

University of Strathclyde
Department of Electronic and Electrical Engineering
Communications Division

**Genetic algorithms for large
resource allocation problems.**

Claude Muller

April 1995

Submitted to the University of Strathclyde
for the Degree of Doctor of Philosophy.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyrights Acts as qualified by University of Strathclyde Regulation 3.49. Due Acknowledgements must always be made of the use of any material contained in, or derived from, this thesis.

Abstract

Telecommunications management systems must undergo some radical transformations in order to support the next generation of giant gigabit networks. In particular, Artificial Intelligence (AI) will be necessary to automate some aspects of network management. This thesis concentrates on one particular aspect of management, namely resource management.

A large number of telecommunications applications can be viewed as resource allocation problems and the telecommunications operators require efficient techniques to solve these problems. However, due to the complex nature of these scheduling problems, large instances cannot be tackled easily by conventional techniques such as the deterministic and analytical procedures provided by Operational Research. Alternative approaches such as the ones proposed by AI have to be considered.

This thesis examines genetic algorithms (GAs) - one of these AI techniques - applied to vehicle routing problems with time-windows and technological constraints - a special class of resource allocation problem. In particular, this thesis proposes a new class of GAs. By working directly on the schedules rather than using an indirect representation, these GAs overcome some important limitations shown by traditional GA-based systems.

These direct GAs are examined over a wide range of problems. In a first set of experiments, the direct GAs totally dominate the traditional GA-based systems. In a second study, which compares the performances of these new GAs with the performances of random search, hill-climbing, simulated annealing and tabu search, it is shown that there is no universally best algorithm. However, under certain conditions some algorithms should be preferred to others. In particular, when time is available, the direct GAs should be chosen for the solution of under-resourced problems.

Acknowledgements

I wish to thank Professor D.G.Smith for his support and encouragements during the past four years. Also, I would like to thank Dr Evan Magill and Dr Patrick Prosser whose advice, knowledge and critical supervision have been invaluable. Special thanks also go to my two partners in this project, Chris and Craig. Thank you for your help and friendship. Thanks, also, to Professor J.R. McDonald, and to Professor P. Mars of Durham University, for acting as internal and external examiners respectively.

I would also like to thank David, Lee and Grace - Glasgow would have been a different experience without your early friendship. Just a shame that I have to support United now!!! Also, I have been lucky enough to be part of a very friendly group, thank you guys. I am so sorry, Peter, that 5-a-side is not our strong point. May be, we should practice a little bit more, and go to Caffé Qui a little bit less. Thanks are also due to Alasdair, Sheona, Iain, Neil, Lois, and Lorna's Mum and Dad - 35 Kessington Road has been a great place to relax and be comforted.

Je voudrais également remercier mes parents pour leur nombreux encouragements et leur assistance tout au long de ces 28 années. Merci à vous pour vos efforts. Finally, I would like to thank my wife, Lorna, for all her help, support and patience. What an understatement !!! Thank you very much, Lorna. Promise, I will try not to be so impatient in the future ...

Contents

Acknowledgements	i
Contents	ii
List of Figures	vii
List of Tables	x
Acronyms	xii
1 Introduction.	1
1.1 Objective.	1
1.2 The case for Genetic Algorithms.	3
1.3 Thesis outline.	5
2 AI & Telecommunications.	7
2.1 The Information Society.	7
2.2 Network Management.	13

2.2.1	Functional Division.	14
2.2.2	Telecommunications Management Networks.	16
2.2.3	Network and Service management.	18
2.3	AI in Network Management.	20
2.3.1	Expert Systems.	22
2.3.2	Model-based Reasoning.	23
2.3.3	Constraint Satisfaction.	24
2.3.4	Neural Networks.	24
2.3.5	Distributed Artificial Intelligence.	25
2.4	Conclusion.	26
3	Novel Approaches to Resource Allocation.	28
3.1	The resource allocation problem.	29
3.2	Genetic Algorithms.	33
3.2.1	Introduction.	33
3.2.2	Genetic algorithms: a brief description.	34
3.2.3	Genetic algorithms: the theory.	38
3.2.4	The driving forces.	41
3.2.5	GAs for resource allocation.	46
3.2.6	GAs for Vehicle Routing.	56
3.2.7	GAs for telecommunications.	59

3.3	Simulated Annealing.	62
3.3.1	Introduction.	62
3.3.2	SA: A brief description.	64
3.3.3	The driving forces.	67
3.3.4	Comparative studies.	69
3.4	Tabu Search.	72
3.4.1	Introduction.	72
3.4.2	TS: a brief description.	73
3.4.3	The driving forces.	74
3.4.4	Applications of TS.	78
3.5	Brief comparison of GA, SA and TS.	79
4	Preliminary Studies.	82
4.1	A simple prototype.	82
4.2	A particular VRP.	88
4.3	Conclusion.	93
5	Experimental Set-up.	97
5.1	The set of problems.	97
5.2	An active constrained-based model.	100
5.3	The cost function.	101
5.4	Direct GAs.	102

5.4.1	Crossover #1	104
5.4.2	Crossover #2.	107
5.4.3	Crossover #3.	108
6	Experiments and Discussions.	112
6.1	Genetic Algorithms.	113
6.1.1	Comparison of different genetic algorithms.	114
6.1.2	The effect of population size.	133
6.1.3	GA with repair algorithm.	146
6.2	The 10 minute Beauty Contest.	154
6.3	CPU Intensive Experiments	165
7	Conclusion.	171
7.1	Summary.	171
7.2	Future work.	174
7.2.1	Stochastic techniques and other problem-domains.	175
7.2.2	The objective function is a parameter.	177
7.2.3	Dynamic problems.	177
7.2.4	A distributed and cooperative approach.	177
7.2.5	Integration with a graphical user interface.	179
7.3	Concluding remark.	181

List of Figures

2.1	The Information Society.	8
2.2	Network management.	15
2.3	TMN.	17
2.4	A layered view of network management.	19
2.5	A selection of AI techniques.	22
3.1	A VRP and one of its possible solutions.	31
3.2	A chromosome.	35
3.3	Flow of a genetic algorithm.	36
3.4	The stochastic nature of GAs.	38
3.5	A chromosome and its schemata.	39
3.6	Building blocks.	40
3.7	Representations for Resource Allocation Problems.	47
3.8	Indirect representation.	48
3.9	Implementation of simulated annealing.	66
3.10	Tabu search - Short-term memory module.	73

4.1	Current domains and durations of the operations.	83
4.2	One optimal solution found by the genetic algorithm.	84
4.3	Object representation of a genetic algorithm.	85
4.4	Object representation of a chromosome.	87
4.5	A particular VRP.	89
4.6	Job-centred and engineer-centred perspectives.	91
4.7	Evolution of performances with time.	93
5.1	The 36 problems.	98
5.2	Pseudo-code for Crossover #1.	105
5.3	Crossover #1, step by step.	106
5.4	Pseudo-code for Crossover #3.	110
5.5	Crossover #3, step by step.	111
6.1	Selection function - an elitist strategy.	116
6.2	Ranking of GAs over 36 problems.	119
6.3	GAs - When the best solution was found (% of 10 minutes).	122
6.4	Population free of duplicates: the larger, the harder.	135
6.5	Different population sizes, different ranking.	138
6.6	Ranking of GAs with or without repair algorithm.	147
6.7	Ranking of the search techniques for 10 minute runs.	157
6.8	Ranking of the search techniques for 4 hour runs.	165

6.9	Computational effort - Comparison between short and intensive runs.	166
7.1	Management system: observation, control, and presentation.	180

List of Tables

1.1	An example of exponential complexity.	3
3.1	Some resource allocation problems.	29
3.2	Operators for the sequencing representation.	50
3.3	Sequencing crossovers - A comparison.	52
3.4	GA, SA, and TS: A comparison.	80
4.1	Description of a job set.	82
4.2	Optimal solution found by the genetic algorithm.	85
4.3	Performances of the two scheduling functions.	92
6.1	GAs performance - Average final cost.	125
6.2	GAs performance - Average amount of work done in minutes.	126
6.3	GAs performance - Average amount of travel time in minutes.	127
6.4	GAs performance - Average amount of work not done in minutes.	128
6.5	GAs performance - Ranking of GAs (versus cost).	129
6.6	GAs performance - Amount of iterations in 10 minutes.	130

6.7	GAs performance - Amount of solutions rejected in 10 minutes. . .	131
6.8	GAs performance - When the best solution was found (% of 10 minutes).	132
6.9	Average final cost for different population sizes.	140
6.10	Amount of iterations for different population sizes.	141
6.11	% of chromosomes accepted for different population sizes.	142
6.12	Average final cost for different population sizes and acceptance policies.	143
6.13	Ranking for different population sizes and acceptance policies. . . .	144
6.14	When the best solution is found (in % of 10 minutes) for different population sizes and acceptance policies.	145
6.15	Effect of repair algorithm - Average final cost.	150
6.16	Effect of repair algorithm - Ranking of the algorithms.	151
6.17	Effect of repair algorithm - Amount of iteration in 10 minutes. . . .	152
6.18	Effect of repair algorithm - When the best solution was found (in % of 10 minutes).	153
6.19	10 minute runs - Average final cost.	161
6.20	10 minute runs - Ranking of the search techniques.	162
6.21	10 minute runs - Amount of exploration in 10 minutes.	163
6.22	10 minute runs - When the best solution is found (in % of 10 minutes).	164
6.23	4 hour runs - Average final cost.	168
6.24	4 hour runs - Ranking of the search techniques.	169
6.25	4 hour runs - Amount of exploration.	170

Acronyms

AI	Artificial Intelligence
ATM	Asynchronous Transfer Mode
AT&T	American Telephone and Telegraph Corporation
BISDN	Broadband Integrated Services Digital Network
CAP	Channel Assignment Problem
CCITT	Comite Consultatif International Telegraphique et Telephonique
CIM	Computer Integrated Manufacturing
CNA-M	Cooperative Networking Architecture-Management
CSP	Constraint Satisfaction Problem
DAI	Distributed Artificial Intelligence
DAR	Dynamic Anticipatory Routing
DCF	Data Communications Function
EDS	Electronic Data Systems
FIFO	First In - First Out
GA	Genetic Algorithm
GUI	Graphical User Interface
HC	Hill-climbing
IN	Intelligent Network
ISO	International Standards Organisation
ITU	International Telecommunications Union
ITU-TSS	International Telecommunications Union - Telecommunications Standards Sector
MF	Mediation Function
NEF	Network Element Function
NT&T	Nippon Telephone and Telegraph Corporation

OSF	Operations System Function
PMX	Partially Mapped Crossover operator
PTT	(Ministry of) Posts, Telegraphy and Telephony (a general term for a public telecommunications operator)
QAP	Quadratic Assignment Problem
QoS	Quality of Services
RACE	Research and development in Advanced Communications for Europe
RBOCs	Regional Bell Operating Companies
RS	Random Search
SA	Simulated Annealing
SDH	Synchronous Digital Hierarchy
SIP	Societa Italiana per l'Esercizio delle Telecomunicazioni
TINA-C	Telecommunications Information Network Architecture Consortium
TMN	Telecommunications Management Network
TS	Tabu Search
TSP	Travelling Salesman Problem
VC	Virtual Channel
VLSI	Very Large Scale Integration
VP	Virtual Path
VRP	Vehicle Routing Problem
WSF	Work Station Function

Chapter 1

Introduction.

1.1 Objective.

This thesis investigates Genetic Algorithms (GAs) for solving large-scale resource allocation problems, with special interest for problems in the field of telecommunications.

A large number of telecommunications applications can be viewed as resource allocation problems; for instance:

- Work force management - this is concerned with the optimal allocation of a work force to a set of construction or maintenance operations.
- Frequency allocation - the bandwidth employed for mobile phone is a limited resource and, therefore the frequencies must be allocated efficiently in order to maximise the number of channels.
- Allocation of bandwidth for video broadcasting networks.
- Design and dimensioning - given a description of network requirements and

a description of the available resources, it is possible to ensure that proposed designs meet a specified level of service.

An automated and efficient management (or allocation) of both the workforce and of the elements composing the network offers the opportunity to:

- reduce the cost of operating the network,
- improve the quality of service,
- provide service to a larger number of customers,
- compete in an open environment,
- and finally, increase profits.

Therefore, an efficient solution of this class of problems, both in terms of quality and running time, is of major concern to telecommunications operators.

In this field of research, a major difficulty is that large problems are intractable. Although conventional techniques, such as the analytical procedures proposed by operational research, can now tackle problems of respectable size, there is still an important need for approximate techniques.

Basically, these conventional techniques might require exponential time in the worst case, and hence, might fail to provide a solution - for large instances - in reasonable time.

Table 1.1, taken from [GARE79], illustrates this point and this table shows how long it would take a computer to explore m^n states for different small values of m and n . It is assumed that the computer performs an instruction in a micro-second.

m	n=10	n=20	n=30	n=40	n=50	n=60
2	0.001 second	1.0 second	17.9 minutes	12.7 days	35.7 years	36.6 centuries
3	0.059 second	58 minutes	6.5 years	3855 centuries	$2 * 10^7$ millennia	$1.3 * 10^{12}$ millennia

Table 1.1: An example of exponential complexity.

By contrast, some instances of the problem used for the experiments (see Chapters 4-5-6) have an estimated search space of more than 10^{500} states¹.

Clearly, an exhaustive exploration of such large search spaces cannot be performed in reasonable time. Systematic or deterministic algorithms cannot, realistically, guarantee that they will find the optimal solution of large instances. Observe that, even a significant acceleration (2 or more orders of magnitude) of computer speed would not be sufficient. These problems are simply intractable by nature.

Therefore, alternative approaches such as the ones proposed by Artificial Intelligence (AI) have to be considered - approaches which sacrifice optimality for efficiency.

1.2 The case for Genetic Algorithms.

A Genetic Algorithm ² is an AI search technique. More precisely, it is a stochastic procedure, i.e. random-based decisions take an important part in the search process. This particular property has been one of the principal motivations for using GAs to solve resource allocation problems.

Indeed, their stochastic nature might permit GAs to be more efficient than deter-

¹By comparison, the universe has an estimated size of 10^{80} atoms.

²Section 3.2.2 provides a brief tutorial on genetic algorithms.

ministic algorithms for handling the complex nature of some resource allocation problems (see Section 3.1 for further explanation) and generally, GAs are regarded as being an efficient means for the rapid exploration of large and complex search spaces [GLOV87] [DAVI91a].

The application of GAs to resource allocation per se is not novel. However, this research proposes a GA which overcomes some important limitations of many existing GA-based systems. In such systems, the GA works with an intermediate representation and is used in association with a solution-builder (see Section 3.2.5). The GA manipulates order-based chromosomes and the solution-builder uses these order-based chromosomes to create schedules. These systems have a number of weaknesses:

1. The indirect representation might not allow total exploration of the search space, i.e. the intermediate representation might mask some parts of the real search space.
2. This indirect approach has been applied successfully to a wide range of sequencing problems. However, this technique is likely to be disruptive when applied to problems which are not pure sequencing problems.
3. This indirect approach means that, for each new chromosome, a new solution is entirely created from scratch and this new solution is also entirely evaluated from scratch. These two operations are extremely expensive in terms of computational effort.

This research proposes and investigates a class of GAs which work directly on the candidate-solutions, i.e. the solutions are the chromosomes. During the reproduction process, part of two solutions are mixed together to produce a new solution. Therefore, there is no distortion due to an intermediate representation and the search process is also accelerated as the GA is no longer generating and evaluating

a new solution entirely from scratch for each new chromosome.

The performances of these new GAs are studied over a wide range of vehicle routing problems³ (VRPs) - a special class of resource allocation problem, and compared against the performances of other search techniques, namely random search (RS), hill-climbing (HC), simulated annealing (SA), tabu search (TS) and also some GAs using traditional (i.e. indirect) crossovers. VRPs have been used for these experiments because the sponsor of this research project, British Telecom, is particularly interested in finding efficient tools to solve these problems.

In this particular study, direct GAs totally dominate the indirect candidate over the entire range of problems. Furthermore, the study shows that there is no universally best algorithm in the set {GA,HC,RS,SA,TS}, but it is shown that under certain conditions, some algorithms should be preferred to others.

When time is available, GA should be preferred for the solution of under-resourced VRPs. However, when the algorithm is only given a limited amount of time or when the VRP is over-resourced, then SA should be preferred.

1.3 Thesis outline.

This thesis contains 7 chapters (including this introduction). Chapters 2 and 3 both present some background material that readers familiar with telecommunications and/or AI may wish to skip. Chapters 4 to 6 concentrate on the different studies and Chapter 7 concludes this thesis. The remainder of this section now gives a brief overview of each chapter:

- Chapter 2, "*AI & Telecommunications*", examines the relationship between AI and telecommunications and highlights the necessity for the use of AI in

³These problems are introduced in Chapter 3.

network management systems.

- Chapter 3, "*Novel Approaches to Resource Allocation*", focuses on three search techniques: genetic algorithms, simulated annealing and tabu search and reviews some of their applications to large resource allocation problems.
- Chapter 4, "*Preliminary Studies*", presents two early projects and draws some initial conclusions. In the first project, an indirect GA is applied to a job-shop scheduling problem and in the second study, an indirect GA is used to solve a particular instance of the vehicle routing problem.
- Chapter 5, "*Experimental Set-up*", gives some detailed information (a) on the 36 problems used for the experiments, (b) on the model used to encode the solutions, and (c) on the cost function. Finally, this chapter also introduces and explains a new class of direct genetic algorithms especially designed for vehicle routing problems.
- Chapter 6, "*Experiments and Discussions*", examines the performance of the direct GAs introduced in the previous chapter. This chapter reports three sets of experiments.
- Chapter 7, "*Conclusion*", summarises the work that has been presented in the previous chapters, discusses areas of future work, and concludes this thesis.

Chapter 2

AI & Telecommunications.

This chapter examines the relationship between two rapidly evolving technologies, Artificial Intelligence and Telecommunications, and highlights the necessity for AI in the next generation of management systems.

The Information Revolution has already started and the pace of this revolution will accelerate in the next decade. The key component of this change, the telecommunications network, will have to undergo some rapid and profound transformations. Novel approaches such as AI will be necessary to automate some aspects of network management.

2.1 The Information Society.

This section highlights the importance of information for the social and economic developments of our society, and describes the rapid changes faced by the telecommunications industry as a result of this information revolution.

The information revolution started 150 years ago with the invention of telegraph

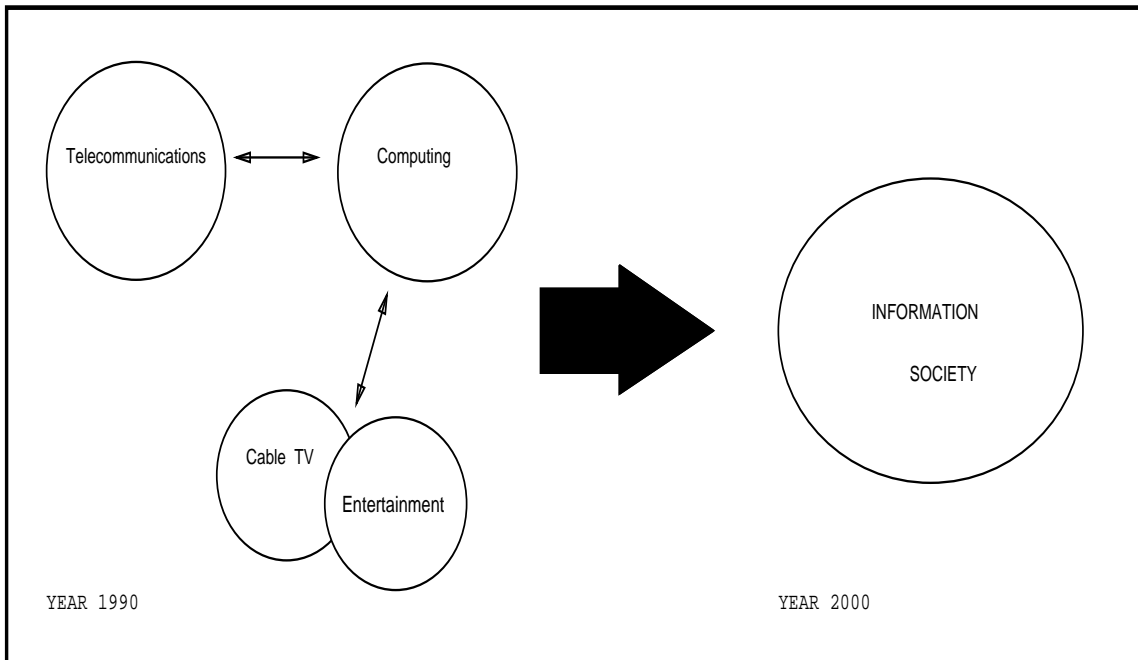


Figure 2.1: The Information Society.

and has known three main periods: the Wire Age (1844-1900), the Wireless Age (1900-1970) and we are now entering the Age of the Information Society [DIZA89].

As illustrated in Figure 2.1, this information society can be described as the coming together of the telecommunications, computing, cable television and entertainment industries [RUDG93]. More precisely, the principal characteristic of this new Age is the convergence (or synergy) between telecommunications and a number of computing-related industries.

The Intelligent Network (IN) with its intensive use of databases and its flexible architecture is a fine example for illustrating this convergence. In fact, the entire network is becoming more and more like an enormous software package [MAYO90]. Even at the lowest level, i.e. in the digital switch itself, the intelligence has moved from the hardware platform to the software. Indeed, some telephones exchanges are now using over 12 million lines of code [JOHN89a].

In order to accelerate the pace of this convergence, different industries (e.g. telecom-

munications, computing and cable TV companies) have come closer together in the recent past. Indeed, there have been many cross alliances, for instance, between (1) telecommunications and software companies and (2) telecommunications and cable TV companies ¹:

- France Telecom and Cap Gemini Sogeti (Europe's biggest software house)
- SIP (Italy's operator) and Finsiel (Europe's second biggest software house) [ECON93a]
- McCaw Cellular Communications (America's largest cell-phone operator and a subsidiary of AT&T) and Microsoft
- Sprint (America's third-biggest long-distance telephone operator) and Electronic Data Systems (EDS, a subsidiary of General Motors and the world's biggest computer services group) [ECON94].

Another example of this convergence is the consortium, named TINA-C (Telecommunications Information Network Architecture-Consortium), recently launched by BT², Bellcore and NT&T, which aims at regrouping equipment suppliers, computer makers, and software houses, and transforming telecommunications networks into information networks [COMM92a].

The origins of this convergence lies in the fact that the importance of information (and its processing) in our society is increasing rapidly; information is becoming a strategic resource as vital as energy. It is likely that, in the near future, a large section of the workforce will either create, process or disseminate information [RACE90]. This means that the weight of the information-based industries - and this is especially true for the telecommunications industry - within the economy is

¹The merger of Bell Atlantic with Tele-Communications International (TCI) has been widely publicised, but this merger collapsed in April 1994.

²formerly British Telecom.

increasing rapidly. Already, in 1993, the world's two most valuable companies were NT&T and AT&T with a respective market capitalisation of \$146 billion and \$85 billion [DUBR93].

In Europe alone, broadband services could generate annual revenue for the carriers of about \$13.1 billion by the year 2000 [RACE90]. By that time, the European Union expects the telecommunications industry to account for more than 7% of its Gross National Product [ECON92a].

In order to meet these expectations, the telecommunications industry must undergo some radical transformations. Indeed, while the computer industry has enjoyed an extremely rapid growth, the pace for the telecommunications industry so far has been slower [RACE90] and the lack of efficient means of telecommunications still represents an obstacle in the development of the information society.

However, this situation is changing rapidly for political and economical reasons. North America has lost its leadership in consumer electronics while Europe has lost its position both in consumer electronics and in the computer industry. Therefore, these two players cannot afford to lose their leading positions in the telecommunications industry.

Europe traditionally has had an important share of the telecommunications market with successful companies such as Alcatel, Siemens, Ericsson, GPT and Philips. In 1991, these European companies sold nearly 55% of all the telecommunications equipment worldwide with the remaining 45% shared between Japan (15%) and North America (30%) [BELL93].

Europe must maintain this strong position despite the intensification of the competition. As a consequence, the European Union launched in 1987 a massive research initiative, RACE (Research and Development in Advanced Communications for Europe) for the purpose of developing technological solutions (e.g. optical switching systems, digital mobile systems) which aimed at introducing cost-effective broad-

band services by the year 1996 [RACE90].

There is a second force which motivates the development of a successful telecommunications industry: corporate users have indeed discovered how to use information technologies as strategic tools. In other words, the ability to process and transfer information quickly and accurately is becoming essential in order to maintain the competitiveness of a business [WILL91]. For instance, the credit-card group VISA implemented recently, a telecommunications system which demonstrates just how important telecommunications services can be as corporate tools. This system, called *PS 2000*, is described as "*a vast information super-highway*" with two computer centres and 9 million miles of optical fibre connecting 18 400 banks around the world. The total cost of this system was around \$40 million; however, it is expected that the *PS 2000* system will reduce the cost resulting of counterfeiting from \$125 million to \$20 million in just one year. Overall, the banks were expected to save some \$270 million in 1993 [ECON93c].

This example shows that private networks, and the associated services and information management facilities that they provide, are becoming increasingly important in the management of a company. Hence, from a company's point of view, the effective management of telecommunications services represents a strategic issue affecting profitability [WILL91].

Corporate users want the telecommunications industry to answer their needs, i.e. they want a telecommunications market driven by the users. Many believe that this will only be possible through the creation of a competitive market. Hence, many large companies are putting pressure on governments and PTTs to increase competition³. Until recently, a real competitive market was not a possible option. In the USA, AT&T was broken up in 1984 and it still had 66% of the long distance

³For instance, Fiat, Europe's second biggest car maker, is lobbying the European Union; it claims that it could cut its \$100 million-a year phone bill by at least \$43 million if the European telecommunications market was liberalised.

traffic in 1992. In the same year, the regional Bell operating companies (RBOCs) controlled, on average 75% of the local traffic [COMM92b]. In the United Kingdom, the duopoly between BT and Mercury was also started in 1984 and BT still had a 92% share of the market in 1992 [ECON92b]. Basically, the investment required for the realisation of a physical network is immense and represents a considerable barrier for any new competitor.

However, the deployment of new technologies such as cable TV, satellites and cellular services allow new entrants to overcome this obstacle. For instance, in [ECON92b], the author claims that a high capacity digital wireless network covering the whole of Britain could cost under \$2 billion to build (around one order of magnitude less than the original copper network).

To illustrate this new situation, AT&T has recently made a strategic alliance with McCaw Cellular Communications which has a cellular network covering more than 40% of the US population. Such an alliance permits AT&T to re-enter the local loop market and compete with the RBOCs [COMM92b] [ECON92c]. AT&T, MCI and Bell Atlantic have contracts with cable television companies that will allow them to attack very lucrative sectors of telecommunications; video and data communications [BELL93] [ECON93b]. In the United Kingdom, a number of companies will rapidly challenge the duopolists (i.e. BT and Mercury); for instance, Energis (a telecommunications company owned by the National Grid company), Ionica, British Rail Telecommunications, and British Waterways. On top of this, 39 cable television companies have been granted a licence by the British government to offer telecommunications services [FAGA93].

Thus, the telecommunications market will, very soon, be driven by the users' requirements and the principal demands will be service quality, open access, flexibility, mobility and lower cost. More precisely, Nicholas and Yeomans, in [NICH93], give the following list of basic requirements:

- availability: same systems at any point of requirement, same operator at any point of requirement, single point of accountability.
- open access: ability to inter-operate between systems, ability to inter-operate between providers, ability to choose between systems or providers.
- open use: no utilisation restrictions, application independent, service independent, technology independent.
- rational prices: transparent, non discriminating, commercially realistic.
- accountability: contractual, commercial, defined conditions, dispute process.

These demands introduce major challenges to all strands of the telecommunications industry [FITC91] and in order to meet these demands, the operators must transform their networks radically. The three branches of telecommunications, i.e. transmission, switching, and management, will have to undergo significant transformation. In the near future, SDH ⁴(for transmission) and ATM ⁵(for switching) will provide a network infrastructure capable of offering a wide range of services with diverse traffic characteristics (e.g. bandwidth requirements, delay and loss sensitivity, and burstiness level). In order to support this infrastructure, the way the operators manage their networks will need to be revised entirely [HELL91].

2.2 Network Management.

To quote from [GUID90], "*network management is about the people, procedures and tools required effectively to manage a telecommunications network at each stage in its life cycle (i.e. pre-service, in-service and future service).*" Network management can also be defined as the tasks of:

⁴SDH stands for Synchronous Digital Hierarchy.

⁵ATM stands for Asynchronous Transfer Mode.

1. maintaining an efficient and reliable network,
2. increasing performance both in terms of quality and quantity,
3. while using a minimum of network resources.

Network management may, in fact, be viewed from different perspectives. For the owner, the network is a business, and the primary objective is to provide as much service whilst using the smallest possible investment (i.e. maximise profits); for the operator, the network has to be robust, easy to maintain, analyse and expand and finally, for the user the network must provide (as seen before) service quality, open access, flexibility, mobility and low cost [LIEB88].

Network management must accommodate these different demands.

2.2.1 Functional Division.

The problem with network management is that it is an enormous and complex task which cannot be tackled by a single agent; there has to be a division of the tasks and ISO (International Organisation for Standardisation) has separated the different management functions into five classes [KLER88]:

- **Accounting and commercial management** determines and allocates costs, and charges for the use of the network elements.
- **Performance management** monitors, evaluates and optimises the performance of the network.
- **Configuration management** maintains precise information about the identity and status of every network element.
- **Security management** provides access protection and control over the network elements.

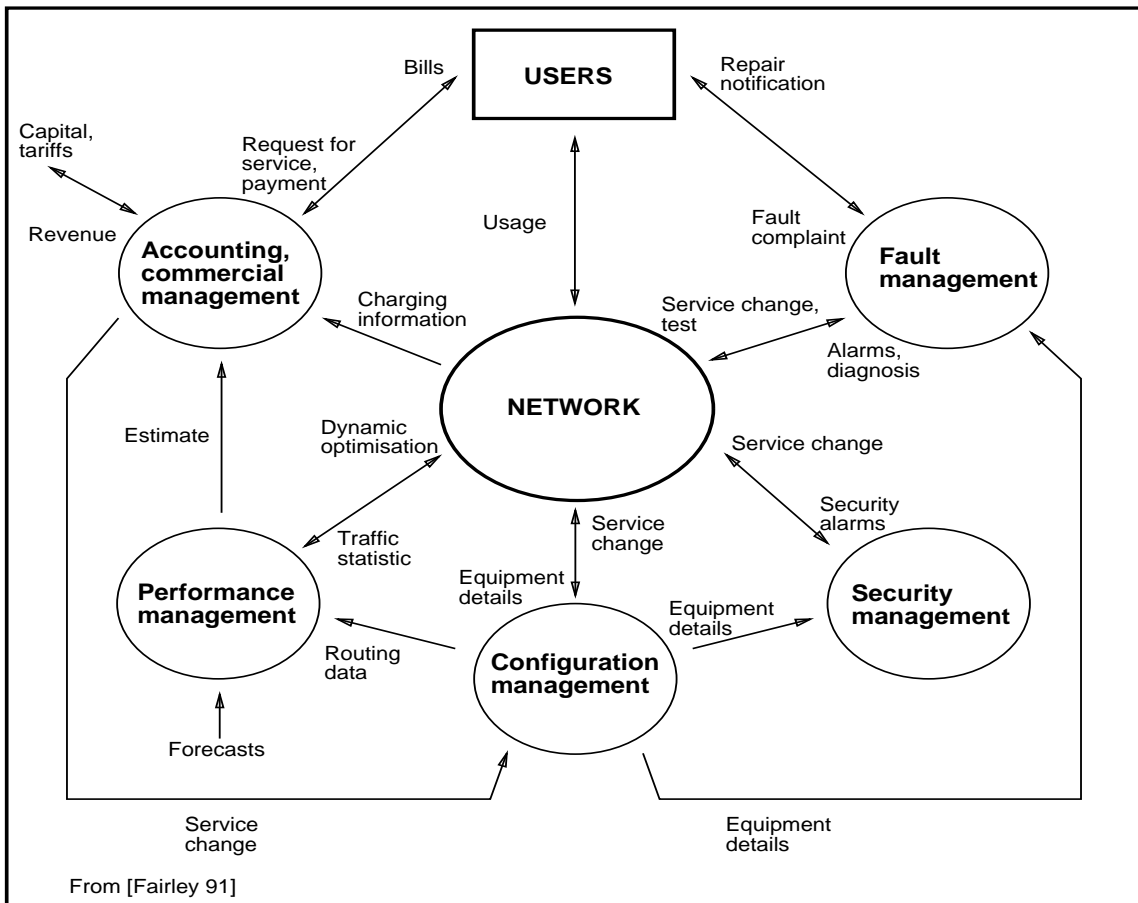


Figure 2.2: Network management.

- **Fault management** is concerned with detection, diagnosis, isolation and repair of faults occurring in the network.

Figure 2.2, from [FAIR91], provides a global view of network management; this diagram also shows the relationships and the flow of data between these different functions.

2.2.2 Telecommunications Management Networks.

The management functions discussed in the previous section are performed by individual operations systems. In order to achieve a global, coherent and integrated management, there is also a need for a management network which would provide the ability for the different operations systems to exchange information and interact.

The CCITT⁶ (Comité International Télégraphique et Téléphonique) Recommendation M.30 proposes a telecommunication management network (TMN) to facilitate the exchange of information between operations systems and network elements [CCIT88].

TMN will enable an operator to manage its own network efficiently but it will also enable different operators to exchange information, to support interaction and to encourage mutual co-operation. This network will have a crucial part to play in the successful deployment of a competitive market, where several companies will operate, where network control will take place in an open co-operative environment (across geographical and technological boundaries) and finally, where the different actors will have to interact and exchange information.

M.30 proposes five functional blocks:

- **Operations System Function (OSF):** performs the different management functions.
- **Mediation Function (MF):** provides the links required between the network elements and the operations systems.
- **Data Communications Function (DCF):** transports management mes-

⁶Following a reorganisation of the ITU (an agency of the United Nations dealing with telecommunications) in 1993, the CCITT has now been replaced by the ITU-TSS (international Telecommunications Union - Telecom Standards Sector).

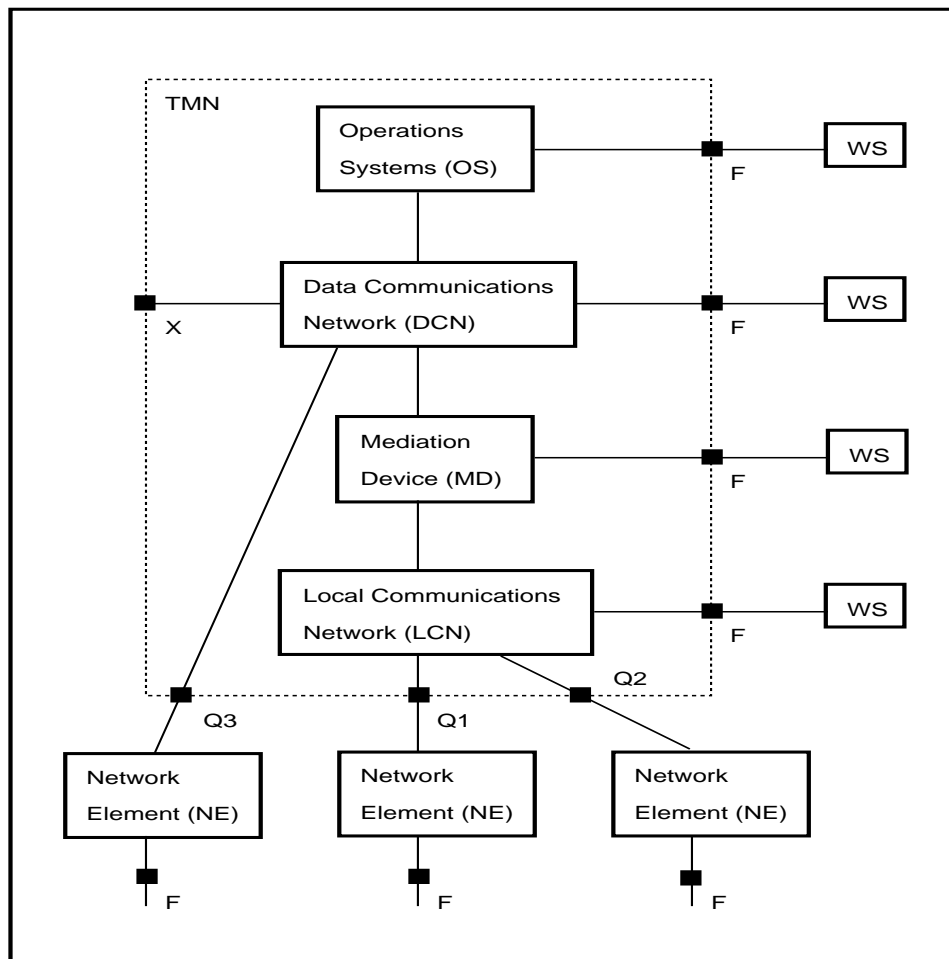


Figure 2.3: TMN.

sages between the different elements of the TMN.

- **Network Element Function (NEF):** the different elements which make up the network.
- **Work Station Function (WSF):** provides an interface between the operator and the TMN.

The interconnections between the various TMN elements (i.e. NEF, MF, and OSF) are made possible by the Q.1, Q.2, and Q.3 interfaces, either directly or via the DCF. The F interfaces link the OSFs either to the WSFs or to the MFs. The G interfaces (not displayed on Figure 2.3) provide the interfaces between the operator

and the TMN, more precisely between the users and the WSFs. Finally, the X interfaces support the dialogue between several separate TMNs [GUID90].

In addition to TMN, the ISO/NM Forum also proposes a framework for the cooperative management of a full range of networks [EMBR90]. Both architectures provide a communications infrastructure to allow for the transmission of monitoring and control data.

2.2.3 Network and Service management.

A layered view of network management is also required: operators can no longer simply maintain their networks and must also manage efficiently the numerous services provided by both themselves and other parties. BT's CNA-M architecture [WILL91] and RACE [SMIT91] have introduced the concept of service management, which is expressed as a distinct layer above network management. A telecommunications service may involve several networks (e.g. voice and data) and several operators (e.g. due to national boundaries). The service management layer will integrate the traffic, manage the underlying networks and perform the interface between customers and operators.

The CNA-M (Cooperative Networking Architecture-Management) is shown in Figure 2.4. Essentially, CNA-M divides network management into a hierarchy of management layers: business, service, network, and element management. The peak of the pyramid indicates minimum operator involvement, while the base of the pyramid indicates maximum operator involvement.

Each layer performs a well-defined class of functions. As seen earlier, the "service layer" is responsible for the interface between the customers and the operator. In [SINN91], this layer is defined as "the layer which must provide, create, monitor, inform, maintain and account for the services of each customer". The service layer




Figure 2.4: A layered view of network management.

requests its information from the network layer. It also provides data to the business layer, which monitors and manages the overall business.

This section has presented several propositions which should permit operators to manage their networks in an efficient manner. However, the increasing complexity of the network and its associated services means that a reduction of cost and the improvement of quality can only be achieved through a new level of automation [MAMD91]. Modern systems can be monitored comprehensively, but the resulting large volume of data must be filtered to identify relevant information. Conventional computer applications provide a degree of automation to this process but, since much of the data is incomplete and conflicting, human interaction remains essential.

In principle, Artificial Intelligence could limit the need for human intervention, and so increase the level of automation [HASK90]. Hence, the application of AI to management has appealed to many operators. The next section will focus on this issue, and discuss the use of AI to help automate aspects of management (at each

layer in the management hierarchy).

2.3 AI in Network Management.

”Artificial Intelligence is the science (and art) of developing computer programs that emulate human-like thinking and reasoning. Unlike conventional software that primarily manipulates data, AI programs depend critically on the manipulation of knowledge.” [CEBU89]

To quote from [BODE90], *”Artificial intelligence was conceived in the mid-1940s, born in the early 1950s, and christened in 1956 (at a small meeting of computer scientists, psychologist, and physiologists)”*. Since then, AI has generated an important debate and a large section of the scientific community does not perceive AI as ”real science”, they would rather speak of the myths of AI (see [WEIZ90], [PENR89]).

A possible reason for this, is that the expectations have always been far too unrealistic. For instance, Herbert Simon and Allen Newell wrote in 1958:

”There are now in the world machines that think, that learn, and that create. Moreover, their ability to do these things is going to increase rapidly until - in the visible future - the range of problems they can handle will be coextensive with the range to which the human mind has been applied.” [SIMO58]

In the middle of the 90’s, this assertion is still unrealistic; indeed, not many 1994-AI systems would satisfy this statement. In fact, at one time, the critics were very important; so much so, that, in 1972, Sir James Lighthill nearly killed off British research by writing in his review for the British government:

”In no part of the field have the discoveries made so far produced the major impact that was then promised (in the 1950s).” [INDE93]

Despite these problems, AI is not just a technology for the future; in fact, it can offer very real benefits today and some products using AI technologies, are now emerging. In 1993, the AI industry represented a market of over \$600 million in America and has reached a mature position [HARM94]. Various AI techniques are now considered for the new generation of management systems. For instance, RACE had 6 projects investigating the applicability of advanced information processing techniques (including AI techniques) to network management in the areas of maintenance, network and customer administration systems, and traffic and quality of services (QoS) management [SMIT91].

AI techniques offer a number of advantages over traditional methods; in particular, they can address ill-defined problems (e.g. missing and/or incoherent data), solve problems that conventional algorithms find intractable (e.g. due to combinatorial explosion) and, in some systems, can learn from past experience.

AI encompasses a wide range of techniques (e.g. knowledge acquisition and representation, case-based reasoning, machine learning, robotics, simulation, virtual reality, etc.). Figure 2.5 illustrates a limited subset of these techniques. In this figure, AI has been divided into a number of classes. The first class is concerned with high-level tasks such as planning, design or diagnosis; it includes expert systems, model based reasoning, constraint satisfaction, and also, natural language and human-computer interfaces⁷. The second class, neural networks, is less procedural and is more concerned with low level tasks such as perception and pattern recognition. The third class, Distributed Artificial Intelligence (DAI), attempts to solve a problem with a collection of agents rather than with a single agent. The last class, Meta-heuristics, includes genetic algorithms, simulated annealing and tabu search. These approaches are now considered separately.

⁷Although these last two techniques will play an important part in the successful development of the next generation of network interfaces, they will not be discussed further.

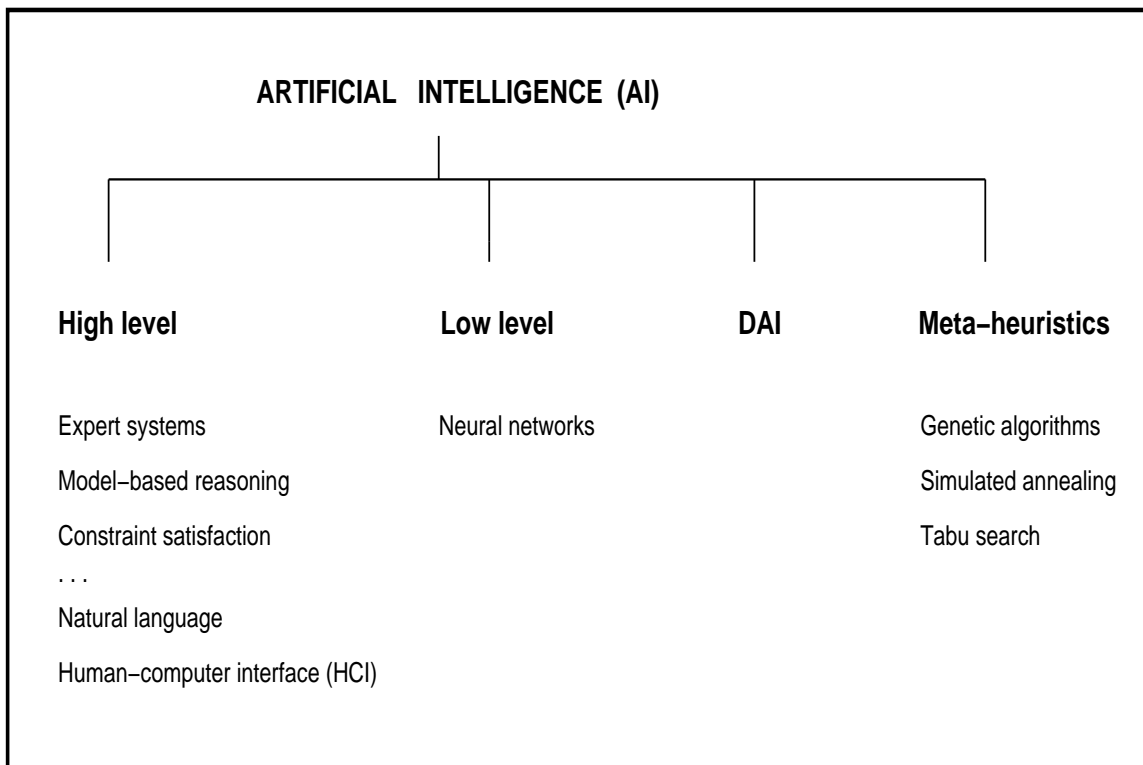


Figure 2.5: A selection of AI techniques.

2.3.1 Expert Systems.

For many companies, expert systems have represented the first step towards AI. They are employed to solve complex tasks and are only knowledgeable in a narrow domain. Under this paradigm, knowledge was originally represented as a set of **“if-then”** rules; they represent the rules of thumb, heuristics employed by the human field-experts to solve their problems.

The principal advantages of expert systems are their simplicity and the flexibility of their internal architectures which permit the rapid development of prototypes. However, the acquisition and validation of the rules represent two major problems [LIPP91]. Also, expert systems have often been developed as prototypes or stand-alone systems and they suffer from a lack of integration in their environments ⁸

⁸This might also be the case for all the other AI-based systems.

[WRIG90]. Finally, their complexity means that expert systems cannot be considered for time-critical operations. All these limitations reduce the applicability of these products to real world problems.

However, ranging from diagnosis through planning to traffic management applications, a large number of expert systems have been developed. In [LIEB88] and in [MALC90], the authors provide general surveys, while design applications and maintenance applications are reviewed in [LIRO91a] and [LIRO91b] respectively.

2.3.2 Model-based Reasoning.

Model-based systems [LIPP91], [MANN90] represent the second generation of AI systems and attempt to overcome the limitations encountered by expert systems. Currently, these systems can mainly be found in maintenance and diagnostic applications [LOES90].

These systems reason with a model of the network under control. This model contains all the necessary information comprising information about the internal behaviour of the components and the dependencies between them. For example, given a set of fault symptoms reported by the network elements, the model-based system reasons with the model and generates a set of fault hypotheses. These fault hypotheses are then used to predict a behaviour (i.e. a set of symptoms). The predicted and the observed symptoms are then compared in order to refine and minimise the set of fault hypotheses.

Model-based systems are not limited by knowledge acquisition and their models can be modified easily. Their principal limitations are (1) the definition of the correct levels of abstraction for the different models [LIPP91] and (2) their computational complexity. Time-critical applications (e.g. traffic management) will require an important speed-up.

2.3.3 Constraint Satisfaction.

Constraint satisfaction views a problem as a set of variables, where each variable has a domain of values and a set of constraints that act between subsets of the variables. The problem is then to find an assignment of values to variables, from their respective domains, such that all the constraints acting between the variables are satisfied [TSAN93]. Constraint satisfaction can be used to solve problems which have a large number of constraints, for example, in machine vision, belief maintenance, temporal reasoning, graph problems, floor plan design [KUMA92] and also in design, planning [LUSH90] or scheduling applications [BURK91].

Unfortunately, the general CSP is NP-complete⁹ and in the worst case, deterministic algorithms may require exponential time. Hence, for large instances, these algorithms may fail to provide an answer in reasonable time.

2.3.4 Neural Networks.

A neural network is based on a massively parallel architecture where many primitive processing elements interact with each other. Its principal task is to perform pattern recognition and it can be trained to recognise a set of patterns [MASS90]. After training, a neural network should be able to handle imperfect or incomplete data, thus providing a degree of fault tolerance.

A number of neural network models are currently under investigation for telecommunications applications. For instance, multi-layer perceptions have been applied mainly to speech recognition and vision [MYER90] [MCCU88]. Feed-backward networks (e.g. Hopfield networks) have been applied to combinatorial optimisation problems, large resource allocation problems [JOHN92] and traffic management problems [MAMD91] [MORR91]. Neural networks can be used on-line for real-

⁹This term is explained in chapter 3.

time applications and a number of VLSI (Very Large Scale Integration) and optical neural networks are now available [HOWA88] [PSAL90].

2.3.5 Distributed Artificial Intelligence.

DAI is the branch of AI that deals with the interaction of intelligent agents [GASS89] [ERCA91]. In other words, a DAI system consists of a group of agents that may interact by cooperation, by coexistence or by competition [CHAI92]. Traditionally, each agent is in charge of a sub-part of the problem. Another alternative is to let each agent solve the entire problem, i.e. the same problem is given to n agents. In both cases, the agents communicate and exchange their findings during the search; different levels of communication, co-operation and control among the agents can be used to achieve a solution.

DAI represents a very promising technology for telecommunications. Several reasons (e.g. complexity of the tasks, competitive environment, and geographical distribution) will force the distribution of management responsibilities among a large number of intelligent agents [GRIF91]. Thus, the capacity to reason with incomplete, inconsistent, and distributed knowledge, the exchange of information and co-operation between local agents will be essential in order to achieve a coherent and global management. Some telecommunications applications are already using DAI techniques; Lebouc and Stern propose a general architecture for distributed network management [LEBO91], Lirov and Melaned solve design problems in a distributed manner [LIRO91a], and Garijo and Hoffman use this technique to solve operations and maintenance problems [GARI92].

2.4 Conclusion.

This section now attempts to associate the four management layers presented in Section 2.2.3 with the different AI techniques introduced in Section 2.3, i.e. decide which AI techniques might be used for which management layer.

At the lowest level (e.g. element management), the time-response becomes critical. The main emphasis is not on providing the optimal answer, but rather it is on providing a *good solution* quickly. The environment changes extremely rapidly, and if the time-response is too slow, the solution becomes irrelevant as the problem changes before the solution can be implemented. AI-based systems have to operate on-line, and they should preferably be hardware-based (like neural networks) rather than software-based. These systems have to process the information available (often incomplete and/or incoherent) and make their decision as quickly as possible, hence iterative methods are not suitable for this kind of problems. Here, the main objective of AI is to accelerate the control process, try to automate that process, and if possible remove any human intervention from the control loop.

At the highest layers of the hierarchy, the main problem for the operator is that he or she has to cope with an overwhelming amount of data. At this level, AI-based systems should be used to assist the user, as decision-support or consultative tools. They can operate off-line; the time-response is important but not critical. The main challenge is to present the user with the right information at the right time. At this level, it will be possible to use AI techniques for processing the raw data and present only the relevant information. For example, model-based expert systems can be used to hide the complexity of the network behind several layers of abstraction.

The next chapter focuses on one particular aspect of management where AI can also be particularly effective: Resource Management (e.g. network design, frequency assignment, work force management). Three meta-heuristics, namely genetic al-

gorithms, simulated annealing and tabu search are presented and some of their applications to resource allocation problems are reviewed.

Chapter 3

Novel Approaches to Resource Allocation.

This chapter investigates three iterative search techniques which offer some promising results for the solution of large resource allocation problems. These techniques are genetic algorithms (GAs), simulated annealing (SA) and tabu search (TS). In contrast with constructive techniques such as constraint satisfaction where the algorithm works with an incomplete solution and constructs the final solution piece by piece, these three methods work as follows: the algorithm starts with one (several for GA) candidate solution and proceeds by transforming this solution into another one, in an attempt to improve its quality. This cycle is repeated until certain termination criteria are satisfied.

This chapter starts by introducing the general resource allocation problem. The different techniques are then presented separately.

3.1 The resource allocation problem.

In a paper entitled "*Why scheduling is hard (and how to do it anyway)?*", Parunak defines a schedule, i.e. one feasible solution of a resource allocation problem, in the following terms:

"A schedule is a subset of a cartesian product of three sets. There is a set of tasks (**What**) that need to be done, there is a set of time-periods (**When**) during which a task might be performed and finally, there is a set of resources (**Where**) that tasks occupy as they execute."
[PARU87]

In other words, Parunak defines the resource allocation problem as the task of deciding **What** happens **When** and **Where**.

	TSP	VRP	Scheduling	Telecommunications
What	visits	deliveries	jobs	calls
When	time-windows	time-windows	time-windows	time-frames
Where	N cities	locations of customers	machines	network elements

Table 3.1: Some resource allocation problems.

Several classes of resource allocation problems ¹ are now considered and are mapped into the above definition (as illustrated in Table 3.1):

- A travelling salesman problem (TSP). In this classical problem, a salesman has to visit N cities and returns to the city of origin. Each city is to be visited once and only once, and the route is to be made as short as possible. In this case, "What" is the set of visits, "When" is the set of time-windows when the different cities can be visited and "Where" is the set of N cities.

¹Applications of GA, SA and TS, in these 4 different problem-domains, are reviewed later in this chapter.

- A vehicle routing problem (VRP) consists of a workforce of engineers operating from one base, and a set of customers requesting a service (e.g. repair). One possible objective might be to perform all jobs while minimising the total distance travelled by the different engineers. Furthermore, this basic problem might be complicated by several side-constraints, e.g. the engineers might operate from several bases located in different parts of the map; the duration of a job might be dependent on the skill of the engineer allocated to this particular job. In this problem-domain, "What" represents the different services performed by the engineers, "When" is the set of time-windows and "Where" is the location of the customers.

To illustrate this class of problem, Figure 3.1 displays one possible solution of a particular VRP (studied in more detail in Section 4.2). The squares represent 11 bases and the circles represent 250 jobs (geographically distributed). This figure also shows the tours of the different engineers and a set of conflicting jobs which have not been allocated.

- Scheduling tasks such as job shop scheduling, flow shop scheduling, wafer fabrication and flexible manufacturing systems; "What" may be a set of jobs to be done, "When" (as usual) may be a set of time-windows during which the jobs can be done and "Where" may be a set of machines that perform the jobs.
- As discussed previously, a large number of telecommunications applications can be regarded as resource allocation problems. In a telecommunications environment, "What" may be a set of calls with diverse traffic characteristics (e.g. bandwidth requirements, delay and loss sensitivity, and burstiness level), "When" may be a set of time-frames when the calls can be carried over the network, and "Where" may be a set of network elements - physical links, switches, multiplexers - required for the completion of the different calls.




Figure 3.1: A VRP and one of its possible solutions.

Resource allocation problems have attracted a lot of attention because their efficient solution can generate massive benefits, e.g. reduction of cost and improvement of quality. However, such problems can be extremely hard to solve; typically, they are NP-hard [PARU87]. NP-hard problems are the hardest problems in the classification of Garey and Johnson [GARE79]. In this classification, combinatorial problems are divided into four classes, from P to NP-hard problems:

1. a **P problem** can be solved by a deterministic algorithm in polynomial time; given the same input, a deterministic algorithm will always follow the same sequence of instructions.

2. an **NP problem** can be solved by a non-deterministic algorithm in polynomial time; a non-deterministic algorithm can be viewed as a two-stage algorithm, first the algorithm guesses a possible solution to the problem under study and in a second stage, the algorithm verifies that guess.
3. an **NP-complete problem** is NP and is at least as hard as every other NP problem.
4. an **NP-hard problem** is NP complete, and may be harder than NP, i.e. it may require exponential time, even for a non deterministic algorithm. In this case, the algorithm may fail to provide a solution (for large instances) in reasonable time.

Observe that this decomposition is a worst case analysis for a whole class of problems but it says nothing about the complexity of a particular instance. It may be the case that particular instances are easier to solve.

Furthermore, the solution of these resource allocation problems is rendered more complicated by two other difficult tasks:

1. implementing a suitable model - or representation - of the problem.
2. designing an appropriate objective function.

The model must give a correct representation of the "real-world" problem and must allow a total exploration of the search space, i.e. the set of all possible solutions. If it provides an over-constrained ² view of the problem, some possible solutions might not be accessible. On the other hand, if the representation is under-constrained, the search technique might have to explore an unnecessarily large search space or might create solutions which are not feasible.

²A problem is said to be over-constrained when too many constraints are acting on the problem and a problem is under-constrained when only a few constraints need to be satisfied.

The objective function - or evaluation function - plays a crucial part in the search process by judging and comparing the quality of the candidate solutions, and hence by guiding the search technique in its exploration of the search space.

However, resource allocation problems have many conflicting objectives and the objective function must reflect and balance the relative importance of these different objectives. For instance, in a telecommunications environment, the evaluation function may have to take into account the following objectives: minimise the worst blocking probability of a set of call classes with diverse traffic characteristics, maintain fairness among these classes, balance load on the network, maximise throughput, maintain robustness, minimise delays.

The remainder of this chapter examines three search techniques, GA, SA and TS. Each technique is described in some detail, and a number of applications are reviewed for each technique. The last section compares the three techniques and concludes this chapter.

3.2 Genetic Algorithms.

3.2.1 Introduction.

In the real world, species can adapt to changing and complex environments mainly because their evolution is dictated by two powerful mechanisms:

1. selection, i.e. survival of the fittest,
2. recombination, i.e. inheritance of genetic material.

Holland [HOLL92] used these two principles to develop GAs. In a few words, a GA is a stochastic and iterative search procedure which works with a pool of candidate

solutions. Iteration after iteration, the algorithms selects above-average solutions for reproduction and recombines them. The ultimate objective of this iterative process is to build a high quality - or optimal - solution from the information present in the pool of solutions.

Holland's original work is reported in a book entitled "*Adaptation in natural and artificial systems*" [HOLL75]. His principal claims were (1) that, given some conditions on the problem-domain, simple data structures, i.e. bit strings, could be used to represent solutions of complex problems and (2) that a repeated cycle of appropriate selections and recombinations could rapidly improve the average quality of a population of solutions.

These characteristics have advocated GAs for the solution of complex problems that conventional algorithms, e.g. hill climbing, exhaustive tree search, cannot handle easily. In particular, this technique is viewed as a promising search technique which gives the opportunity to explore large and complex search spaces in an efficient manner [GLOV87] [DAVI91a]. For instance, GAs have been applied to the following problem-domains: function optimisation [DAVI91b], travelling salesman [GOLD85] [GREF87], vehicle routing [THAN91], bin-packing [DAVI85a], scheduling [LAWT92], graph colouring [DAVI85a], [DAVI91b], routing for telecommunications networks [Cox et al 91] and multiple-fault diagnosis [LIEP91].

3.2.2 Genetic algorithms: a brief description.

How does a GA work?

A genetic algorithm works with a population of chromosomes (or strings) where each chromosome (see Figure 3.2) represents one possible solution to the problem under study; each element (or locus) of a chromosome describes one feature of the problem, and each locus can be assigned different values (or alleles):




Figure 3.2: A chromosome.

Originally, the different elements could only be assigned a binary³ value (0 or 1); the feature was either there or not. Other chromosomal representations have been designed in order to apply GAs to problem-domains where the binary representation is not suitable. Subsections 3.2.5 to 3.2.7 illustrate this last point, and in particular Section 3.2.5 discusses the applications of GAs within the resource allocation domain.

There are many different ways in which to implement a GA; however, they all share the same general principles (see Figure 3.3 for the basic GA cycle). Two possible models - the steady-state model and the generational model - are now introduced.

The steady-state model works as follows: an initial population of chromosomes is created randomly. Each chromosome is given a fitness value which is directly related to the quality of the solution that its binary string represents.

Once this initial stage is finished the algorithm enters a loop. Two chromosomes are selected with a probability which relates to their fitness; the chromosomes with the best fitness being more likely to be selected. Copies of their strings are then manipulated and recombined with different genetic-like operators, e.g. crossover, mutation and inversion, in order to create a new chromosome. Once its fitness has

³The binary representation was preferred because it offered the finest granularity and greatest flexibility, especially in identifying similarities between chromosomes.

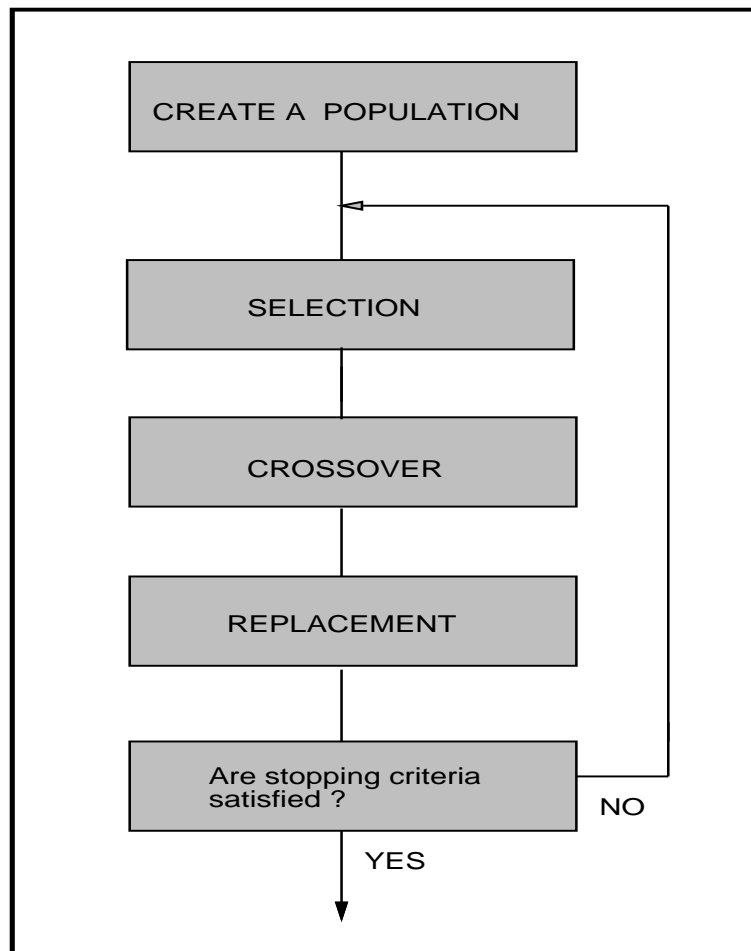


Figure 3.3: Flow of a genetic algorithm.

been evaluated, the "child" is inserted in the population and the worst chromosome of the population, i.e. chromosome with the worst fitness, is eliminated. This cycle "selection, reproduction, replacement" is repeated until a certain termination criterion is satisfied.

The generational model is another, more popular alternative. Under this model, the GA works with two populations: the old one and the new one. Parents (in the old population) are selected for reproduction and they generate offsprings in the new population; an entire population is generated at each population. In some cases, a certain percentage of the old population might be copied - or cloned - in the new population; the term of *generational gap* is used to describe this phenomenon.

Comparison with other search techniques.

GAs differ from other search techniques in several ways:

1. Random-based decisions play an important part in the search process; basically, the final solution delivered by a GA depends on both the initial random starting points, i.e. the initial population, and the random-directed actions performed by different operators during the search. Stochastic techniques such as GAs provide an alternative for the solution of large combinatorial problems as deterministic techniques cannot tackle these problems properly [GARE79].

The stochastic nature of GAs may give rise to quite different computational behaviours, and this has been the principal motivation for advocating GAs for the solution of NP-complete problems [PROS91].

Figure 3.4 illustrates this stochastic nature of GAs. Here, a GA has solved a problem 10 times and on each occasion, it has reached a different solution. This graph shows the reduction in terms of cost over 100 iterations for each of the 10 runs. The shape of these curves is typical, most of the improvements come at the start, and the end of the run only sees small improvements - or no improvement at all. This is a common situation for most optimisation algorithms.

2. Typically, search techniques move from point to point in the search space; instead, GAs work with a population of solutions. For this reason, a GA can be said to have a more global view of the search space than a conventional algorithm and is less likely to be trapped in some inferior region of the search space and, as a consequence, it is less likely to converge towards a poor local optimum [GREF87].




Figure 3.4: The stochastic nature of GAs.

3. The two previous features are integrated into an iterative process. Generation after generation, using random-directed decisions, a GA first selects above-average solutions for reproduction and then recombines them to produce a new set of solutions. If a piece of information is present in many above-average solutions, it will to be propagated to future generations. As stated previously, the ultimate objective of this iterative process is to "build" the optimal solution from the information present in the pool of solutions.

The following subsection introduces the theory behind GAs.

3.2.3 Genetic algorithms: the theory.

This subsection presents the schema (pl: schemata) theorem which provides a formal platform for GAs and can be used to explain their internal mechanisms.

What is a schema?

A schema is the generalisation or abstraction of a binary string; each element of a schema can take the values: 0,1 or #, where # means that the value is irrelevant (0 or 1). For instance, the strings "1000" or "0000" both contain the schema "#000".




Figure 3.5: A chromosome and its schemata.

Figure 3.5 shows a chromosome and its schemata. A chromosome of length L contains 2^L schemata. It means that the evaluation of one chromosome implies the evaluation of 2^L schemata. At each generation, a GA recombines the schemata present in the different chromosomes of its population.

The schema theorem presented in the next paragraph demonstrates that the strength of GAs lies in the efficient processing of these schemata [DAVI91a].

The Schema Theorem.

The Schema Theorem [HOLL75] states that *"Assuming that selection chances are proportional to chromosome fitness, a schema occurring in chromosomes with above-average fitness will tend to occur more frequently in the next generation, and one occurring in chromosomes with below-average fitness will tend to occur less frequently (ignoring the disruptive effects of crossover and mutation)"*. This theorem immediately provides some answers to some important questions:

- What are the basic assumptions behind GAs?
- Why does a GA work with a population of solutions?
- What is the importance of one string?




Figure 3.6: Building blocks.

Assuming a suitable problem-domain, the basic assumptions are that:

1. the optimal solution and the above-average solutions of a problem have some similarities; they have some common schemata.
2. the recombination of schemata present in above-average solutions should lead to the identification of the optimal solution.

Schemata can be viewed as building blocks. Given the assumption that a solution consists of a series of these building blocks (see Figure 3.6), the difference between an above-average solution and the optimal solution is that the optimal solution has all the right building blocks whereas an above-average solution has some of these building blocks but also some poor building blocks. Therefore, if a building block is present in several above-average solutions, that building block is likely to be part of the optimal solution.

In order to identify similarities among above-average chromosomes, a GA must work with a population of solutions. In fact, a chromosome by itself has no real importance. The search is guided by the similarities, i.e. if a schema is present in a large number of above-average chromosomes, this schema is likely to be propagated to the next generation. However, what happens to an individual solution, i.e. the chromosome, is not relevant [GREF87] [GLOV87].

3.2.4 The driving forces.

This subsection analyses a number of mechanisms and parameters which influence and guide the search process of a GA. As stated earlier, a GA has to (1) select above-average solutions for reproduction and (2) recombine them to produce a new set of solutions. Therefore, the different driving forces have two principal objectives:

1. Make sure that the recombination of similarities from above-average solutions is fruitful, i.e. the recombination of two above-average solutions tends to generate an above-average solution.
2. Maintain the best possible balance between solution quality and rapid convergence. Different applications may have different needs: time-critical applications (e.g. traffic, routing of telecommunications networks) will sacrifice optimality for a rapid convergence towards a satisfactory solution, whereas applications with no (or minimal) time-constraints (e.g. design and dimensioning applications) will favour optimality rather than rapid convergence.

Each driving force is now considered in turn:

Representation/operators.

The data structure used to represent candidate solutions, i.e. the representation, and the operators, i.e. crossover and mutation, cannot be considered separately; the operators must be able to work with the representation implemented and therefore, the choice of the data structure dictates the design of the operators. As a result, these two elements can be viewed as a united driving force which must ensure that:

1. Genetic operations are not disruptive⁴ and also children should inherit a maximum of information from their parents.
2. The resolution of the representation is adequate. In fact, the representation can be seen as a discretisation of the search space. If the representation is too restrictive (i.e. over-constrained), some regions of the 'real' search space might not be accessible, while if the representation is too relaxed (i.e. under-constrained), the GA may be forced to explore an unnecessary large search space. In both cases, the GA may not be able to converge towards the optimal solution.

The internal algorithms - crossover and mutation - used to encode the operators, are strongly dependent on the chromosomal representation adopted to solve a specific problem. However, these problem-specific algorithms perform similar functions, i.e. a crossover operator is always a crossover operator whatever representation is being used. The following paragraphs analyse these generic functions.

Crossover and mutation.

The roles of the crossover and mutation operators can be described as follows:

⁴A crossover operation is said to be disruptive if in general, the recombination of two chromosomes tend to generate a solution with a worse fitness than its parents.

The crossover operator acts by combining parts of two parents. Clearly, it recombines groups of bits from two chromosomes and produces two new chromosomes. It does not create any new genetic material but simply manipulates the information already existing in the population. It is analogous to mating in biological organisms.

Mutation operates on a single parent by randomly changing some part of it. It is a background operator, i.e. it occurs rarely, and the main task of mutation is to re-introduce diversity into the population during the search process in order to avoid premature convergence⁵. In genetic terms, this operation is analogous to the mutation of genes, in which the code for one amino acid changes to the code of a different amino acid [SING94].

Crossover is the key-operator for GAs, mutation alone could not generate the recombination process achieved by the crossover operator [SCHA92].

However, the actions of the crossover might be disruptive and these disruptions represent the main difficulty faced by GAs. The search performed by a GA is reduced to a basic random search if the children cannot inherit sufficient information from their parents and, therefore, the disruptions caused by the crossover must be kept to a reasonable level.

For certain classes of problems, it might not be possible to design a suitable crossover operator. For instance, crossover might be especially disruptive for problems with a high level of epistasis⁶ where the various elements of the chromosome are severely constraining each other; the recombination of two above-average solutions is not likely to produce an above-average solution. Davidor advocates pure random search for a solution to these problems [DAVI91a].

⁵Premature convergence can be defined as the loss of diversity in the population before the optimal solution has been found.

⁶In this context, the term "epistasis" is related to the level of dependence existing among the different elements of a chromosome. If many elements are heavily dependent on the values assigned to other elements, the level of epistasis is said to be high.

It is also worth mentioning that GAs might not be the appropriate tool for the solution of problems with a low level of epistasis. In this case, the various elements of a chromosome are not constraining each other; the values they are assigned are more or less independent from each other and a hill-climbing procedure, focusing on each individual element separately, will be preferred. It is simpler to implement and, in this case, more efficient.

The initial population.

The initial population must supply a suitable sample of the search space to the genetic algorithm. The diversity of the initial population must be sufficient so that there is sufficient variety in the schemata to be processed.

In order to give a smart start to the GA, the initial population can be produced with a heuristic rather than using a random method. However, Grefenstette, in [GREF87], reports that the seeding of the initial population with above-average solutions generated by a heuristic can lead to a premature convergence. The author concludes that this heuristic-seeding must be done with care. Our results confirm this point (see Subsection 6.1.2).

The size of the population can also affect the efficiency of GAs. Very small populations may not create enough variety whereas large populations may be prohibitive in terms of computational effort.

The selection function.

The selection function guides the search by determining which chromosomes are used during the reproduction process; it also has a key role to play when deciding between optimality and running time.

For (near) real time applications where a rapid solution is desired, Davidor advo-

cates an exponential selection function: the best solutions are given an exponential chance of being selected. This strategy leads to a rapid convergence towards a good, but probably sub-optimal solution [DAVI91a].

For applications with no time-constraints, the primary concerns of the selection function are (1) to guide the search efficiently and also (2) to avoid a premature loss of diversity in the population [DAVI91b]. In particular, two critical situations might emerge and the selection function must handle them.

1. One chromosome, i.e. a super individual, might have a much better cost than the other members of the population. In this case, the selection function has to attenuate the selection chance of this particular chromosome. If the probability of being selected is directly related to the fitness of the chromosomes, it is likely that the super-individual will always be selected; the population will then lose its diversity and a rapid convergence towards a local optimum will happen. One possible solution is to normalise the situation and to make the chance of selection of each chromosome directly proportional to its ranking in the population, rather than to its individual fitness. This normalisation technique limits the risk of domination by one super-individual.
2. All the chromosomes have a similar fitness. In this case, chromosomes with the highest fitness must be given a better chance. In fact, in order to guide the search efficiently, the selection function must introduce some pressure and exaggerate the selection chances of the fittest chromosomes.

In summary, this subsection has discussed the importance of the different driving forces. These different ingredients all have a critical part to play in the success or failure of a GA. Firstly, assuming a suitable problem-domain, a representation and its associated operators must be designed so that:

- the action of the crossover creates as little disruption as possible and tends

to generate better solutions from one generation to the other,

- the resolution of the representation is adequate. It is necessary for the representation to provide a correct and complete view of the search space.

Moreover, the operators, the initial population and the selection function must be designed so that the best possible balance between optimality and running time is kept.

The remainder of this section on GAs contains three subsections describing the applications of GAs to the following problem-domains:

- Resource allocation (Subsection 3.2.5),
- Vehicle routing (Subsection 3.2.6),
- Telecommunications (Subsection 3.2.7).

3.2.5 GAs for resource allocation.

This subsection reviews a number of applications which use GAs for the solution of resource allocation, i.e. scheduling, problems.

The chromosomal representation that these systems use influences the design of the entire GA, that is, the design of the operators and the resolution of the search space. Therefore, these applications are reviewed with respect to how they represent the problems they try to solve and also, we present a study of the alternative representation schemes, and comment on their effectiveness, or otherwise.

Figure 3.7 presents a number of possible representations which can be divided into two broad classes:

- indirect representation,

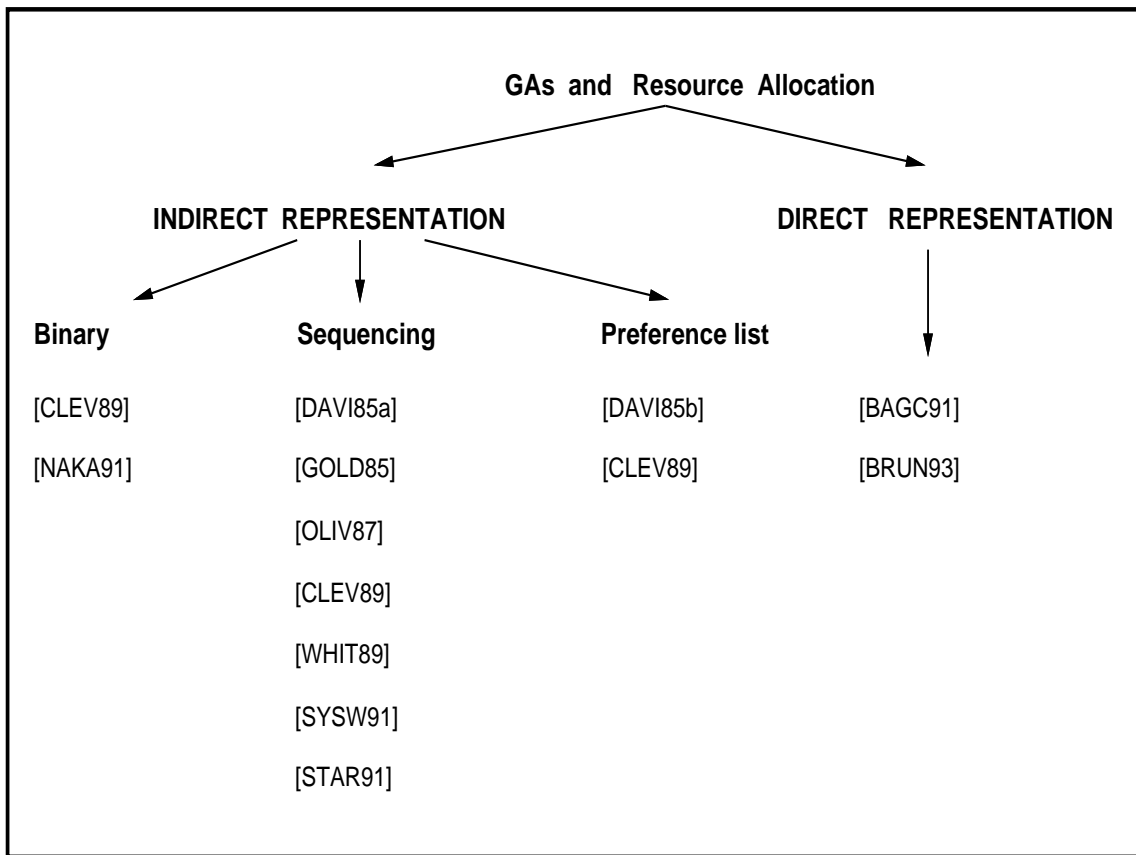


Figure 3.7: Representations for Resource Allocation Problems.

- and direct representation.

The next two paragraphs examine these two styles of representation.

The indirect representation.

The indirect representation is the most widely used. Davis was the first to propose the two-fold "GA-solution builder"⁷ architecture [DAVI85a]. His claims were that for certain classes of epistatic problems such as bin-packing and graph colouring, the conventional, i.e. binary, representation is not suitable. The recombination of two above-average solutions is not likely to produce above-average solutions, rather

⁷or schedule-builder when solving scheduling problems.

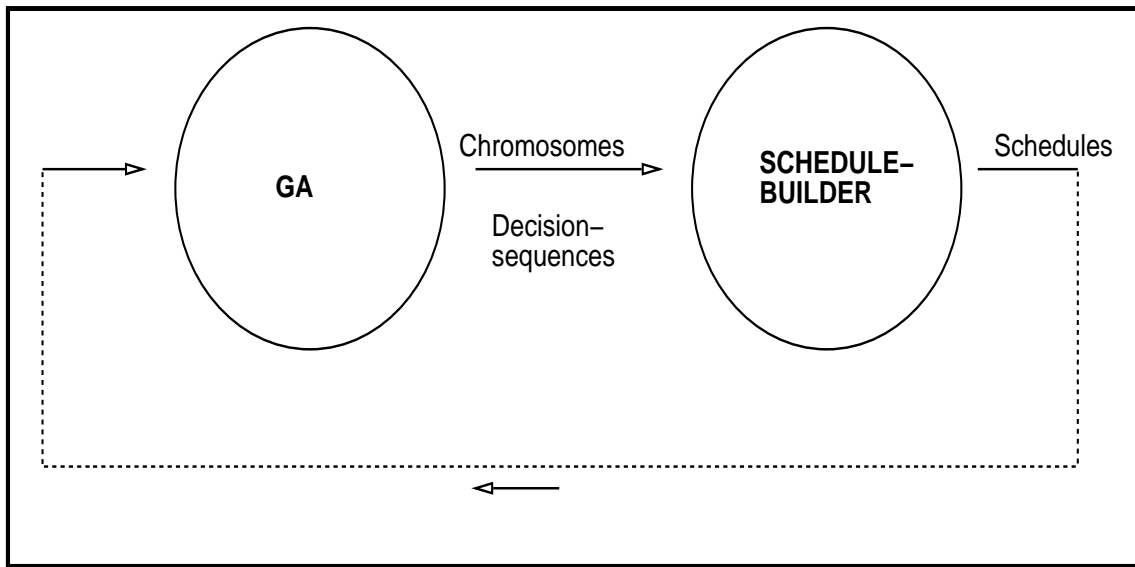


Figure 3.8: Indirect representation.

it is (highly) likely to produce illegal solutions. He, thus, presented an alternative technique, where the GA was only doing part of the search and a solution-builder was decoding the intermediate solutions manipulated by the GA.

Figure 3.8 illustrates how GA-based systems, using this representation, work.

The search process is shared by a GA and a schedule-builder. The GA uses an intermediate representation for its chromosomes. These chromosomes are viewed as decision sequences by the schedule-builder which uses them to generate legal solutions. Thus, each chromosome is associated with a legal schedule and the fitness of the chromosome is defined by the quality of the schedule.

Three different classes of indirect representations can be distinguished:

- binary representation,
- sequencing representation,
- and preference list.

These three sub-classes are now considered in turn.

Binary representation. Nakano applied a conventional GA to three job-shop problems; the largest problem had 20 jobs and 5 machines [NAKA91]. There were some sequencing constraints, and each job had to be processed on different machines following a given sequence.

Each bit of the chromosome was associated with a pair of jobs and a machine; the value of the bit (0 or 1) defined the processing priority between the two jobs on that particular machine.

As the search space covered by the binary representation was much larger than the real search space, most of the solutions produced by the crossover operator were illegal. As a consequence, a schedule-builder was needed to generate legal schedules. Nakano claimed that the solution-quality achieved by the GA approached those obtained by branch and bound methods. However, it is not clear what information is transmitted from generation to generation and how children can inherit that valuable information from their parents.

Cleveland and Smith compared a binary representation with some sequencing and preference list representations on a series of problems on scheduling flow-shop releases [CLEV89]. Their results are discussed later (see paragraph on preference lists).

Sequencing representations. This approach has been the most popular to date with the GA community for the solution of scheduling problems. Here, the problem is reduced to a sequencing problem. Basically, the original problem is transformed into the problem of finding an ordering (or sequence) such that the schedule generated from this ordering has a minimal cost.

Table 3.2 presents a collection of crossover operators which work with a sequencing representation:

The principal difference between these operators is the information that each of

Type	Order-based	PMX (see below)	Cycle crossover	Edge recombination	Position-based
Reference Pressure(s)	[DAVI85a] 1. Order 2. Position	[GOLD85] 1. Order 2. Position	[OLIV87] 1. Position	[WHIT89] 1. Adjacency	[SYSW91] 1. Position

Table 3.2: Operators for the sequencing representation.

them attempts to preserve during the reproduction operation. Under "pressure" is identified the objective(s) for each operator. These are amplified below.

The GA developed by Davis [DAVI85a] was the first to integrate a sequencing representation and an order-based crossover. In this case, the objective of the crossover was to preserve the relative order of the elements contained in the sequences. During crossover, one part of the sequence is inherited from one parent and the other part is filled with the remaining elements, following the relative order of the other parent.

The partially mapped crossover operator (PMX) has been proposed in [GOLD85]; this operator is an order-based crossover which also attempts to preserve the absolute positions for one part of the child and the relative order for the other part. However, Starweather et al. claimed that the PMX operator is more disruptive than the order-based crossover [STAR91].

The cycle crossover has been designed by Oliver [OLIV87]. The objective of this operator is to preserve positional information. Each position of the sequence is guaranteed to inherit the value assigned to the same position in one of the two parents.

The edge-recombination operator was specifically designed for the solution of TSPs [WHIT89]. The authors claimed that the important information, i.e. the information that children must inherit from their parents, for this particular class of problems, is not the absolute positions or the relative order of the elements in the

sequence. Rather, it is the fact that some elements are next to each other in the sequence. Basically, the important information in a solution of a TSP is the fact that city a is visited just before or just after city b , i.e. the fact that cities a and b are adjacent in the tour. Hence, preservation of adjacency information is the objective for this operator.

The position-based operator [SYSW91] had the same objective as the cycle crossover; it attempts to preserve positional information. However, Starkweather et al. claimed that this operator is more disruptive than the cycle crossover and is in fact a variant of the order-based crossover [STAR91].

The performances of these different operators vary from one problem-domain to another, and are, in fact, strongly dependent on the specific nature of the problem in hand. For instance, in [OLIV87], the authors compared the performance of three operators: the order-based crossover, the PMX and the cycle crossover on a single 30 cities TSP. The order-based crossover did 11% better than the PMX and 15% better than the cycle crossover. They also compared (1) the tours achieved by a GA with an order-based crossover and (2) the optimal solution; the best results were within 0.25% of the optimal solution.

Syswerda, in [SYSW91], compared three crossovers (order-based, position-based and edge recombination) on a job-sequencing problem with resource, time and set-up time constraints, and with 90 tasks to be scheduled on 30 resources. The performances of the position-based and the order-based crossover mechanisms were quite similar whereas the edge recombination operator produced rather poor results and was, in fact, doing little better than random search. This operator, designed especially for the solution of TSPs, was obviously not suitable for this type of scheduling problems.

Starkweather et al., in [STAR91], compared six crossover operators on two different problems: a 30 city TSP and a real world warehouse/shipping scheduling problem

with 195 jobs. The six operators were: the edge recombination crossover, two variants of the order-based crossover, the position crossover, the PMX and the cycle crossover (see Table 3.3).

	TSP	Scheduling problem
1	edge recombination	position-based
2	order-based (1 and 2)	order-based (1)
3	position-based	PMX
4	PMX	order-based (2)
5	cycle	edge recombination
6	-	cycle

Table 3.3: Sequencing crossovers - A comparison.

For the TSP, the ranking of the operators was as follows: (1) edge recombination operator, (2) order-based crossovers, (3) position-based crossover, (4) PMX and (5) cycle crossover. The edge recombination, the two order-based crossovers and the position-based achieve similar results; they were all within 1 % of the optimal solution. Compared to the previous operators, both the PMX and the cycle crossover performed poorly on this problem and they were respectively 10% and 20% away from the optimal solution.

For the second problem, the scheduling problem, the ranking of the operators was rather different. The edge recombination operator performed poorly compared to its competitors. The position-based crossover and one of the two order-based crossovers achieved the best result, followed by the PMX operator and the remaining order-based crossover. The cycle crossover was again quite disappointing.

In summary, all these comparative studies show how much the performance of the different operators depends on the specific nature of the problem in hand. The edge recombination operator attempts to preserve adjacency and, thus, is suitable for the solution of TSPs; however, it might not be the best choice for scheduling problems where the notion of adjacency is not so important. The order-

based crossover and the position-based crossover appear to be a suitable choice, for instance, for problems with precedence constraints where the relative order might be the relevant information. None of these studies has been favourable to the cycle crossover; however, this operator might be required to apply GA to sequencing problems where the absolute positions of the various elements in the sequence is of prime importance.

Preference lists. This representation, proposed by Davis [DAVI85b] for the solution of a small job shop scheduling problem is rather rich and complex. The schedule-builder is, in fact, a simulator which must take decisions for each machine at regular intervals; the decision can be to allocate a task to a machine or, the machine may have either to wait for a job or be idle for a given time period.

The chromosome (here, preference list) dictates, for each of these decision points, which decision the simulator should prefer. Clearly, for each decision point, the chromosome provides an ordered list of possible decisions; each of these is attempted by the simulator until one can be accomplished.

Davis [DAVI85b] designed three operators for this representation:

1. "Run-idle" attempted to limit the amount of waiting time experienced by the different machines.
2. "scramble" "scrambled" (or rearranged) the elements of a preference list.
3. "crossover" exchanged preference lists between solutions.

The principal limitation of this representation is that it is not a practical approach for large problems. The preference list must provide an ordered list for each decision-point and for each machine, and therefore can become very large very rapidly when the size of the problem increases.

Cleveland and Smith undertook a comparative study of binary representations, sequencing representations and preference lists [CLEV89]. Two complex single-machine scheduling problems were used as the basis for the comparison. In fact, the same problem was viewed from two different angles. In the first case, the scheduling problem was reduced to a sequencing problem where the order in which the jobs were released was the only concern. In the other case, the problem was viewed as a more general timing problem and, absolute release times for each job were required.

As expected, sequencing approaches performed much better than both the binary representation and the preference list on the sequencing problem. The principal difference between these three techniques was that both the binary representation and the preference list introduced the notion of time in their representation whereas the sequencing approaches were only concerned with the order of the jobs in the decision-sequences. The authors concluded that the binary representation and the preference list were, in a sense, too complex for this problem and could not guide the search in an efficient manner.

The sequencing approaches showed their weakness on the second timing problem. Indeed, in this problem where absolute release times were part of the solution, both the binary representation and the preference list provided significantly better results than the sequencing approaches.

In summary, three classes of indirect representations have been examined. Until recently, indirect representations were the most popular choice for the solution of scheduling problems. The principal reason for this was the relative simplicity of the operators involved.

Direct representations.

Direct representation has been proposed as a promising alternative for the solution of resource allocation problems [BAGC91] [BRUN93]. The genetic operators are more complex. However, they overcome most of the limitations encountered by their indirect counterparts, with which they differ in two principal ways:

1. The GA works directly on the original search space, i.e. the schedules are the chromosomes. There are no distortions introduced by an intermediate search space.
2. The schedule builder is eliminated. More precisely, the scheduling function is performed by the crossover operator. The GA performs as much of the search as possible. Basically, all the variables are encoded in the representation and hence, the schedule builder might not be needed at all or used at worst to refine marginally the solutions proposed by the recombination process.

The definition of suitable operators, crossover and mutation, represent the principal problems with this class of representation. In comparison with the sequencing representations, where the crossover operator only recombines two sequences of elements, the task is here much more complex. The crossover must mix and recombine two schedules and guarantee that, during this operation, the schedules have inherited the important information from their parents.

Bagchi et al., in [BAGC91], used a job shop scheduling problem to compare three different representations; the principal difference between these three representations being the amount of information encoded in the representation.

For the simplest representation, the original problem was reduced to a sequencing problem and the chromosomes were lists, i.e. permutations, of jobs. The schedule-builder had an important role to play in the search process. At the other end of

the spectrum, the third representation integrated all variables and hence, the entire search process was performed by the GA.

In this comparative study, the performance of the GA was directly related to the amount of knowledge encoded in the representation; the more information that was encoded the better was the performance.

In [BRUN93], the author also puts forward a strong case for the direct approach. A large real-world problem was used to compare one sequencing representation and one direct representation. Experiments were performed with different settings, i.e. different sizes of populations, initialisation heuristics, number of iterations, selection methods. The direct approach with its "knowledge-augmented" representation consistently outperformed the sequencing approach.

In conclusion, this section has presented two styles of representations, indirect and direct. Until recently, the indirect style was favoured because of the simplicity of its operators. However, direct approaches are now emerging. They do represent a promising alternative to the more traditional indirect approaches and merit further investigation.

3.2.6 GAs for Vehicle Routing.

This subsection reviews a number of applications which use GAs to solve VRPs. All of the systems presented in this section use an indirect representation; each of these systems is now presented in turn:

Thangiah et al.[THAN91] developed *Gideon* which addressed VRPs with time windows and capacity constraints; the objective of this system was to "minimise the number of vehicles and the distance travelled for servicing the set of customers without being tardy or exceeding the capacity or travel time of the vehicles".

In *Gideon*, a three-stage strategy was adopted. In a first step, the region covered by the different vehicles was divided into sectors; all the customers located in one sector were serviced by the same vehicle. Then, a heuristic was used to build the routes of the different engineers, i.e. given a cluster of jobs, the heuristic was used to find the best possible route among these customers. Finally, a local search (improvement) procedure was employed to refine the final answer.

Only the first step of the search involved a GA; each chromosome was a binary string which represented a list of seed angles used to divide customers into clusters. The objective of the GA was to find a set of seed angles minimising the total distance travelled by the engineers.

The authors reported two comparative studies of *Gideon* with other techniques. In [THAN91], *Gideon* was compared to two local procedures developed by Solomon [SOLO83]. Three different classes of VRPs were used:

- Class I: VRP where locations of customers were generated randomly,
- Class II: VRPs with clusters of customers,
- Class III: VRPs with semi-clusters of customers.

These problems also varied in terms of fleet size, vehicle capacity, travel times, time-windows, and customer service times. In 41 out of 56 problems, *Gideon* performed better than the local search procedures; the average reduction in terms of fleet size was 3.9% and in terms of distance travelled by the vehicles, the average reduction was 4.4%. In general, the improvements achieved by the GA were greater for problems in Classes I and III than for problems in Class II.

In [THAN93], a set of 25 large problems (200 customers for each problem) was used to compare the performance of *Gideon* with the performance of a greedy algorithm. Again, the geographical distribution of the customers was used as a parameter to divide the VRPs into 5 categories, each containing 5 problems.

In the first class, the customers were uniformly distributed across the entire area, while in the other four classes the 200 customers were regrouped respectively into 1,2,3 or 4 clusters. Within each class, the problems varied in terms of time deadlines, and varied from loose to tight constraints.

In this comparative study, three sets of experiments were performed. In the first set, the different problems were solved solely by using the greedy algorithm. In the second set, the greedy algorithm was augmented by a post-optimisation procedure and finally, *Gideon* was used in the third set of experiments.

In terms of solution quality, *Gideon* performed well for problems where the customers were distributed uniformly and/or with tight time deadline constraints. On the other hand, the association, greedy algorithm and post-optimisation, achieved good results for problems in which customers were tightly clustered and with loose time deadlines.

In terms of computational effort, the comparison is more delicate. In most cases, both versions of the greedy algorithm, i.e. with and without the post-optimisation, outperformed *Gideon*. However, in some cases, the greedy algorithm augmented by the post-optimisation procedure was extremely slow in providing an answer.

Blanton and Wainwright also used a GA to solve VRPs with time-windows and capacity constraints [BLAN93]. Their principal claim was that traditional operators usually associated with the sequencing representation (e.g. order-based, position-based, cycle crossovers) were not suitable for the solution of problems with multiple constraints. In order to overcome this limitation, a new class of crossovers was developed, namely the merge cross operators, which use global precedence, i.e. priority, information during the recombination process.

Basically, a global precedence arrangement indicates, for each pair of customers, the order in which the customers should be processed. Different criteria (e.g. time-windows of the different customers, distances or capacities) might be used to gen-

erate different arrangement and hence, different crossovers.

Twelve merge cross operators, using different arrangement criterion (or combination of criterion) were compared with three traditional crossovers, PMX, cycle crossover and the edge recombination operator, on four different problems with 15, 30, 75 and 99 customers respectively. In each case, the best solution was achieved by one of the merge cross operators. As a group, these operators outperformed the traditional operators. However, no merge cross operator was consistently better than the rest of its competitors. This raises some serious doubts concerning the robustness of these operators. In fact, the success or failure of a given merge cross operator appears to be greatly dependent on the specific characteristics of the problem in hand.

All the systems presented in this subsection used an indirect representation. No GA with a direct representation has yet been implemented to solve VRPs. The comparison in Section 3.2.5, where systems developed by Bagchi et al. and Bruns outperformed their competitors, motivates the investigation of this class of representation for the solution of VRPs.

3.2.7 GAs for telecommunications.

This short section presents three recent applications of GAs in the field of telecommunications.

Many features of typical telecommunications problems motivates a GA solution. Indeed, a large number of telecommunications applications have large and complex search spaces and/or have severe time-constraints. In both cases, conventional techniques might fail to provide an answer in reasonable time, and non-deterministic search techniques such as GAs may be well-suited for the solution of these problems.

Clitherow and Fisher, from Bellcore, developed a GA to solve small design problems

(from 5 to 30 nodes) [CLIT89]. In fact, two GAs were implemented. The first GA used conventional operators whereas the second was equipped with problem-specific operators. These problem-specific operators were two mutation mechanisms whose actions were dictated by a rule base, a set of heuristics used usually by human experts to design a network manually.

In both GAs, each element of the chromosome represented the capacity of a particular link between two nodes, i.e. there was one element for each link of the network.

The GAs were given a smart start; the initial population was not purely random. For each element of the chromosome, the probability distribution used to define the initial capacity was centred around an approximate solution.

The principal results of the study were that:

1. The GA equipped with the two problem-specific mutation operators consistently outperformed the conventional GA,
2. The number of iterations required for convergence increased as the network grew in size; hardly a surprising outcome !

Cox proposed an algorithm for the dynamic anticipatory routing (DAR) of calls in circuit-switched telecommunications networks [COX91]; this scheduling problem was concerned with the capacity reservation for services such as video broadcasting where demand can be planned in advance.

The elements of the problem were:

1. The configuration of the network, e.g. number of nodes, links and capacities,
2. The traffic pattern which was described in probabilistic terms,

3. A call table describing the individual calls requested (e.g. source node, destination node, desired start time, duration)
4. A loss function defining a penalty cost for each call type.

The objective was to find a dynamic routing policy minimising the average loss per unit of time resulting from blocked calls.

The DAR algorithm incorporated a two-fold procedure:

1. In the first part, the algorithm was using the traffic pattern only to solve the problem in a probabilistic manner and provide a path assignment rule to each call. This path assignment rule was, in fact, a small set of path selection probabilities mapping call i into path j with probability $p(i, j)$. The objective was to minimize the expected penalty cost, generated by blocked calls, per unit time.
2. The objective of the second part, called the bandwidth packing algorithm, was to improve the initial assignments made in the first part, by using information relevant to the current situation: (a) the current state of the network and (b) a list of call requests with their characteristics.

Four different algorithms, namely a greedy algorithm, a pairwise exchange heuristic, uniform random search, and a GA were implemented to perform the bandwidth-packing task. The performance of the four algorithms was compared using a 50 node test problem. All perform reasonably well. The worst one, uniform random search, was only 8% away from the optimal solution. In comparison to its competitors, the GA performed well, but was outperformed by a hybrid algorithm: GA followed by the pairwise exchange procedure.

Finally, Liepins and Potter reported a GA approach for diagnosing multiple faults in a microwave communications system [LIEP91]. The authors defined multiple

diagnosis as the identification of a set of faults that best corresponds to a set of alarms. First, the original diagnosis problem was mapped into a set covering problem and the GA was used to find the optimal answer to this constrained problem.

In conclusion, two important observations can be made from the reviews presented here:

1. The performances achieved by the different GAs, presented in the three previous sections, compared generally favourably with those achieved by their competitors.
2. Subsection 3.2.5, "*GAs for resource allocation*", has highlighted the reason why a direct representation is an appropriate choice. This style of representation should be investigated for the solution of VRPs.

3.3 Simulated Annealing.

3.3.1 Introduction.

SA is an iterative and stochastic search procedure which was first introduced by Kirkpatrick et al. [KIRK83] and Cerny [CERN85]. The principal characteristic of SA is that it provides a mechanism which (potentially) prevents a local search algorithm from being trapped in a local minimum.

Local search algorithms are often used to solve large combinatorial problems for which exact algorithms do not exist or cannot provide an answer in reasonable time.

Usually, a local search algorithm works as follows: the algorithm starts with an initial solution which might be either chosen at random or provided by another search procedure.

A rearrangement operator is used to generate a neighbour⁸ of the current solution; a neighbour can be defined as a similar solution which differs (from the original solution) only by a single change which is performed by the rearrangement operator.

Then, the difference of quality between the current solution and its newly generated neighbour is calculated. If the neighbour is of better quality, it becomes the new current solution. If not, the original current solution is kept. Thus, a local search algorithm can be viewed as a descent algorithm. Only moves which decrease the cost, i.e. downhill moves, are accepted. Moves which increase the cost, i.e. uphill moves, are always rejected.

These two steps of neighbour generation and comparison are repeated until no improvement can be detected in the neighbourhood of the current solution. Thus, this final solution is guaranteed to be a local minimum.

However, this local minimum might be distant from the optimal solution. In fact, the principal limitation of local search algorithms is that they do not provide a mechanism which would allow an escape from local minima. If the algorithm has reached a local minimum, the search can go no further. There is no exception to this rule, even if the local minimum found is of poor quality in global terms.

A simple answer to this problem is to repeat the local search from different starting points and keep the best solution found. In this case, different parts of the search space are explored and, intuitively, the algorithm is more likely to find a final solution of better quality. However, this naive approach does not always overcome the underlying problem.

SA is an alternative approach. Basically, this procedure accepts uphill moves with a controlled probability and is therefore less likely than a local search algorithm to be trapped in a poor local minimum.

⁸and therefore, the neighbourhood is the set of all the neighbours.

SA has been used to tackle some classical combinatorial problems, e.g. traveling salesman [GOLD86], graph partitioning [JOHN89b], graph colouring problems [JOHN91] and has given some mixed results.

The remainder of this section consists of five subsections. Subsection 3.3.2 introduces SA and describes its internal algorithm. Subsection 3.3.3 examines the role played by the different parameters. Subsection 3.3.4 reports the results of some comparative studies of SA with conventional techniques when applied to different classes of combinatorial problems.

3.3.2 SA: A brief description.

SA originated from an analogy drawn between thermodynamics and combinatorial optimisation. In this context, the search for an optimal solution in a combinatorial optimisation problem is analogous to the production of a crystal using an annealing technique.

The annealing technique works as follows: first, a solid is melted at a high temperature and then, the temperature is slowly and gradually lowered with most of the time spent at temperatures near freezing.

At high temperatures, the atoms of the system can move about freely. In the ground states, the particles are arranged in a highly structured manner and the energy of the system is minimal; the annealing technique has produced a crystal.

These ground states and this pure crystal can only be obtained if the temperature is reduced in slow stages. Otherwise, if the cooling process is too quick⁹, the system will not reach its minimum energy state and the crystal will have many defects and only locally optimal structures. Basically, the period of time at each temperature

⁹The process in which the temperature is instantaneously lowered is defined as quenching. In the context of simulated annealing, local search is analogous to quenching.

must be sufficiently long to reach equilibrium.

The Metropolis algorithm [METR53] developed for the simulation of a collection of atoms in equilibrium at a given temperature is the central part of a SA algorithm. This algorithm works as follows:

- An atom is given a random displacement and the resulting change in the energy level of the system, ΔE , is calculated.
- If ΔE is negative, the newly generated configuration is accepted as the new configuration.
- If ΔE is positive, the case is treated in a probabilistic manner. The probability of acceptance is given by:

$$P(\Delta E) = \exp(-\Delta E / K_b T)$$

where K_b is the Boltzmann constant and T is the temperature.

- If $P(\Delta E)$ is greater than a random number between 0 and 1 then the newly generated configuration is accepted, otherwise it is rejected and the original configuration remains the current configuration.

In summary, the Metropolis algorithm allows both uphill and downhill moves to occur. However, there is an important difference: a downhill move is always accepted, whereas an uphill move is only accepted with a certain probability which decreases with the temperature, i.e. the lower the temperature, the less likely an uphill move will occur.

The SA algorithm is illustrated in pseudo-code in Figure 3.9. The inner loop of this algorithm is a Metropolis algorithm and the outer loop determines the changes in temperature. A SA algorithm works as follows:


```
1  S := initial-solution;
2  T := initial-temperature;
3  WHILE (not frozen)
4  DO BEGIN
5      WHILE (not equilibrium)
6      DO BEGIN
7          S' := nearby-solution(S);
8          Δ := change-in-cost(S, S');
9          IF (Δ > 0)
10         THEN accept S' with probability  $\exp(-\Delta/T)$  as S
11         ELSE accept S' as S
12         END
13     T := reduce-temperature(T);
14     END
15 END;
```

Figure 3.9: Implementation of simulated annealing.

First, an initial solution and an initial temperature are chosen; this initial temperature is set high, so that the term $\exp(-\Delta/T)$ ¹⁰ tends to be large and large uphill moves are allowed.

The SA algorithm attempts a certain number of moves at each temperature and the Metropolis algorithm is used to define the validity of these moves. At each temperature, the algorithm must proceed long enough for the solution to reach a state of equilibrium.

The temperature is gradually dropped, and as T decreases, the term $\exp(-\Delta/T)$ decreases, minimising uphill changes to the system. Eventually, as T tends to zero, $\exp(-\Delta/T)$ also tends to zero, and no uphill moves can take place, i.e. freezing occurs.

¹⁰Note that T is a control parameter using the same units than the cost function and therefore, the Boltzmann constant K_b is not required anymore.

3.3.3 The driving forces.

This subsection analyses the different parameters which guide and influence the search performed by a SA algorithm. These parameters are:

1. A neighbourhood and a set of operators which define (1) the search space the algorithm has to explore and (2) how a neighbour solution can be generated.
2. An initial temperature selected so that it allows most uphill moves.
3. A temperature function which defines how the temperature is reduced during the search. The rate of reduction in T affects the granularity of the schedule. If T decreases too fast then quenching would occur and it would generate a non-optimal solution (i.e. a glass instead of the desired crystal).
4. The number of iterations to be performed at each temperature.

Neighbourhood/operators.

The design of a suitable neighbourhood is critical to the success or failure of SA. Eglese [EGLE90] claims that a search space with a "smooth" topology where the local optima are shallow should be preferred to a search space with a "bumpy" topology where there are many deep local minima. Eglese, in the same paper, also covered constrained problems. The two traditional techniques for solving such problems were proposed: (1) restrict the search space to solutions which conform to all the constraints or (2) allow solutions which do not satisfy all constraints but which are then given a penalty. The author claimed that the second technique is likely to generate a simpler search space with a smoother topology.

The annealing schedule.

The initial temperature, the temperature function and the number of iterations performed at each temperature define the annealing schedule which is also called the cooling schedule. Eglese reported several classes of annealing schedules [EGLE90]:

1. The original annealing schedules [KIRK83]. The initial temperature is set high enough to allow most uphill moves. A proportional temperature function (i.e. $T(t+1) = \alpha T(t)$) is implemented and the decrements becomes smaller and smaller as t increases so that an important part of the search is performed at low temperatures. At a given temperature, equilibrium is reached and the algorithm moves to a lower temperature when a sufficient number of transitions have been accepted, subject to a constant upper bound. This original annealing schedule can be modified in several ways: for instance, the number of iterations performed at each temperature might be kept constant, or be proportional to the size of the problem in hand, and so on.
2. Hajek [HAJE88] proposed an important theoretical result, stating that, in order to reach an optimal solution, there was no requirement to attain equilibrium at a succession of reducing temperatures. Rather, the main condition was for the cooling to be carried out sufficiently slowly. This result meant that there is a balance between a large reduction in temperature and a small number of iterations at a fixed temperature. It also generated a spectrum of alternatives for an annealing schedule, with the two extremes being:
 - The annealing schedule proposed by Lundy and Mees [LUND86] where only one iteration is performed at each temperature and the following temperature function is used:

$$T(t+1) = T(t)/(1 + BT(t))$$

where B is a constant. This function provides a slower temperature reduction than the proportional function used by Kirkpatrick et al. [KIRK83].

- The annealing schedule proposed by Connolly [CONN88] where the temperature is kept constant during the search. This can be viewed as a random search which accepts uphill moves with the probability $\exp(-C\Delta)$ and C is kept constant during the entire search process. Connolly used this method for the solution of quadratic assignment problems¹¹ (QAP). This work achieved results which outperformed the results achieved by the annealing schedule proposed by Kirkpatrick et al. [KIRK83]. The author also claimed that, for this class of problems, a sequential generation of neighbours was superior to a random generation.

3.3.4 Comparative studies.

This subsection reports the results of several papers which compared the performances of SA with the performance of other search procedures on different classes of combinatorial problems. For each paper, this study (1) presents the class of problems being examined and (2) examines and compares the results achieved by the SA algorithm and its competitors.

Golden and Skiscim used SA to solve routing and location problems, in particular travelling salesman problems (TSPs) and p-median location problems¹² [GOLD86].

A set of TSPs of varying sizes (between 25 and 318 cities) were used to compare the results achieved by SA¹³ with those of:

¹¹The quadratic assignment problem can be defined as the assignment of inter-communicating objects to locations such that the total cost of communicating is minimised.

¹²The p-median problem is to locate p facilities at various nodes of a network in such a way that the sum of the distances from each node in the network to its nearest facility is minimised.

¹³In these tests, the SA algorithm implemented used a 2-opt procedure as its main operator.

1. A specialised heuristic for the TSP; the CCAO procedure. The authors described the CCAO heuristic as a *"hybrid procedure that use the convex hull of points as the starting sub-tour and insert nodes using a combination of the greatest angle method and the cheapest insertion criteria."* Also, a branch exchange heuristic was used as a post-optimisation procedure.
2. The classic 2-opt procedure of Lin and Kernighan [LIN73].

SA was consistently outperformed by the CCAO procedure, both in terms of quality and running time, and could only provide better results than the 2-opt procedure on a single problem.

For the p-median problems, SA was compared with a local search procedure, specifically designed for this class of problems. In general, this procedure required considerably less computational effort than SA and provided results of comparable quality.

Golden and Skiscim also studied the relationship between problem size and running time and concluded that, in practice, SA required a computational effort which grows faster than cubically.

Johnson et al. used SA for the solution of graph partitioning problems and compared its performance with standard local search algorithms [JOHN89b]. The graph partitioning problem is the problem of partitioning the vertices of a graph into two equal size sets in order to minimise the number of edges with end points in both sets. The authors used two classes of graphs for their comparison: (1) random and (2) geometric graphs.

For random graphs, SA outperformed its competitors both in terms of quality and running time, particularly, when the density and/or the size of the graph increased. However, for geometric graphs, the situation was reversed and local search procedures performed better than SA. By way of summary, the authors

offered a few general observations from their extensive series of tests and it appeared that:

- SA required long running times,
- the best way to add time was to add it uniformly at each temperature,
- there was no gain spending much time at high temperatures,
- the temperature function provided by Kirkpatrick et al. [KIRK83] was as good as other techniques, e.g. logarithmic cooling, linear cooling,
- there may be an advantage in starting with a good solution rather than a randomly generated one,
- replacing $\exp(-\Delta/T)$ by an approximation or a look-up table could provide an important speed-up without degrading the quality of the solutions achieved.

Finally, Johnson et al. also examined the performances of SA for the solution of graph colouring and number partitioning problems [JOHN91]. Given a graph and n colours, the graph colouring problem is to assign a colour to each node of the graph such that no pair of nodes connected by a link have the same colour. In the sequence partitioning problem a sequence of real numbers a_1, a_1, \dots, a_n in the interval $[0, 1]$ must be partitioned into two sets A_1 and A_2 such that:

$$\left| \sum_{a_i \in A_1} a_i - \sum_{a_i \in A_2} a_i \right|$$

is minimised.

For graph colouring problems, three implementations of SA, with different internal representations, were tested and in general, achieved good final solutions at the expense of long calculation times. The authors viewed number partitioning problems as a class of deceptive problems for local search heuristic, and hence, SA. The tests

confirmed their view and SA was outperformed by its competitors both in terms of solution quality and running time.

3.4 Tabu Search.

3.4.1 Introduction.

Tabu Search ¹⁴ (TS) is a neighbourhood search technique which was first developed by Glover [GLOV86] [GLOV88] [GLOV89b]. Like SA, TS can escape from a local optimum by accepting uphill moves but in contrast with SA (and GA for that matter), most implementations of TS are wholly deterministic; the major exception being the probabilistic Tabu Search proposed by Faigle and Kern [FAIG92] which selects moves based on probabilities. The term *tabu* refers to the fact that the algorithm prevents the occurrence of some moves during the search; these moves are *tabu*. The objectives of TS can be stated as follows:

- To avoid being trapped in a local optimum.
- To avoid cycling, i.e. avoid wasting computational effort by re-visiting points in the search space.
- To maintain a correct balance between intensification and diversification, i.e. balancing between focusing the search effort in a specific (promising) area of the search space and encouraging the exploration of new regions.

In order to achieve these aims, the algorithm simulates a number of memory modules with different time spans (short-, intermediate- and long-term). These modules

¹⁴This technique was originally developed in North America. This explains why it is called *tabu search* instead of *taboo search*.

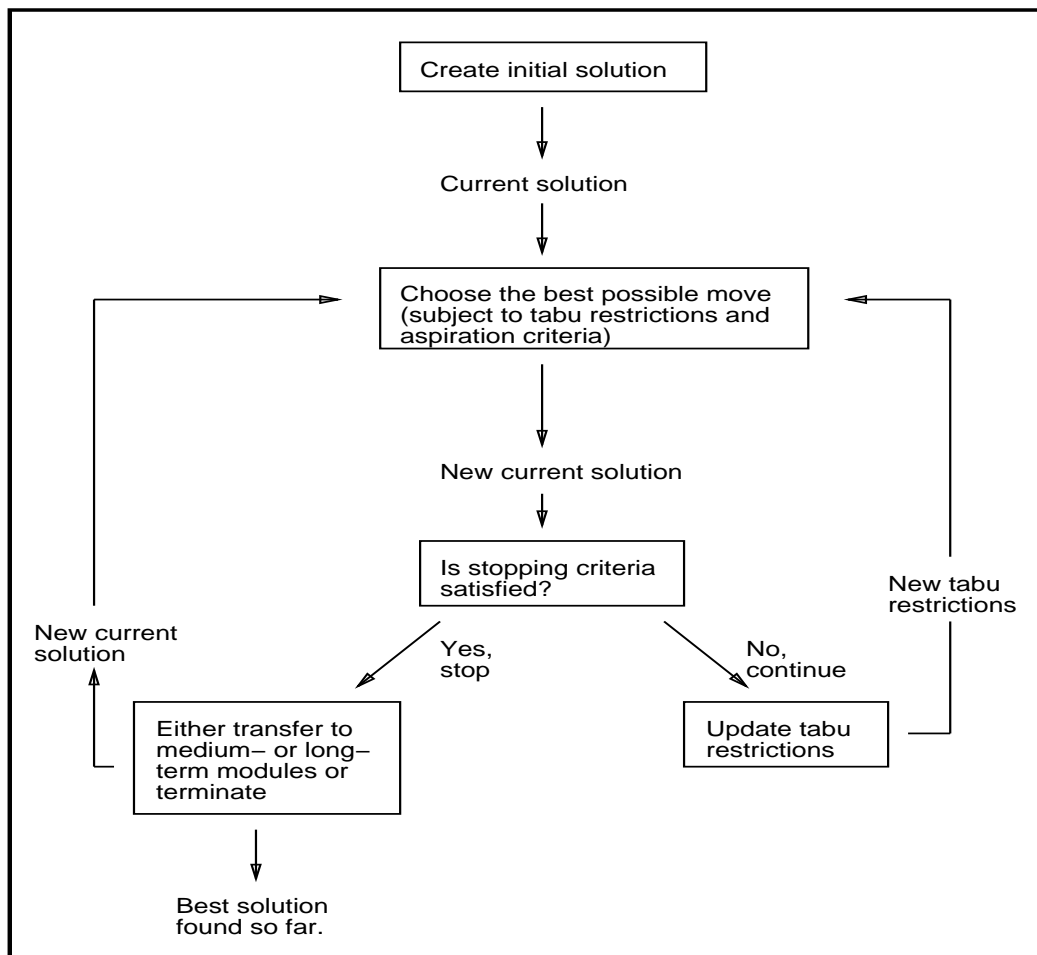


Figure 3.10: Tabu search - Short-term memory module.

constrain the search in different ways, and their interactions provide the means for answering the above objectives.

The remainder of this section consists of 4 subsections. In Subsection 3.4.2, the central part of TS, i.e. the short-term memory module, is described step by step. Subsection 3.4.3 examines the roles played by the different driving forces of TS while Subsection 3.4.4 reviews a number of applications.

3.4.2 TS: a brief description.

TS can be implemented in many ways and this paragraph describes the central part of most implementations (see Figure 3.10); this module corresponds to the short-term memory of TS. Given an initial solution S , the algorithm explores the entire neighbourhood $N(S)$ of S and consider all possible moves in both uphill and downhill direction. All moves are considered except for a certain number of tabu moves (tabu moves which meet the aspiration criteria ¹⁵ are also considered). The algorithm then continues its search by moving to S' (and S' becomes S); S' being the solution with the highest evaluation in the neighbourhood even if it implies an uphill move. With this mechanism ¹⁶, TS avoids being trapped; it does not terminate when it reaches a local optimum.

This cycle "exploration, selection" is repeated until a certain termination criterion is satisfied. At this point and depending on the implementation, either the algorithm terminates or the solution is passed on to the medium- and long-term memory modules.

3.4.3 The driving forces.

Tabu Search is composed of three modules:

- a short-term memory module which, at each iteration, seeks to make the best possible move. This module is constrained by the tabu restrictions and the aspiration criteria (this module was described in some detail in the previous section).
- a medium-term memory module which intensifies the search.

¹⁵Basically, a move is admissible either if it is not tabu or if its tabu status can be overridden by the aspiration criteria.

¹⁶With this mechanism, at the end of the search, the current solution S might not be the best solution encountered during the search. In order to overcome this problem, TS keeps a record of its best solution.

- a long-term memory module which diversifies the search.

Each driving force, namely the tabu restrictions, the aspiration criteria, the medium-term memory and the long-term memory modules, are now considered in turn:

Tabu restrictions.

In order to avoid being trapped in a local optimum, TS allows non-improving moves. However, this means that the algorithm could cycle around the same solutions¹⁷ unnecessarily. The aim of the tabu restrictions is to prevent this cycling behaviour. The tabu restrictions can be implemented as follows: each time a move is made, its "inverse"¹⁸ is added to a FIFO (First In - First Out) list called the tabu list. This list has a maximum length of T (called the tabu tenure). This limit means that a move only remains tabu for a limited amount of time¹⁹.

In order to avoid cycling, the ideal situation would be to record each solution already visited during the past search and to forbid the algorithm to re-visit these points. Unfortunately, this approach is not acceptable for large-scale applications. For such applications, the storage of all previous solutions would require a prohibitive amount of memory and the cost of comparing the current solution against (potentially) all previous solutions would also be prohibitive. This explains why TS only records attributes of a limited number of past solutions.

It should be noted that an algorithm may have several tabu lists, with each list dealing with a specific set of attributes. In this case, Glover recommends that each tabu tenure (i.e. lengths of the tabu lists) should be different. His point is that,

¹⁷For instance, the following situation may arise, the algorithm takes a non-improving move to escape from a local optimum, only to return to that point in the very next move because the algorithm is designed so that it always choose the best possible move.

¹⁸Tabu restrictions are also often used to prevent repetitions as well as reversals.

¹⁹In many applications, T is set to 7. In this case, once a move is added to the tabu list, it remains tabu for 7 iterations.

some attributes might contribute more strongly to a tabu restriction than others and, hence they should be given a briefer tabu tenure in order to avoid making the restrictions too severe [GLOV89a].

This "attribute approach" is not without problems. An attribute is not specific to one single solution. Therefore, once an attribute is pushed into the tabu list, it prevents the algorithm from visiting a number of solutions rather than masking one single point. Some of these solutions have never been visited before and are therefore perfectly valid; the algorithm should be able to visit these points. In this sense, it can be said that the tabu restrictions over-constrain the search. These non-desirable effects of the tabu restrictions are counter-balanced by the aspiration criteria.

The aspiration criteria.

The aspiration criteria is used to override the tabu restrictions. As these restrictions over-constrain the search, it is indeed necessary for the algorithm to have a mechanism which, on specific occasions, relaxes these constraints and allows a tabu move to take place. For instance, an algorithm may decide that if a tabu move results in a solution better than any solution visited so far, the tabu restrictions should be relaxed and the tabu move should be accepted [DEWE89]. In [LAGU91], the authors have observed that, in a run of 1000 iterations, an average of 7 tabu moves are executed as a result of meeting the aspiration criteria.

The medium-term memory module.

This module is used within TS "*to achieve regional intensification*" [GLOV89a]. In other words, this module tries to detect features such as values received by particular variables which are common to above-average solutions. In order to achieve this, the best solutions found by the short-term memory module are recorded and

compared. For instance, in a traveling salesman problem, the algorithm might then notice that some edges are often part of good solutions. Once this learning stage is completed, the algorithm uses this new knowledge to focus the search and tries to generate solutions that exhibit these features. This can be achieved by restricting the set of possible moves during the *period of regional search intensification* [GLOV89a].

The long-term memory module.

This module is used within TS *"to achieve global diversification"* [GLOV89a]. While the previous module tries to intensify the search in some promising regions of the search space, this module tries to direct the search towards regions which have not been explored yet; this technique gives a global view of the search space to the algorithm.

In this sense, TS is similar to GAs. Indeed, with their pools of solutions, GAs also have this global view. However, while GAs give a random sampling of the search space, TS (more precisely the long-term memory) tries to generate solutions which are definitely different from the past solutions. Glover describes the objective of this long-term memory as follows:

"The objective is to create evaluation criteria that can be used by a heuristic search process which is specifically designed to produce a new starting point, thus generating such a point by purposeful instead of random means. These evaluation criteria penalize the features that long-term memory finds to be prevalent in previous executions of the search process." [GLOV89a].

3.4.4 Applications of TS.

Compared with GA and SA, introduced in 1975 and 1983 respectively, TS is still in its infancy and there are few applications of TS available. The same can be said for detailed or comparative studies of TS. Therefore, this subsection only reviews a limited number of applications.

In [MALE89], the authors performed two sets of experiments. In the first one, SA and TS are compared over a number of travelling salesman problems (up to 100 cities) and TS consistently outperforms SA. In the second set of experiments, the authors implemented parallel versions of TS and SA; in both cases, the overall system consists of a number of agents and each agent is solving the entire problem. In parallel TS, each agent performs a tabu search with different parameters, i.e. different attributes and aspiration criteria. Periodically, the agents are stopped and then the best solution is communicated to all the agents and the agents re-start from this new initial solution. In parallel SA, all the agents have a different annealing schedule. The main observation in this set of experiments is that both parallel TS and SA achieve a super-linear speed-up when compared with the original SA and TS. Since then, a similar phenomenon was reported in [CLEA91] where it is claimed that this super-linear speed-up was due to the fact that the different agents use the extra information (gained from cooperation with the other agents) to guide and focus the search in the most promising regions of the search space.

Perry also implemented a parallel TS for the solution of a time-constrained scheduling problem ²⁰ [PERR90]. The problem is the allocation of surface-to-air-missiles to counter a number of threats. TS was used mainly because it is an iterative approach and hence, a solution can be provided (if needed) at any time. An initial solution is provided very rapidly, then, TS is used to refine this original solution.

²⁰In this case, the main point of using a parallel algorithm is to increase the number of feasible schedules that can be generated in a given period of time.

Unfortunately, the authors provides few details concerning its application.

In [ADEN92], the authors compared the performance of TS against the performances of some greedy heuristics over 480 flow shop problems of different sizes. TS dominates its opponents in most cases. In [LAGU91], TS is used to solve a single machine scheduling problem. This problem is a sequencing problem and the objective is to minimise the sum of the set-up costs and delay penalties of the schedule. In this paper, the authors compare two moves operator for TS: swap and insert. The different algorithms are compared over some (relatively) small problems: 20 jobs, 35 jobs problems where the optimal solution can be found in less than 30 seconds. This study shows that insert is better than swap. Finally, a hybrid TS allowing both swap and insert moves was implemented. A second set of experiments showed that this hybrid version performed better than any one of the two original algorithms.

3.5 Brief comparison of GA, SA and TS.

GA, SA and TS originate from three different metaphors. GA replicates the principles of the theory of Evolution, i.e. selection and inheritance. SA is drawn from an analogy between thermodynamics and combinatorial problems and TS simulates three different types of memory (short-, medium-, and long-term).

Table 3.4 presents the key-features and the driving forces of these three techniques. First, they all are iterative and do not construct their final solution piece by piece. Rather, they start their search with an initial solution (or an initial population for GA) and gradually, iteration after iteration, improve the quality of that solution. Both GA and SA are stochastic and TS is, in most implementations, deterministic. SA is a neighbourhood search technique; it works with only one solution and at

	GA	SA	TS
Type	global stochastic iterative	neighbourhood stochastic iterative	neighbourhood (and possibly global) deterministic iterative
Solution-candidate(s)	pool	one	one
Driving forces	initial population, crossover, mutation, selection function.	equilibrium, cooling rate, initial temperature.	tabu restrictions, aspiration criteria, medium-term memory, long-term memory.

Table 3.4: GA, SA, and TS: A comparison.

each iteration, this solution is only slightly altered. On the other hand, GA works with a pool of solutions and what happens to a particular chromosome is of little importance. This global view of the search space gives GA the opportunity to detect similarities among above-average solutions and then, GA can favour the propagation of these similarities during the search process.

When comparing SA and GA, it could be said that GA uses its population to *leap* from solution to solution in the search space whereas SA only goes from one solution to its neighbour. Our results in Chapter 6 show that these two algorithms have totally different computational behaviour.

Primarily, TS is a neighbourhood search technique; like SA, it uses a move operator to jump from neighbour to neighbour. However, the long-term memory module ²¹ provides TS with a global view of the search space. In contrast to GA, this *global view* is not provided by a random sampling of the search space; rather TS tries to generate (using a deterministic heuristic) new starting points which are radically different from past ones.

In conclusion, GA, SA, and TS are three AI search techniques which can be used

²¹Observe that this module is absent from most TS implementations.

to solve large combinatorial problems. They use their intelligence to explore large search spaces and find optimal or, more realistically, near-optimal solutions.

Chapter 4

Preliminary Studies.

This chapter describes two preliminary studies. First, a GA using an indirect representation is used to solve a simple scheduling problem. Then, a similar indirect GA is applied to a particular instance of the VRP. These two studies highlight the limitations of indirect approaches and advocate the adoption of a direct representation.

4.1 A simple prototype.

This prototype was used initially as an introduction to GAs. The objective of the GA was to find a schedule minimising the total tardiness ¹ of a job set.

operations	A	B	C	D	E	F	G	H	I	J
duration	10	10	10	10	40	20	10	40	25	25
domain	20-50	20-40	30-40	20-99	50-200	80-200	100-200	0-200	0-150	0-150

Table 4.1: Description of a job set.

¹Tardiness is the amount of time between the moment when a task should have been completed and the moment when the task is actually completed.



Figure 4.1: Current domains and durations of the operations.

Table 4.1 details the operations which make up the job set. The domain of an operation represents the set of possible values for the start-time of the operation, i.e. the time-window when the operation can be started. The different operations form a job set which must be performed on a single resource and this resource can only handle one operation at a time.

Figure 4.1 displays a Gantt chart which gives a graphical representation of this problem. The domains of the different operations are represented by a large horizontal white box. Within this large white box, a smaller box represents the duration of the operation.

The GA uses a steady-state model (described in section 3.2.2) and the representation employed for the chromosome is as follows:

$$o_1, o_2, \dots, o_N$$

This is a simple order-based representation; o_i refers to the i th operation to be

instantiated. Each chromosome of the initial population is created randomly and represents one possible ordering of the job set.




Figure 4.2: One optimal solution found by the genetic algorithm.

A schedule-builder uses the chromosome to generate a schedule; the order imposed by the chromosome is strictly followed and the schedule-builder assigns the earliest possible start-time to each operation. Each start-time has to be compatible with the previous decisions. If a start time cannot be found in the domain of the operation, the schedule-builder relaxes the problem and allocates a late start-time which is set to be as early as possible.

The different chromosomes are ranked ² with respect to the total tardiness of their schedules and the probability of a chromosome being selected for reproduction is directly proportional to its rank in the population. The search is pursued until either an optimal solution, i.e. a solution with no tardiness, has been found or the maximum number of iterations has been reached.

²The best chromosome is the one which has the smallest tardiness.

The following table presents one optimal solution found by the GA with a population of 20 chromosomes and 500 iterations:

operations	A	B	C	D	E	F	G	H	I	J
start-time	50	40	30	60	175	80	165	100	0	140

Table 4.2: Optimal solution found by the genetic algorithm.

This optimal solution is also displayed in Figure 4.2; the solid black bars show the scheduling decisions, the length of the bar corresponds to the duration of the operation and its position represents the start-time of the operation.

During this project, an object-oriented approach has been used with Sun Common Lisp 4.0, CLOS (Common Lisp Object System), SPE (Symbolic Programming Environment 1.2) and Lispview 1.0 running on a SPARCstation IPC under Unix. Figures 4.3 and 4.4 show object representations of both a GA and a chromosome.

```

SPE: Listener<2>
Help Edit Text Search Show Stop Package File Module

#<Genetic-Agent #X27F5E76>
  is an instance of the class GENETIC-AGENT:
The following slots have allocation : INSTANCE:
NAME          "agent1"
TYPE          GENETIC_ALGORITHM
OBJECTIVE     NIL
STATUS       ACTIVE
FITTEST      #<Chromosome #X27F624E>
SIZE        20
*CHROMO*    (#<Chromosome #X2A83AEE> #<Chromosome #X281CCEE> #<Chrom\
osome #X2A99F9E> #<Chromosome #X282A28E> #<Chromosome #X2A9A08E> #<Chromosome #X\
2A9A0CE> #<Chromosome #X282F4CE> #<Chromosome #X281079E> #<Chromosome #X2A99EFE>\
#<Chromosome #X2A83956> ...)
INIT_POPULATION_FCT  INIT
SCHEDULING_FCT       SCHEDULING
EVALUATION_FCT       EVALUATION
SELECTION_FCT        SELECTION
CROSSOVER_FCT        CROSSOVER
MUTATION_FCT         NIL
INVERSION_FCT        NIL
DATA_COLLECTION_FCT  NIL

```

Figure 4.3: Object representation of a genetic algorithm.

Each instance of the class **genetic-agent** (see Figure 4.3) has the following slots:

name: unique name given by the user, when the GA is created, for manipulation purposes.

type: type of the object, i.e. genetic algorithm.

objective: documentation string describing the objective of the GA, e.g. *minimise total tardiness*.

status: status of the process associated with the object (i.e. init, active, inactive or terminated).

fittest: fittest element of the population.

size: size of the population, i.e. number of chromosomes in the population.

***chromo*:** list of chromosomes present in the population.

init_population_fct: function called by the process to create the population.

scheduling_fct: function called by the process to translate, i.e. map, order-based chromosomes into schedules.

evaluation_fct: function called by the process to evaluate the schedules.

selection_fct: function called by the process to select two parents from the population.

crossover_fct: function called by the process to perform the crossover operation.

mutation_fct: function called by the process to perform the mutation operation.

inversion_fct: function called by the process to perform the inversion operation.

data_collection_fct: function called by the process to collect data during the search.

```

SPE: Listener<2>
Help Edit Text Search Show Stop Package File Module
;;; SPE: Lisp Listener
> (describe (first (*chromo* (first *ga*))))

#<Chromosome #X247219E>
  is an instance of the class CHROMOSOME:
  The following slots have allocation :INSTANCE:
  G-STRING      ((("i" 0) ("c" 30) ("g" 165) ("j" 175) ("b" 40) ("f" 80) ("a" 50\
) ("e" 200) ("d" 60) ("h" 100))
  OBJECTIVE     NIL
  TARDINESS     60
  FITNESS       0
  NB_JOBS_LATE  2
  HISTORY       ("agent1")
  CHILDREN      0
  BIRTH         88
> █

```

Figure 4.4: Object representation of a chromosome.

Each instance of the class chromosome (see Figure 4.4) has the following slots:

g-string: An order-based representation of the problem. Each couple of this string represents an operation and its allocated start-time.

tardiness: total tardiness of the jobs set.

fitness: fitness of the chromosome.

nb_jobs_late: number of operations that the scheduling function has not scheduled properly, i.e. number of operations with a late start time.

children: number of children of the chromosome.

birth: date of birth of the chromosome, i.e. number of cycles between the creation of the population (i.e. date 0) and the birth of the chromosome.

4.2 A particular VRP.

In this project, a GA is used to find near optimal solutions to one particular instance of the VRP. This VRP has the following parameters:

1. a workforce of 118 engineers operating from 11 bases,
2. a set of 250 jobs geographically distributed,
3. the objective of performing all the jobs while minimising the total distance travelled. The following cost function is used:

$$Cost = Travel + \sum_{i \in Conflicts} (Duration_i + 60)$$

where the first term represents the total distance travelled by the different engineers and the second term represents a penalty related to the conflicting jobs; for each job which has not been assigned to an engineer, the cost is increased by the duration of the job and a penalty of 60 points.

This initial problem has been further complicated by a number of side-constraints:

- Technological constraints (I): each job can only be performed by a sub-set of the workforce. On average, an engineer is only qualified to perform 57 jobs out of the 250 possible jobs. A significant number of jobs can only be performed by a few engineer. This may create some *bottleneck* situations.
- Technological constraints (II): the duration of a job depends on the engineer allocated to perform that job. The difference can be as large as 60%, say between an expert and an apprentice.
- Time windows (i.e, temporal constraints): a job can only be performed during a certain period of time. There are 3 classes of time-windows: all day, morning and afternoon.




Figure 4.5: A particular VRP.

Figure 4.5 gives a graphical representation of this problem. The squares represent the 11 bases, where the 118 engineers start and finish their tours, and the circles represents the 250 jobs which have to be performed. All these elements are located within an area of 10 miles by 10 miles. Figure 3.1 in Section 3.1 displays a partial solution to this problem.

The GA works with an indirect representation of the problem; hence, it is used in association with a schedule-builder. As seen in Section 3.2.5 and in the previous prototype, the schedule-builder uses the chromosome as a decision sequence; it follows the order of the chromosome and makes a decision for every element of the chromosome. In fact, two indirect representations have been investigated in this

work:

1. The GA has a job-centred perspective. Each element of a chromosome represents a different job. The schedule-builder follows the ordered list and attempts to attach each job to an engineer. This strategy is similar to the operation-centred perspective adopted by ISIS [FOX84].
2. The GA has an engineer-centred perspective: each element of a chromosome represents a different engineer. The schedule-builder attempts to build a tour for each engineer successively. This strategy is similar to the resource-centred perspective adopted by OPIS [SMIT90].

The schedule-builder works as follows. First, it assumes that all the jobs are conflicting (as soon as a job is allocated to an engineer, it is removed from the set of conflicting jobs). The algorithm attempts to attach the current engineer to all the jobs he/she can do and which are still in the conflict list, i.e. jobs which have not been allocated yet. Then, the algorithm moves to the next engineer in the chromosome.

The two strategies are illustrated in Figure 4.6. Two GAs have been implemented: one with a job-centred perspective and one with an engineer-centred perspective. Both GAs have the following set-up:

- population: 100 chromosomes.
- crossover: The PMX crossover operator (described in Section 3.2.5).
- mutation: The mutation operator works with a single chromosome; two elements of the chromosome swap position. This operation is performed after the crossover operation, on 4% of the children.
- inversion: The inversion operator also works with a single chromosome; it inverts the order of the elements between two randomly-chosen points of the

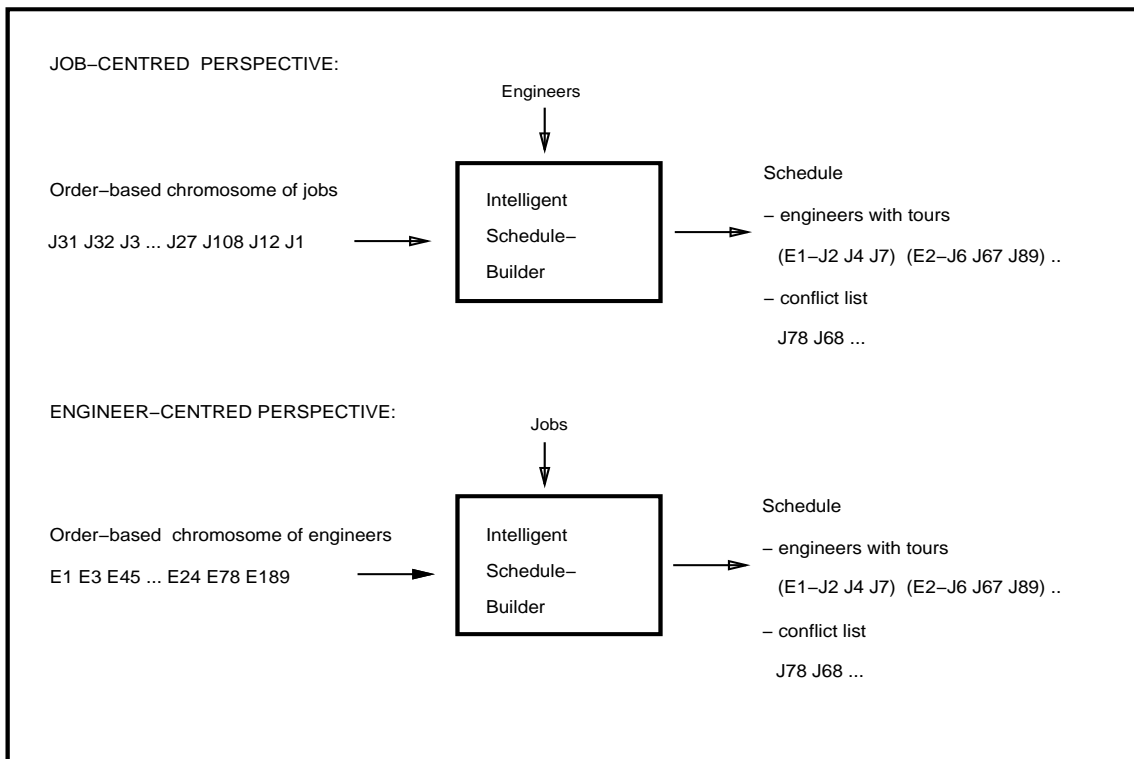


Figure 4.6: Job-centred and engineer-centred perspectives.

chromosome. This operation is performed after the crossover operation, on 4% of the children.

- Termination criteria: The GA is stopped after 100 iterations which is around 20 minutes on a SPARCstation IPC with LISP.
- Model: The two GAs use a steady-state model and the chromosomes are selected according to their rank in the population.
- Due to the stochastic nature of GAs, each algorithm is run 20 times.

Experiments were performed (1) to examine the quality of the schedules produced by the two different schedule-builders, i.e. job- and engineer-centred, and (2) to examine the performances of the two GAs.

Perspective	μ	min	max	σ
JOB	24387	22770	25954	523
ENG	23168	22097	25191	419

Table 4.3: Performances of the two scheduling functions.

For each schedule-builder, a set of experiments has been carried out to check the quality of the schedules produced. In each case, 1000 random strings have been generated and processed through the schedule-builder. In the case of the job-centred function, each string was a random permutation of the list of jobs. For the engineer-centred function, each string was a random permutation of the list of engineers.

Table 4.3 presents the performances of the two schedule-builders. The column μ gives the average cost achieved by the function (out of 1000 tries), min gives the minimum, max gives the maximum, and σ gives the standard deviation.

On average, the engineer-centred function produces better solutions than the job-centred function. The main reason is that the engineer-centred function can handle *bottleneck* engineers in a more efficient manner and, hence, can schedule more jobs than the job-centred function. Observe that the two standard deviations are relatively small compared to the average values. This suggests that the schedules are quite similar to each other. This might limit the performances of both GAs, which might not have enough information (more precisely, enough variety in their respective pool of solutions) to guide their search through the solution space in an efficient manner.

The two curves, displayed in Figure 4.7, show how the GAs performs on this particular problem. As expected, the engineer-centred GA dominates its opponent from start to finish. The two curves are concave and as the search goes on, the average additional improvement found per iteration diminishes. At the end of the search




Figure 4.7: Evolution of performances with time.

(in this case, after 100 iterations), the improvements are marginal. This suggests that both GAs have converged towards a solution. The main reason for such an early convergence is the relatively small size of the populations.

4.3 Conclusion.

The two studies presented in this chapter were used (originally) as an introduction to GAs. In order to solve the two problems under study, a traditional and indirect approach was adopted. The GA, in both cases, was working in association with a schedule-builder.

In the first study, this association (GA and schedule-builder) was capable of reaching optimal solutions, i.e. solutions with no tardiness. This was facilitated by the

fact that the algorithm only had to explore a relatively small search space. In the second study, both GAs (with job-centred and with engineer-centred perspectives) were capable of improving on their initial solutions.

However, the indirect approach has some important drawbacks. First, this type of representation may not allow a total exploration of the search space, i.e. it may compromise *completeness*. This is because the GA works with an indirect view of the problem, e.g. list of jobs, and it may be the case that the optimal solution, or more generally, good quality solutions, cannot be produced by the schedule-builder from *any list of jobs*.

Moreover, despite the fact that indirect GAs have been successfully applied to a wide range of sequencing problems such as the traveling salesman problem, there is no guarantee that this approach will be equally successful when applied to problems that are not pure sequencing problems.

The last point can be explained as follows. Indirect GAs tend to preserve the relative order of the elements in the chromosome during a crossover operation; the main assumption is that the solution-builder will generate good quality solutions, if provided with well-ordered chromosomes. This indirect strategy is likely to be disruptive when applied to problems which are not (100%) sequencing problems. Clearly, indirect GAs might not guide their search properly. In the worst case, a schedule-builder might generate the same solution for any ordered list. This will be the case if the elements within the list are identical. For example, assume an indirect GA working with an engineer-based perspective where all the engineers are identical and are therefore inter-changeable. Any possible permutation of this ordered list will generate the same schedule (or more exactly, schedules with the same cost).

Finally, the schedule-builder generates new solutions entirely from scratch at each crossover. This new solution is then entirely evaluated from scratch. These two

operations are extremely expensive in terms of computational effort.

Again, the above limitations advocate the implementations of a direct approach similar to the ones reported previously by Bagchi et al. [BAGC91] and Bruns [BRUN93].

More generally, these two studies also suffer from a number of important limitations. First, the performances of the algorithms are not compared against the performance of random search; neither are they compared against the performances of other search techniques. Hence, it is not possible to know if the algorithm guides its search in an efficient manner or if the algorithm is just another form of random search; neither is it possible to know how the algorithm compares relatively to other search techniques.

Another important limitation of these two projects is the fact that only one problem is studied. Here, there is the danger of designing an algorithm that is very good at solving a particular problem but nothing else.

In conclusion, an algorithm should be tested against a wide range of problems, each with different characteristics. Its performance should also be compared against those of other techniques. It will then become possible to define the most suitable techniques and parameters for each class of problems. It will also be possible to see how robust a particular technique is, to determine how well it can handle different classes of problems. Finally, such an extensive study should provide enough information about the different techniques and about the problems to anticipate the behaviour of the algorithms and the quality of their answers.

The next two chapters report three sets of experiments which follow the above recommendations. In this study, direct GAs are applied to a wide range of vehicle routing problems. These problems have a substantial range of variation in each of three dimensions, i.e. work load, time constraints and specialisation constraints. Moreover, the performances of these direct GAs are compared against the

performances of other search techniques.

Chapter 5

Experimental Set-up.

This chapter is divided into four sections. In Section 5.1, the set of vehicle routing problems used for the experiments is described. Then, Section 5.2 presents the common model - or data structure - that has been used to represent the solutions of the different problems. Section 5.3 defines the objective function and Section 5.4 introduces a new class of GAs - especially designed for the solution of VRPs - which overcomes many limitations of traditional GA-based systems.

5.1 The set of problems.

Thirty six problems with diverse characteristics have been generated. In order to perform a study as wide ranging as possible, extremes classes of problems have been considered. In particular, this work has examined how three parameters: workload, time constraints (also called time windows) and specialisation constraints (also called technological constraints) can affect the performance of different search techniques.

A problem can be under-resourced, critically resourced or over-resourced, it can have loose or tight time constraints and loose or tight specialisation constraints. Thus, there are twelve different parameter setting and three problems are created at each parameter setting, hence, the thirty six problems. These twelve parameter settings and the associated thirty six problems can be positioned on different points of a cube (see Figure 5.1). On this cube, each point represents three problems and one parameter setting.




Figure 5.1: The 36 problems.

The 6 points on the left face of the cube represent problems with loose time windows. As we move down an edge (any edge) we move from under-resourced problems to over-resourced problems. Moving left to right we increase the tightness of the temporal constraints, and all problems on the right hand face have the tightest temporal constraints. Moving into the page we increase the tightness of the specialisation constraints.

A different data set ¹, i.e. engineers and jobs, is generated randomly for each problem, and this data set respects the following rules:

- All the jobs are located within a 5 kilometres radius of a central location. Hence, all the 36 problems are urban.
- Several jobs can shared the same location.
- All problems have 200 jobs, the durations of these jobs are drawn randomly from the distribution {15,30,45,60,75,90,105,120 minutes} and each duration is equally likely.
- 10% of the jobs are compulsory. If one of these high priority jobs is not allocated, then the search algorithm must create some space in the schedule, for example, by removing one of the low priority jobs.
- The level of time windows (T) is either loose ($T = 0$) or tight ($T = 2$). If $T = 0$, all jobs can be done at any time during the working hours. If $T = 2$, 45% of the jobs can only be done in the morning, 45% can only be done in the afternoon, and 10% can be done at any time - morning and afternoon.
- Under-resourced problems have 30 engineers, critically resourced have 40 and over-resourced have 50.
- The level of specialisation constraints (S) is either loose ($S=0$) or tight ($S = 3$). If $S = 0$, all jobs can be done by any engineer, and if $S = 3$, 80% of the jobs can be done by 20% of the engineers, and 20% of the jobs can be done by 80% of the engineers.
- There is only one base station. All the engineers start their tours from this base station after 9am and come back to this base station for 5pm.

¹The problem generator which was used to create these problems is presented in more detail in Annexe I.

Engineers travel at 12mph. Manhattan distances are assumed and engineers travel up and down, left and right, but not diagonally. Working hours are 9am to 5pm (or 540 to 1020 measured in minutes from midnight). Morning is from 9am to noon (or 540 to 720) and afternoon is from noon to 5pm (or 720 to 1020). No overtime is allowed, engineers must be back at the base station for 5pm.

From now on, the following convention will be used to refer to a problem or to a set of problems:

problem E-T-S-N

where E indicates the number of engineers, T and S indicate the tightness of the time windows and of the specialisation constraints, and N is a number used to distinguish between the three problems generated at each point of the cube. For instance, the expression "*problems 30-*-2-**" refers to the set of problems with 30 engineers and tight specialisation constraints.

5.2 An active constrained-based model.

All the search techniques use a common representation of the VRP. A solution is represented as a list of engineers, and each engineer has his own ordered list of jobs (his tour). For an n engineer problem, there is always the $n + 1$ th engineer (the virtual engineer) that is able to do all work. Therefore, if a job has not yet been allocated to any of the n engineers it will reside in the tour of the $n + 1$ th engineer. This tour is frequently referred to as the "conflict list" or as the set of "jobs not done".

A constraint-based approach ² is taken when representing an engineer's tour. A tour is represented as a sequence of jobs, such that it is possible to perform each job

²This mechanism is required so that the search process does not take an over-constrained

within its time window, travel between jobs in the given sequence, and return to the base before the end of the day (i.e. a legal tour using a least commitment strategy). When a job is inserted into a tour, constraint propagation takes place, updating (restricting) the time windows of the jobs within that tour. Symmetrically, when a job is removed from a tour, constraints are retracted, and the time windows may be relaxed. Therefore, central to the representation of the VRP is a constraint maintenance system that gives an active representation of tours. This constraint maintenance system was implemented by Craig Brind.

5.3 The cost function.

The principal objective is to perform all the jobs and a second objective is to minimise the total distance travelled by the workforce. The following cost function is used to attribute a cost for each solution:

$$cost(s) = C.D(s) + T(s)$$

where C is a constant, $D(s)$ is the sum of the jobs not done in minutes, and $T(s)$ is the sum of the travels in minutes of the engineers. The travel time for an engineer assumes *Manhattan distance*³ covered at a given speed (in this case 12mph), and is measured in minutes.

In some early tests, a small value had been chosen for C ($C = 3$). However, in some instances, the different algorithms were trading work for travel, i.e. algorithms could get a lower cost by doing less work and, hence, much less travel. This situation was undesirable. The objective of performing all the jobs is indeed more important

view of the problem (resulting in lost opportunities) or an under-constrained view of the problem (resulting in infeasible solutions).

³Given two locations $\langle x_1, y_1 \rangle$ and $\langle x_2, y_2 \rangle$, the Manhattan distance is $|x_1 - x_2| + |y_1 - y_2|$.

than the reduction of travel, and the cost function has to reflect this obligation.

By choosing a large enough value for the coefficient C we hope that it is not possible to trade travel for work done. Put another way, given any two solutions s' and s'' , where $cost(s') < cost(s'')$, $D(s') \leq D(s'')$. Therefore, we can make an improvement to s by increasing the amount of work done, i.e. reduce the amount of work *not* done, or by reducing the amount of travel. It should be impossible to reduce the cost of s by reducing the amount of work done in s (increasing the amount of work *not* done in s).

All the experiments were run using a coefficient $C = 600$. Hence, the cost of a solution is 600 times the durations of the jobs not done plus the different travels and, problems with large cost solutions are typically under-resourced, i.e. work is left undone.

5.4 Direct GAs.

The GAs presented in this section are designed with the principal objective of overcoming the many limitations of traditional GA-based systems. This research follows the initial works of Bagchi et al. and Bruns reviewed in Section 3.2.5. These GAs work directly⁴ on the solution-candidates, i.e. the schedules are the chromosomes. There is no longer any schedule-builder. More precisely, the scheduling function is performed directly by the crossover operator. Each element (or locus) of a chromosome represents one engineer, and each one of these loci can be assigned different tours (or alleles). During the reproduction process, parts of two schedules are mixed together to produce one new schedule. Direct crossovers can be split up into two broad categories:

⁴From now on referred to as the "direct" GAs.

- Direct crossovers using a random-based inheritance procedure, where a chromosome inherits random parts of both parents.
- Direct crossovers using a knowledge-based inheritance procedure, where a chromosome inherits the best parts (with respect to one or several given parameter(s)) of its parents. A knowledge-based crossover gives more intelligence to the GA; the operator identifies the best parts, i.e. tours, of a schedule and favours their inheritance during the crossover operation.

Different knowledge-based crossovers can be envisaged, e.g. a crossover may favour the inheritance of tours with a large number of jobs, tours representing a large amount of work, or tours with a low percentage of travel. For instance, when solving under-resourced problems, a GA may attempt to preserve tours with large amounts of work being done. These tours are likely to be well "packed" and, hence, are likely to represent good building blocks for the creation of good-quality solutions.

However, it is expected that these crossovers may not be as robust as random-based crossovers. For instance, in the case of over-resourced problems, the previous crossover, trying to preserve tours with large amounts of work, may not be desirable. Here, all jobs will be performed as the problem is over-resourced and, hence, the primary objective will become the reduction of travel. This crossover will misguide the search. In this case, a random-based crossover ⁵ will be more adaptable or rather, less deterministic and will reach a better solution.

In this study, a number of direct crossovers have been implemented and tested. They are now presented in some detail:

- Crossover #1 and #2 use the same overall algorithm. The principal difference is that Crossover #1 uses a random-based inheritance procedure and

⁵entirely driven by random decisions.

Crossover #2 uses a knowledge-based inheritance procedure.

- Crossover #3 and #4 attempt to overcome some limitations shown by both Crossover #1 and #2. They are both based on the same overall algorithm; Crossover #3 uses a random-based inheritance procedure and Crossover #4 uses a knowledge-based inheritance procedure.

Observe that these crossovers use the tours of two chromosomes, i.e. the *parents*, to produce one chromosome, i.e. the *child*.

5.4.1 Crossover #1

Initially, two parents have been selected for reproduction and the *child* is empty, i.e. all the jobs are in the child's conflict list. This crossover is a three stage process:

1. **Inherit from $parent_1$** : the child inherits half its tours from $parent_1$. For example, if the problem is under-resourced and there are 30 engineers, 15 tours are copied from $parent_1$ to the *child*. All the jobs present in these tours are then removed from the child's conflict list.
2. **Inherit from $parent_2$** : for each job still in the child's conflict list, the crossover tries to allocate the job to the same engineer as in $parent_2$. For example, if job J_6 is done by engineer E_3 in $parent_2$, then the crossover tries to allocate job J_6 to engineer E_3 in the *child*. All jobs which are successfully allocated are removed from the child's conflict list.
3. **Repair**: for each job still in the child's conflict list, the crossover tries to allocate the job to any engineer. Again, all jobs which are successfully allocated are removed from the child's conflict list.

```
1  FUNCTION Crossover1(parent1, parent2)
2  BEGIN
3  child = new-chromosome();
4  child = inherit-half(parent1);
5  FOR job IN conflict-list(child)
6  DO BEGIN
7      decision = nil;
8      eng = who-is-doing(job,parent2);
9      IF eng != CONFLICT-LIST
10     THEN decision = try-to-allocate(job,eng,child);
11     IF decision = true
12     THEN remove-from-conflict-list(job,child);
13  END
14  FOR job IN conflict-list(child)
15  DO BEGIN
16     decision = nil;
17     FOR all engineers While decision = nil
18     DO BEGIN
19         decision = try-to-allocate(job,eng,child);
20         IF decision = true
21         THEN remove-from-conflict-list(job,child);
22  RETURN child;
23  END;
```

Figure 5.2: Pseudo-code for Crossover #1.

The pseudo-code for Crossover#1 is given in Figure 5.2 and Figure 5.3 provides a graphical description of the actions taken during a Crossover#1 operation. In this example, there are 4 engineers, E_1 to E_4 , and 20 jobs to allocate, J_1 to J_{20} .

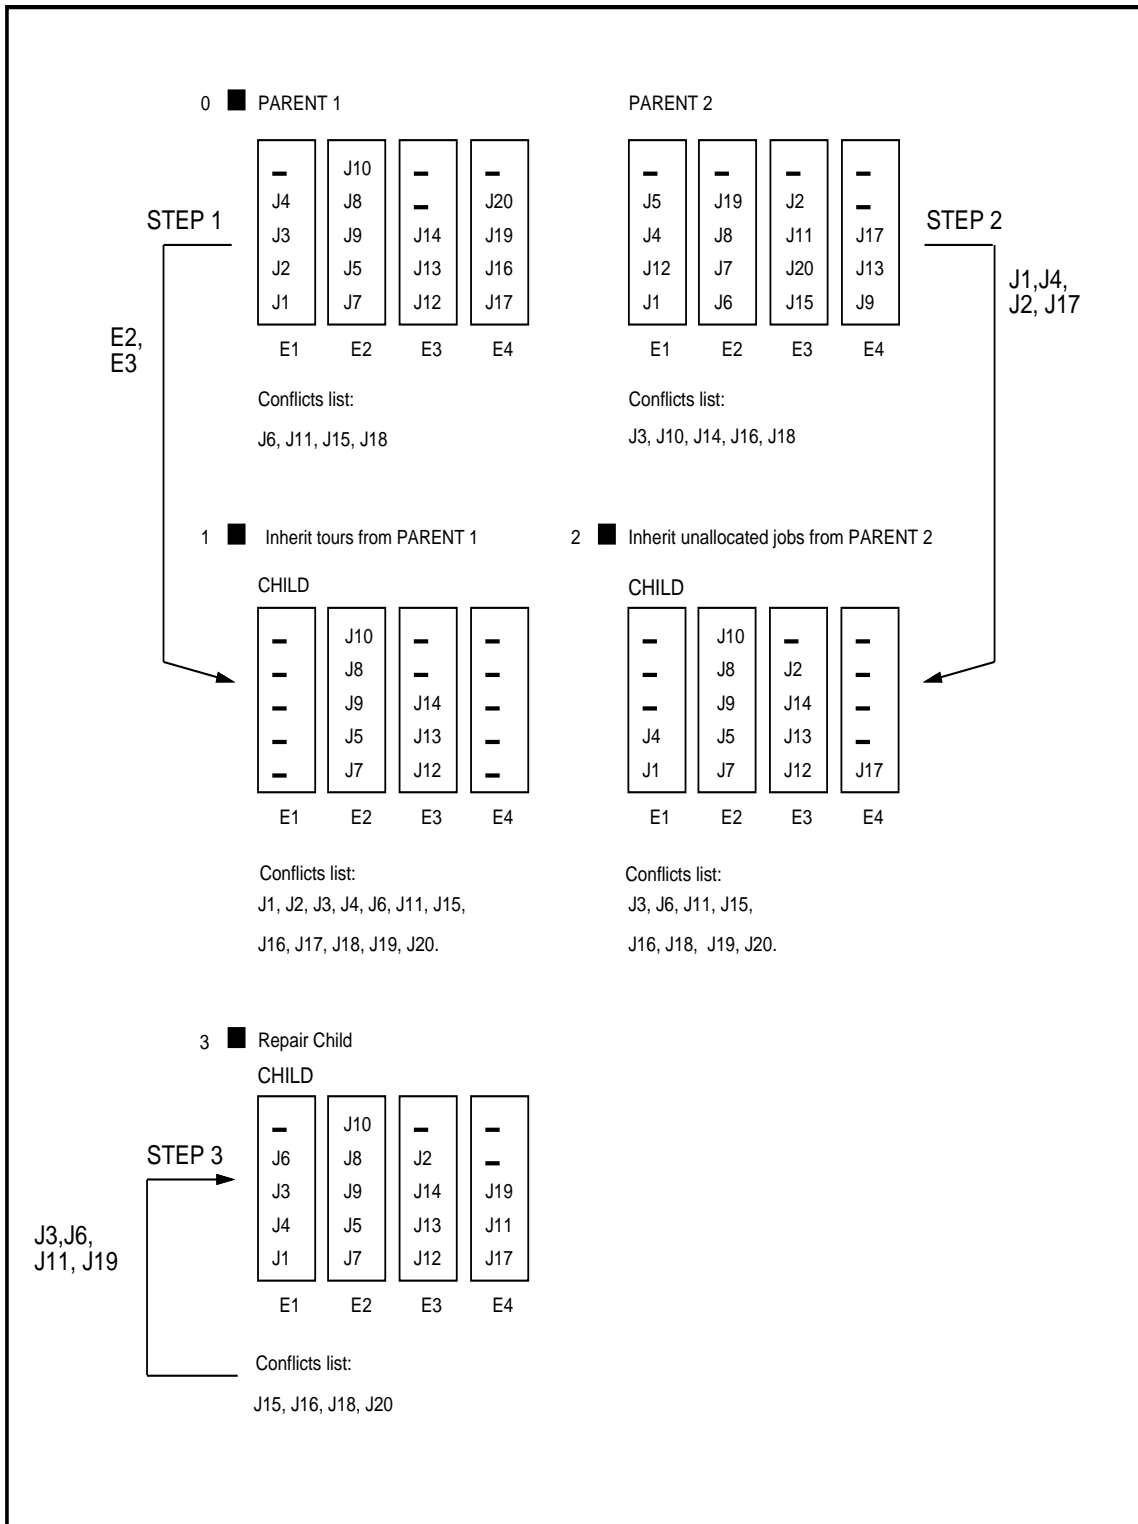


Figure 5.3: Crossover #1, step by step.

In the first stage, the *child* inherits half its tours from *parent*₁. There are four engineers so two tours are copied, tours of engineers E_2 and E_3 have been selected randomly; they are now copied from *parent*₁ into the *child* (step 1). At this point, twelve jobs are still in the child's conflict list. For each of these jobs, the crossover tries to allocate the job to the same engineer as in *parent*₂. In this example, it has been possible to allocate jobs J_1, J_2, J_4 , and J_{17} to the same engineer as in *parent*₂ (step 2). There are still eight jobs remaining in the child's conflict list. The crossover now tries to allocate them to any possible engineer. Jobs J_3, J_6, J_{11} , and J_{19} are inserted during this operation (step 3).

5.4.2 Crossover #2.

This crossover follows the same sequence of steps as Crossover #1, but uses a knowledge-based inheritance procedure. That is, the *child* still inherits half the tours of *parent*₁. However, the inheritance is no longer random-based. It is now knowledge-based and the *child* inherits the best tours of *parent*₁, i.e. the child inherits the best half of *parent*₁. Again, the definition of *best tour* might be related to the number of jobs in a tour, the sum of the durations of the jobs in a tour, or some measure of *profit* (whatever that measure may be). The rest of the algorithm, i.e. steps 2, 3, and 4, is identical to Crossover #1.

Separate tests - done in the early stages of the implementation - showed that the above crossovers (#1 and #2) are extremely disruptive and a GA equipped with either one of these crossovers performs only marginally better than random search.

5.4.3 Crossover #3.

Crossover #3 follows a three stage process:

1. **Inheritance from $parent_1$** : the entire chromosome, i.e. tours and conflict list, is copied from $parent_1$ to the *child*. At this point, the *child* is a clone of $parent_1$.
2. **Inheritance from $parent_2$** :
 - (a) The crossover now decides which tours will be copied from $parent_2$ to the *child*. 50% of the tours in $parent_2$ must be inherited by the child. For instance, if the problem has four engineers, the crossover may decide randomly that the tours of E_2 and E_3 will be copied.
 - (b) Before the crossover can actually copy these tours from $parent_2$ to the *child*, it must work on the child so that future duplications and omissions of jobs in the *child* can be avoided:
 - i. The selected tours (e.g E_2 and E_3) are deleted from the child's schedule and their jobs are pushed into the child's conflict list.
 - ii. Jobs allocated to the selected engineers in $parent_2$ are identified. All these jobs are then deleted from the *child* both in the child's tours and in the child's conflict list.
 - iii. Once this prerequisite work is done, the selected tours are copied from $parent_2$ to the *child*.
3. **Repair *child***: for each job present in the conflict list, the algorithm attempts to allocate the job to any engineer. As usual, all jobs which are successfully allocated are removed from the child's conflict list.

The pseudo-code for Crossover #3 is given in Figure 5.4 and Figure 5.5 provides a graphical description of the actions taken during a Crossover #3 operation.

Figure 5.5 shows how Crossover #3 works. First, the entire schedule is copied from $Parent_1$ to $child$ (step 1). Then, a random decision on which tours are to be copied from $Parent_2$ to $child$ is taken. Here, tours 2 and 3 will be inherited from $Parent_2$ (step 2). At this stage, the algorithm removes all the jobs from the child's tours 2 and 3 and pushes them into the child's conflict list. Jobs $J_5, J_7, J_8, J_9, J_{10}, J_{12}, J_{13},$ and J_{14} are moved to the child's conflict list (step 3). All the jobs allocated to tours 2 and 3 in $Parent_2$ are deleted from the $child$ (both tours and conflict list); jobs $J_2, J_6, J_7, J_8, J_{11}, J_{15}, J_{19},$ and J_{20} are deleted (step 4). $Child$ inherits the selected tours from $Parent_2$; tours 2 and 3 are copied from $parent_2$ (step 5). The final action is to repair the child and attempt to insert unallocated jobs into the schedule; here, jobs J_5, J_9, J_{10}, J_{12} are inserted.

Crossover #4.

This crossover follows the same sequence of steps as Crossover #3, but, like Crossover #2, it uses a knowledge-based inheritance procedure and the child inherits the best half of $Parent_2$. The rest of the algorithm, i.e. steps 2, 3 and 4, is identical to Crossover #3.

The performance of these direct GAs are examined in the next chapter.

```
1  FUNCTION Crossover3(parent1, parent2)
2  BEGIN
3  list-of-tours = tours-selection();
4  FOR tour IN list-of-tours
5  DO BEGIN
6      remove-tour(tour,child);
7  END
8  FOR tour IN list-of-tours
9  DO BEGIN
10     list-of-jobs = push-jobs(tour,parent2)
11  END
12  FOR job IN list-of-jobs
13  DO BEGIN
14  delete-job(job,child)
15  END
16  FOR tour IN list-of-tours
17  DO BEGIN
18     copie-tour(tour,parent2,child)
19  END
20  FOR job IN conflict-list(child)
21  DO BEGIN
22     decision = nil;
23     FOR all engineers While decision = nil
24     DO BEGIN
25         decision = try-to-allocate(job,eng,child);
26         IF decision = true
27         THEN remove-from-conflict-list(job,child);
28  RETURN child;
29  END;
```

Figure 5.4: Pseudo-code for Crossover #3.

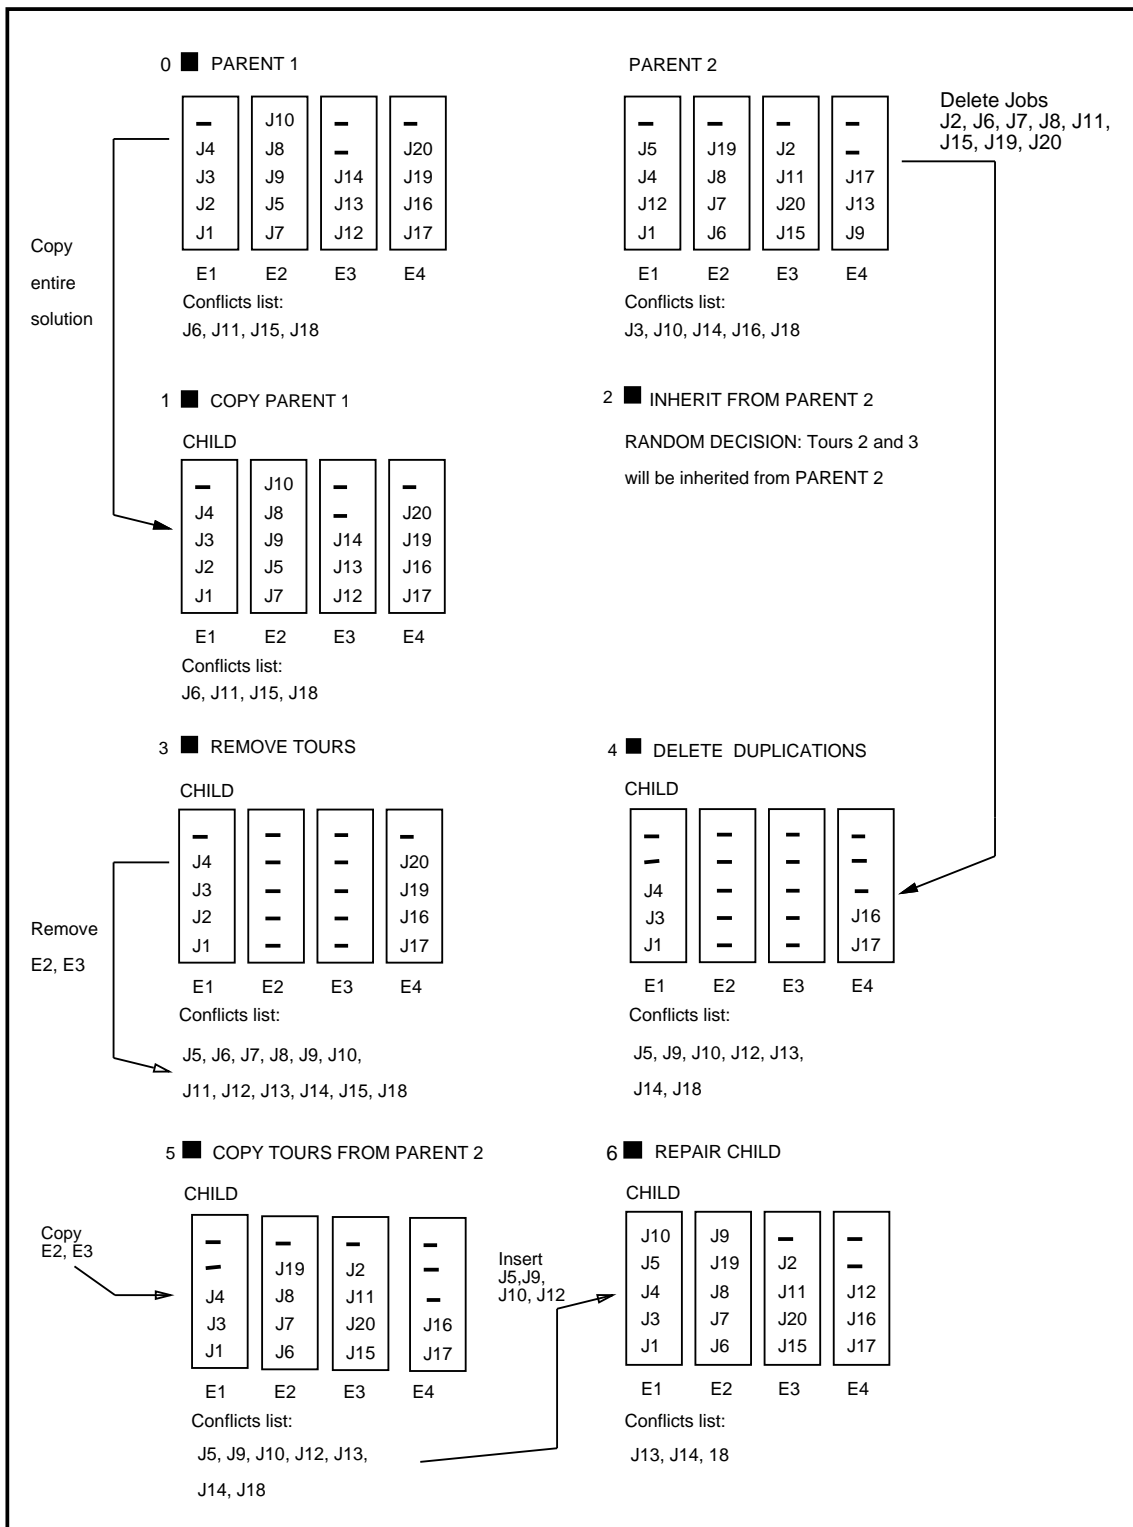


Figure 5.5: Crossover #3, step by step.

Chapter 6

Experiments and Discussions.

This chapter is divided into three sections; each section reports a separate set of experiments:

1. In Section 6.1, the direct GAs introduced in Chapter 5 are compared against an indirect GA, and random search. Two models - steady-state and generational - are implemented. A study is then performed to determine the effect of different population sizes on a given direct GA. Finally, this section reports on the performance of a direct GA associated with a local repair algorithm. Each technique is given ten minutes of CPU time (on a SPARCstation IPC) per problem.
2. In Section 6.2, the best direct GA is compared against random search, hill climbing, simulated annealing and tabu search. Again, each technique is given ten minutes of CPU time per problem.
3. Section 6.3 repeats the previous comparative study. However, in this case, each technique is given the equivalent¹ of four hours of CPU time per problem.

¹In fact, these CPU intensive experiments were executed on a DEC Alpha 3000/300; each algorithm was given one hour of CPU time and the DEC Alpha was roughly four time faster than

The different algorithms were written in C and runs were performed on two SPARC-stations IPC and one DEC Alpha 3000/300 under the UNIX operating system.

6.1 Genetic Algorithms.

This section examines GAs in some detail and addresses a number of issues:

1. Which GA is the best among a set of possible models and crossovers. The behaviour of several GAs are studied over a set of problems and the conditions under which one GA is better than another are defined.
2. The effect of different population sizes on the performance of direct GAs.
3. A repair algorithm is associated with a direct GA and the behaviour of this new hybrid algorithm is studied.

During this study, each algorithm has been applied 5 times to each problem and the start value of each run was chosen randomly, so that an indication of the average behaviour can be obtained. Comparisons between algorithms are based on the average values resulting from these trials. CPU time is limited such that each algorithm is only allowed 10 minutes on each problem. Therefore, to investigate the performance of an algorithm on a single point of the cube takes about 150 minutes - there are 3 problems at each point, and the algorithm is applied 5 time to each problem, 10 minutes for each application. To apply a single algorithm to all 36 problems takes 30 hours and 12 algorithms are examined in this section. Therefore, it took approximately 360 hours on a SPARCstation IPC to perform the set of experiments reported in this section.

a SPARCstation IPC.

6.1.1 Comparison of different genetic algorithms.

During this study, a number of GAs have been implemented. Here, we examine the performances of these different algorithms and compare them over the entire range of problems. Four GAs are compared and RS is used as the reference algorithm:

1. The first GA is the only indirect GA of this comparative study; it uses a generational model and the partially mapped crossover ² (PMX). The chromosome is an ordered list of engineers. The schedule-builder ³ takes the ordered list of engineers and allocates jobs to each engineer in turn. This first GA will be referred to as PMX.
2. The second GA is a direct GA; it uses a generational model and the random-based direct Crossover #3. This GA will be referred to as Direct3.
3. The third GA is also a direct GA; it uses the steady-state model and the random-based Crossover #3. This GA will be referred to as Direct3-s.
4. The fourth GA is again a direct GA; it uses a steady-state model and the intelligent direct Crossover #4. This crossover uses the amount of work done to sort the different tours of a solution and tries to preserve the best tours. This GA will be referred to as Direct4-s.
5. RS performs a random sampling of the search space and keeps a record of the best solution found during the search. RS enters a loop and at each iteration the algorithm creates a random solution via the schedule builder, and then evaluates that solution. The cost of the current solution S is then compared with the cost of the best solution found so far (S_{best}). If $cost(S) <$

²This order-based crossover attempts to preserve the absolute positions of elements for one part of a child and the relative order for the other part. This crossover is described more fully in Subsection 3.2.5

³This schedule-builder is used by RS to generate all its solutions and also by all the direct GAs to create their initial populations.

$cost(S_{best})$ then S_{best} is replaced by S , otherwise S_{best} remains unchanged and S is relinquished. This cycle is repeated until the time limit is reached.

In the early stages of this project, some C libraries (developed by L Corcoran, these libraries are available on the Internet) were used to implement PMX and Direct3. Here, the GA uses a generational model: an entire population is created at each iteration. In fact, this model works with two populations: the old one and the new one. The old population is used to generate the new population. Parents in the old population are selected for reproduction and they generate offspring in the new population. 30% of the old population is copied (or cloned) to the new population at each generation. This corresponds to a generational gap of 30%. A roulette selection technique was adopted where each member of the population is allocated a slice of a roulette wheel; the slice is proportional to the fitness of the chromosome.

Under this model, a GA might converge prematurely. Such a convergence may happen well before the ten minutes limit and when this situation occurs, the GA is stopped. Furthermore, this first model also suffers from some serious limitations in terms of computational efficiency and memory requirement. This prompted the decision to design a new GA. During the design of this new algorithm, the main emphasis was on the design of an efficient algorithm, i.e. this new model had to be *fast and memory efficient*.

The *second-generation GAs*, Direct3-s and Direct4-s, use a steady-state model. Initially, a set of n solutions is generated and sorted by fitness. The selection function works as follows. First, a number, x , is generated randomly in the range 1 to $n - 1$. Then, a number y is generated, where $1 < y \leq x$ and the y th chromosome is selected for reproduction. Once two parents have been selected, the crossover operator generates one child which is then inserted at the correct position in the population and the weakest member of the population is deleted. This cycle of selection, recombination and replacement is repeated until the 10 minute limit is

reached.

Figure 6.1 shows that with this selection function fitter chromosomes have a much higher chance of being selected. To plot this graph, the selection function was run 10 million times. There were 100 chromosomes in the pool, ranked with respect to their cost from the best one to the worst one. The fittest chromosome was selected on 518692 occasions and the worst chromosome was only selected on 902 occasions.




Figure 6.1: Selection function - an elitist strategy.

This selection strategy was adopted because the GA has a strict time-limit of 10 minutes. In this situation, this selection function provides a suitable balance between exploration of the solution space and rapid convergence.

Moreover, in order to prevent a premature convergence, the GA does not accept the insertion of duplicates in the population. Basically, a chromosome cannot be inserted if it is identical, in terms of cost, to a chromosome already present in the population. When this situation occurs, the chromosome is simply rejected and the algorithm generates a new chromosome. Finally, during these first experiments, the

four GAs have a population of 100 chromosomes and none of them uses a mutation operator.

Tables 6.1, 6.2, 6.3, 6.4 present the performances of the five algorithms over the full range of problems. A table entry gives the average value taken over the 5 applications of the algorithm to a given problem. There are five columns for each problem; each column corresponds to a different algorithm. Hence, if we scan across a row, we can compare algorithms on specific problems.

The data is analysed from 4 perspectives. Table 6.1 gives the average final cost reached by the five algorithms for the different problems. Table 6.2 gives the average amount of work done. Table 6.3 gives the average travel time. Table 6.4 gives the average amount of work not done.

Table 6.1 shows that, in general, the five algorithms find solutions with comparable costs, i.e. the solutions are in the same order of magnitude. This is a remarkable result as the cost function used in these tests does not have a smooth topology. Indeed, as discussed earlier, the cost of a solution is 600 times the duration of the jobs not done plus the different travel time and, thus even a marginal improvement in terms of work done leads to a significant reduction of cost. Despite this, the five algorithms generate (in most cases) solutions of similar cost. This signifies that, generally, a problem is equally difficult for the five algorithms. If one algorithm finds a given problem difficult, then all the others will find this particular problem difficult as well.

However, there are some noticeable exceptions. For instance, on the 40-2-0-3 problem, the average costs for RS, PMX, Direct3 and Direct3-s are 56334, 47417, 12963 and 12861 respectively but Direct4-s manages an average cost of only 1984.

This results from the fact that only Direct4-s can allocate all the work, while RS, PMX, Direct3 and Direct3-s still have 20, 75, 18 and 18 minutes respectively of work not done. As this example shows, massive differences (in terms of cost) can

be expected between two algorithms if one manages to allocate all the jobs and the other one does not. In this case, it can be said that the problem under study is easy for Direct4-s but hard for the other techniques. The situation where the different algorithms do not have the same behaviour only occurs in the 40-* problems with time-windows and technological constraints.

These differences in performance suggest that these particular problems are critical, i.e. some of the algorithms can allocate all the work, some cannot. Clearly, it appears that the load of the engineers is a key parameter that determines whether all the work is allocated or causes some to remain undone.

Typically, the 30-* problems are under-resourced and all the algorithms generate solutions with large costs. The best solution is simply the one which minimises the amount of work not done. Travel is only a secondary objective here.

The 50-* problems are over-resourced and all the jobs are done. Hence, the reduction of travel becomes the primary concern and the best solution is the one that minimises travel. Since all jobs are allocated, travel offers the only scope for improvement.

For the 40-* problems in the middle of the spectrum, the situation is more delicate. Some problems have solutions in which all the work is done, and some have solutions in which some jobs remain unallocated. Clearly, these problems are critically resourced and they are the most interesting problems. It is of little interest to an enterprise to solve a problem by continually having excessive resources, or to continually fail to meet those objectives due to under-resourcing.

Tables 6.1 and 6.2 show that the solution with the largest amount of work done is always the best solution. This was guaranteed by the cost function. Furthermore, the best solution also tends to be the one which minimises travel. For instance, on the 30-0-0-1 problem, Direct4-s allocates the largest amount of work (12972) but it also has the smallest amount of travel (1286). Clearly, in this case, the best

solution is doing more work and less travel than the other solutions. This is a typical situation and suggests that reduction of travel is of prime importance. This observation will be used in Section 6.1.3 to create a repair algorithm.

Table 6.5 ranks the algorithms. Given a problem instance P_i , the performance of algorithm A_i is compared against the others. If an algorithm A_i finds a better cost than an algorithm A_j , then the algorithm A_i is awarded one point. If the two algorithms find solutions of equal value, then both get zero points. In total, an algorithm can score a maximum of four points and a minimum of 0 points for a given problem instance.

Figure 6.2 summarises the average, the minimum and the maximum scores for each algorithm over the entire set of 36 problems.

Algorithm	average	min	max
RS	0.42	0	1
PMX	0.58	0	1
Direct3	2.86	2	4
Direct3-s	2.97	2	4
Direct4-s	3.16	2	4

Figure 6.2: Ranking of GAs over 36 problems.

From Table 6.5 and Figure 6.2, it can be seen that the three direct GAs totally dominate RS and the indirect candidate, PMX. In this particular set of experiments, there is no single example where PMX or RS provide a better solution than Direct3, Direct3-s or Direct4-s. Furthermore, in most cases, there is a significant gap between the solutions of RS and PMX, and the solutions of the three direct GAs. The average costs achieved by the PMX algorithm are remarkably similar to those achieved by RS; these two algorithms achieve average scores of 0.42 (for RS) and 0.58 (for PMX). These are very low scores compared with the scores of the three direct GAs. Clearly, the association of order-based GA and schedule-builder

is not suitable for this type of problem, i.e. the PMX algorithm does not guide its search properly.

Among the three direct GAs, there is no overall champion. Direct3, Direct3-s and Direct4-s have comparable average scores of 2.86, 2.97 and 3.16.

The GA using the knowledge-based crossover, Direct4-s, achieves the best score of 3.16. As expected, this crossover is especially effective on the 30-* problems and on the critical 40-* problems. On this particular set of problems, i.e. 12 under-resourced 30-* problems and 6 critical 40-* problems, Direct4-s reaches the best solution on 17 occasions and is only once beaten by Direct3-s. However, on the 50-* problems, Direct4-s is outperformed by Direct3 and Direct3-s; both of these GAs use the same random-based direct crossover.

These results come as no surprise. Direct4-s uses its intelligence to inherit the best tours during the crossover operation; in this case, the best tours are the ones which have the largest amount of work done. This knowledge-based crossover is particularly effective on under- or critically-resourced problems where the algorithm must use all its intelligence in order to reduce the amount of work not done. However, on the 50-* problems, where travel is the only scope for improvement, this technique does not guide the search properly and the two random-based direct GAs (Direct3 and Direct3-s) are more effective.

Direct3-s is particularly effective on the 50-* problems, where out of 12 instances it scores the maximum 4 points on 9 occasions. In general, Direct3-s is marginally better than Direct3. As these two algorithms use the same crossover, this difference suggests that, on average, the steady state model is more effective than the generational model.

Table 6.6 gives a guide to the amount of exploration performed by a given technique over a given problem in 10 minutes. The RS column shows the number of schedules created. The four other columns show the number of crossover operations. The two

algorithms, PMX and Direct3, experienced premature convergence in a number of instances, hence the two columns PMX and Direct3, also show (between brackets) the real amount of time in seconds spent by the algorithm before convergence.

In general, RS and PMX perform fewer iterations than the three direct GAs. Also, PMX performs slightly fewer iterations than RS. These results were expected. RS is a blind search, i.e. a new solution is generated and evaluated from scratch at each operation. These two operations are extremely expensive in terms of computational effort. The indirect GA, PMX, has the same behaviour, i.e. generation and evaluation from scratch for each crossover, but on top of this, PMX also has to generate order-based chromosomes and this small overhead explains the slightly bigger value of RS over PMX.

In the case of the direct GAs, the crossover operator combines two solutions so a direct GA only has to reschedule part of a solution. Therefore, the exploration of a single point in the search space will be substantially less expensive than for RS or PMX.

Looking at the RS column, it can be seen that more exploration takes place over problems with specialisation. Generally, RS examines 3 times as many points in the 30-*-3 problems as in the 30-*-0 problems. This also holds for the 40-* and the 50-* problems. PMX and the three direct GAs also experience this phenomenon.

Moving from under-resourced to over-resourced problems, direct GAs perform more exploration. This is because, in under-resourced problems, the computational cost of these GAs is dominated by the repair of chromosomes trying to insert as many jobs in the schedule as possible. In over-resourced problems, however, less repair is required and more exploration takes place.

Direct3-s always performs more iterations than Direct3. This is another indication that the steady-state model is more efficient than the generational model. Table 6.6 also shows that PMX and Direct3 converge extremely rapidly in some cases. For

instance, on the 40.2.3.1 problem, the two GAs converge and stop after only 202 and 56 seconds respectively. This suggests that the generational model is not capable of maintaining sufficient diversity within the population during the search. This might be due to the selection technique which gives too much emphasis to the fittest elements, but it might also be due to some other factors such as the population being too small, or the lack of a mutation operator.

Total convergence cannot happen under the steady state model. Indeed, this model does not accept duplicates and thus, the population cannot converge towards a uniform population. However, if the level of diversity is not kept at a proper level, the search might stagnate and the algorithm might well be spending most of its time creating solutions which will be rejected because they are already present in the population. This problem becomes more acute with large populations.

Table 6.7 shows the number of solutions rejected, in 10 minutes, by the two GAs, Direct3-s and Direct4-s. These are the only algorithms which use the steady state model. When compared with Table 6.6, which presents the number of solutions visited in 10 minutes, this table suggests that both Direct 3-s and Direct 4-s are spending most of their time creating solutions which are rejected immediately.

Algorithm	average	min	max
RS	-	-	-
PMX	-	-	-
Direct3	-	-	-
Direct3-s	42.5%	11%	83%
Direct4-s	27.6%	1%	78%

Figure 6.3: GAs - When the best solution was found (% of 10 minutes).

Table 6.8 shows, for each problem, when the best solution was found during the search, and Figure 6.3 summarises this information. Both tables show the consequence of the lack of diversity in the population. In most cases, the two GAs

find their best solutions relatively early during the search, i.e. they do not use the full 10 minutes, and this situation is even more acute when solving problems with tight technological constraints. This might be due to the fact that in these types of problems, the different GAs perform far more iterations, and therefore, the loss of diversity happens earlier in the 10 minutes.

Several options can be suggested in order to overcome this loss of diversity:

1. Incorporate a mutation operator. Here, a possible mutation would be to randomly move a job from one engineer to another. However, this type of action is already performed during the crossover operation by the internal repair algorithm, i.e. *implicit mutation*, and hence, this operator is not an option worth considering here.
2. Adopt a new selection function. Rather than using an exponential selection function such as the one displayed in Figure 6.1, one possible option would be to normalise the situation and to make the selection probability of each chromosome directly proportional to its ranking in the population, rather than to its individual cost. This normalisation technique would limit the risk of domination by one single super-individual. However, this technique was implemented in the generational model and did not improve the situation.
3. Increase the size of the population. Very small populations might not create enough variety. However, larger populations might be prohibitive in terms of computational effort and this must be considered in a time-constrained environment.

The following section will investigate this last option.

In conclusion, the direct GAs totally dominate PMX over the entire set of problems. In this particular study, the indirect GA is not directing its search properly and

appears to be just another form of random search. Crossover #3 (the random-based crossover) should be preferred for the solution of over-resourced problems and Crossover #4 (the knowledge-based crossover) should be preferred for both under-resourced and critical problems. The steady-state model is overall the best option.

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	975696	982923	690824	660190	588086
30-0-0-2	775956	768757	467641	426173	384688
30-0-0-3	617570	622975	293034	293033	251543
30-0-3-1	763409	765237	529049	516432	435282
30-0-3-2	815645	803022	584880	530766	487509
30-0-3-3	1092715	1078299	849329	824090	793437
30-2-0-1	909188	898370	786606	703645	649560
30-2-0-2	772461	767042	700241	651570	608226
30-2-0-3	1499448	1499493	1421860	1423532	1409089
30-2-3-1	1364732	1339495	1213344	1225979	1180816
30-2-3-2	1154086	1148643	1054976	993695	919782
30-2-3-3	909303	885900	781294	784890	765036
40-0-0-1	2246	2260	1367	1333	1335
40-0-0-2	2231	2236	1325	1335	1338
40-0-0-3	2186	2173	1299	1307	1322
40-0-3-1	2315	2298	1481	1555	1582
40-0-3-2	2200	2192	1476	1460	1527
40-0-3-3	2382	2405	1529	1562	1640
40-2-0-1	445075	454055	417804	387155	392486
40-2-0-2	54407	63443	54321	23583	5436
40-2-0-3	56334	47417	12963	12861	1984
40-2-3-1	22329	13401	2453	2273	2176
40-2-3-2	2279	2325	2083	1989	1954
40-2-3-3	191368	184196	142638	120974	108409
50-0-0-1	2313	2297	1415	1359	1374
50-0-0-2	2235	2233	1325	1402	1422
50-0-0-3	2234	2217	1317	1320	1280
50-0-3-1	2276	2277	1522	1575	1625
50-0-3-2	2365	2313	1521	1545	1654
50-0-3-3	2325	2317	1530	1610	1633
50-2-0-1	2184	2225	1684	1689	1776
50-2-0-2	2107	2101	1493	1560	1560
50-2-0-3	2119	2107	1549	1554	1581
50-2-3-1	2184	2189	1843	1867	1891
50-2-3-2	2354	2390	2112	2106	2049
50-2-3-3	2385	2416	1950	1988	2058

Table 6.1: GAs performance - Average final cost.

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	12327	12315	12801	12852	12972
30-0-0-2	12285	12297	12798	12867	12936
30-0-0-3	12279	12270	12819	12819	12888
30-0-3-1	12156	12153	12546	12567	12702
30-0-3-2	12144	12165	12528	12618	12690
30-0-3-3	12267	12291	12672	12714	12765
30-2-0-1	12183	12201	12387	12525	12615
30-2-0-2	12111	12120	12231	12312	12384
30-2-0-3	11889	11889	12018	12015	12039
30-2-3-1	11424	11466	11676	11655	11730
30-2-3-2	11835	11844	12000	12102	12225
30-2-3-3	11103	11142	11316	11310	11343
40-0-0-1	13155	13155	13155	13155	13155
40-0-0-2	12780	12780	12780	12780	12780
40-0-0-3	13410	13410	13410	13410	13410
40-0-3-1	12990	12990	12990	12990	12990
40-0-3-2	13575	13575	13575	13575	13575
40-0-3-3	13230	13230	13230	13230	13230
40-2-0-1	13287	13272	13332	13383	13374
40-2-0-2	13293	13278	13293	13344	13374
40-2-0-3	13560	13575	13632	13632	13650
40-2-3-1	13422	13437	13455	13455	13455
40-2-3-2	13635	13635	13635	13635	13635
40-2-3-3	13560	13572	13641	13677	13698
50-0-0-1	13860	13860	13860	13860	13860
50-0-0-2	13575	13575	13575	13575	13575
50-0-0-3	13215	13215	13215	13215	13215
50-0-3-1	13335	13335	13335	13335	13335
50-0-3-2	13740	13740	13740	13740	13740
50-0-3-3	13830	13830	13830	13830	13830
50-2-0-1	13185	13185	13185	13185	13185
50-2-0-2	13290	13290	13290	13290	13290
50-2-0-3	13710	13710	13710	13710	13710
50-2-3-1	13470	13470	13470	13470	13470
50-2-3-2	13215	13215	13215	13215	13215
50-2-3-3	13335	13335	13335	13335	13335

Table 6.2: GAs performance - Average amount of work done in minutes.

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	1896	1923	1424	1390	1286
30-0-0-2	1956	1957	1441	1373	1288
30-0-0-3	1970	1975	1434	1433	1343
30-0-3-1	2009	2037	1649	1632	1482
30-0-3-2	2045	2022	1680	1566	1509
30-0-3-3	1915	1899	1529	1490	1437
30-2-0-1	1988	1970	1806	1645	1560
30-2-0-2	2061	2042	1841	1770	1626
30-2-0-3	1848	1893	1660	1532	1489
30-2-3-1	2132	2095	1944	1979	1816
30-2-3-2	2086	2043	1976	1895	1782
30-2-3-3	2103	2100	1894	1890	1836
40-0-0-1	2246	2260	1367	1333	1335
40-0-0-2	2231	2236	1325	1335	1338
40-0-0-3	2186	2173	1299	1307	1322
40-0-3-1	2315	2298	1481	1555	1582
40-0-3-2	2200	2192	1476	1460	1527
40-0-3-3	2382	2405	1529	1562	1640
40-2-0-1	2275	2255	2004	1955	1886
40-2-0-2	2207	2243	2121	1983	1836
40-2-0-3	2334	2417	2163	2061	1984
40-2-3-1	2529	2601	2453	2273	2176
40-2-3-2	2279	2325	2083	1989	1954
40-2-3-3	2368	2396	2238	2174	2209
50-0-0-1	2313	2297	1415	1359	1374
50-0-0-2	2235	2233	1325	1402	1422
50-0-0-3	2234	2217	1317	1320	1280
50-0-3-1	2276	2277	1522	1575	1625
50-0-3-2	2365	2313	1521	1545	1654
50-0-3-3	2325	2317	1530	1610	1633
50-2-0-1	2184	2225	1684	1689	1776
50-2-0-2	2107	2101	1493	1560	1560
50-2-0-3	2119	2107	1549	1554	1581
50-2-3-1	2184	2189	1843	1867	1891
50-2-3-2	2354	2390	2112	2106	2049
50-2-3-3	2385	2416	1950	1988	2058

Table 6.3: GAs performance - Average amount of travel time in minutes.

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	1623	1635	1149	1098	978
30-0-0-2	1290	1278	777	708	639
30-0-0-3	1026	1035	486	486	417
30-0-3-1	1269	1272	879	858	723
30-0-3-2	1356	1335	972	882	810
30-0-3-3	1818	1794	1413	1371	1320
30-2-0-1	1512	1494	1308	1170	1080
30-2-0-2	1284	1275	1164	1083	1011
30-2-0-3	2496	2496	2367	2370	2346
30-2-3-1	2271	2229	2019	2040	1965
30-2-3-2	1920	1911	1755	1653	1530
30-2-3-3	1512	1473	1299	1305	1272
40-0-0-1	0	0	0	0	0
40-0-0-2	0	0	0	0	0
40-0-0-3	0	0	0	0	0
40-0-3-1	0	0	0	0	0
40-0-3-2	0	0	0	0	0
40-0-3-3	0	0	0	0	0
40-2-0-1	738	753	693	642	651
40-2-0-2	87	102	87	36	6
40-2-0-3	90	75	18	18	0
40-2-3-1	33	18	0	0	0
40-2-3-2	0	0	0	0	0
40-2-3-3	315	303	234	198	177
50-0-0-1	0	0	0	0	0
50-0-0-2	0	0	0	0	0
50-0-0-3	0	0	0	0	0
50-0-3-1	0	0	0	0	0
50-0-3-2	0	0	0	0	0
50-0-3-3	0	0	0	0	0
50-2-0-1	0	0	0	0	0
50-2-0-2	0	0	0	0	0
50-2-0-3	0	0	0	0	0
50-2-3-1	0	0	0	0	0
50-2-3-2	0	0	0	0	0
50-2-3-3	0	0	0	0	0

Table 6.4: GAs performance - Average amount of work not done in minutes.

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	1	0	2	3	4
30-0-0-2	0	1	2	3	4
30-0-0-3	1	0	2	3	4
30-0-3-1	1	0	2	3	4
30-0-3-2	0	1	2	3	4
30-0-3-3	0	1	2	3	4
30-2-0-1	0	1	2	3	4
30-2-0-2	0	1	2	3	4
30-2-0-3	1	0	3	2	4
30-2-3-1	0	1	3	2	4
30-2-3-2	0	1	2	3	4
30-2-3-3	0	1	3	2	4
40-0-0-1	1	0	2	4	3
40-0-0-2	1	0	4	3	2
40-0-0-3	0	1	4	3	2
40-0-3-1	0	1	4	3	2
40-0-3-2	0	1	3	4	2
40-0-3-3	1	0	4	3	2
40-2-0-1	1	0	2	4	3
40-2-0-2	1	0	2	3	4
40-2-0-3	0	1	2	3	4
40-2-3-1	0	1	2	3	4
40-2-3-2	1	0	2	3	4
40-2-3-3	0	1	2	3	4
50-0-0-1	0	1	2	4	3
50-0-0-2	0	1	4	3	2
50-0-0-3	0	1	3	2	4
50-0-3-1	1	0	4	3	2
50-0-3-2	0	1	4	3	2
50-0-3-3	0	1	4	3	2
50-2-0-1	1	0	4	3	2
50-2-0-2	0	1	4	2	3
50-2-0-3	0	1	4	3	2
50-2-3-1	1	0	4	3	2
50-2-3-2	1	0	2	3	4
50-2-3-3	1	0	4	3	2

Table 6.5: GAs performance - Ranking of GAs (versus cost).

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	1400	1292 (616)	5603 (602)	8190	9700
30-0-0-2	1400	1292 (615)	6296 (570)	11100	13830
30-0-0-3	1400	1292 (618)	7548 (552)	14990	18880
30-0-3-1	4400	4175 (603)	5657 (179)	36560	42050
30-0-3-2	4400	3944 (572)	5032 (169)	36210	38900
30-0-3-3	4600	4352 (601)	6337 (234)	26640	29130
30-2-0-1	1400	1346 (620)	5086 (529)	7460	7970
30-2-0-2	1420	1360 (622)	5344 (551)	7400	8590
30-2-0-3	1500	1428 (616)	4188 (603)	4820	4910
30-2-3-1	4680	4420 (605)	5208 (221)	18360	17650
30-2-3-2	4800	3903 (520)	4120 (163)	21830	24470
30-2-3-3	5200	4896 (604)	5820 (194)	24030	25660
40-0-0-1	1400	1292 (623)	7656 (321)	64000	69340
40-0-0-2	1400	1292 (623)	8690 (346)	60500	60270
40-0-0-3	1400	1292 (625)	7765 (334)	70140	50050
40-0-3-1	4500	4243 (605)	7969 (149)	157340	113300
40-0-3-2	4500	4243 (603)	7629 (151)	149060	113670
40-0-3-3	4700	4080 (552)	8309 (147)	171480	130280
40-2-0-1	1500	1360 (608)	6174 (364)	15160	15340
40-2-0-2	1500	1278 (576)	3427 (133)	81730	78210
40-2-0-3	1500	1237 (560)	3998 (117)	91110	82450
40-2-3-1	4600	1373 (202)	3359 (56)	186310	127170
40-2-3-2	4600	3576 (498)	4896 (74)	177130	154750
40-2-3-3	4200	3359 (518)	3930 (89)	85420	87750
50-0-0-1	1400	1292 (626)	8187 (330)	66540	65010
50-0-0-2	1400	1278 (620)	8826 (352)	73420	64250
50-0-0-3	1400	1292 (624)	8064 (306)	75420	49720
50-0-3-1	4200	4025 (606)	6963 (138)	142350	111230
50-0-3-2	4300	4080 (604)	7534 (146)	143190	125800
50-0-3-3	4400	4148 (604)	7316 (145)	163200	114920
50-2-0-1	1500	1360 (611)	7956 (281)	67400	75030
50-2-0-2	1400	1319 (610)	8078 (297)	74340	72610
50-2-0-3	1500	1360 (614)	8404 (312)	66440	78000
50-2-3-1	4500	4066 (584)	6596 (122)	174430	116160
50-2-3-2	4000	3753 (605)	5521 (96)	164700	129830
50-2-3-3	4800	4528 (603)	6922 (121)	172670	130680

Table 6.6: GAs performance - Amount of iterations in 10 minutes.

	RS	generational model		steady state model	
		PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	-	-	-	5267	7599
30-0-0-2	-	-	-	8046	11174
30-0-0-3	-	-	-	11013	15438
30-0-3-1	-	-	-	31190	29854
30-0-3-2	-	-	-	27821	30805
30-0-3-3	-	-	-	22864	25715
30-2-0-1	-	-	-	4703	6079
30-2-0-2	-	-	-	4343	6693
30-2-0-3	-	-	-	3287	3968
30-2-3-1	-	-	-	16256	15049
30-2-3-2	-	-	-	18238	20237
30-2-3-3	-	-	-	21745	22994
40-0-0-1	-	-	-	62868	68570
40-0-0-2	-	-	-	59203	59222
40-0-0-3	-	-	-	69049	49167
40-0-3-1	-	-	-	155762	111778
40-0-3-2	-	-	-	146885	111803
40-0-3-3	-	-	-	169267	128575
40-2-0-1	-	-	-	13719	14178
40-2-0-2	-	-	-	78269	75347
40-2-0-3	-	-	-	87596	80357
40-2-3-1	-	-	-	184238	122417
40-2-3-2	-	-	-	171470	152128
40-2-3-3	-	-	-	82859	82768
50-0-0-1	-	-	-	65246	63942
50-0-0-2	-	-	-	72024	62893
50-0-0-3	-	-	-	74018	48799
50-0-3-1	-	-	-	141087	110428
50-0-3-2	-	-	-	141676	124813
50-0-3-3	-	-	-	161978	113756
50-2-0-1	-	-	-	66178	74005
50-2-0-2	-	-	-	73108	71386
50-2-0-3	-	-	-	65288	77084
50-2-3-1	-	-	-	172891	114196
50-2-3-2	-	-	-	160497	126058
50-2-3-3	-	-	-	170053	129504

Table 6.7: GAs performance - Amount of solutions rejected in 10 minutes.

	generational model			steady state model	
	RS	PMX	Direct3	Direct3-s	Direct4-s
30-0-0-1	-	-	-	74	78
30-0-0-2	-	-	-	37	39
30-0-0-3	-	-	-	44	37
30-0-3-1	-	-	-	23	9
30-0-3-2	-	-	-	11	19
30-0-3-3	-	-	-	22	22
30-2-0-1	-	-	-	58	41
30-2-0-2	-	-	-	67	46
30-2-0-3	-	-	-	83	65
30-2-3-1	-	-	-	36	23
30-2-3-2	-	-	-	39	8
30-2-3-3	-	-	-	25	23
40-0-0-1	-	-	-	45	24
40-0-0-2	-	-	-	54	12
40-0-0-3	-	-	-	55	35
40-0-3-1	-	-	-	50	42
40-0-3-2	-	-	-	34	40
40-0-3-3	-	-	-	25	12
40-2-0-1	-	-	-	50	33
40-2-0-2	-	-	-	66	18
40-2-0-3	-	-	-	50	27
40-2-3-1	-	-	-	24	41
40-2-3-2	-	-	-	15	1
40-2-3-3	-	-	-	39	14
50-0-0-1	-	-	-	48	41
50-0-0-2	-	-	-	47	17
50-0-0-3	-	-	-	63	24
50-0-3-1	-	-	-	30	28
50-0-3-2	-	-	-	46	24
50-0-3-3	-	-	-	42	34
50-2-0-1	-	-	-	56	8
50-2-0-2	-	-	-	34	30
50-2-0-3	-	-	-	55	62
50-2-3-1	-	-	-	27	6
50-2-3-2	-	-	-	28	5
50-2-3-3	-	-	-	27	6

Table 6.8: GAs performance - When the best solution was found (% of 10 minutes).

6.1.2 The effect of population size.

This section investigates the effects of different population sizes on the performance of GAs. The principle objective here is to reduce the loss of diversity by increasing the size of the population. Only one single direct GA is studied and, again, RS is used as the reference algorithm. Direct4-s has been chosen for these tests because this algorithm achieves the best performance on the hard problems where some work remains undone, i.e. under-resourced 30-* and critical 40-* problems.

This algorithm was run with three different population sizes: 100, 200 and 400 chromosomes. Therefore, the performance of four different runs are reported in the next tables: RS, Direct4-s:100, Direct4-s:200 and Direct4-s:400.

Analysis of the results in the previous section show that there is no instance where work is being traded for travel, i.e. a low cost solution is always better than a high cost solution (never does less work). Consequently, work done, travel or work not done are no longer reported. Hence, the data is only analysed from one perspective: final cost. Table 6.9 gives the average final cost reached by the algorithms on the different problems.

Direct4-s:200 always performs better than Direct4-s:100 and, in general, Direct4-s:400 performs better than Direct4-s:200. However, there are some noticeable exceptions where Direct4-s:200 achieves better solutions than Direct4-s:400. For instance, on problems with no technological or temporal constraints, i.e. 30-0-0-*, 40-0-0-* and 50-0-0-*, Direct4-s:200 always performs significantly better. On the last two groups (40-0-0-* and 50-0-0-*), even the “smallest” GA, Direct4-s:100, is reaching far better solutions than direct4-s:400. Intriguingly, in these two examples, the performances of Direct4-s:400 are comparable to the performances of RS.

An explanation for these peculiar results can be found in Table 6.10 which shows the amount of exploration performed by the four algorithms in 10 minutes. As

expected, Direct4-s:100 performs more iterations than Direct4-s:200 which in turn performs more iterations than Direct4-s:400. Clearly, larger populations require more time to create the initial population, and thus, less time is available for the actual search.

However, a linear relationship should be expected, such that if the population size is doubled, the time required to create the initial population should double as well. Unfortunately, in some cases, the relationship seems to be much worse than linear. In some instances, 10 minutes are not even sufficient for Direct4-s:400 to create its initial population. For instance, this is the case for all the 40-0-0-* and all the 50-0-0-* problems. Figure 6.4 shows the time in seconds that the schedule-builder needs in order to create an initial population free of duplicates for the 30-0-0-1 problem. The situation is even more acute for any of the 40-0-0-* or 50-0-0-* problems.

The last point “*free of duplicates*” is the main problem here. If duplicates were allowed, then the relationship would be linear. However, in order to limit the risk of premature convergence, duplicates are not accepted, and as the search goes on, it seems harder and harder for the schedule-builder to create a new, not yet visited, solution and this explains the exponential behaviour. After a certain stage, the schedule-builder is spending most of its time generating solutions which are rejected immediately.

Table 6.11 illustrates this problem; this table reports the percentage of schedules accepted when 100, 500, 1000, 2000 or 5000 schedules are created. For instance, when the schedule-builder creates 100 solutions for a 30-0-0-1 problem, 98% are accepted and 2% are rejected because they are duplicates. when the schedule-builder creates 5000 schedules only 33.9%, i.e. 1693 schedules are accepted.

Table 6.11 shows that the situation is particularly serious for problems with no constraints, i.e. 30-0-0-*, 40-0-0-*, and 50-0-0-* problems, and also for most of the 50-* problems with the noticeable exception of the 50-2-3-* problems.

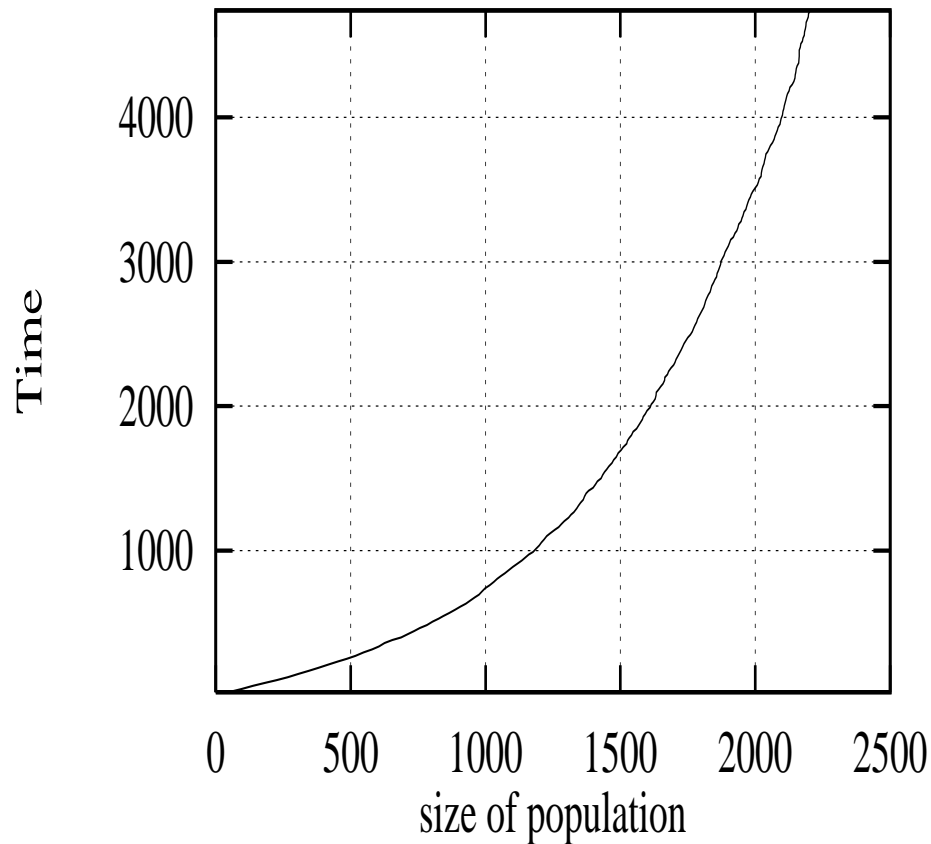


Figure 6.4: Population free of duplicates: the larger, the harder.

This phenomenon raises an important issue. It suggests a negative answer to the question “Could a GA benefit from a *smart* start?” In other words, if a GA uses an intelligent schedule-builder to generate its initial population, will the algorithm find better solutions?

For instance, several methods could be envisaged for the development of an intelligent schedule-builder:

- Jobs could be allocated to engineers in a round-robin manner, resulting in a more equitable distribution of jobs to engineers.
- Geographical information could be taken into account, such that tours radiate from the base, resulting in natural clusters of jobs.

- Some look-ahead techniques might be employed such that critical jobs and engineers can be identified and scheduled before all others.

In this study, a rather naive schedule-builder was used. As seen earlier, it takes an ordered list of engineers and allocate jobs to each engineer in turn. During this operation, not a single heuristic is used. However, this basic schedule-builder is already generating a large amount of duplicates and unfortunately, one should expect a more intelligent schedule-builder to be even more deterministic ⁴ and generate even more duplicates.

In these circumstances, it seems unlikely that a GA could benefit from a good start, i.e. if the *smart* schedule-builder is too deterministic. Clearly, with an intelligent schedule-builder, the different solution-candidates in the initial population tend to be similar - good quality solutions but similar solutions. The GA may converge rapidly towards a sub-optimal solution, thus losing the benefits of a good starting point. Obviously, the best option would be an intelligent schedule-builder which creates radically different solutions and thus maintains the level of diversity in the population.

Nevertheless, there might be the possibility of seeding the initial population; i.e., only a limited section of the population is generated by the intelligent method, the rest (say 80% for instance) is generated by the normal, i.e. naive, method. This approach might provide the best compromise, since the GA would have access to a limited set of good quality solutions but at the same time, the population would have a sufficient level of diversity.

In conclusion, there are two conflicting observations. First, the presence of duplicates in the population has to be avoided because it leads to a loss of diversity which, in turn, leads to a premature convergence towards a sub-optimal solution.

⁴Ultimately, the "*best*" schedule-builder would always generate the same solution, i.e. this solution would be the "optimal" solution.

On the other hand, if duplicates are rejected, then the GA might spend most of its time generating solutions which are rejected straight away.

In consequence, a GA might improve its performances by accepting duplicates despite the possible loss of diversity.

Two GAs which accept duplicates are studied in the remainder of this section:

1. The first GA only accepts duplicates when the initial population is being created, but rejects them during the rest of the search. This GA will be referred to as Dupli-1.
2. The second GA accepts duplicates at any time during the search. This GA will be referred to as Dupli-2.

Both Dupli-1 and Dupli-2 use the steady-state model and Crossover #4, they also both have a population of 400 chromosomes. Hence, direct comparisons with Direct4-s:100, Direct4-s:200 and Direct4-s:400 are possible.

Table 6.12 gives the average final cost reached by RS, Direct4-s:100, Direct4-s:200, Direct4-s:400, Dupli-1 and Dupli-2. Table 6.13 gives the ranking of these algorithms for the different problems and Figure 6.5 summarises this information. It gives the average score, the minimum and the maximum score for each algorithm over the entire set of problems; in this case, each algorithm can score a minimum of 0 point and a maximum of 5 points.

On average, Dupli-1 and Dupli-2 find the best solutions and, hence, have the best scores of all the algorithms: 4.05 for Dupli-1 and 3.05 for Dupli-2. Dupli-1 seems to be a more robust algorithm.

Of the 50-* problems, the 50-2-3-1 problem is the only instance (out of 12) where the best solution is not found by either Dupli-1 or Dupli-2. In the same manner, on the easy 40-* problems (40-* problems where all the work is done), the best

Algorithm	average	min	max
R S	0.0	0	0
Direct4-s:100	1.27	1	2
Direct4-s:200	3.11	2	5
Direct4-s:400	3.05	1	5
Dupli-1	4.05	2	5
Dupli-2	3.5	1	5

Figure 6.5: Different population sizes, different ranking.

solution is typically found either by Dupli-1 or by Dupli-2 as well. Therefore, “*accept duplicates*” appears to be a good policy on over-resourced problems.

The situation is more delicate for under-resourced problems. For instance, on the 30-* problems, Dupli-1 and Dupli-2 together can only manage 4 *pole-positions*. Generally, Dupli-2 is performing rather poorly on these problems. Direct4-s:200 is particularly effective on problems with no technological constraints and Direct4-s:400 works well on problems with tight technological constraints.

Table 6.7 (in the previous section) showed that the level of rejection was especially important in the case of over-resourced problems. In these situations, the GA was spending most of its search rejecting schedules. By accepting duplicates, the GA converges more rapidly, but at least, the algorithm is moving around in the search space. The level of rejection is not so important in under-resourced problems and these “*accept duplicates*” policies might not be desirable.

Table 6.14 shows (for the 6 algorithms) when the best solutions were found (in %) on the different problems. Observe that the results for Dupli-2 might be slightly misleading. Dupli-2 accepts duplicates at any time during the search, hence there can be several copies of the best solution in the population. In the worst case, if the population has converged, each member of the population is a copy of the best solution. Despite this, this table only indicates when the last copy of the best solution was created.

Table 6.14 shows that, typically, when the population size is increased, the algorithm finds its best solution at a later stage. This behaviour was expected. First, the generation of a larger initial population requires more computational effort, so the genetic search starts later. More importantly, a larger population helps to delay premature convergence and hence, the algorithm can use the extra-time to produce better solutions.

Dupli-1 exhibits a typical behaviour, it finds its best solutions earlier when solving problems with tight technological constraints. On problems with no such constraints, Dupli-1 finds its best solutions in the latter stages of the search, i.e. after 90-95% of the search. This suggests that this algorithm might benefit from longer runs; there is still scope for significant improvement at the end of the 10 minutes.

In conclusion, larger populations should provide a higher level of variety in the pool and, hence, help the GA to reach better solutions by preventing premature convergence. However, the limitations of the schedule-builder mean that the algorithm spends a prohibitive amount of time creating large populations unless duplicates are allowed in the population. This study also clearly suggests that a GA might not be able to benefit from a *smart* schedule-builder which affects the level of diversity in the population.

	RS	Direct4-s:100	Direct4-s:200	Direct4-s:400
30-0-0-1	975696	588086	526785	553823
30-0-0-2	775956	384688	319818	325189
30-0-0-3	617570	251543	175831	186656
30-0-3-1	763409	435282	374021	325324
30-0-3-2	815645	487509	451450	354138
30-0-3-3	1092715	793437	712334	651047
30-2-0-1	909188	649560	582873	642311
30-2-0-2	772461	608226	554164	584830
30-2-0-3	1499448	1409089	1380200	1396575
30-2-3-1	1364732	1180816	1123114	1083411
30-2-3-2	1154086	919782	851285	804407
30-2-3-3	909303	765036	750542	654992
40-0-0-1	2246	1335	1249	2201
40-0-0-2	2231	1338	1256	2217
40-0-0-3	2186	1322	1215	2157
40-0-3-1	2315	1582	1450	1406
40-0-3-2	2200	1527	1394	1338
40-0-3-3	2382	1640	1512	1495
40-2-0-1	445075	392486	370864	352839
40-2-0-2	54407	5436	1750	1841
40-2-0-3	56334	1984	1717	1719
40-2-3-1	22329	2176	1960	1860
40-2-3-2	2279	1954	1838	1708
40-2-3-3	191368	108409	46982	9122
50-0-0-1	2313	1374	1276	2276
50-0-0-2	2235	1422	1252	2189
50-0-0-3	2234	1280	1210	2183
50-0-3-1	2276	1625	1513	1465
50-0-3-2	2365	1654	1519	1469
50-0-3-3	2325	1633	1477	1452
50-2-0-1	2184	1776	1628	2165
50-2-0-2	2107	1560	1487	2049
50-2-0-3	2119	1581	1498	2075
50-2-3-1	2184	1891	1771	1642
50-2-3-2	2354	2049	1965	1876
50-2-3-3	2385	2058	1882	1835

Table 6.9: Average final cost for different population sizes.

	RS	Direct4-s:100	Direct4-s:200	Direct4-s:400
30-0-0-1	1400	9700	8150	4170
30-0-0-2	1400	13830	11450	4410
30-0-0-3	1400	18880	15900	6030
30-0-3-1	4400	42050	37800	32960
30-0-3-2	4400	38900	38340	30450
30-0-3-3	4600	29130	27140	22470
30-2-0-1	1400	7970	6660	3570
30-2-0-2	1420	8590	6970	3700
30-2-0-3	1500	4910	4100	2930
30-2-3-1	4680	17650	18470	15940
30-2-3-2	4800	24470	20830	19410
30-2-3-3	5200	25660	22380	21280
40-0-0-1	1400	69340	27660	0
40-0-0-2	1400	60270	25080	0
40-0-0-3	1400	50050	26550	0
40-0-3-1	4500	113300	89390	45640
40-0-3-2	4500	113670	85160	44770
40-0-3-3	4700	130280	83880	47090
40-2-0-1	1500	15340	11950	5930
40-2-0-2	1500	78210	54100	7570
40-2-0-3	1500	82450	42580	13220
40-2-3-1	4600	127170	99340	59430
40-2-3-2	4600	154750	92290	50350
40-2-3-3	4200	87750	80080	71420
50-0-0-1	1400	65010	30140	0
50-0-0-2	1400	64250	25650	0
50-0-0-3	1400	49720	25570	0
50-0-3-1	4200	111230	76650	28260
50-0-3-2	4300	125800	81400	34650
50-0-3-3	4400	114920	76610	34730
50-2-0-1	1500	75030	41920	0
50-2-0-2	1400	72610	31950	0
50-2-0-3	1500	78000	30160	0
50-2-3-1	4500	116160	104020	55650
50-2-3-2	4000	129830	99450	55770
50-2-3-3	4800	130680	93960	54070

Table 6.10: Amount of iterations for different population sizes.

	100	500	1000	2000	5000
30-0-0-1	98	82.4	70.4	55.15	33.86
30-0-0-2	97	83.4	71.2	54.55	34.2
30-0-0-3	95	83	71.6	55.3	33.86
30-0-3-1	99	92.4	88.1	79.15	61.66
30-0-3-2	99	92.4	89.2	79.05	60.64
30-0-3-3	98	91.6	85.5	76.35	55.58
30-2-0-1	99	90.2	84.1	72.8	50.7
30-2-0-2	96	89.6	82.7	70.5	49.92
30-2-0-3	98	90.4	84	72.7	52.92
30-2-3-1	99	96.6	94.8	90.55	78.42
30-2-3-2	99	97.8	94.7	90	77.06
30-2-3-3	98	96.2	94.3	89.05	77.26
40-0-0-1	82	45.4	28.3	16.25	7.42
40-0-0-2	81	48.4	29.8	17.4	7.96
40-0-0-3	84	45.2	27.9	15.8	7.48
40-0-3-1	97	87	79.5	68.1	50.94
40-0-3-2	97	91.6	85.5	74.5	56.64
40-0-3-3	96	90.6	84.4	73.8	56.34
40-2-0-1	99	95.8	90.7	84.5	69.34
40-2-0-2	99	94.4	90	81.15	62.8
40-2-0-3	97	95.2	91	84.45	68.9
40-2-3-1	98	96.8	94.6	89.9	78.4
40-2-3-2	97	97	94.2	87.4	75.42
40-2-3-3	99	97	93.8	88.25	73.9
50-0-0-1	84	45	28	16.7	8.04
50-0-0-2	88	46.8	29.3	17.7	8
50-0-0-3	84	45	29.1	16.5	7.48
50-0-3-1	86	50	32.3	20	10.96
50-0-3-2	85	52.2	35.6	22.4	13.3
50-0-3-3	87	53.8	35.9	23.25	13.26
50-2-0-1	85	47.8	27.9	16.25	7.7
50-2-0-2	81	44.6	27.9	16.5	7.74
50-2-0-3	79	43	27	15.7	7.12
50-2-3-1	97	90.6	84.1	71.4	52.24
50-2-3-2	98	95.4	93.2	84.5	67.24
50-2-3-3	98	90.8	84.3	73.95	56.14

Table 6.11: % of chromosomes accepted for different population sizes.

	RS	Direct4-s:100	Direct4-s:200	Direct4-s:400	Dupli-1	Dupli-2
30-0-0-1	975696	588086	526785	553823	535779	559222
30-0-0-2	775956	384688	319818	325189	308984	341434
30-0-0-3	617570	251543	175831	186656	193840	201057
30-0-3-1	763409	435282	374021	325324	301907	413671
30-0-3-2	815645	487509	451450	354138	379346	438829
30-0-3-3	1092715	793437	712334	651047	690690	744747
30-2-0-1	909188	649560	582873	642311	645932	636926
30-2-0-2	772461	608226	554164	584830	577598	550580
30-2-0-3	1499448	1409089	1380200	1396575	1387475	1387506
30-2-3-1	1364732	1180816	1123114	1083411	1088868	1126788
30-2-3-2	1154086	919782	851285	804407	802645	854868
30-2-3-3	909303	765036	750542	654992	691023	701913
40-0-0-1	2246	1335	1249	2201	1240	1150
40-0-0-2	2231	1338	1256	2217	1199	1123
40-0-0-3	2186	1322	1215	2157	1212	1135
40-0-3-1	2315	1582	1450	1406	1416	1402
40-0-3-2	2200	1527	1394	1338	1329	1397
40-0-3-3	2382	1640	1512	1495	1458	1424
40-2-0-1	445075	392486	370864	352839	345682	342030
40-2-0-2	54407	5436	1750	1841	1824	9033
40-2-0-3	56334	1984	1717	1719	1739	1832
40-2-3-1	22329	2176	1960	1860	1845	2042
40-2-3-2	2279	1954	1838	1708	1691	1839
40-2-3-3	191368	108409	46982	9122	27135	61440
50-0-0-1	2313	1374	1276	2276	1289	1139
50-0-0-2	2235	1422	1252	2189	1252	1174
50-0-0-3	2234	1280	1210	2183	1182	1115
50-0-3-1	2276	1625	1513	1465	1447	1423
50-0-3-2	2365	1654	1519	1469	1420	1443
50-0-3-3	2325	1633	1477	1452	1443	1442
50-2-0-1	2184	1776	1628	2165	1589	1527
50-2-0-2	2107	1560	1487	2049	1479	1352
50-2-0-3	2119	1581	1498	2075	1454	1436
50-2-3-1	2184	1891	1771	1642	1700	1725
50-2-3-2	2354	2049	1965	1876	1836	1912
50-2-3-3	2385	2058	1882	1835	1789	1851

Table 6.12: Average final cost for different population sizes and acceptance policies.

	RS	Direct4-s:100	Direct4-s:200	Direct4-s:400	Dupli-1	Dupli-2
30-0-0-1	0	1	5	3	4	2
30-0-0-2	0	1	4	3	5	2
30-0-0-3	0	1	5	4	3	2
30-0-3-1	0	1	3	4	5	2
30-0-3-2	0	1	2	5	4	3
30-0-3-3	0	1	3	5	4	2
30-2-0-1	0	1	5	3	2	4
30-2-0-2	0	1	4	2	3	5
30-2-0-3	0	1	5	2	4	3
30-2-3-1	0	1	3	5	4	2
30-2-3-2	0	1	3	4	5	2
30-2-3-3	0	1	2	5	4	3
40-0-0-1	0	2	3	1	4	5
40-0-0-2	0	2	3	1	4	5
40-0-0-3	0	2	3	1	4	5
40-0-3-1	0	1	2	4	3	5
40-0-3-2	0	1	3	4	5	2
40-0-3-3	0	1	2	3	4	5
40-2-0-1	0	1	2	3	4	5
40-2-0-2	0	2	5	3	4	1
40-2-0-3	0	1	5	4	3	2
40-2-3-1	0	1	3	4	5	2
40-2-3-2	0	1	3	4	5	2
40-2-3-3	0	1	3	5	4	2
50-0-0-1	0	2	4	1	3	5
50-0-0-2	0	2	3	1	4	5
50-0-0-3	0	2	3	1	4	5
50-0-3-1	0	1	2	3	4	5
50-0-3-2	0	1	2	3	5	4
50-0-3-3	0	1	2	3	4	5
50-2-0-1	0	2	3	1	4	5
50-2-0-2	0	2	3	1	4	5
50-2-0-3	0	2	3	1	4	5
50-2-3-1	0	1	2	5	4	3
50-2-3-2	0	1	2	4	5	3
50-2-3-3	0	1	2	4	5	3

Table 6.13: Ranking for different population sizes and acceptance policies.

	RS	Direct4-s:100	Direct4-s:200	Direct4-s:400	Dupli-1	Dupli-2
30-0-0-1	-	78	62	91	96	99
30-0-0-2	-	39	54	94	91	99
30-0-0-3	-	37	76	92	97	99
30-0-3-1	-	9	61	63	64	99
30-0-3-2	-	19	40	43	67	99
30-0-3-3	-	22	41	79	62	99
30-2-0-1	-	41	89	95	93	97
30-2-0-2	-	46	80	92	94	98
30-2-0-3	-	65	82	95	92	98
30-2-3-1	-	23	52	82	73	83
30-2-3-2	-	8	64	78	62	99
30-2-3-3	-	23	63	74	90	99
40-0-0-1	-	24	42	0	94	99
40-0-0-2	-	12	64	0	97	99
40-0-0-3	-	35	59	0	90	99
40-0-3-1	-	42	45	66	73	99
40-0-3-2	-	40	34	64	56	99
40-0-3-3	-	12	34	63	67	99
40-2-0-1	-	33	73	88	96	99
40-2-0-2	-	18	41	94	88	99
40-2-0-3	-	27	48	86	96	99
40-2-3-1	-	41	27	74	76	99
40-2-3-2	-	1	24	67	52	99
40-2-3-3	-	14	18	64	64	99
50-0-0-1	-	41	58	0	81	99
50-0-0-2	-	17	53	0	94	99
50-0-0-3	-	24	40	0	93	99
50-0-3-1	-	28	47	67	75	99
50-0-3-2	-	24	56	68	68	99
50-0-3-3	-	34	35	67	74	99
50-2-0-1	-	8	48	0	87	99
50-2-0-2	-	30	44	0	94	99
50-2-0-3	-	62	39	0	84	99
50-2-3-1	-	6	46	62	55	99
50-2-3-2	-	5	47	59	59	99
50-2-3-3	-	6	47	71	65	99

Table 6.14: When the best solution is found (in % of 10 minutes) for different population sizes and acceptance policies.

6.1.3 GA with repair algorithm.

In this section, a GA is coupled with a repair algorithm ⁵. In their original form, GAs are regarded as an efficient tool for detecting promising regions in large solution spaces. However, they might not be well equipped for climbing these promising peaks. Rather than merely exploring the search space, an efficient algorithm also needs to perform some sort of local improvement or repair. The association “*repair algorithm and GA*” might provide a better balance between exploration and exploitation; the GA explores the search space and the repair algorithm performs a local search in the neighbourhood of the solutions provided by the GA.

This hybrid algorithm follows the same cycle as the previous steady-state GAs. However, once the crossover operator has generated a new solution, this solution is transferred to the repair algorithm. Once the repair algorithm has completed its task, the newly modified solution is finally inserted in the population.

The main objective of the repair algorithm is to improve the quality of the schedules. This can be achieved in two different ways:

1. Reduce the total distance travelled by the engineers.
2. Reduce the amount of work not done. This option provides the best scope for improvement. A reduction in work not done will generate a significant reduction of cost.

The repair algorithm, adopted in this study, works as follows: first, the VRP is viewed as a set of independent TSPs and a tour improvement heuristic [BRIN94] is used to rearrange the different tours. The main motivation here, is not a mere reduction of travel, rather the main concern is to rearrange the tours, so that each engineer spends a minimum time travelling and a maximum time working. Once

⁵Observe that a repair algorithm is already implemented in each direct crossover. The repair algorithm, presented in this section, is a more complex algorithm.

this first stage is completed, the repair algorithm attempts to insert unallocated jobs in the schedule where the tour improvement heuristic has created some empty slots.

In this section, Direct4-s:200, Dupli-1 and Dupli-2 are coupled with the repair algorithm. Table 6.15 presents the final costs of six algorithms: Direct4-s:200, Direct4-s:200 with repair, Dupli-1, Dupli-1 with repair, Dupli-2 and finally Dupli-2 with repair.

Table 6.16 gives the ranking of these algorithms over the different problems; this information is summarised in Figure 6.6 which gives the average, the minimum and the maximum scores for each algorithm over the entire set of problems. In this case, one algorithm can score a maximum of 5 points and a minimum of 0 points for a given problem.

Algorithm	average	min	max
Direct4-s:200	1.38	0	5
Direct4-s:200 + repair	2.63	0	5
Dupli-1	2.75	1	5
Dupli-1 + repair	2.66	0	5
Dupli-2	2.22	0	5
Dupli-2 + repair	3.33	0	5

Figure 6.6: Ranking of GAs with or without repair algorithm.

Overall, the addition of the repair algorithm appears to be beneficial. The average score of Direct4-s:200 jumps from 1.38 to 2.63 when the repair algorithm is added (and from 2.22 to 3.33 for Dupli-2). However, Dupli-1 does not seem to benefit from this association, its average score drops from 2.75 down to 2.66. These last results should be considered carefully as the difference between the two scores is minimal.

Before the addition of the repair algorithm, the ranking between the algorithms was

Dupli-1 > Dupli-2 > Direct4-s:200 (where the relation “>” is read as “generally achieves a better score than”).

The repair algorithm modifies the previous ranking, which now becomes Dupli-2 with repair > Dupli-1 with repair > Direct4-s:200. This modification is due mainly to the fact that Dupli-1 does not take advantage of the repair algorithm.

Dupli-2 with repair has the best score overall and it performs well on 50-* problems. However, this algorithm finds relatively poor solutions for under-resourced problems.

With under-resourced and critical problems, i.e. all 30-* and some 40-* problems, there is no clear champion. Direct4-s:200 manages 3 pole-positions, Direct4-s:200 with repair also manages 3, Dupli- 1 manages 4, Dupli-1 with repair manages 6, Dupli-2 manages 1 and finally Dupli-2 with repair also manages 1 pole- position.

The repair algorithm is especially effective on the 50-* problems. Furthermore, Table 6.15 indicates that the repair algorithm is especially useful when solving problems with tight technological constraints. This result is quite unexpected. Indeed, technological constraints limit the possible rearrangements of the different tours and, thus a high level of technological constraint should reduce the scope for improvement.

Table 6.17 shows the amount of exploration performed by the different algorithms in 10 minutes. As expected, the addition of the repair algorithm drastically alters the computational behaviour of the different algorithms: the number of iterations performed in 10 minutes drops considerably when the repair algorithm is added. At each iteration, the repair algorithm rearranges all the tours of the current solutions and then tries to insert unallocated jobs in the schedule. These two operations require a significant amount of time, and this overhead means that the GA is performing fewer iterations. In other words, the hybrid algorithm is doing more exploitation and less exploration. As a direct consequence, a smaller number of

solutions are visited. However, it also means that each one of these iterations is likely to generate a better solution.

Finally, Table 6.18 shows for the 6 algorithms when the best solutions were found (in %) for the different problems. The principal observation is that, when coupled with a repair algorithm, a GA tends to find its best solution at a later stage. This is especially true for problems with tight technological constraints. This result suggests that these hybrid algorithms would also benefit from longer runs. Each iteration demands more computational effort and 10 minutes is apparently not a sufficient for the GA to perform sufficient exploration.

In general, the association GA-repair algorithm is fruitful and this is especially true for over-resourced problems. This hybrid algorithm has a different behaviour from a “*pure 100 %*” GA; it performs fewer iterations, i.e. less exploration and more exploitation. Results suggest that longer runs could be beneficial.

	Direct4-s 200	Direct4-s 200 repair	Dupli-1	Dupli-1 repair	Dupli-2	Dupli-2 repair
30-0-0-1	526785	553813	535779	615070	559222	609660
30-0-0-2	319818	321592	308984	393661	341434	364844
30-0-0-3	175831	173961	193840	251505	201057	231668
30-0-3-1	374021	348769	301907	312679	413671	377560
30-0-3-2	451450	373937	379346	361335	438829	424366
30-0-3-3	712334	712298	690690	685255	744747	706887
30-2-0-1	582873	584639	645932	683744	636926	696327
30-2-0-2	554164	550484	577598	595691	550580	601013
30-2-0-3	1380200	1396475	1387475	1419982	1387506	1430746
30-2-3-1	1123114	1124798	1088868	1097834	1126788	1106827
30-2-3-2	851285	806175	802645	788203	854868	851209
30-2-3-3	750542	718002	691023	669351	701913	716188
40-0-0-1	1249	1177	1240	1224	1150	1100
40-0-0-2	1256	1172	1199	1176	1123	1061
40-0-0-3	1215	1166	1212	1204	1135	1026
40-0-3-1	1450	1380	1416	1378	1402	1352
40-0-3-2	1394	1332	1329	1326	1397	1297
40-0-3-3	1512	1436	1458	1414	1424	1392
40-2-0-1	370864	352730	345682	378092	342030	370829
40-2-0-2	1750	1739	1824	1806	9033	1730
40-2-0-3	1717	1679	1739	1821	1832	1724
40-2-3-1	1960	2003	1845	1863	2042	1950
40-2-3-2	1838	1705	1691	1668	1839	1745
40-2-3-3	46982	64915	27135	23461	61440	46922
50-0-0-1	1276	1246	1289	1297	1139	1179
50-0-0-2	1252	1171	1252	1255	1174	1067
50-0-0-3	1210	1152	1182	1220	1115	1044
50-0-3-1	1513	1429	1447	1422	1423	1360
50-0-3-2	1519	1454	1420	1386	1443	1354
50-0-3-3	1477	1469	1443	1443	1442	1369
50-2-0-1	1628	1566	1589	1598	1527	1534
50-2-0-2	1487	1386	1479	1446	1352	1326
50-2-0-3	1498	1457	1454	1425	1436	1400
50-2-3-1	1771	1675	1700	1642	1725	1583
50-2-3-2	1965	1886	1836	1824	1912	1862
50-2-3-3	1882	1877	1789	1793	1851	1756

Table 6.15: Effect of repair algorithm - Average final cost.

	Direct4-s 200	Direct4-s 200 repair	Dupli-1	Dupli-1 repair	Dupli-2	Dupli-2 repair
30-0-0-1	5	3	4	0	2	1
30-0-0-2	4	3	5	0	2	1
30-0-0-3	4	5	3	0	2	1
30-0-3-1	2	3	5	4	0	1
30-0-3-2	0	4	3	5	1	2
30-0-3-3	1	2	4	5	0	3
30-2-0-1	5	4	2	1	3	0
30-2-0-2	3	5	2	1	4	0
30-2-0-3	5	2	4	1	3	0
30-2-3-1	2	1	5	4	0	3
30-2-3-2	1	3	4	5	0	2
30-2-3-3	0	1	4	5	3	2
40-0-0-1	0	3	1	2	4	5
40-0-0-2	0	3	1	2	4	5
40-0-0-3	0	3	1	2	4	5
40-0-3-1	0	3	1	4	2	5
40-0-3-2	1	2	3	4	0	5
40-0-3-3	0	2	1	4	3	5
40-2-0-1	1	3	4	0	5	2
40-2-0-2	3	4	1	2	0	5
40-2-0-3	4	5	2	1	0	3
40-2-3-1	2	1	5	4	0	3
40-2-3-2	1	3	4	5	0	2
40-2-3-3	2	0	4	5	1	3
50-0-0-1	2	3	1	0	5	4
50-0-0-2	1	4	2	0	3	5
50-0-0-3	1	3	2	0	4	5
50-0-3-1	0	2	1	4	3	5
50-0-3-2	0	1	3	4	2	5
50-0-3-3	0	1	2	3	4	5
50-2-0-1	0	3	2	1	5	4
50-2-0-2	0	3	1	2	4	5
50-2-0-3	0	1	2	4	3	5
50-2-3-1	0	3	2	4	1	5
50-2-3-2	0	2	4	5	1	3
50-2-3-3	0	1	4	3	2	5

Table 6.16: Effect of repair algorithm - Ranking of the algorithms.

	Direct4-s 200	Direct4-s 200 repair	Dupli-1	Dupli-1 repair	Dupli-2	Dupli-2 repair
30-0-0-1	8150	3170	4720	2020	4830	2050
30-0-0-2	11450	4040	5540	2280	6770	2550
30-0-0-3	15900	4830	5870	2530	9260	2810
30-0-3-1	37800	9990	32540	8110	37560	8900
30-0-3-2	38340	9770	34510	8230	35260	8750
30-0-3-3	27140	8760	24410	7280	25350	7560
30-2-0-1	6660	2860	3720	1930	4150	1940
30-2-0-2	6970	2970	3880	2030	4060	2090
30-2-0-3	4100	2010	3000	1560	2990	1560
30-2-3-1	18470	7080	17540	6210	17910	6640
30-2-3-2	20830	8120	19680	6830	21140	7300
30-2-3-3	22380	8920	21990	7350	23280	7890
40-0-0-1	27660	7770	9270	4790	72060	7000
40-0-0-2	25080	7590	7650	4490	60810	6740
40-0-0-3	26550	8040	8550	4640	65450	7260
40-0-3-1	89390	14280	45250	11830	150930	14630
40-0-3-2	85160	13780	42860	10940	153620	14570
40-0-3-3	83880	14260	49660	11900	153150	15100
40-2-0-1	11950	4480	6290	3020	8980	3020
40-2-0-2	54100	9890	10760	4680	50350	6530
40-2-0-3	42580	10280	8830	4590	63510	7630
40-2-3-1	99340	17590	62880	13850	151400	17240
40-2-3-2	92290	16530	55920	12460	149370	16340
40-2-3-3	80080	15340	72390	12280	83900	14230
50-0-0-1	30140	8160	10140	4810	60620	7140
50-0-0-2	25650	7740	8040	4650	64610	6530
50-0-0-3	25570	6940	9900	4790	57250	7310
50-0-3-1	76650	13770	44190	12020	132980	14510
50-0-3-2	81400	13730	45200	11510	132790	14510
50-0-3-3	76610	13860	44030	11810	134910	14800
50-2-0-1	41920	10290	14120	6420	71640	10990
50-2-0-2	31950	9800	13300	6100	74980	10270
50-2-0-3	30160	9420	14100	6710	84770	10660
50-2-3-1	104020	16950	55640	13410	140940	16620
50-2-3-2	99450	16230	48810	12710	126290	16210
50-2-3-3	93960	17190	56280	14010	137980	17000

Table 6.17: Effect of repair algorithm - Amount of iteration in 10 minutes.

	Direct4-s 200	Direct4-s 200 repair	Dupli-1	Dupli-1 repair	Dupli-2	Dupli-2 repair
30-0-0-1	62	94	96	78	99	93
30-0-0-2	54	87	91	82	99	97
30-0-0-3	76	94	97	88	99	95
30-0-3-1	61	84	64	97	99	99
30-0-3-2	40	56	67	93	99	99
30-0-3-3	41	57	62	95	99	99
30-2-0-1	89	92	93	92	97	89
30-2-0-2	80	95	94	92	98	96
30-2-0-3	82	94	92	95	98	77
30-2-3-1	52	65	73	91	83	99
30-2-3-2	64	61	62	93	99	99
30-2-3-3	63	69	90	86	99	99
40-0-0-1	42	87	94	91	99	99
40-0-0-2	64	86	97	93	99	96
40-0-0-3	59	71	90	90	99	99
40-0-3-1	45	60	73	92	99	99
40-0-3-2	34	72	56	91	99	99
40-0-3-3	34	73	67	91	99	99
40-2-0-1	73	94	96	96	99	96
40-2-0-2	41	77	88	95	99	99
40-2-0-3	48	70	96	97	99	99
40-2-3-1	27	67	76	83	99	99
40-2-3-2	24	59	52	90	99	99
40-2-3-3	18	43	64	90	99	99
50-0-0-1	58	75	81	93	99	99
50-0-0-2	53	84	94	91	99	99
50-0-0-3	40	76	93	95	99	99
50-0-3-1	47	74	75	86	99	99
50-0-3-2	56	56	68	94	99	99
50-0-3-3	35	71	74	85	99	99
50-2-0-1	48	80	87	91	99	99
50-2-0-2	44	65	94	93	99	99
50-2-0-3	39	85	84	88	99	99
50-2-3-1	46	75	55	80	99	99
50-2-3-2	47	65	59	89	99	99
50-2-3-3	47	73	65	86	99	99

Table 6.18: Effect of repair algorithm - When the best solution was found (in % of 10 minutes).

6.2 The 10 minute Beauty Contest.

This section examines and compares the performances of five algorithms; namely, genetic algorithms (GAs), hill-climbing (HC), random search (RS), simulated annealing (SA) and tabu search (TS) over the entire range of problems. In all trials the computational resource was limited to 10 SPARCstation IPC minutes per algorithm per problem. Hence, the term "*10 minute beauty contest*". The experiments reported in this section required over 150 hours of computational effort.

The five algorithms were implemented with the following parameters:

- RS used, again, as the reference algorithm. The results for this algorithm are the same as the results presented in the previous sections.
- Dupli-1 was chosen as the "*GA candidate*". Among all the GAs, Dupli-2 with repair had the best score overall; however, it found relatively poor solutions for the 30-*-* problems. Hence, Dupli-1 was chosen primarily because it finds slightly better solutions than the other GAs on under-resourced and critical problems. For information, this GA used the knowledge-based Crossover #4, the steady-state model and the exponential selection function; it had no mutation operator and operated with a population of 400 chromosomes. It accepted duplicates but only during the initial building of the population. Thereafter, duplicates were rejected. Finally, Dupli-1 was not coupled with the repair algorithm because of the length of the runs and to allow enough exploration.
- TS starts its search with an initial random solution created with the schedule-builder (this is also true for HC and SA). Then, the algorithm iterates. The move operator, described in the next paragraph, used by TS is common to the three neighbourhood search techniques. TS explores an entire neighbourhood at each iteration. Here, the size of a neighbourhood is $n + 1$ solutions

where there are n engineers and one conflict-list. In order to create the n th neighbour, TS takes the current solution and randomly selects a job J_i from the n th engineer's tour. Then, TS repeatedly selects engineers randomly until J_i fits in the selected engineer's tour. Once the $n + 1$ neighbours have been generated, the best solution in the neighbourhood (S_{best}) becomes the current solution. However, before being accepted, S_{best} must satisfy the restriction criteria which is to forbid moving jobs which have already been moved within the tabu tenure; here, 125 moves. This obligation is waived when S_{best} meets the aspiration criteria, i.e. the cost of S_{best} has to be less than the cost of the best solution found so far. Further information concerning the tuning of TS is available in [PROS94] and [BRIN94].

- In a similar manner to TS, HC generates $n + 1$ neighbours at each iteration around the current solution S . Once this has been completed, the best solution S_{best} in the neighbourhood is compared with S . The difference in cost between S and S_{best} is calculated. If S_{best} has a lower cost, the move is accepted and S_{best} becomes S , otherwise S remains unchanged. This cycle is then repeated until the time limit is reached.
- SA attempts a given number of moves at a temperature $Temp$, then reduces the temperature by a given amount. That is, $Temp = Temp \times R$, where R is the cooling rate. In the experiments performed, the initial temperature was set to $Temp = 10.0$, the cooling rate $R = 0.992$ and the number of successful moves at a given temperature was 500. These values provided a very smooth cooling schedule. The search process is terminated after 10 minutes. SA uses the following rules to accept or reject a neighbourhood solution S' :
 - if $cost(S') \leq cost(S)$ then S' is accepted as the new current solution.
 - if $cost(S') > cost(S)$ then the case is treated in a probabilistic manner.

In this case, the probability of acceptance is given by:

$$P(\Delta cost) = e^{-Temp/\Delta} \leq random(1.0)$$

where $Temp$ is the control parameter temperature, $\Delta cost = cost(S') - cost(S)$, and $random(1.0)$ generates a random number in the range 0.0 to 1.0, drawn from the uniform distribution. If $P(\Delta cost)$ is greater than the random number then S' is accepted as the new current solution, otherwise it is rejected and the original solution S remains the current solution.

The move used by TS, HC, and SA works as follows. A Job J_i is randomly selected. This job is then removed from the tour of the relevant engineer E_x and an attempt is made to insert the job into the tour of some other randomly selected engineer E_y , where E_y is technically able to do J_i . Note that E_y may be the conflict-list, i.e. E_y may be the virtual engineer, so long as J_i is not compulsory. The move may:

1. increase the number of jobs being done if E_x is the conflict-list.
2. modify, i.e. increase or decrease, the total travel by reducing the tour length of E_x and increasing the tour length of E_y .
3. decrease the amount of work done if E_y is the conflict list (the virtual engineer).

Moreover, if E_y is the conflict list, i.e. if the algorithm is moving a job J_i from E_x into the conflict-list, the algorithm immediately attempts to move a job J_j from the conflict-list back to engineer E_x . Moving J_i into the conflict-list will always generate a large increase in cost, and moving J_j back in the schedule will always generate a large reduction in cost. If job J_i has a longer duration than job J_j , the moves will most likely be rejected. Otherwise, the moves will (again, most likely) be accepted.

Table 6.19 presents the performances of the 5 algorithms over the entire range of problems and Table 6.20 presents the ranking of the algorithms. Figure 6.7 provides the average score, the minimum and the maximum for the five algorithms. In this case, each algorithm can score a maximum of 4 points and a minimum of 0 points.

Algorithm	average	min	max
GA	1.58	1	4
HC	1.91	0	3
RS	0.03	0	1
SA	3.88	2	4
TS	2.58	2	4

Figure 6.7: Ranking of the search techniques for 10 minute runs.

Figure 6.7 shows that SA, with an impressive average score of 3.88, dominates its opponents; SA is in fact only beaten on two occasions: on the 30-0-3-1 problem by GA and on the 30-2-3-3 problem by TS. GA performs relatively well on under-resourced problems; however, it achieves some rather poor results when applied to over-resourced problems. Typically, GA is only able to beat RS on these problems. This last remark explains the low score (1.58) of GA. In other words, GA is quite capable of packing jobs into an under-resourced schedule; however, it cannot reduce travel as well as the three neighbourhood search techniques. This means that when the reduction of travel becomes the only scope for improvement, i.e. on over-resourced problem, GA cannot compete with SA, TS, and HC.

The poor performance of GA on over-resourced problems has a number of causes. First, the GA adopted for this comparative study, Dupli-1, uses Crossover #4. This knowledge-based crossover is not recommended for the solution of over-resourced problems as it might misguide the search (see Subsection 6.1.1). Clearly, a GA using a random-based crossover (such as Crossover #3) would have achieved better results for these particular problems. Secondly, in an over-resourced problem, a significant proportion of the engineers are idle, and have empty tours. In a direct

representation when performing crossover, the GA may inherit many empty tours and this is of no real value, i.e. each time this happens, the GA is losing one opportunity of mixing good tours. Finally, once the inheritance phase is completed, the GA tries to bind each unallocated job to any possible engineer. It might be a better idea to allocate jobs first to existing tours and then if any job remains unallocated, consider idle engineers.

While HC is dominated by SA, GA and TS on under-resourced problems, it performed relatively well on over-resourced problems. However, the algorithm finds it difficult to maintain its level of performance when the level of constraints (both temporal and technological) increases, and on the 40-0-3-*, 40-2-0-*, 40-2-3-*, 50-2-3-*, typically HC only manages to beat RS and GA. TS has the second highest score (2.58) behind SA. Like the other neighbourhood search techniques, TS performs well on over-resourced problems and also performs fairly well on under-resourced problems, especially when the level of constraint is significant.

Table 6.21 shows the amount of exploration performed by the five algorithms. The GA column shows the number of crossovers and the RS column shows the number of schedules created in 10 minutes. The HC and TS column give the number of neighbourhoods examined; where a neighbourhood is the number of engineers in the problem plus one (for the conflict list). Therefore, the 50-* problems have a larger neighbourhood than the 30-* problems. The SA column gives the number of temperature reductions. There must be 500 successful moves before the temperature is reduced. Comparisons should be made downwards through the table not across rows.

The five algorithms fall into 3 categories: blind, neighbourhood and knowledge intensive techniques. RS is a blind technique, which does not use the history of the search in any way ⁶ and creates and evaluates a new solution from scratch at each

⁶The exception being that RS compares the current solution with the best solution found so far.

iteration. These two operations are extremely expensive in terms of computational effort and hence, RS can only visit a limited number of points in the search space. HC, SA, and TS are neighbourhood techniques. When generating new solutions these techniques only need to generate a simple move, i.e. the deletion of a job from an engineer's tour and the addition of that same job to the tour of another engineer, and evaluate the effect of that simple move. These two operations are relatively inexpensive and HC, SA, and TS can explore a large number of points in the search space (these points might be quite similar to each other).

Compared to the four other techniques, GA is a knowledge intensive technique which attempts to exploit the information, i.e. the genetic material, present in a pool of candidate-solutions. The algorithm combines two solutions into a child and evaluates that child. These two operations are a few orders of magnitude more expensive than the operations performed by HC, SA, and TS. Therefore, GA performs far fewer iterations, i.e. less exploration but more exploitation, than the three neighbourhood techniques.

When creating the child, GA does not have to re-schedule the entire solution, parts of both parents are inherited without any alteration. Therefore, GA performs more iterations than RS. Moreover, more exploration takes places over problems with high specialisation. This phenomenon is observed with the five algorithms. Also, when moving from under-resourced to over-resourced, the five algorithms performed significantly more iterations.

Table 6.22 shows when the different algorithms find their best solution (in %) over the entire set of problems. When moving from under-resourced to over-resourced problems, the different algorithms find their best solution earlier in the 10 minutes; they also find their best solution earlier when applied to problems with high specialisation. Clearly, on these two classes of problems, the algorithms are performing more iterations and reach their best solutions more rapidly. In general, GA finds

its best solution at a later stage than the other algorithms ⁷. This again suggests that GA could benefit from longer runs.

In conclusion, given 10 minutes, SA is the best technique overall. In an environment where the algorithm is given only a limited amount of time, SA is a robust technique which performs well both on over-resourced and under-resourced problems. GA works fairly well on under-resourced problems; however, it is outperformed by the neighbourhood search techniques (SA, HC, and TS) on over-resourced problems. This is due mainly to the fact that GA does not *reduce travel* as well as SA, HC, and TS.

⁷This might be partly due to the fact that GA has to create an initial population before starting its search

	GA	HC	RS	SA	TS
30-0-0-1	535779	580803	975696	440246	559202
30-0-0-2	308984	399073	775956	258495	348610
30-0-0-3	193840	244296	617570	107333	206448
30-0-3-1	301907	449678	763409	337981	337968
30-0-3-2	379346	498310	815645	350548	379354
30-0-3-3	690690	778948	1092715	638426	676248
30-2-0-1	645932	604391	909188	505374	550357
30-2-0-2	577598	561250	772461	465438	492380
30-2-0-3	1387475	1399526	1499448	1354369	1368904
30-2-3-1	1088868	1225485	1364732	1052343	1077584
30-2-3-2	802645	924950	1154086	770189	793561
30-2-3-3	691023	760957	909303	697780	672561
40-0-0-1	1240	920	2246	768	1157
40-0-0-2	1199	897	2231	785	1176
40-0-0-3	1212	934	2186	761	1079
40-0-3-1	1416	1310	2315	963	1172
40-0-3-2	1329	1166	2200	889	1078
40-0-3-3	1458	1305	2382	998	1199
40-2-0-1	345682	335799	445075	254771	258469
40-2-0-2	1824	11894	54407	898	1040
40-2-0-3	1739	1109	56334	925	1072
40-2-3-1	1845	55402	22329	1155	1218
40-2-3-2	1691	1320	2279	1021	1056
40-2-3-3	27135	157962	191368	13802	46245
50-0-0-1	1289	990	2313	818	1248
50-0-0-2	1252	900	2235	817	1170
50-0-0-3	1182	825	2234	770	1097
50-0-3-1	1447	1176	2276	941	1216
50-0-3-2	1420	1218	2365	974	1199
50-0-3-3	1443	1166	2325	972	1172
50-2-0-1	1589	1067	2184	897	1080
50-2-0-2	1479	875	2107	718	971
50-2-0-3	1454	967	2119	826	1010
50-2-3-1	1700	1244	2184	1000	1079
50-2-3-2	1836	1374	2354	1136	1189
50-2-3-3	1789	1304	2385	1109	1154

Table 6.19: 10 minute runs - Average final cost.

	GA	HC	RS	SA	TS
30-0-0-1	3	1	0	4	2
30-0-0-2	3	1	0	4	2
30-0-0-3	3	1	0	4	2
30-0-3-1	4	1	0	2	3
30-0-3-2	3	1	0	4	2
30-0-3-3	2	1	0	4	3
30-2-0-1	1	2	0	4	3
30-2-0-2	1	2	0	4	3
30-2-0-3	2	1	0	4	3
30-2-3-1	2	1	0	4	3
30-2-3-2	2	1	0	4	3
30-2-3-3	3	1	0	2	4
40-0-0-1	1	3	0	4	2
40-0-0-2	1	3	0	4	2
40-0-0-3	1	3	0	4	2
40-0-3-1	1	2	0	4	3
40-0-3-2	1	2	0	4	3
40-0-3-3	1	2	0	4	3
40-2-0-1	1	2	0	4	3
40-2-0-2	2	1	0	4	3
40-2-0-3	1	2	0	4	3
40-2-3-1	2	0	1	4	3
40-2-3-2	1	2	0	4	3
40-2-3-3	3	1	0	4	2
50-0-0-1	1	3	0	4	2
50-0-0-2	1	3	0	4	2
50-0-0-3	1	3	0	4	2
50-0-3-1	1	3	0	4	2
50-0-3-2	1	2	0	4	3
50-0-3-3	1	3	0	4	2
50-2-0-1	1	3	0	4	2
50-2-0-2	1	3	0	4	2
50-2-0-3	1	3	0	4	2
50-2-3-1	1	2	0	4	3
50-2-3-2	1	2	0	4	3
50-2-3-3	1	2	0	4	3

Table 6.20: 10 minute runs - Ranking of the search techniques.

	GA	HC	RS	SA	TS
30-0-0-1	4720	3240	1400	190	3140
30-0-0-2	5540	3280	1400	190	3100
30-0-0-3	5870	3200	1400	190	3100
30-0-3-1	32540	10840	4400	660	10400
30-0-3-2	34510	10520	4400	652	10120
30-0-3-3	24410	11040	4600	648	10360
30-2-0-1	3720	3560	1400	190	3200
30-2-0-2	3880	4040	1420	230	5120
30-2-0-3	3000	7440	1500	458	6760
30-2-3-1	17540	12600	4680	758	12360
30-2-3-2	19680	11240	4800	626	10240
30-2-3-3	21990	14960	5200	930	14600
40-0-0-1	9270	17840	1400	1228	20180
40-0-0-2	7650	19120	1400	1272	21260
40-0-0-3	8550	17320	1400	1200	19600
40-0-3-1	45250	22560	4500	1548	22040
40-0-3-2	42860	22080	4500	1506	21120
40-0-3-3	49660	23240	4700	1544	21540
40-2-0-1	6290	11560	1500	1162	10720
40-2-0-2	10760	11640	1500	1240	12940
40-2-0-3	8830	12600	1500	1186	13060
40-2-3-1	62880	18600	4600	1592	18340
40-2-3-2	55920	19000	4600	1556	17920
40-2-3-3	72390	16320	4200	1386	15220
50-0-0-1	10140	23200	1400	1586	22540
50-0-0-2	8040	24400	1400	1588	22920
50-0-0-3	9900	25640	1400	1610	23700
50-0-3-1	44190	28040	4200	1842	23940
50-0-3-2	45200	27120	4300	1858	23300
50-0-3-3	44030	26240	4400	1914	23100
50-2-0-1	14120	18720	1500	1758	16000
50-2-0-2	13300	24560	1400	1802	19960
50-2-0-3	14100	20720	1500	1748	17320
50-2-3-1	55640	22520	4500	2038	19640
50-2-3-2	48810	21040	4000	2006	18700
50-2-3-3	56280	22600	4800	2050	20240

Table 6.21: 10 minute runs - Amount of exploration in 10 minutes.

	GA	HC	RS	SA	TS
30-0-0-1	96	85	-	95	86
30-0-0-2	91	72	-	93	83
30-0-0-3	97	72	-	91	70
30-0-3-1	64	21	-	46	58
30-0-3-2	67	7	-	50	65
30-0-3-3	62	21	-	52	81
30-2-0-1	93	51	-	72	74
30-2-0-2	94	20	-	100	81
30-2-0-3	92	53	-	85	69
30-2-3-1	73	13	-	46	81
30-2-3-2	62	6	-	45	60
30-2-3-3	90	7	-	38	64
40-0-0-1	94	41	-	61	15
40-0-0-2	97	65	-	50	24
40-0-0-3	90	73	-	47	39
40-0-3-1	73	23	-	35	38
40-0-3-2	56	50	-	31	46
40-0-3-3	67	70	-	28	54
40-2-0-1	96	94	-	47	48
40-2-0-2	88	15	-	41	47
40-2-0-3	96	27	-	51	54
40-2-3-1	76	23	-	26	70
40-2-3-2	52	36	-	28	66
40-2-3-3	64	16	-	31	86
50-0-0-1	81	59	-	37	15
50-0-0-2	94	63	-	43	22
50-0-0-3	93	63	-	70	20
50-0-3-1	75	49	-	29	24
50-0-3-2	68	77	-	26	33
50-0-3-3	74	74	-	27	36
50-2-0-1	87	19	-	30	24
50-2-0-2	94	24	-	43	17
50-2-0-3	84	16	-	38	16
50-2-3-1	55	27	-	22	46
50-2-3-2	59	31	-	26	51
50-2-3-3	65	45	-	21	41

Table 6.22: 10 minute runs - When the best solution is found (in % of 10 minutes).

6.3 CPU Intensive Experiments

This section analyses the performance of GA, HC, RS, SA, and TS when these algorithms are given significantly more time to explore the search space. The main motivation here is to check if the algorithms can benefit from these longer runs. Furthermore, these intensive runs provide the opportunity to verify the quality of the solutions provided by the short runs, i.e. is the difference between the solutions provided by the short runs and the intensive runs significant? Is the extra computational effort worthwhile?

RS and TS remain unchanged for these runs; their parameters keep the same values as in Section 6.2. The *GA candidate* is still Dupli-1 but it has a larger population of 2000 chromosomes rather than 400. For SA, the number of successful moves at a given temperature is set to 12000 moves instead of 500.

These intensive runs were performed on a DEC Alpha 3000/300 and each algorithm was given one hour on each problem. As previously, the algorithm was applied 5 times to each problem. Hence, to apply a single algorithm to all 36 problems took approximately 180 hours. These intensive runs were performed with 5 algorithms and took approximately 900 hours which is equivalent to 3600 hours on a SPARC-station IPC.

Algorithm	average	min	max	average - 10 min
GA	2.27	1	4	1.58
HC	1.55	0	3	1.91
RS	0.30	0	1	0.03
SA	3.6	2	4	3.88
TS	2.25	1	4	2.58

Figure 6.8: Ranking of the search techniques for 4 hour runs.

Table 6.23 presents the performances of the 5 algorithms over the entire range of

problems and Table 6.24 presents the ranking of the algorithms. Figure 6.8 provides the average, the minimum and the maximum scores for the five algorithms. In this case, each algorithm can score a maximum of 4 points and a minimum of 0 points.

Table 6.23 shows that GA, RS, SA, and TS use the extra time to improve their results significantly. For the short runs, the ranking between the different algorithms was $SA > TS > HC > GA > RS$. With the intensive runs, this ranking is modified and becomes $SA > GA \geq TS > HC > RS$.

SA remains the best algorithm overall. However, GA now dominates the other competitors on under-resourced problems and reduces the gap with (SA, HC, and TS) on over-resourced problems considerably. In fact, GA and RS are the only two techniques which improve their average score. GA goes from 1.58 up to 2.27 and RS goes from 0.03 up to 0.3. HC goes from 1.91 down to 1.55, SA goes from 3.88 down to 3.6 and TS goes from 2.58 down to 2.25.

Table 6.25 shows the amount of exploration performed by the different algorithms. A comparison between Table 6.25 and Table 6.21 (which reports the same data for the short runs) underlines the considerable effort behind each intensive run. For instance, on the 30-0-0-1 problem, RS visited 33030 solutions whereas in the short run, it only had time to generate 1400 solutions. GA performed 109256 crossovers (only 4720 for the short run), HC performed 79200 cycles (only 3240 for the short run) and TS did 60600 cycles (only 3140 for the short run). This information is summarised in Figure 6.9.

Runs	GA	HC	RS	TS
short	4720	3240	1400	3140
intensive	109256	79200	33030	60600

Figure 6.9: Computational effort - Comparison between short and intensive runs.

The situation is more delicate for SA. The algorithm is performing fewer iterations

during these intensive runs. For instance, on the 30-0-0-1 problem, SA performs 142 iterations whereas in the short run, it managed 190 iterations. This is due to the fact that the number of successful moves at a given temperature went from 500 for the short runs up to 12000 for the intensive runs.

In other words, the main difference for SA between short and intensive runs is that the algorithm is given much more time to reach equilibrium at each temperature; the function used to reduce the temperature is identical in both sets of runs. Some theoretical results (see Subsection 3.3.3) suggest that this annealing schedule might not be the best option, i.e. there is no requirement to attain equilibrium at a succession of reducing temperatures, rather the main condition for a good annealing schedule is for the cooling to be carried out sufficiently slowly.

A number of alternatives for the annealing schedule can be envisaged. For instance, SA could perform only one iteration at each temperature and reduce the temperature very slowly or, at the other end of the spectrum, SA could keep the temperature constant during the search. In this case, SA would be a random search which accepts uphill moves with the probability $\exp(-C\Delta)$ and C is kept constant during the search.

In conclusion, there is no overall champion among the five algorithms which all benefit from the longer runs. Typically, GA takes advantage of the extra computational effort to provide the best solutions on under-resourced problems. Again, SA is a robust technique which performs well both on under-resourced and over-resourced problems.

	GA	HC	RS	SA	TS
30-0-0-1	418611	744850	950471	389783	433014
30-0-0-2	191802	521610	747099	193632	240439
30-0-0-3	91073	372232	585131	107356	114528
30-0-3-1	255015	597360	720178	314549	273084
30-0-3-2	283818	618848	770597	337931	312673
30-0-3-3	611376	924910	1060289	609596	605968
30-2-0-1	483638	757418	862319	489156	499902
30-2-0-2	485557	649270	740012	449042	459864
30-2-0-3	1327704	1520217	1483255	1339915	1345469
30-2-3-1	1016668	1382046	1314405	1041505	1054138
30-2-3-2	725090	1185889	1112688	750395	739479
30-2-3-3	608013	915757	846271	641898	645515
40-0-0-1	914	857	2203	696	1095
40-0-0-2	872	811	2186	687	1146
40-0-0-3	843	738	2136	671	1053
40-0-3-1	1122	1006	2245	868	1114
40-0-3-2	1090	1063	2148	821	1022
40-0-3-3	1156	1042	2341	913	1150
40-2-0-1	271441	557247	416225	238547	233278
40-2-0-2	1350	74777	27472	807	924
40-2-0-3	1312	91005	20370	853	1004
40-2-3-1	1556	55287	2494	1063	1085
40-2-3-2	1484	13784	2196	953	981
40-2-3-3	1733	278488	148155	1084	1164
50-0-0-1	945	849	2248	727	1175
50-0-0-2	895	750	2163	714	1163
50-0-0-3	862	719	2171	673	1122
50-0-3-1	1161	1030	2234	876	1176
50-0-3-2	1139	1033	2284	905	1158
50-0-3-3	1150	965	2260	889	1116
50-2-0-1	1260	919	2151	883	1063
50-2-0-2	1073	784	2032	699	962
50-2-0-3	1169	904	2074	818	999
50-2-3-1	1377	1092	2141	939	971
50-2-3-2	1561	10286	2333	1055	1099
50-2-3-3	1498	1271	2337	1028	1061

Table 6.23: 4 hour runs - Average final cost.

	GA	HC	RS	SA	TS
30-0-0-1	3	1	0	4	2
30-0-0-2	4	1	0	3	2
30-0-0-3	4	1	0	3	2
30-0-3-1	4	1	0	2	3
30-0-3-2	4	1	0	2	3
30-0-3-3	2	1	0	3	4
30-2-0-1	4	1	0	3	2
30-2-0-2	2	1	0	4	3
30-2-0-3	4	0	1	3	2
30-2-3-1	4	0	1	3	2
30-2-3-2	4	0	1	2	3
30-2-3-3	4	0	1	3	2
40-0-0-1	2	3	0	4	1
40-0-0-2	2	3	0	4	1
40-0-0-3	2	3	0	4	1
40-0-3-1	1	3	0	4	2
40-0-3-2	1	2	0	4	3
40-0-3-3	1	3	0	4	2
40-2-0-1	2	0	1	3	4
40-2-0-2	2	0	1	4	3
40-2-0-3	2	0	1	4	3
40-2-3-1	2	0	1	4	3
40-2-3-2	2	0	1	4	3
40-2-3-3	2	0	1	4	3
50-0-0-1	2	3	0	4	1
50-0-0-2	2	3	0	4	1
50-0-0-3	2	3	0	4	1
50-0-3-1	2	3	0	4	1
50-0-3-2	2	3	0	4	1
50-0-3-3	1	3	0	4	2
50-2-0-1	1	3	0	4	2
50-2-0-2	1	3	0	4	2
50-2-0-3	1	3	0	4	2
50-2-3-1	1	2	0	4	3
50-2-3-2	2	0	1	4	3
50-2-3-3	1	2	0	4	3

Table 6.24: 4 hour runs - Ranking of the search techniques.

	GA	HC	RS	SA	TS
30-0-0-1	109256	79200	33030	142	60600
30-0-0-2	157880	79120	33170	138	66200
30-0-0-3	229604	75920	33020	138	65200
30-0-3-1	262948	278480	108310	458	230400
30-0-3-2	291840	281040	107060	516	214800
30-0-3-3	244428	282000	111510	412	231200
30-2-0-1	97960	90560	33650	134	69600
30-2-0-2	103164	127600	34430	224	126800
30-2-0-3	54596	210080	36640	370	149000
30-2-3-1	193360	415840	114750	654	256000
30-2-3-2	228748	371840	117620	530	232600
30-2-3-3	224460	432960	128630	762	318600
40-0-0-1	365912	395680	32500	908	535800
40-0-0-2	362880	418560	33060	1036	565200
40-0-0-3	356580	374000	32750	932	527000
40-0-3-1	439100	512480	109030	1196	589200
40-0-3-2	447996	468080	109530	1180	563800
40-0-3-3	432684	533840	115310	1210	574600
40-2-0-1	174048	260160	35370	930	285000
40-2-0-2	409540	265760	35040	1000	347400
40-2-0-3	441248	240000	35090	952	332600
40-2-3-1	559396	444240	112230	1176	491200
40-2-3-2	515452	408720	112340	1132	471800
40-2-3-3	537916	409600	102640	1052	414800
50-0-0-1	361508	502400	32170	1216	583000
50-0-0-2	351576	531840	32140	1242	589200
50-0-0-3	326192	550880	32500	1278	612800
50-0-3-1	435360	622560	102480	1374	609000
50-0-3-2	442340	599040	104440	1382	622800
50-0-3-3	446072	603680	107210	1334	628400
50-2-0-1	477620	414800	34920	1324	426400
50-2-0-2	483448	540000	33800	1362	496600
50-2-0-3	479948	398560	34690	1306	425400
50-2-3-1	543000	513520	107820	1466	522200
50-2-3-2	546856	453040	96760	1456	489200
50-2-3-3	523348	521280	116650	1426	537800

Table 6.25: 4 hour runs - Amount of exploration.

Chapter 7

Conclusion.

This chapter summarises the work that has been presented, discusses areas of future work, and finally concludes this thesis.

7.1 Summary.

The growing complexity of telecommunications networks means that their management is also becoming more and more complicated. Due to several factors, e.g. amount of data, inconsistent and incoherent data, and time constraints, Artificial Intelligence techniques will be essential in order to carry out some of the management functions.

This thesis focuses on one particular aspect of network management: resource management. The key-motivation behind this work is two-fold:

- It is essential for the telecommunications operators to solve their resource allocations problems in an efficient manner.
- Many classical techniques are inadequate because, due to the sheer size of the

problems encountered, they cannot provide an answer in reasonable time.

Genetic algorithms have been proposed as a possible alternative. However, traditional indirect GAs are not suitable in many cases because they are too slow or too disruptive. To answer these limitations, this thesis proposes a new class of GAs which works directly on the chromosomes. In other words, the solutions are the chromosomes, the GA is no longer working with an indirect representation.

This thesis offers three large empirical studies which observe the performance of these direct GAs, and compare their performance against those of other search techniques in the context of workforce management. A set of thirty six problems is used for these observations and comparisons; each problem is radically different from the rest. More precisely, the effects of three parameters have been studied: work load, time windows, and specialisation constraints. A problem can be under-resourced, critically-resourced, or over-resourced, time windows can either be loose or tight, and specialisation constraints can also be either loose or tight. These three parameters and their possible values mean that there are twelve different configurations, and for each configuration, there are three randomly generated problems. Hence, the thirty six problems. Empirical evidence shows that *work load* is a key-parameter, which decides how the algorithms perform, i.e. determines whether all the work will be done or some will remain undone.

The two chromosomal representations, direct and indirect, are compared, and it is shown that the direct GAs totally dominate the indirect GA, both in terms of quality of solution and amount of exploration. The average costs achieved by the indirect GA are in fact remarkably similar to those achieved by RS and clearly, the association of *order-based GA and schedule-builder* is not suitable for this type of problem; the indirect GA is not guiding its search properly. In general, the steady state model performs better than the generational model.

When the problem under study is over-resourced, a more adaptable random-based

crossover operator should be chosen; however, for an under- or critically-resourced problem, the knowledge-based crossover which tries to maximise the amount of work done should be preferred. This study also shows that increasing population size and diversity helps the search process, but maintaining diversity within the population is no easy task. As the population size increases, the probability of producing duplicates increases, and the efforts to produce novel chromosomal material can be expensive. This also suggests that producing the initial population via a *smart* schedule-builder may lead to a loss of diversity and a poor GA performance. A possible alternative may be to seed the initial population, the algorithm will then have access to a limited set of above-average solutions and at the same time, the level of diversity will not be too severely affected.

Then, the GA is equipped with a repair algorithm. For each child created, the repair algorithm uses a TSP heuristic to improve the different tours and then tries to insert un-allocated jobs in the schedule, where the heuristic has created some empty slots. In general, this association of GA and repair algorithm is fruitful, particularly on over-resourced problems. Due to the increased computational effort, some exploration is lost. This suggests that GAs may benefit from extended runs.

In the 10 minute beauty contest, SA dominates the other competitors, performing well on all problems. GA works fairly well on under-resourced problems; however, it is outperformed by HC, SA, and TS on over-resourced problems; mainly because GA cannot "reduce travel" as well as its three opponents. Again, the results suggest that GA would benefit from extended runs.

The CPU intensive experiments (4 hour CPU instead of 10 minutes) show that indeed GA benefits from more CPU time, and on under-resourced problems, GA generally outperforms all of the other techniques. Again, SA is a robust technique which performs well both with under-resourced and over-resourced problems. The fact that SA is more adaptable than GA should not be unexpected. The internal algorithm of GA has been designed to maximise the amount of work done, this

complex internal algorithm is not as adaptable as the simpler SA.

On a more general note, when comparing search techniques, the amount of time spent implementing them should be taken into consideration. In this study, the implementation and the tuning of SA took less time - several weeks less - than for GA. Hence, with respect to that particular criterion, SA is again dominant. Moreover, HC, RS, SA and TS all require much less memory than GA. While these four algorithms can work with a minimum of two solutions, i.e. the current solution and the best solution, GA requires an entire population. This situation can only become worse when the size of the problem increases.

Therefore, from this study, it appears that in many cases, SA should be preferred. For instance, this study advocates SA when the algorithm is only given a limited amount of time, i.e. in this case, 10 minutes, or when the VRP is over-resourced.

On the other hand, GA with a knowledge-based crossover should be adopted - when CPU time is available - for the solution of under-resourced VRPs. Out of the five search techniques studied in this thesis, direct GA is the most knowledge-based. On hard problems, and when given enough time, GA uses this extra information to provide the best solutions. This suggests that direct GA, using the correct parameters (e.g. cost, profit) to preserve the best parts of the solutions, could be used to solve highly constrained problems. SA, and also TS, would be preferable for lightly constrained problems.

7.2 Future work.

This work can now proceed in a number of ways:

- Apply stochastic techniques to other more telecommunications-oriented problem-domains

- Consider several objective functions rather than only one
- Consider dynamic problems rather than static problems
- Consider a distributed and cooperative approach
- Integrate the search algorithms with a graphical user interface

7.2.1 Stochastic techniques and other problem-domains.

First, it may be preferable to mention where these stochastic techniques will not be of any use. At the lowest levels in the management hierarchy, AI techniques will be used on-line in order to accelerate the control process and reduce human intervention. This will not be achieved with GA, HC, RS, SA or TS. These techniques are iterative and thus, they cannot be considered for such fast control tasks. Rather, neural networks may be the answer with their parallel architecture and hardware implementations.

Stochastic techniques will be in demand at the highest levels in the management hierarchy, where they will be used to solve different classes of resource allocation problems: not only work force management problems but also frequency allocation problems, design and dimensioning problems, or virtual path management problems in ATM networks.

For instance, a heuristic algorithm has been proposed to solve the virtual path management problem [LOGO92]. This suggests that stochastic techniques such as GA or SA could also be applied to this particular problem.

The next paragraphs define the virtual path management problem and show how stochastic techniques could be used to tackle this particular problem.

ATM is the multiplexing and switching technology adopted by ITU-TSS for the future Broadband Integrated Services Digital Network (BISDN). ATM is a connection-

oriented technology, i.e. no information may be transferred between the users until a connection is established. Such a connection in ATM is referred to as a virtual channel (VC). In turn, these VCs are grouped into virtual paths (VPs). To be precise, a VP is a bundle of VCs that have the same end-points. Virtual paths provides a flexible logical network on top of the physical network.

Virtual path management is responsible for creating, maintaining, releasing, and modifying the network of virtual paths. It becomes necessary to modify the existing virtual paths when some of them experience unacceptably high levels of call rejection. Typically, two strategies can be used:

- If only a small modification is required, then only the bandwidths of the VPs are modified. This method is called bandwidth re-allocation. The virtual path manager re-allocates dynamically the bandwidth among the different paths.
- If a larger modification is necessary, the routes of the VPs also have to be modified. It then becomes a routing problem.

These two functions - bandwidth re-allocation and routing - can be performed at the network level by a central operations system. This operations system collects traffic load and VP (bandwidths and routes) information for the entire network, applies a heuristic to this information and generates the appropriate control actions. Then, the new routing tables and/or the new bandwidth allocations are down-loaded to the different nodes of the network. Stochastic techniques working as meta-heuristics could be associated with this heuristic in order to guide its search and improve its performance.

7.2.2 The objective function is a parameter.

As already mentioned in Section 3.1, The objective function plays an essential part in the search by judging the quality of the candidate solutions and hence, by guiding the search process in its exploration. Management problems can have many - possibly conflicting - objectives: minimise the worst blocking probability, maintain fairness, balance load, maximise throughput, maintain robustness, minimise delays. Therefore, a future study should consider the objective function as a parameter and study its effects on the different algorithms.

7.2.3 Dynamic problems.

The techniques presented in this thesis were designed to tackle static problems; they do not respond to changes in the environment. Different events may occur within a dynamic environment, such as the modification or deletion of a job, the addition of a resource. A scheduling system must address these different situations.

Solving a dynamic problem can be regarded as solving a static problem of the same class, many times in succession. Hence, a dynamic algorithm would have to solve a succession of complex resource allocation problems. Rather than starting from scratch all the time, the ideal technique would be to repair the current solution and make it consistent with the new problem. Moreover, the solutions should be made *robust* so that, when a modification in the environment occurs, the solution requires little repair [DRUM92].

7.2.4 A distributed and cooperative approach.

Here, the problem is no longer solved by a single agent, rather the problem is given to a number of agents; each agent is able to solve the problem entirely on

its own. Agents can communicate during the search. Once an agent has found a good solution, the solution is transferred to the other agents which can use this information to guide and direct their search in the most promising areas of the search space¹. The decision to adopt such a distributed and cooperative approach could bring many benefits:

1. Clearwater et al. [CLEA91] report that a set of cooperative agents solving a common cryptarithmic problem² can achieve a combinatorial implosion, i.e. a super-linear speed-up with respect to the number of agents. This speed-up is due to the cooperation between the agents and the specific nature of cryptarithmic problems. In the context of VRPs, a distributed system might not achieve such a formidable speed-up. An important objective for future experiments will be to determine if (as expected) cooperation among, say, n agents enables a faster search than n individual agents working separately. Future experiments should also examine the effects of different communications strategies. For instance, a distributed system could have a star topology with a central server controlling the flow of messages between all the agents. On the other hand, the different agents could be fully connected - or they could use a ring topology - and be able to communicate without the assistance of a central controller.
2. A problem can be solved from different perspectives. Agents may employ different algorithms, focus on different objectives³ and hence, come up with

¹For example, when a GA receives a solution sent by another agent, this solution is evaluated and inserted in the local population. If the solution manages an above-average score, it is likely that this solution will participate to a number of crossover operations. In a sense, this technique performs the same function as the search intensification module of TS.

²A cryptarithmic problem can be viewed as a constraint satisfaction problem. A typical example of such problems is given by the sum $DONALD + GERALD = ROBERT$. The task is to find unique digit assignments to each of the letters so that the numbers add up correctly. For this example, the only solution is $A = 4, B = 3, D = 5, E = 9, G = 1, L = 8, N = 6, O = 2, R = 7, T = 0$.

³Observe that communications between two agents using totally different perspectives may be meaningless.

radically different answers at the end. Either the system can attempt to manipulate and integrate these answers or leave the user decide which solution best suits the problem. Such a system would be more adaptable. In a single-agent system, if the agent finds a problem hard to solve, there is no alternative whereas if an agent fails to solve the problem properly in a multi-agent system, the other agents may be able to provide a better answer.

3. It can also provide an *any-time* system, i.e. a system which can deliver a solution at any-time. The agents can solve the problem at different levels of granularity. Some agents, i.e. the shallow agents, may attempt to provide fast non-optimal solutions whereas other agents, i.e. the deep agents, may require more processing but deliver solutions of better quality. With these multiple levels of reasoning, the overall system can provide an answer at any time. The quality of this answer increases with time as deep agents terminate their search. For instance, for the solution of under-resourced problems, such a system could incorporate GA-based and SA-based agents, the SA-based agents would provide good quality solutions at the early stages of the search, and later on, the GA-based agents could provide better solutions.

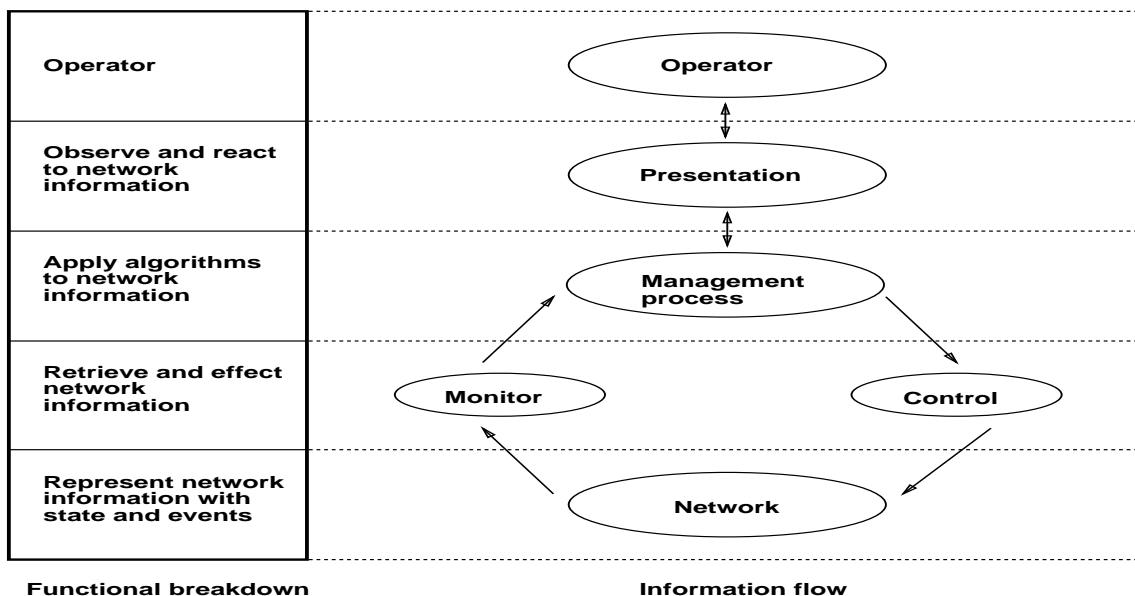
7.2.5 Integration with a graphical user interface.

Finally, the last step in the development of a complete management system may be the integration of these different search techniques with a graphical user interface (GUI).

The management process would be broken into two modules (see Figure 7.1) [CRUT93]:

1. The process module would monitor the state of the network and generate the appropriate actions.
2. The presentation module would play an essential part by filtering the infor-

mation presented to the operator and by allowing the operator to interact directly with the network.



From: "Management and control for giant gigabit networks",
by L. Crutcher and A.A. Lazar, IEEE Network, November 1993.

Figure 7.1: Management system: observation, control, and presentation.

The GUI could follow the Seeheim model⁴ [GREE85], already used in a previous project [MULL92]. This model divides the interface into three layers: the presentation layer defines the presentation components, the dialogue layer monitors the dialogue between the user and the application and checks for consistency, and the application layer receives messages from the dialogue layer and activates the corresponding function of the internal application using a message-passing mechanism.

The operator will be able to view the management problem at different levels of abstraction, e.g. call, path, physical link and network levels. The use of object-oriented techniques will give the opportunity to implement a number of sophisticated mechanisms such as direct manipulation and direct observation.

⁴This model has been employed by the ESPRIT project "VITAMIN" (VISualisation sTANDARD tools in Manufacturing INDustry) for the development of industrial application-interfaces in Computer Integrated Manufacturing (CIM).

Direct manipulation means that the user is able to manipulate any network element directly from his/her terminal by manipulating the graphical object which represents this network element. The manipulations are translated automatically into internal function calls by the interface itself and direct observation means that the information on the screen is automatically updated whenever the internal data is altered, thus allowing a continuous representation of the data of interest and also allowing the user to observe immediately the consequences of his/her actions on the screen.

This GUI equipped with its sophisticated mechanisms will give the opportunity to the operator to interact in an efficient manner with the internal application and the physical network.

7.3 Concluding remark.

This thesis has reported and discussed the results of three large empirical studies. The main objective here was to study how a new class of GAs perform. The breadth of this study was essential if any useful knowledge was to be gained about the applicability of stochastic techniques - in, particular, direct GAs - to vehicle routing problems. A study on a smaller scale would only have distorted the results.

More generally, many AI techniques are still in their infancy and large empirical studies such as the one provided in this thesis are essential in order to gain a better understanding of these techniques and how they can be applied to telecommunications problems.

It cannot be expected that one AI technique will dominate all the other techniques on all the problems. More likely, each technique will have its own *territory*. Therefore, such large empirical studies are essential to define which technique is best for which class of problem and to be able - for a given problem - to anticipate the

behaviour of the different techniques and the quality of their answers.

Bibliography

- [ADEN92] B. Adenso-Diaz. Restricted neighbourhood in the tabu search for the flowshop problem. *European Journal of Operational Research*, 62, pages 27-37, 1992.
- [BAGC91] S. Bagchi, S. Uckum, Y. Miyabe, and K. Kawamura. Exploring problem-specific recombination operators for job shop scheduling. In *[ICGA91]*.
- [BELL93] E. Bell. Telecommunications: telephony merges with video. *IEEE Spectrum*, January 1993.
- [BLAN93] J. L. Blanton and R.L. Wainwright. Multiple vehicle routing with time and capacity constraints using genetic algorithms. In *[ICGA93]*.
- [BODE90] A. Boden. *The social impact of thinking machines*, pages 95–103. Blackwell, 1990.
- [BRIN94] C. Brind, C. Muller, and P. Prosser. A study of stochastic search techniques applied to vehicle routing problems - Part 2. *BT report*, University of Strathclyde, June 1994.
- [BRUN93] R. Bruns. Direct chromosome representation and advanced genetic operators for production scheduling. In *[ICGA93]*.

- [BURK91] P. Burke and P. Prosser. A distributed asynchronous system for predictive and reactive scheduling. *The International Journal for Artificial Intelligence in Engineering*, 6(3), July 1991.
- [CCIT88] CCITT. Recommendation M.30, principles for a telecommunications management network. In Blue book, general maintenance principles. Maintenance of international transmissions systems and telephone circuits, Volume IV.1, pages 22–61, 1988.
- [CEBU89] K. Cebulka, M. Muller, and C. Riley. Applications of artificial intelligence for meeting network management challenges in the 1990's. *IEEE Globecom'89*.
- [CERN85] V. Cerny. Thermodynamical approach to the traveling salesman problem: an efficient simulation algorithm. *Journal of Optimisation Theory and Applications*, 45:41–51, 1985.
- [CHAI92] B. Chaib-Draa, J. Moulin, R. Mandiau, and G. Millot. Trends in distributed artificial intelligence. *Artificial Intelligence Review*, 6(1), 1992.
- [CLEA91] S.C. Clearwater, B.A. Huberman, and T. Hogg. Cooperative solution of constraint satisfaction problems. *Science*, 254:1181-1183, 22 November 1991.
- [CLEV89] A. Cleveland and S.F. Smith. Using genetic algorithms to schedule flow shop releases. In *[ICGA89]*.
- [CLIT89] P. Clitherow and G. Fisher. Knowledge based assistance of genetic search in large design spaces. In *IEA/AIE-89: Proceedings of the Second International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, volume 2, pages 729–734, ACM.

- [COMM92a] ISS generate info network consortium agreement. *Communications International*, December 1992.
- [COMM92b] AT&T back in the local loop. *Communications International*, December 1992.
- [CONN88] D.T. Connolly. An improved annealing scheme for the QAP. *European Journal of Operational Research*, 46:93–100, 1988.
- [COX91] L.A. Cox, L. Davis, and Y. Qiu. Dynamic anticipatory routing in circuit-switched telecommunications networks. In [DAVI91b].
- [CRUT93] L. Crutcher and A. Lazar. Management and control for giant gigabit networks. *IEEE Network*, pages 62-71, November 1993.
- [DAVI91a] Y. Davidor. *Genetic algorithms and robotics: a heuristic strategy for optimization*. World Scientific Series in Robotics and Automated Systems (Volume 1). Singapore: World Scientific, 1991.
- [DAVI85a] L. Davis. Applying adaptive algorithms to epistatic domains. In *Proc International Joint Conference on Artificial Intelligence*, 1985.
- [DAVI85b] L. Davis. Job shop scheduling with genetic algorithms. In [ICGA85].
- [DAVI87] L. Davis, editor. *Genetic algorithms and simulated annealing*. Morgan Kaufmann, 1987.
- [DAVI91b] L. Davis, editor. *Handbook of genetic algorithms*. Van Nostrand Reinhold, 1991.
- [DEWE89] D. de Werra and A. Hertz. Tabu search techniques: a tutorial and an application to neural networks. *OR Spektrum*, 11:131-141, 1989.
- [DIZA89] W.P. Dizard. *The coming information age: an overview of technology, economics and politics*. Longman, third edition, 1989.

- [DRUM92] M. Drummond. Practical approaches to scheduling and planning. AAAI 1992 Spring Symposium Series Reports, page 26, AI magazine, Volume 13, No 3, Fall 1992.
- [DUBR93] S. Dubreuil. Les 100 qui font les telecoms dans le monde. *Telecoms Magazine*, November 1993 (in French).
- [ECON92a] Developing cross-border European telecommunications. *The Economist*, (14th, 20th), November 1992 (advertisement feature for Northern Telecom).
- [ECON92b] The turning of telecoms. *The Economist*, (14th, 20th), November 1992.
- [ECON92c] The big break. *The Economist*, (14th, 20th), November 1992.
- [ECON93a] British Telecom and EDS. *The Economist*, (6th-12th), February 1993.
- [ECON93b] Baby Bells and cable TV. *The Economist*, (6th, 12th), February 1993.
- [ECON93c] VISA international, digital credit. *The Economist*, (25th, 2nd), September 1993.
- [ECON94] SPRINT and EDS, party line. *The Economist*, (21st, 27th), May 1994.
- [EGLE90] R.W. Eglese. Simulated annealing: a tool for operational research. *European Journal of Operational Research*, 46:271–281, 1990.
- [EMBR90] C. Embry, A. Manson, and F. Milham. An open management architecture: OSI/NM forum architecture and concepts. *IEEE Network Magazine*, July 1990.

- [ERCA91] J. Ercau and J. Ferber. L'intelligence artificielle distribuée. *La Recherche*, 22(233), June 1991 (in French).
- [FAGA93] M. Fagan. Electricity firms make a threatening connection. *The Independent*, Wednesday 10th November 1993.
- [FAIG92] U. Faigle and W. Kern. Some convergence results for probabilistic tabu search. *ORSA Journal on Computing*, 4:1,32-37, 1992.
- [FAIR91] M. Fairley. Network management (a tutorial). *Third IEE Conference on Telecommunications, Edinburgh*, March 1991, IEE.
- [FITC91] A. Fitchett and H. Reeve. Future challenges for network management. *British Telecommunications Engineering Journal*, 10, October 1991.
- [FOX84] M. S. Fox and S. .F. Smith. ISIS, a knowledge-based system for facturing scheduling. *Expert systems* , Vol 1, Part 1, pages 25-49, 1984.
- [GARE79] M. Garey and D. Johnson. *Computers and intractability: a guide to the theory of NP-completeness*. W.H. Freeman, 1979.
- [GARI92] J. Garijo and D.A. Hoffman. A multi-agent architecture for operation and maintenance of telecommunication networks. In *12th International Conference on Expert Systems and their Applications, Avignon'92*, 1992.
- [GASS89] L. Gasser. Distributed artificial intelligence. *AI Expert*, July 1989.
- [GLOV87] D.E. Glover. Solving a complex keyboard configuration problem through generalized adaptive search. In *[DAVI87]*.
- [GLOV86] F. Glover. Future paths for Integer Programming and links to Artificial Intelligence. *Comp Oper Res*, 13:533-549, 1986.

- [GLOV88] F. Glover. Tabu Search. *CAAI Report 88-3*, University of Colorado, Boulder, 1988.
- [GLOV89a] F. Glover. Tabu Search, Part I. *ORSA Journal on Computing*, 1:3, 190-206, 1989.
- [GLOV89b] F. Glover and H. J. Greenberg. New approaches for heuristic search: a new bilateral linkage with artificial intelligence. *European Journal of Operational Research*, 39:119-130, 1989.
- [GOLD85] D. E. Goldberg and R. Lingle. Alleles, loci, and the traveling salesman problem. In *[ICGA85]*.
- [GOLD86] B.L. Golden and C.C. Skiscim. Using simulated annealing to solve routing and location problems. *Naval Research Logistics Quarterly*, 33:261-279, 1986.
- [GREE85] M.Green. Report on dialogue specification tools. In *User-interface Management System*, G.E.Pfaff (editor), Springer-Verlag, pages 9-20, 1985.
- [GREF87] J. J. Grefenstette. Incorporating problem specific knowledge into genetic algorithms. In *[DAVI87]*.
- [GRIF91] D. Griffiths and C. Whitney. The role of intelligent software agents in integrated communications management. *British Telecom Technology Journal*, 9(3), July 1991.
- [GUID90] Project Guideline. An implementation architecture for the telecommunications management network for the CEC RACE programme. Technical report, RACE project R1003, March 1990.
- [HAJE88] B. Hajek. Cooling schedules for optimal annealing. *Mathematics of Operations Research*, 13:311-329, 1988.

- [HARM94] P. Harmon. The size of the commercial market in the US. *Intelligent Software Strategies*, X(1), January 1994.
- [HASK90] L. Haskins. Artificial intelligence and corporate knowledge. *British Telecommunications Engineering Journal*, 9, April 1990.
- [HELL91] R.J. Helleur, and N.R. Milway. Network management systems: introductory overview. *British Telecommunications Engineering Journal*, 10, October 1991.
- [HOLL75] J.J. Holland. *Adaptation in natural and artificial systems*. University of Michigan Press, 1975.
- [HOLL92] J.J. Holland. Genetic algorithms. *Scientific American*, 267(1), 1992.
- [HOWA88] E. Howard, D. Jackel, and P. Graf, H. Electronic neural networks. *AT&T Technical Journal*, 67(1), January 1988.
- [ICGA85] J. Grefenstette (editor). Proceedings of the First International Conference on Genetic Algorithms and their Applications. Morgan Kaufmann, 1985.
- [ICGA87] J. Grefenstette (editor). Proceedings of the Second International Conference on Genetic Algorithms and their Applications. Morgan Kaufmann, 1987.
- [ICGA89] J. Shaeffer (editor). Proceedings of the Third International Conference on Genetic Algorithms and their Applications. Morgan Kaufmann, 1989.
- [ICGA91] R. Belew, and L. Booker (editors). Proceedings of the Fourth International Conference on Genetic Algorithms and their Applications. Morgan Kaufmann, 1991.

- [ICGA93] S. Forrest (editor). Proceedings of the Fifth International Conference on Genetic Algorithms and their Applications. Morgan Kaufmann, 1993.
- [INDE93] The brain machines. *The Independent on Sunday*, 1993.
- [JOHN89a] K. Johns. Managing large software systems. *Telesis*, 1989.
- [JOHN89b] D.S. Johnson, C.R. Aragon, L.A. McGeoch, and C. Schevon. Optimization by simulated annealing: an experimental evaluation; Part 1, graph partitioning. *Operations Research*, 37(6):865–892, November 1989.
- [JOHN91] D.S. Johnson, C.R. Aragon, L.A. McGeoch, and C. Schevon. Optimization by simulated annealing: an experimental evaluation; Part 2, graph coloring and number partitioning. *Operations Research*, 39(3):378–406, May 1991.
- [JOHN92] D. Johnson and M. Adorf. Scheduling with neural networks: the case of Hubble space telescope. *International Journal on Computers and Operations Research, special issue on neural networks*, 1992.
- [KIRK83] S. Kirkpatrick, C.D. Gelatt, and M.P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [KLER88] H. Klerer. The OSI management architecture: an overview. *IEEE Spectrum*, 2(2), February 1988.
- [KUMA92] V. Kumar. Algorithms for constraint satisfaction problems: a survey. *AI Magazine*, Spring, 1992.
- [LAGU91] M. Laguna, J.W. Barnes and F. Glover. Tabu search methods for a single machine scheduling problem. *Journal of Intelligent Manufacturing*, 2, pages 63-74, 1991.

- [LAWT92] G. Lawton. Genetic algorithms for schedule optimization. *AI Expert*, May 1992.
- [LEBO91] P. Lebouc and E. Stern. Distributed problem solving in broadband telecommunications network management. In *11th International Conference on Expert Systems and their Applications*, Avignon'91, 1991.
- [LIEB88] J. Liebowitz (editor). *Expert systems applications in telecommunications*. Wiley Series in Telecommunications. 1988.
- [LIEP91] G.E. Liepins and W.D. Potter. A genetic algorithm approach to multi-fault diagnosis. In *[DAVI91b]*.
- [LIN73] S. Lin and B.W. Kernighan. An effective heuristic algorithm for the travelling salesman problem. *Operations Research*, 21:498–516, 1973.
- [LIPP91] B. Lippolt, H. Velthuisen, and E. Kwast. Evaluation of applying model-based diagnosis to telecommunications networks. In *11th International Conference on Expert Systems and their Applications*. Avignon'91, 1991.
- [LIRO91a] Y. Lirov and B. Melaned. Expert design systems for telecommunications. *Expert Systems with Applications*, 2:219–228, 1991.
- [LIRO91b] Y. Lirov and O. Yue. Expert maintenance systems in telecommunication networks. *Journal of Intelligent and Robotic Systems*, (4):303–319, 1991.
- [LOES90] R. Loesener, C.A. Azmoodeh, J. Bigham, and B. Stahl. The use of advanced information processing techniques in maintenance systems for integrated broadband networks. In *Fourth Race TMN Conference*, Dublin, November 1990.

- [LOGO92] M. Logothetis, and M. Shioda. " Centralised virtual path bandwidth allocation scheme for ATM networks. *IEICE Trans Commun*, Vol E75-B, No 10, pages 1071-1080, October 1992.
- [LUND86] M. Lundy and A. Mees. Convergence of an annealing algorithm. *Mathematical Programming*, 34:111–124, 1986.
- [LUSH90] E.P. Lusher. AI and communications network design. *AI Expert*, August 1990.
- [MALC90] J. Malcom and T. Wooding. IKBs in network management. *Computer Communications*, November 1990.
- [MALE89] M. Malek, M. Guruswamy, M. Pandya, and H.Owens. Serial and parallel simulated annealing and tabu search algorithms for the traveling salesman problem. *Annals of Operations Research*, 21,59-84, 1989.
- [MAMD91] H. Mamdani and R. Smith. Advanced information processing for network management. *British Telecommunications Technology Journal*, 9(3), July 1991.
- [MANN90] K. Manning and D. Spencer. Model-based network management. In *Fourth RACE TMN Conference*, Dublin, November 1990.
- [MASS90] E. Masson and Y. Wang. Introduction to computation and learning in artificial neural networks. *European Journal of Operational Research*, 47:1–28, 1990.
- [MAYO90] J. Mayo. The telecommunications explosion: evolution of the I.N. In *The Information Technology Revolution, third edition*, T. Forester (editor), Blackwell, 1990.

- [MCCU88] N. McCulloch, W. Ainsworth, and A. Linggard. Multi-layers perceptrons applied to speech technology. *British Telecom Technology Journal*, 6(2), April 1988.
- [METR53] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E Teller. Equation of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.
- [MORR91] J. T. Morris and B. Samadi. Neural networks in communications: admission control and switch control. In *International Conference in Communications (ICC'91)*, Denver, 1991.
- [MULL92] C.Muller, E.H.Magill, and D.G.Smith. Multi-thread interface for telecommunications applications. Eighth International Conference on Software Engineering for Telecommunications Systems and Services, IEE (conference publication no 352).
- [MYER90] J. Myers. The British Telecom Connectionist Project. In *The market place, neural computing; Third International Seminar*, IBL Technical Services, 1990.
- [NAKA91] R. Nakano. Conventional genetic algorithm for job shop problems. In *[ICGA91]*.
- [NICH93] D. Nicholas and J. Yeomans. 1992 - The impact of the EC on telecommunications (a tutorial). In *The 4th IEE Conference on Telecommunications, Manchester.*, April 1993.
- [OLIV87] I. M. Oliver, D. J. Smith, and J.R.C. Holland. A study of permutation crossover operators on the traveling salesman problem. In *[ICGA87]*.

- [PARU87] H. Van Dyke. Parunak. Why scheduling is hard (and how to do it anyway)? In *Proceedings of the 1987 Material Handling Focus*. Georgia Institute of Technology, September 1987.
- [PENR89] R. Penrose. *The Emperor's new mind*. Oxford University Press, 1989.
- [PERR90] E.L. Perry. Solution of time constrained scheduling problems with parallel tabu search. *Proc 1990 DARPA workshop, Innovative approaches to planning, scheduling, and control*, 231-239, 1990.
- [PROS91] P. Prosser and C. Muller. A design for an intelligent problem solving agent. Technical Report, University of Strathclyde, November 1991.
- [PROS94] P. Prosser, C. Muller, and C. Brind. A preliminary study of stochastic search techniques applied to vehicle routing problems. *BT report*, University of Strathclyde, February 1994.
- [PSAL90] D. Psaltis, D. Brady, G. Xiang-Guang, and S. Lin. Holography in artificial neural networks. *Nature*, 343, January 1990.
- [RACE90] *RACE II, workplan, background material, rationale and overview, definition of scope, and task description, RA3027-2*. 1990.
- [RUDG93] W. Rudge. I'll be seeing you. *IEE Review*, November 1993.
- [SCHA92] J.D. Schaffer and L.J. Eshelman. On crossover as an evolutionary viable strategy. In *[ICGA91]*.
- [SIMO58] A. Simon and A. Newel. Heuristic problem solving: the next advance in operations research. *Operations Research*, page 8, January 1958.
- [SING94] A. Singleton. Genetic programming with C++. *Byte*, pages 171-176, February 1994.

- [SINN91] N. Sinnadurai and G. Morrow. Service management systems. *British Telecommunications Engineering Journal*, 10, October 1991.
- [SMIT91] R. Smith and I. Williamson. Network management for RACE. *British Telecommunications Engineering Journal*, 10, October 1991.
- [SMIT90] S. F. Smith, P. S. Ow, and J. Y. Potvin. OPIS: an opportunistic factory scheduling system. *IEA/AIE 90: industrial and engineering applications of artificial intelligence and expert systems*, Proceedings of the third international conference, Charleston, South Carolina, July 15-18, 1990, vol 1.
- [SOLO83] M. Solomon. *Vehicle routing with time window constraints: models and algorithms.* PhD thesis, Department of Decision Sciences, University of Pennsylvania, 1983.
- [STAR91] T. Starkweather, S. McDaniel, K. Mathias, D. Whitley, and S. Whitley. A comparison of genetic sequencing operators. In *[ICGA91]*.
- [SYSW91] G. Syswerda. Schedule optimization using genetic algorithms. In *[DAVI91b]*.
- [THAN91] S. R. Thangiah, K. E. Nygard, and P. L. Juell. Gideon: a genetic algorithm system for vehicle routing with time windows. In IEEE Computer Society Press, *CAIA-91: Proceedings of the Seventh Conference on Artificial Intelligence Applications*, 1991.
- [THAN93] S. R. Thangiah, R. Vinayagamorthy, and A. V. Gubbi. Vehicle routing with time deadlines using genetic and local algorithms. In *[ICGA93]*.
- [TSAN93] E. Tsang. *Foundations of constraint satisfaction*. Academic Press, 1993.

- [WEIZ90] J. Weizenbaum. *The myths of artificial intelligence*. In *The Information Technology Revolution, third edition*, T. Forester (editor), Blackwell, 1990.
- [WHIT89] D. Whitley, T. Starweather, and D. Fuquay. Scheduling problems and traveling salesman: the genetic edge recombination operator. In *[ICGA89]*.
- [WILL91] K. Willetts. Co-operative management, the key to managing customers networks. *British Telecom Engineering Journal*, 10, October 1991.
- [WRIG90] R. Wright and T. Vesonder. Expert systems in telecommunications. *Expert Systems with Applications*, 1:127–136, 1990.

Annexe A: The problem generator.

Annexe A describes the software which generated the 36 problems used for the experiments. This problem generator is mainly the result of some lengthy discussions between Dr Prosser, Mr Brind and myself. Dr Prosser implemented this problem generator. The following description is taken from a previous technical report: "A preliminary study of stochastic search techniques applied to vehicle routing problems", by P.Prosser, C.Muller, and C.Brind, University of Strathclyde, February 1994.

The VRP Generator

In this section we describe the vehicle routing problem (vrp) generator. We describe the parameters to the generator, the output from the generator, and the characteristics of the problems created. Finally, the database of vrp's for this preliminary study is described.

Parameters

The vrp generator takes the following parameters

Output file: The output file name typically follows the convention C-E-J-D-T-S-N where C is the name of the coordinates file, E is the number of engineers, J is the number of jobs, D is the *bag* of durations (typically $d1$ or $d2$, and see below), T is the *bag* of time windows (typically $t0$, $t1$, or $t2$, again see below), S is the amount of specialisation in the workforce (typically $s0$, $s1$, $s2$, or $s3$, and again see below), and N is a sequence number (the number of the problem generated with these characteristics).

Coordinates file (C) A file of postal codes with their corresponding $x - y$ coordinates, where both x and y are measured in decametres. At present we have 7 data sets, namely Glasgow, G5k (a subset of the Glasgow postal codes within a 5 kilometres of a central location), G7k (as G5k but a 7k radius), Birmingham, B5k (extracted from Birmingham), B7k, and Uxbridge. When generating a problem, the $x - y$ locations of the jobs are randomly drawn from the specified coordinates file *without removal*. That is, it is possible to produce two jobs with the same postal codes and (consequently) the same $x - y$ location.

It should be noted that a problem generated from *5k (where * may be G or B) will tend to involve less travel (tend to be more *urban*) than a similar problem generated from *7k (tends to be more *rural*)⁵. Problems generated with the Uxbridge data set tend to be predominantly *rural*, involving considerable travel.

Speed This is the average speed of an engineer, in miles per hour. Typically we assume 12mph for our urban problems. This is then converted to metres per

⁵... and might even be considered to be sub-urban

minute.

Number of engineers (E) The number of engineers for the problem (typically 30, 40, or 50).

Number of jobs (J) The number of jobs in the problem (typically 200).

Bag of durations (D) A list of possible durations for jobs, measured in minutes.

When a job is created a duration for that job is drawn randomly from this bag.

⁶ A random number is then generated (drawn from the uniform distribution) in the range 1 to n , where n is the number of elements in the bag. The n th element of the bag is then selected. For example, if we wish to generate a set of jobs such that the probability that a job has a duration of 60 minutes is 0.1 (ie. $p(60) = 0.1$), $p(45) = 0.3$, and $p(30) = 0.6$ we would use the bag $\{60, 45, 45, 45, 30, 30, 30, 30, 30, 30\}$.

Bag of time windows (T) A bag of temporal intervals, where a temporal interval has a start time s and an end time e . When a job is created a time window is then selected from this bag (using the same technique described above), such that the duration of the job is less than or equal to $e - s$ of the selected temporal interval. Typically a temporal interval would be (540 1020) (9 o'clock to 5 o'clock, measured in minutes from midnight, also called *all day*), (540 720) the morning, and (720 1020) the afternoon. Therefore, we may have a bag $\{(540 720), (540 720), (540 1020), (720 1020), (720 1020)\}$ which will result in the probability of a job being done in the morning equal to 0.4 ($p(\text{morning}) = 0.4$), $p(\text{all day}) = 0.2$, and $p(\text{afternoon}) = 0.4$. ⁷

Bag of specialisations (S) A bag of specialisations, where a specialisation is a percentage of the population of engineers that can do a job. When a job is

⁶A bag can be thought of as a kind of set, except elements can occur many times. The bag is then used as an *explicit* representation of a distribution.

⁷Note that with a bag $\{(540 1020)\}$ all jobs may be done any time within the working day, and corresponds to no temporal constraints

created a value S is drawn from this bag (using the same technique as above). $E \times S$ engineers are then randomly chosen from the set of engineers, and marked as being capable of doing that job. For example, the bag $s_0 = \{1.0\}$ implies any job can be done by any engineer (consequently no technological constraints), $s_1 = \{0.8, 0.8, 0.8, 0.8, 0.2\}$ implies 80% of the jobs can be done by 80% of the engineers and 20% of the jobs can be done by 20% of the engineers (a weak technological constraint), and $s_3 = \{0.2, 0.2, 0.2, 0.2, 0.8\}$ implies 80% of the jobs can be done by 20% of the engineers and 20% of the jobs can be done by 80% of the engineers (a tight technological constraint).

Probability of compulsory jobs The probability that a job must be done. Typically $p(\text{compulsory})$ is 0.1

Normalised distances If set to true then all $x - y$ coordinates are normalised. That is, the minimum x location (x_{min}) is found from the job set, and the minimum y location is also found (y_{min}). The $x - y$ locations of all jobs are then reduced by x_{min} and y_{min} .

All jobs do-able If true then any job that is created must be feasible. A job is feasible if it is possible to travel to that job, start and complete the job within its time window, and return to the base before end of day (and all problems have been created with this flag set to true).

Generator Output

Given the output file name *outf* (for example) the generator outputs 4 files in the following directories.

vrp/outf The actual machine readable vehicle routing problem.

gnu/outf.jobs The $x - y$ coordinates of the jobs. These can then be plotted (using gnuplot) to show graphically the distribution of the jobs.

gnu/outf.base The $x - y$ location of the base

parms/outf A summary on the problem generated. This file records the actual parameters used when generating the problem, and gives a summary analysis of the problem. This is described in more detail below.

Analysis of Problems

As noted above, the parms file for a problem (parms/outf) records the actual parameters used in generating the problem, along with an analysis of the resultant problem. A sample parms file is given below along with its corresponding graphic (the scatter of jobs and the base on the xy plane) in Figure 1.

```
Coordinates file /usr/keg/bt/vrp-prelim/data/g5k
Output files /usr/keg/bt/vrp-prelim/*/g5k-30-200-d1-tw1-s2-1.*
mph=12 n=30 m=200 p-compulsory=0.1
durations (15 30 45 60 75 90 105 120)
time windows ((540 720) (540 720) (540 1020) (720 1020) (720 1020))
specialisations (0.8 0.2)
```

```
Of 200 jobs 200 can be reached
There are 19 compulsory jobs
jobs per engineer 95.2 6.5 83.0 110.0 30 95.0
job durations 68.47 34.13 15.0 120.0 200 60.0
time-windows 277.8 111.5 180.0 480.0 200 300.0
duration over time window 0.28 0.17 0.03 0.67 200 0.25
engineers per job 14.28 8.97 6.0 24.0 200 6.0
distance from base in minutes 9.01 4.25 1.0 16.0 200 9.0
distance between jobs in minutes 13.0 6.54 0.0 31.0 19900 13.0
Average Engineer Load 1.2
coefficients 0.46 0.39 0.01 2.93 20100 0.32
Profit Coefficient = 3
Base Location = (4823 4482)
```




Figure A.1: g5k-30-200-d1-tw1-s2-1

The problem has 30 engineers and 200 jobs. The probability that a job is compulsory is 0.1, and engineers travel at 12 mph. Durations of jobs are in the range 15 minutes to 120 minutes, in 15 minute increments, each duration equally likely (and this bag of durations is named *d1*). The time window for a job can be in the morning, all day, or in the afternoon. The probability of a job having an *all day* time window is 0.2, and the probability of a morning time window is 0.4, the same as the probability of an afternoon time window (NOTE: this bag of time windows is named *tw1*). The bag of specialisations (the *technological* constraints) dictates that 50% of the jobs can be done by 80% of the engineers, and the remaining 50% of the jobs can be done by 20% of the engineers (and this bag of specialisations is named *s2*).

The next block of text refers to the analysis of the problem created. The figures to the right generally fall into six columns: the mean, standard deviation, minimum, maximum, sample size, median. In this case 200 jobs were created and all were feasible, and 19 of those jobs were compulsory.

jobs per engineer shows the number of jobs that each engineer is capable of doing. This can be taken as a measure of the technological constraints.

job durations shows statistics on the durations of the jobs produced. As expected, the durations of the jobs are on average about mid-way between 60 and 75 minutes.

time-windows shows statistics on the durations of time windows for jobs. In this case the average duration of time windows is 277 minutes.

duration over time window gives an indication of the tightness of the temporal constraints. For each job, the duration of that job is divided by the duration of its time window, and this value is collected. If the duration of jobs are high with respect to the duration of the time windows we get a large value,

and this indicates that temporal constraints are tight (ie. the engineers must get to the job close to the start of the time window). Low values correspond to relaxed temporal constraints (ie. engineers can arrive at a job some time after the start of a time window). Note that the tightness of the temporal constraints are controlled by the interplay between the parameters \mathbf{D} (bag of durations) and \mathbf{T} (bag of time windows).

engineers per job is symmetrical to the row *jobs per engineer*, allowing us a job-based perspective of the problem. As the technological constraints become tighter we will see this value fall. In this example we see that on average each job can be done by 14 engineers (out of 30, and is considered a moderately tight technological constraint).

distance from base in minutes For each job, we compute the distance from the base to that job, in minutes, given the location of the base and the job and the speed that engineers travel. For rural problems we expect this value to be larger than its urban counterpart. The distance between locations is taken as the *Manhattan* distance⁸.

distance between jobs in minutes For all pairs of jobs ⁹ compute the travel time between jobs. For rural problems we expect this value to be larger than its urban counterpart. Again, Manhattan distance is assumed.

Average Engineer Load This is an estimate of the loading on engineers, and is computed as follows. Given \mathbf{J} jobs, each of (on average) \mathbf{D} duration, and \mathbf{E} engineers we compute the average number of jobs an engineer is expected to do (\mathbf{J}/\mathbf{E}). Consequently the engineer will be expected to work $\mathbf{D}\cdot\mathbf{J}/\mathbf{E}$ minutes in the day. The engineer must travel from the base to the first job,

⁸Manhattan distance assumes a *grid* layout, such that one travels up and down, and left and right, but not diagonally. This is a typical restriction in an urban environment, and especially so in Manhattan!

⁹If we have n jobs we compute the $n(n-1)/2$ distances between jobs.

travel between the jobs, and travel from the last job to the base. Using the calculations on travel time (from base and inter-job) and job durations we then divide the number of minutes an engineer is expected to work and travel by the number of minutes in the working day (480).

If this value is large (ie greater than 0.9) we anticipate that the problem is under-resourced, and it is unlikely that all jobs will be done. A good solution will then (typically) attempt to minimise jobs not done. When this value is low (ie. less than 0.85) we expect that the problem is over-resourced, and all jobs will be done. Consequently a good solution will be one that minimises travel (since all jobs will probably be done, this is the only scope for differentiating solutions).

coefficients A ratio based on the travel associated with a job, and the duration of that job. For each job we measure the distance that could be traveled in order to perform that job. That is, for job job_j an engineer may have to travel from job job_i to job job_j and take travel time $t_{i,j}$, perform job_j with duration d_j , and then return to the base taking travel time $t_{j,0}$. The coefficient $C_{i,j}$ for this *triangle* of travel and work (from job_i to job_j to the base, job_0) is calculated as $C_{i,j} = \frac{t_{i,j} + t_{j,0}}{d_j}$.

This value is computed for all possible pairs of jobs job_i and job_j , where $i \neq j$, and $0 \leq i \leq J^{10}$

Problems with high values for $C_{i,j}$ tend to have considerable amount of travel, or jobs with relatively low durations. Symmetrically, problems with low values of $C_{i,j}$ tend to have low travel to jobs of relatively high duration. Therefore, a rural problem tends to have high $C_{i,j}$ with respect to its urban counterpart, and problems with short duration jobs have high $C_{i,j}$ compared to similar problems with jobs of large durations.

¹⁰That is we also consider the tour from the base (job_0) to job_j , then returning to the base.

Profit Coefficient The profit coefficient is used by the cost function employed by the search process. The profit coefficient C is taken to be the maximum value of $C_{i,j}$ (rounded up to the nearest integer).

In evaluating a solution we take into consideration the sum of the durations of the jobs that have not been done (\bar{D}), and the sum of the travel time for all engineers (T). The cost of a solution is then $C\bar{D} + T$. The coefficient C guarantees that the term $C\bar{D}$ is always greater than the travel associated with doing the jobs that have not been done. An alternative way of looking at this is, we guarantee that the search process cannot produce a lower cost solution by reducing travel by not doing jobs, consequently there is always a motive to do more work.

Base Location This is the $x - y$ location of the base and is generated such that it corresponds approximately to the centre of gravity of the location of the jobs.

Therefore, it should be apparent that we can produce vrp's that are under-resourced and over-resourced. We can produce an under-resourced vrp by decreasing the number of engineers, or by increasing the number of jobs or the durations of the jobs, or by increasing the travel between jobs. Furthermore, we have a (coarse) measure of this, namely *load*. Generally, we vary load by altering the number of engineers and the durations of jobs.

Similarly, we can vary the tightness of the temporal and technological constraints. We can increase the tightness of the temporal constraints by making the time windows smaller, or by increasing the durations of jobs (but this has the side-effect of increasing load). We can also increase the tightness of the technological constraints (by limiting who can do a job).

Bringing this together, there appears to be (at least) 4 extreme scenarios. First, we may have vrp's that are under-resourced but with no temporal or technological

constraints. Then, we can have under-resourced vrp's with tight temporal and technological constraints. There are then over-resourced problems with no temporal or technological constraints, and over-resourced problems with tight temporal and technological constraints. Therefore, we hope that we can investigate the role of *load* independently of constraint tightness (and vice versa). Further, we have the ability to generate a range of problems within those extremes.

VRP Database

As part of the preliminary study, 216 vrp's have been produced, with the following characteristics.

location All vrp's have been generating using the g5k coordinates file. That is the Glasgow postcodes in a 5k radius of G1. This produces *urban* problems, and all problems tend to have the same distances between jobs, and between jobs and the base.

Number of Engineers The number of engineers is 30, 40, or 50.

Number of Jobs The number of jobs is held constant at 200.

Durations Durations of jobs are drawn randomly from the bags d1 or d2, where

d1 is {15, 30, 45, 60, 75, 90, 105, 120}.

d2 is {15, 15, 15, 30, 30, 30, 45, 45, 60, 60, 120, 180, 240, 300}.

Time Windows The temporal constraints, either t0, t1, or t2

t0 all jobs have a time window running from 9 to 5 (all day)

t1 probability of morning = probability of afternoon = 0.4 and probability of all day = 0.2

t2 probability of morning = probability of afternoon = 0.45 and probability of all day is approximately 0.1

Specialisations Specialisation (the technological constraints), either s0, s1, s2, or s3

s0 all jobs can be done by any engineer

s1 80% of jobs can be done by 80% of the engineers, and 20% of the jobs can be done by 20% of the engineers

s2 50% of the jobs can be done 80% of the engineers, and 50% of the jobs can be done by 20% of the engineers

s3 80% of the jobs can be done by 20% of the engineers, and 20% of the jobs can be done by 80% of the engineers

n A sequence number. Typically 3 problems are generated at each parameter setting

Problems have the following broad characteristics. As the number of engineers increases the load per engineer falls, and the problem moves from being under-resourced to over-resourced. Problems with durations d1 will tend to be less heavily loaded than problems with durations d2. Problems with time windows t0 have no temporal constraints, and problems with t3 have the tightest temporal constraints. Problems with specialisation s0 have no technological constraints and problems with s3 have the tightest technological constraints

At present the following set of problems has been generated. For g5k (urban), number of engineers 30, 40, or 50, all possible combinations of durations (d1 or d2), time windows (tw0, tw1, tw2), and specialisations (s0, s1, s2, s3), 3 problems at each setting (in total $3 \times 2 \times 3 \times 4 \times 3 = 216$ problems).

- Problems g5k-50-* will be over-resourced compared to g5k-40-*, and g5k-30-* is under-resourced.

- Problems *-d2-* will tend to have a greater spread of job durations than problems *-d1-*, and *-d2-* will tend to put a higher loading on engineers
- Problems *-tw0-* have no temporal constraints, and problems *-tw2-* is the most temporal constrained.
- Problems *-s0-* have no technological constraints, and problems *-s3-* have the most technological constraints.

Annexe B: Raw results.

The following tables give the raw experimental results. Each line corresponds to the five applications of an algorithm to a specific problem. A table entry is a triple: cost, travel and work not done.

Problems	1			2			3			4			5		
30-0-0-1	973932	1932	1620	982916	1916	1635	1000927	1927	1665	991944	1944	1650	964899	1899	1605
30-0-0-2	793969	1969	1320	757947	1947	1260	757970	1970	1260	766953	1953	1275	766946	1946	1275
30-0-0-3	622980	1980	1035	604966	1966	1005	622956	1956	1035	631997	1997	1050	631978	1978	1050
30-0-3-1	758017	2017	1260	767052	2052	1275	767057	2057	1275	767043	2043	1275	767020	2020	1275
30-0-3-2	820993	1993	1365	794034	2034	1320	776019	2019	1290	812078	2078	1350	811989	1989	1350
30-0-3-3	1108906	1906	1845	1018777	1777	1695	1099945	1945	1830	1063889	1889	1770	1099981	1981	1830
30-2-0-1	910940	1940	1515	883968	1968	1470	919981	1981	1530	874934	1934	1455	902031	2031	1500
30-2-0-2	766918	1918	1275	758028	2028	1260	776103	2103	1290	776092	2092	1290	758071	2071	1260
30-2-0-3	1504884	1884	2505	1487037	2037	2475	1504834	1834	2505	1495901	1901	2490	1504811	1811	2505
30-2-3-1	1325044	2044	2205	1316123	2123	2190	1370188	2188	2280	1334063	2063	2220	1352060	2060	2250
30-2-3-2	1135987	1987	1890	1153943	1943	1920	1145087	2087	1905	1163168	2168	1935	1145031	2031	1905
30-2-3-3	884097	2097	1470	902048	2048	1500	866125	2125	1440	884132	2132	1470	893102	2102	1485
40-0-0-1	2290	2290	0	2235	2235	0	2269	2269	0	2229	2229	0	2278	2278	0
40-0-0-2	2226	2226	0	2247	2247	0	2185	2185	0	2235	2235	0	2287	2287	0
40-0-0-3	2180	2180	0	2211	2211	0	2190	2190	0	2111	2111	0	2173	2173	0
40-0-3-1	2302	2302	0	2235	2235	0	2335	2335	0	2324	2324	0	2296	2296	0
40-0-3-2	2154	2154	0	2211	2211	0	2229	2229	0	2160	2160	0	2209	2209	0
40-0-3-3	2385	2385	0	2385	2385	0	2441	2441	0	2434	2434	0	2380	2380	0
40-2-0-1	461275	2275	765	461334	2334	765	434268	2268	720	461119	2119	765	452280	2280	750
40-2-0-2	65293	2293	105	56199	2199	90	74199	2199	120	65213	2213	105	56313	2313	90
40-2-0-3	47360	2360	75	47465	2465	75	29442	2442	45	56421	2421	90	56398	2398	90
40-2-3-1	47643	2643	75	11547	2547	15	2510	2510	0	2594	2594	0	2715	2715	0
40-2-3-2	2351	2351	0	2309	2309	0	2423	2423	0	2267	2267	0	2275	2275	0
40-2-3-3	191412	2412	315	182435	2435	300	200383	2383	330	164358	2358	270	182395	2395	300
50-0-0-1	2260	2260	0	2312	2312	0	2332	2332	0	2310	2310	0	2273	2273	0
50-0-0-2	2242	2242	0	2249	2249	0	2231	2231	0	2194	2194	0	2249	2249	0
50-0-0-3	2223	2223	0	2242	2242	0	2181	2181	0	2217	2217	0	2224	2224	0
50-0-3-1	2293	2293	0	2274	2274	0	2234	2234	0	2322	2322	0	2262	2262	0
50-0-3-2	2340	2340	0	2288	2288	0	2299	2299	0	2307	2307	0	2333	2333	0
50-0-3-3	2292	2292	0	2305	2305	0	2359	2359	0	2326	2326	0	2307	2307	0
50-2-0-1	2228	2228	0	2231	2231	0	2208	2208	0	2232	2232	0	2228	2228	0
50-2-0-2	2083	2083	0	2108	2108	0	2106	2106	0	2104	2104	0	2104	2104	0
50-2-0-3	2114	2114	0	2129	2129	0	2070	2070	0	2104	2104	0	2120	2120	0
50-2-3-1	2208	2208	0	2200	2200	0	2173	2173	0	2179	2179	0	2185	2185	0
50-2-3-2	2364	2364	0	2389	2389	0	2393	2393	0	2375	2375	0	2429	2429	0
50-2-3-3	2409	2409	0	2417	2417	0	2440	2440	0	2414	2414	0	2401	2401	0

Table B.1: PMX, raw results.

Problems	1			2			3			4			5		
30-0-0-1	721464	1464	1200	703428	1428	1170	712454	1454	1185	667413	1413	1110	649364	1364	1080
30-0-0-2	469433	1433	780	469409	1409	780	424387	1387	705	523536	1536	870	451444	1444	750
30-0-0-3	253363	1363	420	298448	1448	495	361551	1551	600	271422	1422	450	280387	1387	465
30-0-3-1	514624	1624	855	505629	1629	840	514639	1639	855	541647	1647	900	568708	1708	945
30-0-3-2	586715	1715	975	640774	1774	1065	568615	1615	945	541613	1613	900	586684	1684	975
30-0-3-3	838470	1470	1395	901611	1611	1500	829493	1493	1380	838542	1542	1395	838531	1531	1395
30-2-0-1	766816	1816	1275	748745	1745	1245	775742	1742	1290	820869	1869	1365	820862	1862	1365
30-2-0-2	757938	1938	1260	694747	1747	1155	685854	1854	1140	649783	1783	1080	712883	1883	1185
30-2-0-3	1432638	1638	2385	1396595	1595	2325	1423677	1677	2370	1423701	1701	2370	1432692	1692	2385
30-2-3-1	1162929	1929	1935	1180857	1857	1965	1334049	2049	2220	1226018	2018	2040	1162869	1869	1935
30-2-3-2	1054925	1925	1755	1018919	1919	1695	1046002	2002	1740	1072997	1997	1785	1082039	2039	1800
30-2-3-3	784923	1923	1305	730820	1820	1215	829981	1981	1380	784948	1948	1305	775798	1798	1290
40-0-0-1	1399	1399	0	1346	1346	0	1348	1348	0	1403	1403	0	1343	1343	0
40-0-0-2	1232	1232	0	1342	1342	0	1414	1414	0	1298	1298	0	1340	1340	0
40-0-0-3	1360	1360	0	1207	1207	0	1329	1329	0	1333	1333	0	1268	1268	0
40-0-3-1	1528	1528	0	1465	1465	0	1529	1529	0	1409	1409	0	1476	1476	0
40-0-3-2	1446	1446	0	1509	1509	0	1526	1526	0	1426	1426	0	1477	1477	0
40-0-3-3	1501	1501	0	1505	1505	0	1563	1563	0	1630	1630	0	1448	1448	0
40-2-0-1	442986	1986	735	407033	2033	675	415981	1981	690	389029	2029	645	433993	1993	720
40-2-0-2	29035	2035	45	65210	2210	105	47023	2023	75	74176	2176	120	56165	2165	90
40-2-0-3	56302	2302	90	2166	2166	0	2046	2046	0	2171	2171	0	2134	2134	0
40-2-3-1	2319	2319	0	2427	2427	0	2464	2464	0	2425	2425	0	2630	2630	0
40-2-3-2	2064	2064	0	1943	1943	0	1982	1982	0	1962	1962	0	2466	2466	0
40-2-3-3	164285	2285	270	182349	2349	300	164222	2222	270	20032	2032	30	182305	2305	300
50-0-0-1	1418	1418	0	1393	1393	0	1329	1329	0	1452	1452	0	1485	1485	0
50-0-0-2	1246	1246	0	1363	1363	0	1401	1401	0	1349	1349	0	1268	1268	0
50-0-0-3	1251	1251	0	1288	1288	0	1286	1286	0	1473	1473	0	1287	1287	0
50-0-3-1	1511	1511	0	1459	1459	0	1564	1564	0	1584	1584	0	1493	1493	0
50-0-3-2	1433	1433	0	1570	1570	0	1444	1444	0	1642	1642	0	1517	1517	0
50-0-3-3	1497	1497	0	1527	1527	0	1505	1505	0	1547	1547	0	1576	1576	0
50-2-0-1	1641	1641	0	1659	1659	0	1657	1657	0	1686	1686	0	1778	1778	0
50-2-0-2	1449	1449	0	1523	1523	0	1488	1488	0	1424	1424	0	1585	1585	0
50-2-0-3	1554	1554	0	1579	1579	0	1508	1508	0	1530	1530	0	1578	1578	0
50-2-3-1	1739	1739	0	1933	1933	0	1757	1757	0	1972	1972	0	1818	1818	0
50-2-3-2	2146	2146	0	2128	2128	0	2158	2158	0	2085	2085	0	2043	2043	0
50-2-3-3	2030	2030	0	1933	1933	0	1944	1944	0	1983	1983	0	1860	1860	0

Table B.2: Direct3, raw results.

Problems	1			2			3			4			5		
30-0-0-1	667377	1377	1110	658409	1409	1095	694423	1423	1155	685402	1402	1140	595342	1342	990
30-0-0-2	442412	1412	735	397279	1279	660	478487	1487	795	379294	1294	630	433394	1394	720
30-0-0-3	289403	1403	480	280417	1417	465	280415	1415	465	334530	1530	555	280404	1404	465
30-0-3-1	577730	1730	960	505628	1628	840	523659	1659	870	478584	1584	795	496561	1561	825
30-0-3-2	505577	1577	840	514451	1451	855	496522	1522	825	541624	1624	900	595656	1656	990
30-0-3-3	802444	1444	1335	811540	1540	1350	847469	1469	1410	820502	1502	1365	838497	1497	1395
30-2-0-1	631555	1555	1050	703620	1620	1170	739712	1712	1230	766725	1725	1275	676614	1614	1125
30-2-0-2	685875	1875	1140	622692	1692	1035	685808	1808	1140	586746	1746	975	676731	1731	1125
30-2-0-3	1432546	1546	2385	1405559	1559	2340	1432548	1548	2385	1405435	1435	2340	1441576	1576	2400
30-2-3-1	1171947	1947	1950	1189952	1952	1980	1207981	1981	2010	1253031	2031	2085	1306986	1986	2175
30-2-3-2	1018995	1995	1695	1018924	1924	1695	1009879	1879	1680	946844	1844	1575	973836	1836	1620
30-2-3-3	712875	1875	1185	820799	1799	1365	820922	1922	1365	766816	1816	1275	803040	2040	1335
40-0-0-1	1354	1354	0	1330	1330	0	1388	1388	0	1271	1271	0	1326	1326	0
40-0-0-2	1354	1354	0	1381	1381	0	1278	1278	0	1271	1271	0	1392	1392	0
40-0-0-3	1284	1284	0	1309	1309	0	1372	1372	0	1276	1276	0	1295	1295	0
40-0-3-1	1630	1630	0	1582	1582	0	1629	1629	0	1478	1478	0	1456	1456	0
40-0-3-2	1442	1442	0	1394	1394	0	1448	1448	0	1505	1505	0	1513	1513	0
40-0-3-3	1556	1556	0	1490	1490	0	1646	1646	0	1589	1589	0	1533	1533	0
40-2-0-1	406976	1976	675	370999	1999	615	379890	1890	630	397991	1991	660	379920	1920	630
40-2-0-2	38097	2097	60	73853	1853	120	2083	2083	0	1874	1874	0	2010	2010	0
40-2-0-3	1980	1980	0	56002	2002	90	2057	2057	0	2103	2103	0	2165	2165	0
40-2-3-1	2171	2171	0	2163	2163	0	2465	2465	0	2180	2180	0	2387	2387	0
40-2-3-2	1961	1961	0	2014	2014	0	2022	2022	0	1984	1984	0	1964	1964	0
40-2-3-3	182145	2145	300	110178	2178	180	83158	2158	135	101165	2165	165	128225	2225	210
50-0-0-1	1335	1335	0	1420	1420	0	1277	1277	0	1387	1387	0	1379	1379	0
50-0-0-2	1398	1398	0	1435	1435	0	1472	1472	0	1274	1274	0	1434	1434	0
50-0-0-3	1367	1367	0	1305	1305	0	1308	1308	0	1409	1409	0	1215	1215	0
50-0-3-1	1566	1566	0	1684	1684	0	1500	1500	0	1561	1561	0	1564	1564	0
50-0-3-2	1547	1547	0	1579	1579	0	1576	1576	0	1490	1490	0	1534	1534	0
50-0-3-3	1632	1632	0	1582	1582	0	1565	1565	0	1659	1659	0	1614	1614	0
50-2-0-1	1636	1636	0	1551	1551	0	1776	1776	0	1744	1744	0	1741	1741	0
50-2-0-2	1648	1648	0	1634	1634	0	1491	1491	0	1560	1560	0	1470	1470	0
50-2-0-3	1540	1540	0	1569	1569	0	1513	1513	0	1586	1586	0	1565	1565	0
50-2-3-1	1893	1893	0	1988	1988	0	1768	1768	0	1815	1815	0	1875	1875	0
50-2-3-2	2171	2171	0	2100	2100	0	2212	2212	0	2080	2080	0	1968	1968	0
50-2-3-3	2018	2018	0	2008	2008	0	1907	1907	0	1937	1937	0	2070	2070	0

Table B.3: Direct3-s, raw results.

Problems	1			2			3			4			5		
30-0-0-1	631329	1329	1050	568245	1245	945	595346	1346	990	559239	1239	930	586273	1273	975
30-0-0-2	370299	1299	615	433353	1353	720	433373	1373	720	352207	1207	585	334210	1210	555
30-0-0-3	217317	1317	360	262314	1314	435	208312	1312	345	289387	1387	480	280389	1389	465
30-0-3-1	433407	1407	720	406502	1502	675	442513	1513	735	487554	1554	810	406434	1434	675
30-0-3-2	487494	1494	810	469494	1494	780	532540	1540	885	478514	1514	795	469503	1503	780
30-0-3-3	775406	1406	1290	847559	1559	1410	775368	1368	1290	802476	1476	1335	766379	1379	1275
30-2-0-1	640517	1517	1065	676657	1657	1125	631549	1549	1050	694592	1592	1155	604487	1487	1005
30-2-0-2	658573	1573	1095	595772	1772	990	586601	1601	975	586639	1639	975	613546	1546	1020
30-2-0-3	1432574	1574	2385	1378385	1385	2295	1387526	1526	2310	1432473	1473	2385	1414487	1487	2355
30-2-3-1	1198919	1919	1995	1207905	1905	2010	1162694	1694	1935	1189778	1778	1980	1144785	1785	1905
30-2-3-2	892736	1736	1485	955827	1827	1590	919724	1724	1530	928867	1867	1545	901760	1760	1500
30-2-3-3	757795	1795	1260	784815	1815	1305	766840	1840	1275	739782	1782	1230	775952	1952	1290
40-0-0-1	1354	1354	0	1349	1349	0	1354	1354	0	1314	1314	0	1304	1304	0
40-0-0-2	1339	1339	0	1413	1413	0	1338	1338	0	1253	1253	0	1348	1348	0
40-0-0-3	1299	1299	0	1280	1280	0	1304	1304	0	1300	1300	0	1427	1427	0
40-0-3-1	1652	1652	0	1492	1492	0	1681	1681	0	1518	1518	0	1571	1571	0
40-0-3-2	1494	1494	0	1498	1498	0	1624	1624	0	1517	1517	0	1502	1502	0
40-0-3-3	1623	1623	0	1631	1631	0	1664	1664	0	1597	1597	0	1688	1688	0
40-2-0-1	415982	1982	690	424843	1843	705	352858	1858	585	406869	1869	675	361881	1881	600
40-2-0-2	1798	1798	0	1719	1719	0	1841	1841	0	19978	1978	30	1846	1846	0
40-2-0-3	2054	2054	0	2025	2025	0	1996	1996	0	1895	1895	0	1954	1954	0
40-2-3-1	2187	2187	0	2124	2124	0	2215	2215	0	2115	2115	0	2241	2241	0
40-2-3-2	1967	1967	0	1996	1996	0	1863	1863	0	1965	1965	0	1981	1981	0
40-2-3-3	101160	2160	165	128290	2290	210	83199	2199	135	164209	2209	270	65187	2187	105
50-0-0-1	1381	1381	0	1384	1384	0	1358	1358	0	1320	1320	0	1428	1428	0
50-0-0-2	1353	1353	0	1467	1467	0	1448	1448	0	1412	1412	0	1433	1433	0
50-0-0-3	1270	1270	0	1349	1349	0	1289	1289	0	1288	1288	0	1205	1205	0
50-0-3-1	1679	1679	0	1601	1601	0	1640	1640	0	1599	1599	0	1607	1607	0
50-0-3-2	1770	1770	0	1521	1521	0	1726	1726	0	1612	1612	0	1645	1645	0
50-0-3-3	1576	1576	0	1550	1550	0	1645	1645	0	1718	1718	0	1676	1676	0
50-2-0-1	1774	1774	0	1732	1732	0	1798	1798	0	1821	1821	0	1756	1756	0
50-2-0-2	1547	1547	0	1606	1606	0	1559	1559	0	1528	1528	0	1562	1562	0
50-2-0-3	1608	1608	0	1489	1489	0	1553	1553	0	1568	1568	0	1690	1690	0
50-2-3-1	1872	1872	0	2023	2023	0	1932	1932	0	1797	1797	0	1835	1835	0
50-2-3-2	2119	2119	0	2054	2054	0	2022	2022	0	2168	2168	0	1884	1884	0
50-2-3-3	2119	2119	0	2015	2015	0	2052	2052	0	2123	2123	0	1983	1983	0

Table B.4: Direct4-s, raw results.

Problems	1			2			3			4			5		
30-0-0-1	487160	1160	810	586237	1237	975	514177	1177	855	514154	1154	855	532198	1198	885
30-0-0-2	298185	1185	495	307197	1197	510	361280	1280	600	271144	1144	450	361284	1284	600
30-0-0-3	181246	1246	300	154176	1176	255	181229	1229	300	190272	1272	315	172233	1233	285
30-0-3-1	388419	1419	645	388473	1473	645	370407	1407	615	352395	1395	585	370415	1415	615
30-0-3-2	478476	1476	795	469443	1443	780	424441	1441	705	451424	1424	750	433467	1467	720
30-0-3-3	757378	1378	1260	703291	1291	1170	694325	1325	1155	685299	1299	1140	721377	1377	1200
30-2-0-1	541409	1409	900	622564	1564	1035	586471	1471	975	586454	1454	975	577468	1468	960
30-2-0-2	523588	1588	870	559661	1661	930	532493	1493	885	604590	1590	1005	550489	1489	915
30-2-0-3	1369318	1318	2280	1378382	1382	2295	1360431	1431	2265	1387338	1338	2310	1405532	1532	2340
30-2-3-1	1162664	1664	1935	1108762	1762	1845	1135647	1647	1890	1108723	1723	1845	1099775	1775	1830
30-2-3-2	883660	1660	1470	847737	1737	1410	793634	1634	1320	883695	1695	1470	847703	1703	1410
30-2-3-3	757747	1747	1260	739684	1684	1230	784954	1954	1305	757684	1684	1260	712641	1641	1185
40-0-0-1	1221	1221	0	1257	1257	0	1235	1235	0	1337	1337	0	1196	1196	0
40-0-0-2	1217	1217	0	1298	1298	0	1262	1262	0	1284	1284	0	1221	1221	0
40-0-0-3	1254	1254	0	1220	1220	0	1248	1248	0	1147	1147	0	1209	1209	0
40-0-3-1	1503	1503	0	1463	1463	0	1396	1396	0	1429	1429	0	1463	1463	0
40-0-3-2	1368	1368	0	1436	1436	0	1379	1379	0	1398	1398	0	1393	1393	0
40-0-3-3	1559	1559	0	1482	1482	0	1484	1484	0	1473	1473	0	1562	1562	0
40-2-0-1	398060	2060	660	361888	1888	600	379838	1838	630	370754	1754	615	343783	1783	570
40-2-0-2	1713	1713	0	1675	1675	0	1850	1850	0	1827	1827	0	1686	1686	0
40-2-0-3	1702	1702	0	1785	1785	0	1718	1718	0	1631	1631	0	1749	1749	0
40-2-3-1	1950	1950	0	1950	1950	0	1828	1828	0	2116	2116	0	1960	1960	0
40-2-3-2	1821	1821	0	1780	1780	0	1936	1936	0	1852	1852	0	1805	1805	0
40-2-3-3	64907	1907	105	38053	2053	60	19988	1988	30	56020	2020	90	55943	1943	90
50-0-0-1	1287	1287	0	1209	1209	0	1308	1308	0	1307	1307	0	1273	1273	0
50-0-0-2	1206	1206	0	1232	1232	0	1235	1235	0	1380	1380	0	1211	1211	0
50-0-0-3	1105	1105	0	1131	1131	0	1361	1361	0	1216	1216	0	1239	1239	0
50-0-3-1	1453	1453	0	1466	1466	0	1521	1521	0	1593	1593	0	1533	1533	0
50-0-3-2	1482	1482	0	1501	1501	0	1533	1533	0	1468	1468	0	1612	1612	0
50-0-3-3	1480	1480	0	1490	1490	0	1475	1475	0	1472	1472	0	1469	1469	0
50-2-0-1	1599	1599	0	1756	1756	0	1582	1582	0	1618	1618	0	1588	1588	0
50-2-0-2	1494	1494	0	1552	1552	0	1500	1500	0	1406	1406	0	1487	1487	0
50-2-0-3	1434	1434	0	1542	1542	0	1496	1496	0	1492	1492	0	1529	1529	0
50-2-3-1	1787	1787	0	1675	1675	0	1790	1790	0	1823	1823	0	1783	1783	0
50-2-3-2	1924	1924	0	1989	1989	0	1908	1908	0	2080	2080	0	1927	1927	0
50-2-3-3	1844	1844	0	1875	1875	0	1899	1899	0	1873	1873	0	1920	1920	0

Table B.5: Direct4-s:200, raw results.

Problems	1			2			3			4			5		
30-0-0-1	559220	1220	930	586290	1290	975	568272	1272	945	559205	1205	930	496128	1128	825
30-0-0-2	316162	1162	525	361232	1232	600	361250	1250	600	298187	1187	495	289118	1118	480
30-0-0-3	217296	1296	360	190290	1290	315	199241	1241	330	172257	1257	285	154200	1200	255
30-0-3-1	316292	1292	525	352397	1397	585	316313	1313	525	334316	1316	555	307305	1305	510
30-0-3-2	334330	1330	555	343321	1321	570	379382	1382	630	361338	1338	600	352323	1323	585
30-0-3-3	658284	1284	1095	658246	1246	1095	658242	1242	1095	676281	1281	1125	604183	1183	1005
30-2-0-1	658608	1608	1095	640470	1470	1065	586404	1404	975	631508	1508	1050	694567	1567	1155
30-2-0-2	622635	1635	1035	568601	1601	945	541595	1595	900	604600	1600	1005	586721	1721	975
30-2-0-3	1396614	1614	2325	1405672	1672	2340	1378629	1629	2295	1396414	1414	2325	1405547	1547	2340
30-2-3-1	1072625	1625	1785	1072594	1594	1785	1099527	1527	1830	1090633	1633	1815	1081680	1680	1800
30-2-3-2	811610	1610	1350	811644	1644	1350	838625	1625	1395	802576	1576	1335	757582	1582	1260
30-2-3-3	658604	1604	1095	685621	1621	1140	631550	1550	1050	631605	1605	1050	667582	1582	1110
40-0-0-1	2214	2214	0	2107	2107	0	2220	2220	0	2229	2229	0	2238	2238	0
40-0-0-2	2182	2182	0	2226	2226	0	2225	2225	0	2243	2243	0	2209	2209	0
40-0-0-3	2155	2155	0	2158	2158	0	2177	2177	0	2143	2143	0	2153	2153	0
40-0-3-1	1445	1445	0	1410	1410	0	1361	1361	0	1388	1388	0	1430	1430	0
40-0-3-2	1310	1310	0	1423	1423	0	1318	1318	0	1298	1298	0	1343	1343	0
40-0-3-3	1476	1476	0	1520	1520	0	1479	1479	0	1494	1494	0	1506	1506	0
40-2-0-1	343767	1767	570	352762	1762	585	389203	2203	645	343699	1699	570	334765	1765	555
40-2-0-2	1835	1835	0	1899	1899	0	1782	1782	0	1850	1850	0	1843	1843	0
40-2-0-3	1730	1730	0	1693	1693	0	1784	1784	0	1803	1803	0	1588	1588	0
40-2-3-1	1797	1797	0	1895	1895	0	1870	1870	0	1866	1866	0	1872	1872	0
40-2-3-2	1686	1686	0	1745	1745	0	1663	1663	0	1714	1714	0	1732	1732	0
40-2-3-3	1893	1893	0	1905	1905	0	11020	2020	15	1823	1823	0	28973	1973	45
50-0-0-1	2301	2301	0	2275	2275	0	2293	2293	0	2224	2224	0	2291	2291	0
50-0-0-2	2169	2169	0	2190	2190	0	2216	2216	0	2155	2155	0	2219	2219	0
50-0-0-3	2153	2153	0	2197	2197	0	2222	2222	0	2148	2148	0	2197	2197	0
50-0-3-1	1411	1411	0	1462	1462	0	1500	1500	0	1497	1497	0	1458	1458	0
50-0-3-2	1381	1381	0	1459	1459	0	1412	1412	0	1533	1533	0	1563	1563	0
50-0-3-3	1349	1349	0	1431	1431	0	1419	1419	0	1579	1579	0	1483	1483	0
50-2-0-1	2115	2115	0	2155	2155	0	2190	2190	0	2193	2193	0	2172	2172	0
50-2-0-2	2040	2040	0	2038	2038	0	2068	2068	0	2033	2033	0	2068	2068	0
50-2-0-3	2069	2069	0	2053	2053	0	2098	2098	0	2068	2068	0	2087	2087	0
50-2-3-1	1557	1557	0	1559	1559	0	1672	1672	0	1625	1625	0	1799	1799	0
50-2-3-2	1910	1910	0	1808	1808	0	1933	1933	0	1817	1817	0	1916	1916	0
50-2-3-3	1863	1863	0	1801	1801	0	1863	1863	0	1806	1806	0	1846	1846	0

Table B.6: Direct4-s:400, raw results.

Problems	1			2			3			4			5		
30-0-0-1	541201	1201	900	532151	1151	885	496160	1160	825	550184	1184	915	559203	1203	930
30-0-0-2	325252	1252	540	271149	1149	450	289106	1106	480	289183	1183	480	370230	1230	615
30-0-0-3	154211	1211	255	154157	1157	255	217276	1276	360	199267	1267	330	244293	1293	405
30-0-3-1	271292	1292	450	307310	1310	510	298287	1287	495	298275	1275	495	334372	1372	555
30-0-3-2	343322	1322	570	397356	1356	660	424392	1392	705	379313	1313	630	352349	1349	585
30-0-3-3	694348	1348	1155	685287	1287	1140	676249	1249	1125	712286	1286	1185	685280	1280	1140
30-2-0-1	631471	1471	1050	640552	1552	1065	649528	1528	1080	667526	1526	1110	640583	1583	1065
30-2-0-2	604770	1770	1005	622682	1682	1035	559552	1552	930	604584	1584	1005	496404	1404	825
30-2-0-3	1405457	1457	2340	1378471	1471	2295	1387461	1461	2310	1378531	1531	2295	1387458	1458	2310
30-2-3-1	1045733	1733	1740	1126663	1663	1875	1090660	1660	1815	1108641	1641	1845	1072643	1643	1785
30-2-3-2	802680	1680	1335	802598	1598	1335	784617	1617	1305	802653	1653	1335	820679	1679	1365
30-2-3-3	640587	1587	1065	703575	1575	1170	676641	1641	1125	712679	1679	1185	721635	1635	1200
40-0-0-1	1233	1233	0	1247	1247	0	1204	1204	0	1315	1315	0	1202	1202	0
40-0-0-2	1238	1238	0	1238	1238	0	1242	1242	0	1151	1151	0	1128	1128	0
40-0-0-3	1226	1226	0	1148	1148	0	1225	1225	0	1220	1220	0	1242	1242	0
40-0-3-1	1477	1477	0	1420	1420	0	1485	1485	0	1362	1362	0	1337	1337	0
40-0-3-2	1309	1309	0	1391	1391	0	1323	1323	0	1324	1324	0	1302	1302	0
40-0-3-3	1425	1425	0	1479	1479	0	1490	1490	0	1439	1439	0	1458	1458	0
40-2-0-1	343846	1846	570	335009	2009	555	307720	1720	510	379989	1989	630	361846	1846	600
40-2-0-2	1824	1824	0	1836	1836	0	1793	1793	0	1745	1745	0	1926	1926	0
40-2-0-3	1736	1736	0	1689	1689	0	1633	1633	0	1764	1764	0	1875	1875	0
40-2-3-1	1897	1897	0	1824	1824	0	1788	1788	0	1899	1899	0	1821	1821	0
40-2-3-2	1695	1695	0	1698	1698	0	1623	1623	0	1765	1765	0	1677	1677	0
40-2-3-3	10960	1960	15	1884	1884	0	19974	1974	30	55999	1999	90	46860	1860	75
50-0-0-1	1321	1321	0	1226	1226	0	1332	1332	0	1234	1234	0	1336	1336	0
50-0-0-2	1325	1325	0	1183	1183	0	1269	1269	0	1242	1242	0	1242	1242	0
50-0-0-3	1187	1187	0	1234	1234	0	1172	1172	0	1138	1138	0	1181	1181	0
50-0-3-1	1435	1435	0	1405	1405	0	1373	1373	0	1533	1533	0	1489	1489	0
50-0-3-2	1461	1461	0	1369	1369	0	1348	1348	0	1462	1462	0	1461	1461	0
50-0-3-3	1479	1479	0	1410	1410	0	1391	1391	0	1510	1510	0	1426	1426	0
50-2-0-1	1688	1688	0	1593	1593	0	1575	1575	0	1510	1510	0	1582	1582	0
50-2-0-2	1503	1503	0	1492	1492	0	1477	1477	0	1478	1478	0	1445	1445	0
50-2-0-3	1444	1444	0	1529	1529	0	1384	1384	0	1477	1477	0	1437	1437	0
50-2-3-1	1666	1666	0	1665	1665	0	1797	1797	0	1591	1591	0	1782	1782	0
50-2-3-2	1790	1790	0	1881	1881	0	1883	1883	0	1874	1874	0	1754	1754	0
50-2-3-3	1744	1744	0	1822	1822	0	1773	1773	0	1867	1867	0	1742	1742	0

Table B.7: Dupli-1, raw results.

Problems	1			2			3			4			5		
30-0-0-1	595313	1313	990	550206	1206	915	505166	1166	840	568181	1181	945	577247	1247	960
30-0-0-2	289167	1167	480	361275	1275	600	370258	1258	615	307171	1171	510	379303	1303	630
30-0-0-3	208286	1286	345	190232	1232	315	190221	1221	315	199287	1287	330	217261	1261	360
30-0-3-1	397428	1428	660	388448	1448	645	415438	1438	690	433504	1504	720	433540	1540	720
30-0-3-2	424419	1419	705	469471	1471	780	433399	1399	720	460499	1499	765	406358	1358	675
30-0-3-3	739352	1352	1230	721318	1318	1200	703321	1321	1170	757372	1372	1260	802373	1373	1335
30-2-0-1	649538	1538	1080	613494	1494	1020	667570	1570	1110	640492	1492	1065	613540	1540	1020
30-2-0-2	532510	1510	885	595706	1706	990	559586	1586	930	505504	1504	840	559598	1598	930
30-2-0-3	1378405	1405	2295	1378389	1389	2295	1387659	1659	2310	1387511	1511	2310	1405566	1566	2340
30-2-3-1	1153739	1739	1920	1180740	1740	1965	1081795	1795	1800	1117870	1870	1860	1099796	1796	1830
30-2-3-2	865661	1661	1440	847592	1592	1410	856766	1766	1425	847659	1659	1410	856662	1662	1425
30-2-3-3	685625	1625	1140	685639	1639	1140	712814	1814	1185	748746	1746	1245	676741	1741	1125
40-0-0-1	1103	1103	0	1272	1272	0	1115	1115	0	1136	1136	0	1125	1125	0
40-0-0-2	1132	1132	0	1081	1081	0	1053	1053	0	1208	1208	0	1145	1145	0
40-0-0-3	1114	1114	0	1105	1105	0	1214	1214	0	1093	1093	0	1149	1149	0
40-0-3-1	1406	1406	0	1401	1401	0	1421	1421	0	1422	1422	0	1360	1360	0
40-0-3-2	1440	1440	0	1364	1364	0	1418	1418	0	1364	1364	0	1399	1399	0
40-0-3-3	1376	1376	0	1445	1445	0	1367	1367	0	1452	1452	0	1480	1480	0
40-2-0-1	316876	1876	525	352922	1922	585	343753	1753	570	352771	1771	585	343828	1828	570
40-2-0-2	1775	1775	0	1906	1906	0	1723	1723	0	1852	1852	0	37911	1911	60
40-2-0-3	1826	1826	0	1755	1755	0	1943	1943	0	1897	1897	0	1741	1741	0
40-2-3-1	2082	2082	0	2078	2078	0	2056	2056	0	2095	2095	0	1902	1902	0
40-2-3-2	1903	1903	0	1825	1825	0	1714	1714	0	1847	1847	0	1906	1906	0
40-2-3-3	46894	1894	75	92167	2167	150	83019	2019	135	47064	2064	75	38058	2058	60
50-0-0-1	1103	1103	0	1211	1211	0	1105	1105	0	1138	1138	0	1139	1139	0
50-0-0-2	1210	1210	0	1164	1164	0	1171	1171	0	1108	1108	0	1219	1219	0
50-0-0-3	1125	1125	0	1054	1054	0	1210	1210	0	1090	1090	0	1098	1098	0
50-0-3-1	1451	1451	0	1422	1422	0	1369	1369	0	1467	1467	0	1409	1409	0
50-0-3-2	1482	1482	0	1374	1374	0	1496	1496	0	1481	1481	0	1384	1384	0
50-0-3-3	1455	1455	0	1455	1455	0	1401	1401	0	1453	1453	0	1449	1449	0
50-2-0-1	1575	1575	0	1562	1562	0	1491	1491	0	1461	1461	0	1549	1549	0
50-2-0-2	1291	1291	0	1359	1359	0	1352	1352	0	1461	1461	0	1300	1300	0
50-2-0-3	1479	1479	0	1427	1427	0	1380	1380	0	1475	1475	0	1421	1421	0
50-2-3-1	1752	1752	0	1744	1744	0	1726	1726	0	1651	1651	0	1752	1752	0
50-2-3-2	1825	1825	0	2052	2052	0	1946	1946	0	1910	1910	0	1828	1828	0
50-2-3-3	1874	1874	0	1819	1819	0	1844	1844	0	1862	1862	0	1857	1857	0

Table B.8: Dupli-2, raw results.

Problems	1			2			3			4			5		
30-0-0-1	523158	1158	870	541227	1227	900	577248	1248	960	577247	1247	960	550189	1189	915
30-0-0-2	307145	1145	510	334229	1229	555	316209	1209	525	325192	1192	540	325189	1189	540
30-0-0-3	181106	1106	300	109092	1092	180	190247	1247	315	208225	1225	345	181137	1137	300
30-0-3-1	361402	1402	600	334345	1345	555	379440	1440	630	379394	1394	630	289268	1268	480
30-0-3-2	370350	1350	615	397379	1379	660	343298	1298	570	379322	1322	630	379337	1337	630
30-0-3-3	730296	1296	1215	694293	1293	1155	748334	1334	1245	667240	1240	1110	721327	1327	1200
30-2-0-1	622515	1515	1035	586430	1430	975	577438	1438	960	559421	1421	930	577393	1393	960
30-2-0-2	568531	1531	945	541466	1466	900	523410	1410	870	514433	1433	855	604581	1581	1005
30-2-0-3	1414495	1495	2355	1405462	1462	2340	1378369	1369	2295	1396593	1593	2325	1387458	1458	2310
30-2-3-1	1099570	1570	1830	1090556	1556	1815	1144737	1737	1905	1135582	1582	1890	1153547	1547	1920
30-2-3-2	775523	1523	1290	793558	1558	1320	820581	1581	1365	820634	1634	1365	820580	1580	1365
30-2-3-3	730610	1610	1215	694566	1566	1155	712613	1613	1185	712629	1629	1185	739595	1595	1230
40-0-0-1	1232	1232	0	1175	1175	0	1162	1162	0	1212	1212	0	1104	1104	0
40-0-0-2	1194	1194	0	1202	1202	0	1122	1122	0	1142	1142	0	1201	1201	0
40-0-0-3	1101	1101	0	1236	1236	0	1129	1129	0	1182	1182	0	1186	1186	0
40-0-3-1	1355	1355	0	1400	1400	0	1393	1393	0	1378	1378	0	1377	1377	0
40-0-3-2	1320	1320	0	1340	1340	0	1366	1366	0	1294	1294	0	1342	1342	0
40-0-3-3	1417	1417	0	1342	1342	0	1516	1516	0	1441	1441	0	1467	1467	0
40-2-0-1	361799	1799	600	334685	1685	555	352687	1687	585	361793	1793	600	352686	1686	585
40-2-0-2	1623	1623	0	1755	1755	0	1905	1905	0	1739	1739	0	1675	1675	0
40-2-0-3	1651	1651	0	1769	1769	0	1642	1642	0	1591	1591	0	1744	1744	0
40-2-3-1	2130	2130	0	1944	1944	0	1878	1878	0	2069	2069	0	1996	1996	0
40-2-3-2	1670	1670	0	1782	1782	0	1726	1726	0	1669	1669	0	1682	1682	0
40-2-3-3	1864	1864	0	55910	1910	90	55933	1933	90	91829	1829	150	119042	2042	195
50-0-0-1	1330	1330	0	1193	1193	0	1280	1280	0	1261	1261	0	1168	1168	0
50-0-0-2	1196	1196	0	1210	1210	0	1133	1133	0	1158	1158	0	1159	1159	0
50-0-0-3	1138	1138	0	1177	1177	0	1091	1091	0	1095	1095	0	1260	1260	0
50-0-3-1	1442	1442	0	1376	1376	0	1421	1421	0	1427	1427	0	1482	1482	0
50-0-3-2	1440	1440	0	1503	1503	0	1532	1532	0	1342	1342	0	1453	1453	0
50-0-3-3	1436	1436	0	1429	1429	0	1495	1495	0	1510	1510	0	1477	1477	0
50-2-0-1	1554	1554	0	1513	1513	0	1610	1610	0	1551	1551	0	1603	1603	0
50-2-0-2	1410	1410	0	1445	1445	0	1341	1341	0	1304	1304	0	1431	1431	0
50-2-0-3	1498	1498	0	1529	1529	0	1391	1391	0	1362	1362	0	1509	1509	0
50-2-3-1	1646	1646	0	1672	1672	0	1650	1650	0	1764	1764	0	1647	1647	0
50-2-3-2	1889	1889	0	1898	1898	0	1856	1856	0	1863	1863	0	1926	1926	0
50-2-3-3	1839	1839	0	1872	1872	0	1893	1893	0	1877	1877	0	1906	1906	0

Table B.9: Direct4-s:200+repair, raw results.

Problems	1			2			3			4			5		
30-0-0-1	577195	1195	960	604246	1246	1005	613290	1290	1020	640290	1290	1065	640333	1333	1065
30-0-0-2	397262	1262	660	397239	1239	660	406292	1292	675	379235	1235	630	388279	1279	645
30-0-0-3	244268	1268	405	262316	1316	435	271365	1365	450	235275	1275	390	244304	1304	405
30-0-3-1	298260	1260	495	352353	1353	585	298257	1257	495	280246	1246	465	334282	1282	555
30-0-3-2	388386	1386	645	343255	1255	570	307308	1308	510	379376	1376	630	388350	1350	645
30-0-3-3	703251	1251	1170	613135	1135	1020	694282	1282	1155	730342	1342	1215	685268	1268	1140
30-2-0-1	712632	1632	1185	622519	1519	1035	721609	1609	1200	685475	1475	1140	676487	1487	1125
30-2-0-2	613709	1709	1020	577680	1680	960	604656	1656	1005	586666	1666	975	595746	1746	990
30-2-0-3	1387545	1545	2310	1441623	1623	2400	1423566	1566	2370	1432620	1620	2385	1414558	1558	2355
30-2-3-1	1090588	1588	1815	1072631	1631	1785	1126620	1620	1875	1117663	1663	1860	1081669	1669	1800
30-2-3-2	829634	1634	1380	793600	1600	1320	748600	1600	1245	757601	1601	1260	811581	1581	1350
30-2-3-3	694596	1596	1155	649551	1551	1080	649513	1513	1080	676542	1542	1125	676555	1555	1125
40-0-0-1	1144	1144	0	1252	1252	0	1192	1192	0	1255	1255	0	1278	1278	0
40-0-0-2	1174	1174	0	1123	1123	0	1161	1161	0	1252	1252	0	1172	1172	0
40-0-0-3	1212	1212	0	1303	1303	0	1136	1136	0	1204	1204	0	1169	1169	0
40-0-3-1	1449	1449	0	1323	1323	0	1342	1342	0	1459	1459	0	1319	1319	0
40-0-3-2	1412	1412	0	1232	1232	0	1326	1326	0	1295	1295	0	1369	1369	0
40-0-3-3	1456	1456	0	1414	1414	0	1351	1351	0	1444	1444	0	1407	1407	0
40-2-0-1	361970	1970	600	424899	1899	705	379939	1939	630	343790	1790	570	379865	1865	630
40-2-0-2	1895	1895	0	1655	1655	0	1735	1735	0	2002	2002	0	1745	1745	0
40-2-0-3	1824	1824	0	1911	1911	0	1833	1833	0	1611	1611	0	1926	1926	0
40-2-3-1	1898	1898	0	1859	1859	0	1921	1921	0	1823	1823	0	1815	1815	0
40-2-3-2	1739	1739	0	1738	1738	0	1618	1618	0	1632	1632	0	1614	1614	0
40-2-3-3	19844	1844	30	1820	1820	0	28906	1906	45	46898	1898	75	19837	1837	30
50-0-0-1	1315	1315	0	1297	1297	0	1234	1234	0	1297	1297	0	1342	1342	0
50-0-0-2	1262	1262	0	1288	1288	0	1216	1216	0	1239	1239	0	1271	1271	0
50-0-0-3	1222	1222	0	1206	1206	0	1297	1297	0	1193	1193	0	1185	1185	0
50-0-3-1	1374	1374	0	1383	1383	0	1476	1476	0	1471	1471	0	1410	1410	0
50-0-3-2	1439	1439	0	1394	1394	0	1368	1368	0	1405	1405	0	1327	1327	0
50-0-3-3	1494	1494	0	1370	1370	0	1439	1439	0	1419	1419	0	1493	1493	0
50-2-0-1	1627	1627	0	1597	1597	0	1637	1637	0	1527	1527	0	1604	1604	0
50-2-0-2	1459	1459	0	1440	1440	0	1445	1445	0	1414	1414	0	1472	1472	0
50-2-0-3	1415	1415	0	1390	1390	0	1465	1465	0	1425	1425	0	1433	1433	0
50-2-3-1	1701	1701	0	1762	1762	0	1587	1587	0	1552	1552	0	1609	1609	0
50-2-3-2	1785	1785	0	1861	1861	0	1722	1722	0	1812	1812	0	1942	1942	0
50-2-3-3	1815	1815	0	1811	1811	0	1754	1754	0	1845	1845	0	1742	1742	0

Table B.10: Dupli-1+repair, raw results.

Problems	1			2			3			4			5		
30-0-0-1	622222	1222	1035	631326	1326	1050	631313	1313	1050	641202	1202	900	622240	1240	1035
30-0-0-2	307147	1147	510	379260	1260	630	370282	1282	615	415294	1294	690	352237	1237	585
30-0-0-3	199240	1240	330	271287	1287	450	226256	1256	375	253313	1313	420	208247	1247	345
30-0-3-1	361289	1289	600	370328	1328	615	397421	1421	660	388391	1391	645	370372	1372	615
30-0-3-2	397377	1377	660	379361	1361	630	442327	1327	735	433350	1350	720	469415	1415	780
30-0-3-3	712323	1323	1185	676226	1226	1125	757379	1379	1260	703271	1271	1170	685237	1237	1140
30-2-0-1	685527	1527	1140	703489	1489	1170	694545	1545	1155	676538	1538	1125	721537	1537	1200
30-2-0-2	613584	1584	1020	595639	1639	990	622643	1643	1035	604579	1579	1005	568621	1621	945
30-2-0-3	1432505	1505	2385	1432471	1471	2385	1432507	1507	2385	1432634	1634	2385	1423616	1616	2370
30-2-3-1	1117723	1723	1860	1135678	1678	1890	1081590	1590	1800	1126573	1573	1875	1072575	1575	1785
30-2-3-2	829636	1636	1380	910690	1690	1515	910608	1608	1515	829511	1511	1380	775603	1603	1290
30-2-3-3	676545	1545	1125	721639	1639	1200	748600	1600	1245	721648	1648	1200	712512	1512	1185
40-0-0-1	1085	1085	0	1160	1160	0	1059	1059	0	1128	1128	0	1070	1070	0
40-0-0-2	1061	1061	0	1082	1082	0	1044	1044	0	1060	1060	0	1061	1061	0
40-0-0-3	986	986	0	1104	1104	0	1026	1026	0	978	978	0	1038	1038	0
40-0-3-1	1378	1378	0	1321	1321	0	1335	1335	0	1327	1327	0	1401	1401	0
40-0-3-2	1257	1257	0	1278	1278	0	1343	1343	0	1282	1282	0	1326	1326	0
40-0-3-3	1360	1360	0	1358	1358	0	1471	1471	0	1412	1412	0	1359	1359	0
40-2-0-1	370866	1866	615	352845	1845	585	388882	1882	645	370627	1627	615	370929	1929	615
40-2-0-2	1660	1660	0	1892	1892	0	1714	1714	0	1629	1629	0	1757	1757	0
40-2-0-3	1699	1699	0	1814	1814	0	1679	1679	0	1799	1799	0	1631	1631	0
40-2-3-1	1818	1818	0	1962	1962	0	2096	2096	0	1849	1849	0	2028	2028	0
40-2-3-2	1669	1669	0	1709	1709	0	1868	1868	0	1763	1763	0	1719	1719	0
40-2-3-3	46985	1985	75	55906	1906	90	64824	1824	105	1951	1951	0	64946	1946	105
50-0-0-1	1118	1118	0	1216	1216	0	1143	1143	0	1201	1201	0	1221	1221	0
50-0-0-2	1097	1097	0	1110	1110	0	1039	1039	0	1077	1077	0	1012	1012	0
50-0-0-3	1026	1026	0	1081	1081	0	1040	1040	0	1069	1069	0	1006	1006	0
50-0-3-1	1362	1362	0	1383	1383	0	1424	1424	0	1297	1297	0	1336	1336	0
50-0-3-2	1350	1350	0	1427	1427	0	1372	1372	0	1303	1303	0	1321	1321	0
50-0-3-3	1378	1378	0	1316	1316	0	1369	1369	0	1336	1336	0	1450	1450	0
50-2-0-1	1522	1522	0	1604	1604	0	1599	1599	0	1506	1506	0	1439	1439	0
50-2-0-2	1310	1310	0	1301	1301	0	1230	1230	0	1367	1367	0	1422	1422	0
50-2-0-3	1352	1352	0	1499	1499	0	1348	1348	0	1424	1424	0	1380	1380	0
50-2-3-1	1574	1574	0	1597	1597	0	1620	1620	0	1513	1513	0	1613	1613	0
50-2-3-2	1911	1911	0	1901	1901	0	1790	1790	0	1908	1908	0	1802	1802	0
50-2-3-3	1792	1792	0	1897	1897	0	1724	1724	0	1633	1633	0	1735	1735	0

Table B.11: Dupli-2+repair, raw results.

Problems	1			2			3			4			5		
30-0-0-1	973918	1918	1620	982869	1869	1635	964890	1890	1605	973925	1925	1620	982882	1882	1635
30-0-0-2	748926	1926	1245	784974	1974	1305	793962	1962	1320	766951	1951	1275	784968	1968	1305
30-0-0-3	613955	1955	1020	631968	1968	1050	632010	2010	1050	595929	1929	990	613992	1992	1020
30-0-3-1	757993	1993	1260	758009	2009	1260	748967	1967	1245	767014	2014	1275	785062	2062	1305
30-0-3-2	821011	2011	1365	821068	2068	1365	812072	2072	1350	821055	2055	1365	803023	2023	1335
30-0-3-3	1090903	1903	1815	1090884	1884	1815	1108919	1919	1845	1090926	1926	1815	1081943	1943	1800
30-2-0-1	901974	1974	1500	910967	1967	1515	910998	1998	1515	911029	2029	1515	910974	1974	1515
30-2-0-2	767054	2054	1275	767066	2066	1275	776063	2063	1290	776057	2057	1290	776066	2066	1290
30-2-0-3	1495836	1836	2490	1495845	1845	2490	1504872	1872	2505	1504810	1810	2505	1495877	1877	2490
30-2-3-1	1361186	2186	2265	1370027	2027	2280	1352164	2164	2250	1370176	2176	2280	1370108	2108	2280
30-2-3-2	1163096	2096	1935	1154147	2147	1920	1154056	2056	1920	1163081	2081	1935	1136054	2054	1890
30-2-3-3	901984	1984	1500	884083	2083	1470	920183	2183	1530	911141	2141	1515	929124	2124	1545
40-0-0-1	2272	2272	0	2272	2272	0	2251	2251	0	2265	2265	0	2174	2174	0
40-0-0-2	2183	2183	0	2235	2235	0	2260	2260	0	2245	2245	0	2232	2232	0
40-0-0-3	2166	2166	0	2205	2205	0	2204	2204	0	2168	2168	0	2190	2190	0
40-0-3-1	2306	2306	0	2333	2333	0	2315	2315	0	2313	2313	0	2309	2309	0
40-0-3-2	2205	2205	0	2216	2216	0	2223	2223	0	2175	2175	0	2181	2181	0
40-0-3-3	2356	2356	0	2426	2426	0	2412	2412	0	2380	2380	0	2339	2339	0
40-2-0-1	461248	2248	765	452312	2312	750	443139	2139	735	425304	2304	705	443374	2374	735
40-2-0-2	65155	2155	105	56199	2199	90	56219	2219	90	56219	2219	90	38244	2244	60
40-2-0-3	56272	2272	90	56379	2379	90	56330	2330	90	56369	2369	90	56321	2321	90
40-2-3-1	20561	2561	30	2573	2573	0	47581	2581	75	2480	2480	0	38451	2451	60
40-2-3-2	2320	2320	0	2255	2255	0	2226	2226	0	2310	2310	0	2285	2285	0
40-2-3-3	173288	2288	285	191347	2347	315	218421	2421	360	191351	2351	315	182433	2433	300
50-0-0-1	2317	2317	0	2310	2310	0	2309	2309	0	2300	2300	0	2333	2333	0
50-0-0-2	2236	2236	0	2255	2255	0	2230	2230	0	2251	2251	0	2204	2204	0
50-0-0-3	2241	2241	0	2225	2225	0	2237	2237	0	2258	2258	0	2212	2212	0
50-0-3-1	2279	2279	0	2284	2284	0	2260	2260	0	2296	2296	0	2263	2263	0
50-0-3-2	2335	2335	0	2393	2393	0	2365	2365	0	2398	2398	0	2335	2335	0
50-0-3-3	2338	2338	0	2343	2343	0	2321	2321	0	2323	2323	0	2304	2304	0
50-2-0-1	2156	2156	0	2232	2232	0	2215	2215	0	2209	2209	0	2108	2108	0
50-2-0-2	2116	2116	0	2097	2097	0	2105	2105	0	2111	2111	0	2110	2110	0
50-2-0-3	2133	2133	0	2117	2117	0	2075	2075	0	2141	2141	0	2132	2132	0
50-2-3-1	2220	2220	0	2175	2175	0	2171	2171	0	2185	2185	0	2171	2171	0
50-2-3-2	2313	2313	0	2392	2392	0	2358	2358	0	2397	2397	0	2313	2313	0
50-2-3-3	2412	2412	0	2388	2388	0	2352	2352	0	2390	2390	0	2386	2386	0

Table B.12: Random Search, raw results.

Problems	1			2			3			4			5		
30-0-0-1	577218	1218	960	595220	1220	990	586219	1219	975	586197	1197	975	559161	1161	930
30-0-0-2	388233	1233	645	388252	1252	645	406300	1300	675	442322	1322	735	370259	1259	615
30-0-0-3	271307	1307	450	298393	1393	495	208242	1242	345	235257	1257	390	208283	1283	345
30-0-3-1	451513	1513	750	433475	1475	720	460461	1461	765	460500	1500	765	442445	1445	735
30-0-3-2	514530	1530	855	505524	1524	840	487454	1454	810	523534	1534	870	460510	1510	765
30-0-3-3	784348	1348	1305	802373	1373	1335	811386	1386	1350	739295	1295	1230	757338	1338	1260
30-2-0-1	595393	1393	990	586392	1392	975	622423	1423	1035	595380	1380	990	622371	1371	1035
30-2-0-2	604535	1535	1005	568501	1501	945	514448	1448	855	586493	1493	975	532277	1277	885
30-2-0-3	1386946	946	2310	1395929	929	2325	1404936	936	2340	1395868	868	2325	1413951	951	2355
30-2-3-1	1207505	1505	2010	1279500	1500	2130	1225485	1485	2040	1252498	1498	2085	1162438	1438	1935
30-2-3-2	973518	1518	1620	892558	1558	1485	946487	1487	1575	865629	1629	1440	946561	1561	1575
30-2-3-3	865370	1370	1440	748474	1474	1245	685229	1229	1140	748346	1346	1245	757368	1368	1260
40-0-0-1	965	965	0	895	895	0	990	990	0	820	820	0	932	932	0
40-0-0-2	907	907	0	958	958	0	943	943	0	906	906	0	772	772	0
40-0-0-3	1044	1044	0	991	991	0	795	795	0	950	950	0	891	891	0
40-0-3-1	1303	1303	0	1288	1288	0	1290	1290	0	1305	1305	0	1366	1366	0
40-0-3-2	1109	1109	0	1169	1169	0	1167	1167	0	1119	1119	0	1269	1269	0
40-0-3-3	1278	1278	0	1281	1281	0	1279	1279	0	1353	1353	0	1334	1334	0
40-2-0-1	316020	1020	525	352012	1012	585	369978	978	615	333998	998	555	306990	990	510
40-2-0-2	55093	1093	90	993	993	0	1165	1165	0	1095	1095	0	1125	1125	0
40-2-0-3	1065	1065	0	1187	1187	0	1120	1120	0	1151	1151	0	1026	1026	0
40-2-3-1	46389	1389	75	73490	1490	120	55464	1464	90	100377	1377	165	1294	1294	0
40-2-3-2	1362	1362	0	1402	1402	0	1277	1277	0	1308	1308	0	1253	1253	0
40-2-3-3	235368	1368	390	100320	1320	165	154425	1425	255	163369	1369	270	136329	1329	225
50-0-0-1	950	950	0	1010	1010	0	1031	1031	0	1039	1039	0	924	924	0
50-0-0-2	849	849	0	940	940	0	872	872	0	986	986	0	854	854	0
50-0-0-3	781	781	0	796	796	0	872	872	0	802	802	0	875	875	0
50-0-3-1	1167	1167	0	1231	1231	0	1158	1158	0	1197	1197	0	1129	1129	0
50-0-3-2	1200	1200	0	1230	1230	0	1170	1170	0	1310	1310	0	1184	1184	0
50-0-3-3	1126	1126	0	1250	1250	0	1222	1222	0	1083	1083	0	1153	1153	0
50-2-0-1	1032	1032	0	1037	1037	0	1086	1086	0	1081	1081	0	1101	1101	0
50-2-0-2	845	845	0	880	880	0	809	809	0	934	934	0	909	909	0
50-2-0-3	1012	1012	0	956	956	0	899	899	0	1033	1033	0	936	936	0
50-2-3-1	1158	1158	0	1262	1262	0	1237	1237	0	1293	1293	0	1270	1270	0
50-2-3-2	1377	1377	0	1387	1387	0	1327	1327	0	1414	1414	0	1366	1366	0
50-2-3-3	1249	1249	0	1324	1324	0	1286	1286	0	1314	1314	0	1350	1350	0

Table B.13: Hill-climbing, raw results.

Problems	1			2			3			4			5		
30-0-0-1	433029	1029	720	451038	1038	750	442058	1058	735	451083	1083	750	424025	1025	705
30-0-0-2	235090	1090	390	262087	1087	435	289123	1123	480	253077	1077	420	253101	1101	420
30-0-0-3	118158	1158	195	118144	1144	195	100118	1118	165	100131	1131	165	100115	1115	165
30-0-3-1	325379	1379	540	370414	1414	615	343381	1381	570	307343	1343	510	343390	1390	570
30-0-3-2	415468	1468	690	325298	1298	540	316301	1301	525	343336	1336	570	352340	1340	585
30-0-3-3	622218	1218	1035	622198	1198	1035	658265	1265	1095	640226	1226	1065	649223	1223	1080
30-2-0-1	496342	1342	825	505381	1381	840	523398	1398	870	487335	1335	810	514418	1418	855
30-2-0-2	468982	982	780	450913	913	750	460206	1206	765	487047	1047	810	460042	1042	765
30-2-0-3	1359744	744	2265	1359754	754	2265	1341751	751	2235	1359781	781	2265	1350816	816	2250
30-2-3-1	1054143	1143	1755	1018079	1079	1695	1045124	1124	1740	1054227	1227	1755	1090146	1146	1815
30-2-3-2	748576	1576	1245	811626	1626	1350	757567	1567	1260	811576	1576	1350	721604	1604	1200
30-2-3-3	676221	1221	1125	730157	1157	1215	667180	1180	1110	712105	1105	1185	703237	1237	1170
40-0-0-1	821	821	0	751	751	0	768	768	0	759	759	0	741	741	0
40-0-0-2	824	824	0	747	747	0	780	780	0	772	772	0	802	802	0
40-0-0-3	746	746	0	745	745	0	744	744	0	782	782	0	788	788	0
40-0-3-1	977	977	0	949	949	0	982	982	0	927	927	0	984	984	0
40-0-3-2	884	884	0	854	854	0	892	892	0	931	931	0	887	887	0
40-0-3-3	963	963	0	993	993	0	1045	1045	0	1000	1000	0	989	989	0
40-2-0-1	279991	991	465	261970	970	435	243960	960	405	243968	968	405	243966	966	405
40-2-0-2	921	921	0	869	869	0	864	864	0	884	884	0	955	955	0
40-2-0-3	943	943	0	915	915	0	915	915	0	941	941	0	913	913	0
40-2-3-1	1131	1131	0	1147	1147	0	1192	1192	0	1151	1151	0	1156	1156	0
40-2-3-2	1037	1037	0	940	940	0	1031	1031	0	1090	1090	0	1011	1011	0
40-2-3-3	1273	1273	0	28172	1172	45	1218	1218	0	1158	1158	0	37189	1189	60
50-0-0-1	814	814	0	842	842	0	792	792	0	837	837	0	808	808	0
50-0-0-2	827	827	0	762	762	0	868	868	0	811	811	0	819	819	0
50-0-0-3	791	791	0	762	762	0	750	750	0	774	774	0	774	774	0
50-0-3-1	979	979	0	924	924	0	923	923	0	922	922	0	958	958	0
50-0-3-2	984	984	0	958	958	0	969	969	0	980	980	0	983	983	0
50-0-3-3	974	974	0	939	939	0	1016	1016	0	960	960	0	972	972	0
50-2-0-1	899	899	0	882	882	0	901	901	0	906	906	0	900	900	0
50-2-0-2	713	713	0	722	722	0	729	729	0	728	728	0	698	698	0
50-2-0-3	825	825	0	806	806	0	832	832	0	823	823	0	846	846	0
50-2-3-1	1024	1024	0	1014	1014	0	980	980	0	1002	1002	0	982	982	0
50-2-3-2	1116	1116	0	1135	1135	0	1152	1152	0	1138	1138	0	1141	1141	0
50-2-3-3	1120	1120	0	1094	1094	0	1090	1090	0	1139	1139	0	1103	1103	0

Table B.14: Simulated Annealing, raw results.

Problems	1			2			3			4			5		
30-0-0-1	541177	1177	900	568217	1217	945	568211	1211	945	577238	1238	960	541168	1168	900
30-0-0-2	361255	1255	600	343175	1175	570	397285	1285	660	361240	1240	600	280098	1098	465
30-0-0-3	172185	1185	285	262342	1342	435	208267	1267	345	226274	1274	375	163174	1174	270
30-0-3-1	379425	1425	630	370415	1415	615	343376	1376	570	316325	1325	525	280302	1302	465
30-0-3-2	352323	1323	585	388362	1362	645	406398	1398	675	343285	1285	570	406406	1406	675
30-0-3-3	685259	1259	1140	667265	1265	1110	658215	1215	1095	712293	1293	1185	658210	1210	1095
30-2-0-1	523338	1338	870	559388	1388	930	550290	1290	915	577401	1401	960	541372	1372	900
30-2-0-2	468932	932	780	496031	1031	825	505026	1026	840	495906	906	825	496008	1008	825
30-2-0-3	1350904	904	2250	1386915	915	2310	1386912	912	2310	1368901	901	2280	1350890	890	2250
30-2-3-1	1081172	1172	1800	1036147	1147	1725	1090247	1247	1815	1090128	1128	1815	1090228	1228	1815
30-2-3-2	829685	1685	1380	802548	1548	1335	793490	1490	1320	775462	1462	1290	766621	1621	1275
30-2-3-3	676128	1128	1125	649167	1167	1080	658150	1150	1095	703166	1166	1170	676198	1198	1125
40-0-0-1	1171	1171	0	1187	1187	0	1132	1132	0	1177	1177	0	1121	1121	0
40-0-0-2	1179	1179	0	1143	1143	0	1120	1120	0	1214	1214	0	1225	1225	0
40-0-0-3	1106	1106	0	1023	1023	0	1094	1094	0	1045	1045	0	1128	1128	0
40-0-3-1	1193	1193	0	1157	1157	0	1228	1228	0	1149	1149	0	1133	1133	0
40-0-3-2	1082	1082	0	1035	1035	0	1098	1098	0	1094	1094	0	1081	1081	0
40-0-3-3	1147	1147	0	1201	1201	0	1210	1210	0	1182	1182	0	1258	1258	0
40-2-0-1	280079	1079	465	235098	1098	390	253041	1041	420	244082	1082	405	280048	1048	465
40-2-0-2	980	980	0	1102	1102	0	1047	1047	0	1057	1057	0	1017	1017	0
40-2-0-3	1083	1083	0	1055	1055	0	1052	1052	0	1105	1105	0	1065	1065	0
40-2-3-1	1214	1214	0	1219	1219	0	1203	1203	0	1206	1206	0	1252	1252	0
40-2-3-2	1065	1065	0	1086	1086	0	1020	1020	0	1014	1014	0	1098	1098	0
40-2-3-3	64216	1216	105	64235	1235	105	37264	1264	60	1265	1265	0	64247	1247	105
50-0-0-1	1287	1287	0	1249	1249	0	1290	1290	0	1181	1181	0	1236	1236	0
50-0-0-2	1103	1103	0	1174	1174	0	1269	1269	0	1157	1157	0	1149	1149	0
50-0-0-3	1056	1056	0	1059	1059	0	1127	1127	0	1092	1092	0	1154	1154	0
50-0-3-1	1207	1207	0	1272	1272	0	1227	1227	0	1233	1233	0	1142	1142	0
50-0-3-2	1185	1185	0	1245	1245	0	1225	1225	0	1203	1203	0	1138	1138	0
50-0-3-3	1090	1090	0	1170	1170	0	1228	1228	0	1217	1217	0	1159	1159	0
50-2-0-1	1131	1131	0	1048	1048	0	1071	1071	0	1071	1071	0	1080	1080	0
50-2-0-2	959	959	0	912	912	0	1036	1036	0	978	978	0	970	970	0
50-2-0-3	994	994	0	1043	1043	0	1012	1012	0	1039	1039	0	966	966	0
50-2-3-1	1059	1059	0	1072	1072	0	1080	1080	0	1102	1102	0	1083	1083	0
50-2-3-2	1184	1184	0	1191	1191	0	1209	1209	0	1205	1205	0	1157	1157	0
50-2-3-3	1130	1130	0	1168	1168	0	1166	1166	0	1143	1143	0	1165	1165	0

Table B.15: Tabu Search, raw results.

Problems	1			2			3			4			5		
30-0-0-1	414983	983	690	451092	1092	750	433016	1016	720	424015	1015	705	369951	951	615
30-0-0-2	135888	888	225	190007	1007	315	217045	1045	360	199033	1033	330	217040	1040	360
30-0-0-3	127090	1090	210	145140	1140	240	73049	1049	120	37005	1005	60	73082	1082	120
30-0-3-1	235186	1186	390	271225	1225	450	262271	1271	435	262161	1161	435	244236	1236	405
30-0-3-2	307216	1216	510	298240	1240	495	262217	1217	435	316257	1257	525	235160	1160	390
30-0-3-3	631200	1200	1050	613207	1207	1020	604167	1167	1005	604188	1188	1005	604122	1122	1005
30-2-0-1	523259	1259	870	460174	1174	765	460175	1175	765	460267	1267	765	514318	1318	855
30-2-0-2	469289	1289	780	469411	1411	780	496402	1402	825	487304	1304	810	505381	1381	840
30-2-0-3	1324101	1101	2205	1333121	1121	2220	1324093	1093	2205	1315067	1067	2190	1342138	1138	2235
30-2-3-1	1009417	1417	1680	1045579	1579	1740	1045498	1498	1740	1000493	1493	1665	982357	1357	1635
30-2-3-2	703433	1433	1170	703421	1421	1170	676456	1456	1125	802612	1612	1335	739531	1531	1230
30-2-3-3	613366	1366	1020	604371	1371	1005	613494	1494	1020	613388	1388	1020	595448	1448	990
40-0-0-1	924	924	0	971	971	0	846	846	0	950	950	0	883	883	0
40-0-0-2	876	876	0	848	848	0	878	878	0	886	886	0	875	875	0
40-0-0-3	910	910	0	828	828	0	825	825	0	841	841	0	812	812	0
40-0-3-1	1081	1081	0	1108	1108	0	1155	1155	0	1147	1147	0	1119	1119	0
40-0-3-2	1116	1116	0	1060	1060	0	1173	1173	0	1059	1059	0	1043	1043	0
40-0-3-3	1141	1141	0	1188	1188	0	1088	1088	0	1167	1167	0	1200	1200	0
40-2-0-1	271326	1326	450	280460	1460	465	289509	1509	480	253397	1397	420	262513	1513	435
40-2-0-2	1363	1363	0	1331	1331	0	1362	1362	0	1390	1390	0	1307	1307	0
40-2-0-3	1386	1386	0	1309	1309	0	1368	1368	0	1288	1288	0	1212	1212	0
40-2-3-1	1577	1577	0	1634	1634	0	1490	1490	0	1554	1554	0	1529	1529	0
40-2-3-2	1541	1541	0	1514	1514	0	1385	1385	0	1387	1387	0	1596	1596	0
40-2-3-3	1640	1640	0	1840	1840	0	1744	1744	0	1785	1785	0	1656	1656	0
50-0-0-1	904	904	0	919	919	0	929	929	0	1002	1002	0	973	973	0
50-0-0-2	909	909	0	902	902	0	878	878	0	856	856	0	932	932	0
50-0-0-3	871	871	0	839	839	0	866	866	0	889	889	0	846	846	0
50-0-3-1	1127	1127	0	1202	1202	0	1117	1117	0	1241	1241	0	1121	1121	0
50-0-3-2	1089	1089	0	1122	1122	0	1108	1108	0	1142	1142	0	1237	1237	0
50-0-3-3	1108	1108	0	1180	1180	0	1183	1183	0	1175	1175	0	1108	1108	0
50-2-0-1	1248	1248	0	1195	1195	0	1251	1251	0	1283	1283	0	1323	1323	0
50-2-0-2	989	989	0	1055	1055	0	1011	1011	0	1139	1139	0	1175	1175	0
50-2-0-3	1191	1191	0	1174	1174	0	1105	1105	0	1220	1220	0	1155	1155	0
50-2-3-1	1404	1404	0	1355	1355	0	1402	1402	0	1416	1416	0	1312	1312	0
50-2-3-2	1527	1527	0	1578	1578	0	1558	1558	0	1561	1561	0	1585	1585	0
50-2-3-3	1468	1468	0	1584	1584	0	1438	1438	0	1498	1498	0	1502	1502	0

Table B.16: Genetic Algorithm (CPU intensive runs), raw results.

Problems	1			2			3			4			5		
30-0-0-1	811507	1507	1350	703398	1398	1170	775480	1480	1290	649338	1338	1080	784530	1530	1305
30-0-0-2	550472	1472	915	523399	1399	870	505396	1396	840	541420	1420	900	487364	1364	810
30-0-0-3	325358	1358	540	451571	1571	750	325353	1353	540	415431	1431	690	343449	1449	570
30-0-3-1	613618	1618	1020	622589	1589	1035	667725	1725	1110	523446	1446	870	559425	1425	930
30-0-3-2	640361	1361	1065	640392	1392	1065	631591	1591	1050	604435	1435	1005	577461	1461	960
30-0-3-3	892499	1499	1485	919447	1447	1530	910526	1526	1515	1000628	1628	1665	901454	1454	1500
30-2-0-1	703196	1196	1170	820602	1602	1365	694243	1243	1155	784517	1517	1305	784536	1536	1305
30-2-0-2	631257	1257	1050	640439	1439	1065	676147	1147	1125	685401	1401	1140	613108	1108	1020
30-2-0-3	1495101	1101	2490	1476993	993	2460	1567055	1055	2610	1566915	915	2610	1495022	1022	2490
30-2-3-1	1315346	1346	2190	1414423	1423	2355	1414448	1448	2355	1459534	1534	2430	1306482	1482	2175
30-2-3-2	1198520	1520	1995	1225431	1431	2040	1342378	1378	2235	1144614	1614	1905	1018502	1502	1695
30-2-3-3	982361	1361	1635	910418	1418	1515	856348	1348	1425	982307	1307	1635	847351	1351	1410
40-0-0-1	793	793	0	885	885	0	865	865	0	873	873	0	873	873	0
40-0-0-2	853	853	0	820	820	0	701	701	0	810	810	0	875	875	0
40-0-0-3	779	779	0	707	707	0	715	715	0	785	785	0	707	707	0
40-0-3-1	1037	1037	0	990	990	0	1031	1031	0	974	974	0	998	998	0
40-0-3-2	1128	1128	0	997	997	0	988	988	0	1000	1000	0	1204	1204	0
40-0-3-3	1011	1011	0	1021	1021	0	1144	1144	0	1016	1016	0	1020	1020	0
40-2-0-1	622062	1062	1035	514015	1015	855	577057	1057	960	514049	1049	855	559052	1052	930
40-2-0-2	72961	961	120	117964	964	195	63999	999	105	63986	986	105	54977	977	90
40-2-0-3	136021	1021	225	63995	995	105	63976	976	105	127073	1073	210	63962	962	105
40-2-3-1	109363	1363	180	1207	1207	0	55206	1206	90	1357	1357	0	109304	1304	180
40-2-3-2	1168	1168	0	1127	1127	0	64301	1301	105	1136	1136	0	1188	1188	0
40-2-3-3	208283	1283	345	343291	1291	570	325277	1277	540	190304	1304	315	325288	1288	540
50-0-0-1	820	820	0	799	799	0	873	873	0	862	862	0	892	892	0
50-0-0-2	768	768	0	735	735	0	757	757	0	791	791	0	700	700	0
50-0-0-3	669	669	0	679	679	0	738	738	0	744	744	0	765	765	0
50-0-3-1	1016	1016	0	987	987	0	983	983	0	1199	1199	0	967	967	0
50-0-3-2	1031	1031	0	1097	1097	0	1050	1050	0	970	970	0	1018	1018	0
50-0-3-3	954	954	0	931	931	0	1002	1002	0	959	959	0	981	981	0
50-2-0-1	924	924	0	918	918	0	908	908	0	953	953	0	895	895	0
50-2-0-2	785	785	0	846	846	0	793	793	0	751	751	0	748	748	0
50-2-0-3	893	893	0	916	916	0	898	898	0	949	949	0	864	864	0
50-2-3-1	1050	1050	0	1129	1129	0	1086	1086	0	1084	1084	0	1115	1115	0
50-2-3-2	1322	1322	0	46223	1223	75	1223	1223	0	1319	1319	0	1343	1343	0
50-2-3-3	1338	1338	0	1303	1303	0	1127	1127	0	1212	1212	0	1378	1378	0

Table B.17: Hill-climbing (CPU intensive runs), raw results.

Problems	1			2			3			4			5		
30-0-0-1	955853	1853	1590	946891	1891	1575	937881	1881	1560	955874	1874	1590	955859	1859	1590
30-0-0-2	757914	1914	1260	748901	1901	1245	739902	1902	1230	739889	1889	1230	748892	1892	1245
30-0-0-3	577896	1896	960	586919	1919	975	577921	1921	960	586974	1974	975	595948	1948	990
30-0-3-1	730988	1988	1215	685959	1959	1140	731003	2003	1215	730973	1973	1215	721970	1970	1200
30-0-3-2	775996	1996	1290	776014	2014	1290	748946	1946	1245	776032	2032	1290	775999	1999	1290
30-0-3-3	1045888	1888	1740	1063873	1873	1770	1054893	1893	1755	1054890	1890	1755	1081901	1901	1800
30-2-0-1	865927	1927	1440	865900	1900	1440	865961	1961	1440	874916	1916	1455	838895	1895	1395
30-2-0-2	740024	2024	1230	740035	2035	1230	739986	1986	1230	740003	2003	1230	740016	2016	1230
30-2-0-3	1486842	1842	2475	1486786	1786	2475	1477819	1819	2460	1477949	1949	2460	1486879	1879	2475
30-2-3-1	1298194	2194	2160	1325277	2277	2205	1325197	2197	2205	1325108	2108	2205	1298250	2250	2160
30-2-3-2	1100099	2099	1830	1100008	2008	1830	1118088	2088	1860	1127064	2064	1875	1118185	2185	1860
30-2-3-3	830063	2063	1380	866069	2069	1440	839097	2097	1395	839038	2038	1395	857088	2088	1425
40-0-0-1	2214	2214	0	2174	2174	0	2193	2193	0	2234	2234	0	2200	2200	0
40-0-0-2	2169	2169	0	2203	2203	0	2196	2196	0	2179	2179	0	2185	2185	0
40-0-0-3	2137	2137	0	2125	2125	0	2121	2121	0	2161	2161	0	2137	2137	0
40-0-3-1	2234	2234	0	2246	2246	0	2260	2260	0	2226	2226	0	2261	2261	0
40-0-3-2	2118	2118	0	2158	2158	0	2157	2157	0	2154	2154	0	2156	2156	0
40-0-3-3	2372	2372	0	2328	2328	0	2376	2376	0	2271	2271	0	2358	2358	0
40-2-0-1	416151	2151	690	407332	2332	675	434313	2313	720	434077	2077	720	389253	2253	645
40-2-0-2	38286	2286	60	29371	2371	45	29261	2261	45	2233	2233	0	38213	2213	60
40-2-0-3	29379	2379	45	2380	2380	0	38384	2384	60	29306	2306	45	2401	2401	0
40-2-3-1	2583	2583	0	2515	2515	0	2418	2418	0	2451	2451	0	2507	2507	0
40-2-3-2	2119	2119	0	2184	2184	0	2215	2215	0	2225	2225	0	2238	2238	0
40-2-3-3	146364	2364	240	128271	2271	210	164383	2383	270	155356	2356	255	146403	2403	240
50-0-0-1	2238	2238	0	2248	2248	0	2250	2250	0	2252	2252	0	2252	2252	0
50-0-0-2	2198	2198	0	2144	2144	0	2147	2147	0	2151	2151	0	2178	2178	0
50-0-0-3	2186	2186	0	2158	2158	0	2190	2190	0	2149	2149	0	2174	2174	0
50-0-3-1	2205	2205	0	2243	2243	0	2228	2228	0	2237	2237	0	2257	2257	0
50-0-3-2	2294	2294	0	2295	2295	0	2310	2310	0	2258	2258	0	2265	2265	0
50-0-3-3	2275	2275	0	2245	2245	0	2262	2262	0	2270	2270	0	2249	2249	0
50-2-0-1	2134	2134	0	2185	2185	0	2148	2148	0	2134	2134	0	2156	2156	0
50-2-0-2	2027	2027	0	2039	2039	0	2065	2065	0	2021	2021	0	2009	2009	0
50-2-0-3	2079	2079	0	2085	2085	0	2063	2063	0	2075	2075	0	2072	2072	0
50-2-3-1	2107	2107	0	2152	2152	0	2165	2165	0	2118	2118	0	2167	2167	0
50-2-3-2	2348	2348	0	2321	2321	0	2322	2322	0	2343	2343	0	2332	2332	0
50-2-3-3	2378	2378	0	2317	2317	0	2335	2335	0	2316	2316	0	2342	2342	0

Table B.18: Random Search (CPU intensive runs), raw results.

Problems	1			2			3			4			5		
30-0-0-1	415020	1020	690	397005	1005	660	387985	985	645	351918	918	585	396991	991	660
30-0-0-2	217078	1078	360	190030	1030	315	199014	1014	330	172010	1010	285	190029	1029	315
30-0-0-3	100131	1131	165	100143	1143	165	73113	1113	120	118181	1181	195	145213	1213	240
30-0-3-1	316355	1355	525	289288	1288	480	280269	1269	465	361437	1437	600	325397	1397	540
30-0-3-2	316308	1308	525	316279	1279	525	325318	1318	540	370394	1394	615	361360	1360	600
30-0-3-3	550093	1093	915	649231	1231	1080	604207	1207	1005	649252	1252	1080	595197	1197	990
30-2-0-1	487359	1359	810	460306	1306	765	478329	1329	795	496379	1379	825	523411	1411	870
30-2-0-2	441843	843	735	441852	852	735	450843	843	750	468853	853	780	441821	821	735
30-2-0-3	1341757	757	2235	1332698	698	2220	1341717	717	2235	1350715	715	2250	1332691	691	2220
30-2-3-1	1063154	1154	1770	991022	1022	1650	1018091	1091	1695	1081132	1132	1800	1054127	1127	1755
30-2-3-2	739624	1624	1230	685532	1532	1140	811649	1649	1350	721582	1582	1200	793589	1589	1320
30-2-3-3	667112	1112	1110	586084	1084	975	667139	1139	1110	649123	1123	1080	640036	1036	1065
40-0-0-1	706	706	0	693	693	0	678	678	0	695	695	0	708	708	0
40-0-0-2	703	703	0	668	668	0	692	692	0	682	682	0	692	692	0
40-0-0-3	681	681	0	670	670	0	672	672	0	686	686	0	646	646	0
40-0-3-1	851	851	0	881	881	0	876	876	0	866	866	0	866	866	0
40-0-3-2	823	823	0	832	832	0	862	862	0	808	808	0	783	783	0
40-0-3-3	919	919	0	922	922	0	917	917	0	894	894	0	917	917	0
40-2-0-1	234947	947	390	261962	962	435	225949	949	375	243928	928	405	225949	949	375
40-2-0-2	814	814	0	809	809	0	812	812	0	820	820	0	782	782	0
40-2-0-3	860	860	0	852	852	0	847	847	0	857	857	0	853	853	0
40-2-3-1	1072	1072	0	1042	1042	0	1056	1056	0	1067	1067	0	1078	1078	0
40-2-3-2	953	953	0	959	959	0	941	941	0	975	975	0	938	938	0
40-2-3-3	1143	1143	0	1056	1056	0	1079	1079	0	1087	1087	0	1055	1055	0
50-0-0-1	743	743	0	729	729	0	692	692	0	740	740	0	731	731	0
50-0-0-2	725	725	0	713	713	0	718	718	0	707	707	0	710	710	0
50-0-0-3	689	689	0	681	681	0	664	664	0	673	673	0	662	662	0
50-0-3-1	875	875	0	880	880	0	907	907	0	873	873	0	848	848	0
50-0-3-2	899	899	0	911	911	0	900	900	0	919	919	0	900	900	0
50-0-3-3	882	882	0	877	877	0	909	909	0	877	877	0	901	901	0
50-2-0-1	883	883	0	886	886	0	887	887	0	876	876	0	887	887	0
50-2-0-2	701	701	0	700	700	0	715	715	0	692	692	0	688	688	0
50-2-0-3	812	812	0	826	826	0	817	817	0	812	812	0	823	823	0
50-2-3-1	942	942	0	948	948	0	929	929	0	929	929	0	950	950	0
50-2-3-2	1060	1060	0	1054	1054	0	1064	1064	0	1047	1047	0	1051	1051	0
50-2-3-3	1023	1023	0	1016	1016	0	1052	1052	0	1036	1036	0	1014	1014	0

Table B.19: Simulated Annealing (CPU intensive runs), raw results.

Problems	1			2			3			4			5		
30-0-0-1	460062	1062	765	442028	1028	735	415000	1000	690	414974	974	690	433010	1010	720
30-0-0-2	253085	1085	420	198963	963	330	253038	1038	420	271072	1072	450	226039	1039	375
30-0-0-3	109116	1116	180	127127	1127	210	109134	1134	180	82076	1076	135	145187	1187	240
30-0-3-1	226237	1237	375	271292	1292	450	289314	1314	480	289307	1307	480	289271	1271	480
30-0-3-2	334300	1300	555	280210	1210	465	334286	1286	555	307305	1305	510	307264	1264	510
30-0-3-3	586143	1143	975	604146	1146	1005	604186	1186	1005	604163	1163	1005	631204	1204	1050
30-2-0-1	559414	1414	930	496300	1300	825	451165	1165	750	487299	1299	810	505332	1332	840
30-2-0-2	468877	877	780	450864	864	750	441816	816	735	459887	887	765	477878	878	795
30-2-0-3	1332958	958	2220	1350853	853	2250	1350867	867	2250	1350850	850	2250	1341818	818	2235
30-2-3-1	1054135	1135	1755	1063170	1170	1770	1063148	1148	1770	1063192	1192	1770	1027048	1048	1710
30-2-3-2	703549	1549	1170	757412	1412	1260	712265	1265	1185	730560	1560	1215	793611	1611	1320
30-2-3-3	649141	1141	1080	640132	1132	1065	649118	1118	1080	658108	1108	1095	631079	1079	1050
40-0-0-1	1099	1099	0	1138	1138	0	1236	1236	0	1153	1153	0	1207	1207	0
40-0-0-2	1172	1172	0	1165	1165	0	1167	1167	0	1182	1182	0	1200	1200	0
40-0-0-3	1123	1123	0	1132	1132	0	1168	1168	0	1127	1127	0	1069	1069	0
40-0-3-1	1273	1273	0	1233	1233	0	1250	1250	0	1282	1282	0	1240	1240	0
40-0-3-2	1174	1174	0	1088	1088	0	1131	1131	0	1167	1167	0	1172	1172	0
40-0-3-3	1312	1312	0	1238	1238	0	1262	1262	0	1279	1279	0	1281	1281	0
40-2-0-1	334132	1132	555	352059	1059	585	289126	1126	480	316106	1106	525	352120	1120	585
40-2-0-2	1032	1032	0	989	989	0	1007	1007	0	1015	1015	0	1020	1020	0
40-2-0-3	1120	1120	0	1013	1013	0	1116	1116	0	1149	1149	0	1131	1131	0
40-2-3-1	1148	1148	0	1175	1175	0	1137	1137	0	1165	1165	0	1176	1176	0
40-2-3-2	1054	1054	0	1055	1055	0	1052	1052	0	1042	1042	0	1031	1031	0
40-2-3-3	1066	1066	0	1166	1166	0	1223	1223	0	1148	1148	0	1179	1179	0
50-0-0-1	1231	1231	0	1255	1255	0	1179	1179	0	1237	1237	0	1255	1255	0
50-0-0-2	1108	1108	0	1153	1153	0	1130	1130	0	1128	1128	0	1237	1237	0
50-0-0-3	1180	1180	0	1204	1204	0	1194	1194	0	1086	1086	0	1099	1099	0
50-0-3-1	1205	1205	0	1190	1190	0	1270	1270	0	1225	1225	0	1262	1262	0
50-0-3-2	1291	1291	0	1321	1321	0	1284	1284	0	1213	1213	0	1146	1146	0
50-0-3-3	1298	1298	0	1280	1280	0	1275	1275	0	1271	1271	0	1216	1216	0
50-2-0-1	1139	1139	0	1144	1144	0	1166	1166	0	1089	1089	0	1142	1142	0
50-2-0-2	982	982	0	1017	1017	0	1000	1000	0	986	986	0	1008	1008	0
50-2-0-3	1065	1065	0	1067	1067	0	1038	1038	0	991	991	0	1083	1083	0
50-2-3-1	1059	1059	0	1092	1092	0	1092	1092	0	1057	1057	0	1068	1068	0
50-2-3-2	1203	1203	0	1169	1169	0	1179	1179	0	1194	1194	0	1180	1180	0
50-2-3-3	1136	1136	0	1169	1169	0	1159	1159	0	1171	1171	0	1145	1145	0

Table B.20: Tabu Search (CPU intensive runs), raw results.