42 then gives a comparison of the expected profit and diversity percentage levels for each network.

Table 8-8. The low diversity networks

Network ID	DN-1	DN-2	DN-3
Number of sectors	8	10	5
Expected Income	8257200	8288400	8266800
Required Spend	1729240	1839780	1538260
Expected Profit	6527960	6448620	6728540
Coverage %	98.7	98.7	98.4
Service %	97.1	97.5	97.2
Diversity %	43.3	55.6	35.5

Of the high-diversity networks in described in table 8-9, networks DN-5 and DN-6 were planned with a required diversity constraint of 80%. Network DN-4 achieved its diversity level without any constraints being placed on the optimisation.

Table 8-9. The higher diversity networks.

Network ID	DN-4	DN-5	DN-6
Number of sectors	7	12	12
Expected Income	8179200	8271600	8296800
Required Spend	1698940	1958620	1932860
Expected Profit	6480260	6312980	6363940
Coverage %	98.7	99.5	99.2
Service %	96.2	97.3	97.6
Diversity %	73.9	80.7	80.2




Figure 8-42. A comparison of expected profit and diversity percentage for each network.

This data implies that on the whole the lower diversity networks are better performing. Indeed the best performing network in terms of profit – DN-3 – has the lowest diversity measure by some margin. This can be explained by the increased expenditure required in the higher diversity networks. The figure below displays data from the analysis of the six networks by the NVT. Each network has been analysed at the three different subscription take-up levels – 20%, 25% and 30%.




Figure 8-43. A comparison of profit predicted for each network at three different subscription take-up levels.

These results give no real indication that diversity will help to alleviate the subscription take-up problem. The two best performing networks across the range of subscription take-up levels – DN-3 and DN-4 - were both designed with no particular diversity constraints. The difference in the number of diverse users between the two is considerable. Network DN-3 has a diversity percentage of 35.5% - around 2515 potential users. DN-4 has approximately 5236 diverse users – a level of 73.9%. And yet the two networks produce very similar results. This implies that the overall design of the network is of far greater importance than the number of BS that each individual user can receive service from.

8.7 Known-subscription networks

As has been explained, the networks generated by tools such as AgentOpt and ECHO are produced using statistical probabilities. The NVT tool can then be used to give a further statistical analysis of the networks. In this final investigation in this chapter, we look at the validity of this approach. Firstly, networks are generated from fixed subscription populations. This means that certain users are known to wish to subscribe, so the network can be planned to cater exactly to these people. For a network planning engineer this would be an ideal situation. However, for any sizable commercial network this is an unlikely occurrence. The previous sections of this chapter have been advocating a certain level of over-engineering in order to overcome the uncertainty in subscription. The question is, how do these *certain* networks compare to the *uncertain* networks?

The three user groups detailed in table 8-10 were randomly generated to represent a set of users who have undertaken to subscribe. Each group contains approximately 20% of the users within the Malvern scenario.

Table 8-10. The randomly generated user groups

User Group ID	1	2	3
Number of Users	1462	1416	1415
% of total	20.63514	19.98589	19.97177
Maximum possible revenue	8772000	8496000	8490000

Networks were then planned for these specific user groups. These networks are described in table 8-11 and the cell plans shown in figures 8-44, 8-45 and 8-46.

Table 8-11. The *certain* networks.

Network ID	CN-1	CN-2	CN-3
Expected Income	8682000	8442000	8376000
Expected Profit	6771600	6593600	6533300
Coverage	0.9951	0.9938	0.988
Diversity fraction	0.8207	0.8593	0.6388




Figure 8-44. Network CN-1.




Figure 8-45. Network CN-2.




Figure 8-46. Network CN-3.

Each certain network is then subjected to analysis by the NVT. A comparison of the expected profit for each network, from both the original plan and the subsequent NVT analysis is made in figure 8-47.




Figure 8-47. A comparison of expected profit values.




Figure 8-48. A comparison of invalid user numbers.

The disparity between the original plan and the NVT's predicted profit levels is explained in figure 8-48. This shows that the main cause of invalid users in all three networks is up-link interference. CN-2 and CN-3 suffer to a much greater degree. This gives further proof of the usefulness of the NVT tool even in the case where the subscribers to the network are known in advance of planning.

Figures 8-49, 8-50 and 8-51 show the cell plans for three networks, generated using the standard subscription probability method. The first network, PN-1, was generated with no unusual requirements. PN-2 was generated to provide elevated levels of diversity, whilst PN-3 used a subscription probability of 30% across the range of potential users.




Figure 8-49. Network PN-1




Figure 8-50. Network PN-2




Figure 8-51. Network PN-3

Figures 8-52, 8-53 and 8-54 display results from an NVT analysis of each of the standard networks using the three different user groups as certain subscribers. The equivalent results from the user group specific networks are included for comparison purposes.




Figure 8-52. A comparison of the expected profit across the different groups of users.




Figure 8-53. A comparison of the percentage of possible revenue collected by each network across the different user groups.




Figure 8-54. A comparison of the percentage of possible subscribers served by each network across the different user groups.

These results clearly show that the networks generated by the standard subscription probability method out-perform networks CN-2 and CN-3 on all fronts. Network PN-2 also consistently out-performs CN-1. Although it may seem counter-intuitive that specifically tailored networks would be inferior to more general designs, there are a number of possibilities for why this may be the case. The first factor to consider is the random element within the design

tool. The tailored networks were chosen randomly and there is always the possibility that all three were just bad designs. Also, the optimisation tool that produced the designs may be skewed towards the subscription probability method, although it is difficult to see how this may be the case.

Despite these points, this investigation does strengthen the case for the subscription probability method, especially as in the majority of cases the planning engineer will not have a known population of subscribers waiting patiently for a network to be designed.

8.8 Conclusions

In order to effectively plan BFWA networks for unknown populations of subscribers, methods must be used to increase the inherent flexibility of the network plans. This, in essence, means that a level of over-engineering must be incorporated into the network solution, preferably without having a great effect on the implemented network profit level. This chapter has investigated a number of approaches that should be employed by a planning engineer to this end. There is no guarantee that these methods will produce positive results in every network scenario. It is hoped however, that this chapter has put forward some methods and considerations that will lead to the design and implementation of successful BFWA networks, not least of which is the use of an analytical tool such as the NVT.

Chapter 9

Conclusions and further work

This thesis has presented an ACP tool, called AgentOpt, that employs a novel optimisation scheme based upon the principles of emergent intelligence to design BFWA network solutions. Network designs produced by the tool were evaluated and compared with similar designs that had been produced by other methods. A second tool, the NVT, which performs a statistical analysis of the network designs, was also presented. Finally, investigations were presented to show the usefulness of the NVT in comparing different solutions, and methods were described to enable planning engineers to utilise ACP tools such as AgentOpt to maximum effect. This chapter is divided into two parts discussing the successes and shortcomings of each tool.

9.1 AgentOpt

AgentOpt transforms two essential elements of any potential BFWA network into two distinct populations of autonomous entities – User-Agents and Site-Agents. Through the simple rule-based behaviour that each species of agent exhibit, a self-organizing system is formed. The AgentOpt tool fulfils half it's remit to automatically design networks. It does produce feasible network solutions without any interaction from the planning engineer. However, as shown in chapter 7, the designs are generally slightly inferior to those produced by the ECHO tool, indicating that the optimisation strategy is probably lacking. There are a number of other indications that this is the case. The two most obvious are AgentOpt's reliance on its starting strategy and the

tool's inability to maintain a near-optimal current-solution through out its duration (see section 5.5.1). These two problems seem to be linked. The random-start networks stay much more unstable over the duration of the optimisation. The level of instability appears to be at the root of the problem. The tool was designed to have certain instability levels in order to force the optimisation to cycle through different solutions.

The aim was to have a disarray-rebuild-fine-tune-disarray cycle a number of times during the optimisation run, and this does appear to happen in the random-start networks. However, it appears that with the extra instability from the start, the User-Agents powers of self organization are found wanting. The exact opposite happens with the best path-loss start. In this case the User-Agents quickly find themselves a comfortable configuration and are reluctant to leave it. More investigation is needed on the parameters that govern the behaviour of the system, especially in the areas of User-Agent resentment and distraction.

On the positive side, as described in chapter 7, AgentOpt does produce designs that respond well to improvement by ECHO. The incorporation of the AgentOpt scheme into the ECHO tool as a replacement for the existing second stage is a possibility, leading to a combination that would produce good network designs much more quickly ECHO in isolation.

9.2 NVT

The NVT tool analyses network designs by simulating the take-up of subscription by randomly chosen users within a network region. The repeated generation of these possible network configurations allows statistics to be collected about the networks performance on numerous measures. The NVT

is vital to the strategies laid out in chapter 8 to help a planning engineer choose the best possible network for any situation.

The obvious evolution of the NVT would be towards a more realistic simulation of an implementation of a network over a given time period. Instead of a simple yes/no this user can/cannot be served approach, an extension of the tool would simulate a network from its construction, through user subscription take-up and perhaps over the projected life-time of the network hardware. This could incorporate different roll-out strategies for the hardware implementation, perhaps predict the performance of a fully-subscribed network at different times of day, provide an improved analysis of user interference problems through the use of different duplexing strategies and maybe model the *churn*, or subscriber turn-over that most networks must deal with.

A further approach could be the integration of the present incarnation of the NVT into the AgentOpt or ECHO tool. The analysis of the chosen network could be provided automatically at the end of the optimisation, in which case the integration would simply increase the ease with which the tools are used. Another option is to include the NVT functionality within the optimisation scheme. This would of course increase the computation demand of the scheme by a sizable amount. It would however provide some way of optimising the up-link configuration which has hitherto been ignored by tools such as AgentOpt and ECHO. By including the analytical abilities of the NVT a new tool could also expand upon the strategies outlined in chapter 8 to independently produce a set of networks, optimised to slightly different configurations.

Finally, RPD, the propagation prediction tool described in section 2.4 could also be integrated to produce a single tool that could take as input a

geographical database containing terrain, scenery, potential users and potential base sites, along with list of possible infrastructure options. There would be the potential to convert the graphical interface of the RPD into an interactive 3-D environment to display the geographical network area. The new super-tool could automatically produce, from start to finish, a set of network solutions, a full analysis of each one and a comparison of individual strengths and weaknesses.

Chapter 10

Bibliography

- [1] Craig, K., *et al.*, AY4483 - *Efficient use of Wireless LAN technology for Fixed Wireless Access*. Technical Report: OFCOM, 2004.
- [2] Allen, S.M., *et al.*, *Modelling for Cell Planning Activities in EMBRACE*. Technical Report: Cardiff University, Rutherford Appleton Laboratory, 2001.
- [3] Allen, S.M., *et al.* *Maximising Spectrum Efficiency for BFWA through Automatic Cell Planning*. IEE "Getting the Most Out of the Radio Spectrum" Conference. 10: 1-5, 2002.
- [4] Gamst, A. *Remarks on Radio Network Planning*. 37th IEEE Vehicular Technology Conference. New York, NY, USA. 160-165, 1987.
- [5] Gamst, A., *et al.*, *Cellular Radio Network Planning*. IEEE Aerospace and Electronics Systems Magazine, 1(2): 8-11, 1986.
- [6] Gamst, A., *et al.*, *An Integrated Approach To Cellular Network Planning*. IEEE Journal on Vehicular Technology, 21-25, 1985.
- [7] Martinez-Dalmau, J.L. and Crespo, A. *An Approach for Planning a Cellular Radio Network*. in *Network Planning in the 90's*. 4th International Planning Symposium. North-Holland, Amsterdam, Netherlands. 61-68, 1989.
- [8] Esposito, C. and Willard, C. *PLATON (Planning Tool for Mobile Radio Networks)*. in 2nd International Conference on Universal Personal Communications. New York, NY, USA. 387-391, 1993.
- [9] Fortune and Gay, *Wise Design for Indoor Wireless Systems*. IEEE Computational Science and Engineering, 2(1): 58-68, 1995.
- [10] Calegari, P., *et al.*, *Radio Network Planning with Combinatorial Algorithms*. Proceedings of ACTS Mobile Communications Summit, 707-713, 1996.
- [11] Tutschku, K., *et al.* *ICEPT - An Integrated Cellular Network Planning Tool*. in 47th IEEE Vehicular Technology Conference. New York, NY, USA. 2: 765-769, 1997.

- [12] Wright, M.H. *Optimization Methods for Base Station Placement in Wireless Applications*. 48th IEEE Vehicular Technology Conference (VTC '98). 1: 387- 391, 1998.
- [13] Chamaret, B., *et al.* *Radio Network Optimization with Maximum Independent Set Search*. 47th IEEE Vehicular Technology Conference. 2: 770-774, 1997.
- [14] Calegari, P., *et al.* *Genetic Approach to Radio Network Optimization for Mobile Systems*. 47th IEEE Vehicular Technology Conference. 2: 755-759, 1997.
- [15] Tutschku, K., *et al.*, *An Integrated Approach to Cellular Network Planning*. Proceedings of the 7th International Networks Planning Symposium, 1996.
- [16] Fritsch, T., *et al.* *Field Strength Prediction by Ray-Tracing for Adaptive Base Station Positioning in Mobile Communication Networks*. 2nd ITG Conference on Mobile Communications. Neu Ulm, Germany. 1995.
- [17] Tutschku, K. and Leibnitz, K., *Fast Ray-Tracing for Field Strength Prediction in Cellular Mobile Network Planning*. 46th IEEE Vehicular Technology Conference. 1: 541-545, 1996.
- [18] Vasquez, M. and Hao, J.-K., *A Heuristic Approach for Antenna Positioning in Cellular Networks*. Journal of Heuristics, 7: 443-472, 2001.
- [19] Louta, M.D., *et al.*, *Cost-Efficient Design of Future Broadband Fixed Wireless Access Systems*. Wireless Personal Communications, 2003(27): 57-87, 2003.
- [20] Berry, J., *Radio Planning and the Trunked Network*. Land Mobile. 21- 23, March, 2002.
- [21] Berry, J., *Methodologies in DECT WLL Planning to and Beyond Development*. Technical Report: ATDI Ltd. 2000.
- [22] Svanberg, I. and Morris, M., *Planning for Uncertainty*. Communications International, 19(9): 56-60, 1992.
- [23] Sandrasegaran, K. and Prag, K., *Planning Point-to-Multipoint Rural Radio Access Networks using Expert Systems*. Expert Systems with Applications, 17(3): 145-166, 1999.
- [24] Fruhwirth, *et al.*, *Planning Cordless Business Communications Systems*. IEEE Expert Magazine: 11(1): 50-55, 1996.

- [25] Hao, Q., et al., *A Low Cost Cellular Mobile Communication System: A Hierarchical Optimization Network Resource Planning Approach*. IEEE Journal on Selected Areas in Communications, 15(7): 1315-1325, 1997.
- [26] Siqueira, G.L., et al. *Optimisation of Base Station Antenna Position Based on Propagation*. 47th IEEE Vehicular Technology Conference. 2: 1133-1137, 1997.
- [27] Zhu, Y. *Research on Mobile Wireless Network Planning*. International Conference on Communication Technology, (ICCT96). 1: 342-346, 1996.
- [28] Sherali, H.D., et al., *Optimal Location of Transmitters for Micro-Cellular Radio Communication System Design*. IEEE Journal on Selected Areas in Communications, 14(4): 662-673, 1996.
- [29] Wong, J.K.L. and al, e., *An Improved Integer Programming Formulation for Indoor Base Station Placement*. University of Auckland, New Zealand.
<http://www.esc.auckland.ac.nz/Mason/>
- [30] Ibbetson, L.J. and Lopes, L.B. *An Automatic Base-Site Placement Algorithm*. 47th IEEE Vehicular Technology Conference. 2: 760-764, 1997.
- [31] Krishnamachari, B. and Wicker, S. *Experimental Analysis of Local Search Algorithms for Optimal Base Station Location*. The International Conference on Evolutionary Computing for Computer, Communication, Control and Power (ECCAP). 2: 73-77, 2000.
- [32] Molina, et al. *The Automatic Location of Base-Stations for Optimised Cellular Coverage: A New Combinatorial Approach*. 49th IEEE Vehicular Technology Conference. 1: 606-610, 1999.
- [33] Anderson, H.R. and McGeehan, J.P. *Optimizing Micro-cell Base Station Locations using Simulated Annealing*. 44th IEEE Vehicular Technology Conference. 858-862, 1994.
- [34] Amaldi, E., et al., *Optimisation Models and Algorithms for Downlink UMTS Radio Planning*. IEEE Wireless Communications and Networking Conference (WCNC 2003). 2: 827-831, 2003.
- [35] Allen, S.M., et al. *Frequency Assignment: Methods and Algorithms*. NATO AGARD Conference: Frequency Assignment, Sharing and Conservation in Aerospace Systems. Aalborg, Denmark. K1-K20, 1998.
- [36] Murphy, R., et al., *Frequency Assignment Problems*. Handbook of Combinatorial Optimisation, 16, 1999.

- [37] Allen, S., *et al.*, *Lower Bounding Techniques for Frequency Assignment*. Discrete Mathematics, 1999(197/198): 41-52, 1999.
- [38] Smith, D., *et al.*, *A New Lower Bound for the Channel Assignment Problem*. IEEE Transactions on Vehicular Technology, 49(4): 1265-1272, 2000.
- [39] Smith, D. and Hurley, S., *Bounds for the Frequency Assignment Problem*. Discrete Mathematics, 167: 571-582, 1997.
- [40] Ferland, J., *et al.*, *An Object-Oriented Methodology for Solving Assignment-Type Problems with Neighbourhood Searches*. Journal of Operations Research, 44(2): 347-359, 1996.
- [41] Chu, P. and Beasley, J., *A Genetic Algorithm for the Generalised Assignment Problem*. Computers and Operational Research, 24(1): 17-23, 1997.
- [42] Cuppini, M., *A Genetic Algorithm for Channel Assignment Problems*. European transactions on Telecommunications and Related Technologies. 5(2): 157-167, 1994.
- [43] Kotrotsos, S., *et al.*, *Formulation and Computationally Efficient Algorithms from an Interference-Oriented Version of the Frequency Assignment Problem*. Wireless Personal Communications. 18: 289-317, 2001.
- [44] Anton, *et al.* *Flexible Channel Assignment by Simulated Annealing*. 5th Nordic Seminar on Digital Mobile Radio Communications. Helsinki, Finland. 315-321, 1992.
- [45] Wei-Wang and Rushforth, *An Adaptive Local-Search Algorithm for the Channel Assignment Problem (CAP)*. IEEE Transactions on Vehicular Technology, 45(3): 459-466, 1996.
- [46] Castelino, D., *et al.*, *A Tabu Search Algorithm for Frequency Assignment*. Annals of Operational Research, 63: 301-319, 1996.
- [47] Maniezzo, V. and Carbonaro, A., *An ANTS Heuristic for the Frequency Assignment Problem*. Future Generation Computer Systems, 2000(16): 927-935, 2000.
- [48] Montemmani, R., *et al.*, *An ANTS Algorithm for the Minimum-Span Frequency-Assignment Problem with Multiple Interference*. IEEE Transactions on Vehicular Technology, 51(5): 949-953, 2002.
- [49] Hurley, S., *et al.*, *FASoft: A System for Discrete Channel Frequency Assignment*. Radio Science, 32(5): 1921-1939, 1997.

- [50] DERA, R., University of Bristol, *Technical Working Party on Mobile and Terrestrial Propagation - Proposed Format for 3-D Geographical Datasets (Ver2.1)*. Technical Report: DERA, RAL, University of Bristol, 1999.
- [51] Taplin, R., *et al. Cell Planning Optimisation Tool*. Technical Report: IST-1999-11571 EMBRACE Deliverable D17. Cardiff University, RAL, 2002.
- [52] COST Project 210. Commission of the European Communities Brussels 1992.
- [53] ITU, *ITU-R Recommendation P.526-7, Propagation by diffraction*. 2001.
- [54] ITU, *ITU-R Recommendation P452-10, Prediction procedure for the evaluation of microwave interference between stations on the surface of the Earth at frequencies above about 0.7 GHz*. 2001.
- [55] Allen, S.M., *et al. A Software Tool for Optimised Millimetric Radio Access Cell Planning*. IST Mobile Communications Summit. Thessaloniki. 2002.
- [56] Allen, S.M., *et al. Cost-Effective Cell Planning of Millimetre Broadband Access Systems*. IEE ICAP2003, Exeter, U.K. 1: 280-284, 2003.
- [57] Garey, R.G. and Johnson, D.S., *Computers and Intractability: A Guide to the Theory of NP-Completeness*: WH Freeman & Co., 1979.
- [58] Glasser, C., *et al., The Complexity of Base Station Positioning in Cellular Networks*. Proceedings in Informatics. 8: 167-177. Carleton Scientific, 2000.
- [59] Manber, U., *Introduction to Algorithms: A Creative Approach*. Addison-Wesley, 1989.
- [60] Papadimitrou, C.H. and Steiglitz, K., *Combinatorial Optimisation: Algorithms and Complexity*. Dover, 1982.
- [61] Reeves, C.R., *Modern Heuristic Techniques for Combinatorial Problems*. Advanced Topics in Computer Science, McGraw-Hill, 1995.
- [62] Kirkpatrick, S., *et al., Optimisation by Simulated Annealing*. Science, 220(4598): 671-680, 1983.
- [63] Aarts, E. and Korst, J., *Simulated Annealing and Boltzman Machines*. John Wiley & Sons, 1989.
- [64] Goldberg, D., *Genetic Algorithms in Search, Optimisation and Machine Learning*. Addison-Wesley, 1989.

- [65] Calegari, P., *et al.*, *Combinatorial Optimization Algorithms for Radio Network Planning*. Theoretical Computer Science, 263(1-2): 235-245, 2001.
- [66] Zimmermann, J., *et al.*, *ENCON: An Evolutionary Algorithm for the Antenna Placement Problem*. Computers and Industrial Engineering, 44: 209-236, 2003.
- [67] Glover, F., *Tabu Search- Part I*. ORSA Journal on Computing, 1(3): 190-206, 1989.
- [68] Glover, F., *et al.*, *A User's Guide to Tabu Search*. Annals of Operations Research, 41: 3-28, 1993.
- [69] Corne, D., *et al.*, eds. *New Methods in Optimization*. McGraw-Hill, 1999.
- [70] Demestichas, P., *Computationally Efficient Algorithms for Location Area Planning in future Cellular Systems*. Computer Communications, 2000(23): 1263-1280, 2000.
- [71] Deb, K., *Evolutionary algorithms for multi-criteria optimisation in engineering design*, *Evolutionary Algorithms in Engineering and Computer Science*. John Wiley & Sons, 1999.
- [72] Fonseca, C. and Fleming, P., *An Overview of Evolutionary Algorithms for Multiobjective Optimisation*. Evolutionary Computing, 2(3): 1-16, 1995.
- [73] Meunier, H., *et al.*, *A Multiobjective Genetic Algorithm for Radio Network Optimisation*. Proceedings of the Congress on Evolutionary Computation, 2000. 317-324, 2000.
- [74] Johnson, S., *Emergence*. Penguin, 2001.
- [75] Dorigo, M. and Stutzle, T., *The Ant Colony Optimisation Metaheuristic: Algorithms, Applications and Advances*. 251-285, in *Handbook of Meta-Heuristics*, Kluwer Academic Publishers, 2002.
- [76] Dorigo, M., *et al.*, *The Ant System: Optimisation by a Colony of Cooperating Agents*. IEEE Transactions on Systems, Man and Cybernetics, 26(1): 1-13, 1996.
- [77] Parpinelli, R., *et al.*, *Data Mining with an Ant Colony Optimisation Algorithm*. IEEE Trans. on Evolutionary Computing. 6(4): 321-332, 2002.
- [78] Wagner, I.A., *et al.*, *ANTS: Agents on Networks, Trees and Subgraphs*. Future Generation Computer Systems, 16: 915-926, 2000.

- [79] Di Caro, G. and Dorigo, M., *AntNet: Distributed Stigmergetic Control for Communications Networks*. Journal of Artificial Intelligence Research. 1998(9): 317-365, 1998.
- [80] Dorigo, M. and Gambardella, L., *Ant Colonies for the Travelling Salesman Problem*. BioSystems. 43: 73-81, 1997.
- [81] Levine, J. and Ducatelle, F., *Ant Colony Optimisation and Local Search for Bin Packing and Cutting Stock Problems*. UK Workshop on Computational Intelligence (UKCI-01). Edinburgh, 2001.
- [82] Stutzle, T. and Hoos, H.H., *MAX-MIN Ant System*. Future Generation Computer Systems, 16: 889-914, 2000.
- [83] Trelea, I.C., *The Particle Swarm Optimisation Algorithm: Convergence Analysis and Parameter Selection*. Information Processing Letters, 85: 317-325, 2003.
- [84] Hendtlass, T., *Preserving Diversity in Particle Swarm Optimisation*, P.W.H. Chung, C.J. Hinde, and M. Ali, Editors. Springer-Verlag. 31-40, 2003.
- [85] Salman, A., et al., *Particle Swarm Optimisation for Task Assignment Problem*. Microprocessors and Microsystems, 2002(26): 363-371, 2002.
- [86] Roozmond, D., *Using Intelligent Agents for pro-active, real-time urban intersection control*. European Journal of Operational Research, 2001(131): 293-301, 2001.
- [87] Bocker, J., et al., *Using a Multi-Agent Approach to Optimise the Train Coupling and Sharing System*. European Journal of Operational Research, 2001(134): 242-252, 2001.
- [88] Blum, J. and Eskandarian, A., *Enhancing Intelligent Agent Collaboration for Flow Optimisation of Railroad Traffic*. Transportation Research Part A: Policy and Practice, 36(10): 919-930, 2002.
- [89] Kramer, K., et al., *Tutorial: Mobile Software Agents for Dynamic Routing*. MIT Media Laboratory: Cambridge, MA. 1999.
- [90] Joshi, A., et al., *Multi-Agent Simulation of Complex Heterogeneous Models in Scientific Computing*. Mathematics and Computers in Simulation, 1997(44): 43-59, 1997.
- [91] Cardon, A., et al., *Genetic Algorithms using Multi-Objectives in a Multi-Agent System*. Robotics and Autonomous Systems, 2000(33): 179-190, 2000.

- [92] Gilli, M. and Winker, P., *A Global optimisation heuristic for estimating agent based models*. Computational Statistics & Data Analysis. 42(3): 299-312, 2003.
- [93] Lissajoux, T., et al., *Genetic Algorithms as Prototyping Tools for Multi-Agent Systems: Application to the Antenna Parameter Setting Problem*. Lecture Notes in Computer Science. 1437: 17-28, 1997.
- [94] Socha, K. and Kisiel-Dorohinicki, M., *Agent-Based Multi-Objective Optimisation*. Proceedings of the IEEE 2002 Congress on Evolutionary Computation, CEC '02. 1:109 – 114, 2002.
- [95] Grondalen, O., et al., *TDD versus FDD access schemes*. Technical Report: EMBRACE Deliverable D8, 2002.
- [96] ITU, *ITU-R Recommendation P.530-7: Propagation data and prediction methods required for the design of terrestrial line-of-sight systems*. 2001.
- [97] Usman, I., *Route Diversity Analysis and Modelling for Local Multi-Point Distribution Systems*. RAL, 2002.
- [98] Usman, I., *Personal Communication*. RAL, 2003.

Appendix I

The Malvern Scenario

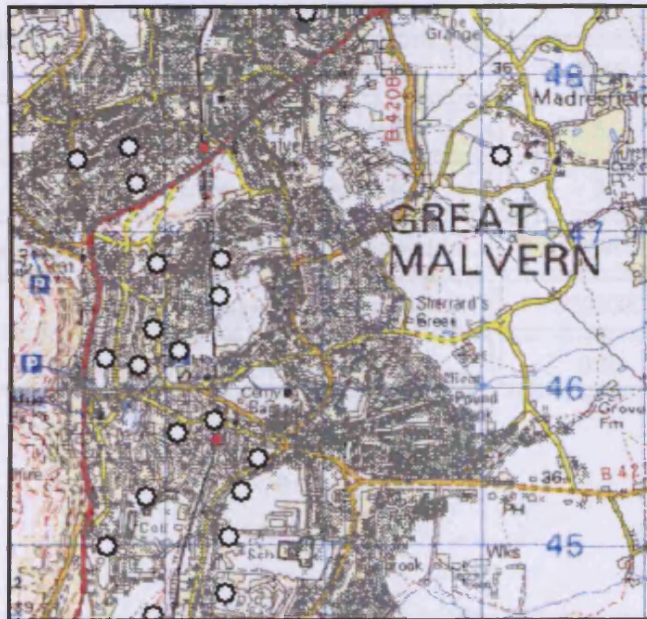


Figure AI-1. A map of the Malvern scenario.
Potential base sites are shown as white dots

Table AI-1. The Malvern scenario in numbers

Required C/I threshold	14
Bandwidth per channel	40000
Signal threshold	-114
Maximum channels per sector	10
Channel back off	3.0
Financial model return period	3.00
Maximum sector power	0.00
Minimum sector power	-30.00
Boundary signal level	1000.00
Number of users	7085
Number of potential sites	21
Number of reception test points	7085
Number of hub antenna models	5
Number of user antenna models	3
Maximal user expected required bandwidth	998985
Maximal user expected revenue	8502000
Average subscription probability	0.20

Appendix II

The Sevenoaks scenario




Figure AII-1. A map of the Sevenoaks scenario.
Potential base sites are shown as white dots.

Table AII-1. The Sevenoaks scenario in numbers.

Required C/I threshold	14
Bandwidth per channel	40000
Signal threshold	-114
Maximum channels per sector	10
Channel back off	3.0
Maximum sector power	0.00
Minimum sector power	-30.00
Boundary signal level	1000.00
Number of users	12966
Number of potential sites	14
Number of reception test points	12966
Number of hub antenna models	5
Number of user antenna models	3
Maximal user expected required bandwidth	285946
Maximal user expected revenue	7519500
Average subscription probability	0.20

Appendix III

AgentOpt initialisation parameters

Table AIII-1. The AgentOpt initialisation parameters.

Parameter Name	Description	Typical Value
AgentOptIterations	The number of iteration in the optimisation run	1000
NeighbourhoodRadius	The size of the radius of each User-Agents neighbourhood, in meters	250
PetitionThreshold	The number of petitioners needed to create a sector	200
CoverageThreshold	The maximum percentage of User-Agents that must be covered for the Site-Agent to concentrate on service	50
MinimumPetitioners	The minimum number of petitioners that needs to be exceeded for the Site-Agent to consider improving a sector.	20
SiteMergePercent	Controls the merging of different site sectors	25
DestroyThreshold	Number of iterations that must be enacted before a sector can be destroyed	20
RandomChoicePercent	Controls the decision to accept a change to a sector configuration	100
DormantThreshold	Number of iterations that a User-Agent has to be inactive before testing to go dormant	50
WakeUpValue	Number of iterations that a User-Agent must stay dormant before testing to wake up	10

Appendix IV

Subscription take-up networks

Further to section 8.1, three networks designed for a 30% subscription take-up are detailed in Table AIV-1 and figures AIV-1, AIV-2 and AIV-3.

Table AIV-1. Networks designed for 30% take-up

Network ID	30%-1	30%-2	30%-3
Number of sites in use	5	4	5
Number of sectors deployed	7	8	11
Expected Profit	9999570	10055910	9161620
Expected Outlay	2227830	2173290	2525780
Expected Income	12227400	12229200	11687400
% of Maximum Income	95.88	95.89	91.6
Coverage	0.99	0.98	0.98
Service	0.96	0.96	0.91
Diversity fraction	0.37	0.45	0.46
Number of channels in use	48	50	39
Bandwidth supplied	1920000	2000000	1560000
Maximum Bandwidth	1500603	1500603	1500603



Figure AIV-1. 30%-1




Figure AIV-2. 30%-2




Figure AIV-3. 30%-3

The following figures are taken from an NVT analysis of the three networks




Figure AIV-4. Expected Profit at different levels of subscription




Figure AIV-5. A measure of invalid users at different subscription levels




Figure AIV-6. The percentage of valid users at different subscription levels

Table AIV-2 and figures AIV-7, AIV-8 and AIV-9 show three networks designed for a 40% subscription probability.

Table AIV-2. Networks designed for 40% take-up

Network ID	40%-1	40%-2	40%-3
Number of sites in use	5	5	6
Number of sectors deployed	11	8	12
Expected Profit	13326320	13622680	12252120
Expected Outlay	2804080	2814920	3182280
Expected Income	16130400	16437600	15434400
% of Maximum Income	94.86	96.67	90.8
Coverage	1.00	0.98	0.996
Service	0.95	0.97	0.908
Diversity fraction	0.57	0.29	0.5735
Number of channels in use	61	58	54
Bandwidth supplied	2440000	2320000	2160000
Maximum Bandwidth Required	2000804	2000804	2000804



Figure AIV-7. 40%-1




Figure AIV-8. 40%-2




Figure AIV-9. 40%-3

As shown in figures AIV-10 and AIV-11, all these networks perform considerably worse than the planning tool predicts. At the 40% take-up level none of the networks can serve the desired number of users, and the profit suffers accordingly.




Figure AIV-10. Expected profit at different subscription levels




Figure AIV-11. Percentage of valid users at different subscription levels




Figure AIV-12. Initial and final errors at a 40% take-up rate

Figure AIV-12 gives an indication of why 40%-1 performs better than 40%-2 and 40%-3. Not only are there 50% fewer initial errors when validating 40%-1, but also the proportion of initial errors that can be solved is large.

Networks designed for a 50% take up rate are shown in Table AIV-3 and figures AIV-13, AIV-14 and AIV-15

Table AIV-3. Networks designed for 50% take-up

Network ID	50%-1	50%-2	50%-3
Number of sites in use	6	5	6
Number of sectors deployed	11	9	12
Expected Profit	16818550	16658750	15056900
Expected Outlay	3437450	3186250	3441100
Expected Income	20256000	19845000	18498000
% of Maximum Income	95.30	93.37	87.1
Coverage	0.99	0.98	0.9975
Service	0.95	0.93	0.87
Diversity fraction	0.51	0.57	0.7327
Number of channels in use	77	80	64
Bandwidth supplied	3080000	3200000	2560000
Maximum Bandwidth Required	2501005	2501005	2501005



Figure AIV-13. 50%-1




Figure AIV-14. 50%-2




Figure AIV-15. 50%-3

The following figures were taken from an NVT analysis of the 50% networks.




Figure AIV-16. A comparison of profit expected by NVT at different subscription levels




Figure AIV-17. A comparison of valid user percentage over different subscription levels




Figure AIV-18. The number of users made invalid through various means at a subscription rate of 50%

These results echo the earlier findings from the 30% and 40% networks.

