Multi-Objective Optimization of Urban Water Resource Systems

By

Seyed Mohammad Mortazavi Naeini

B.Sc., Civil Engineering

M.Sc., Water Engineering

A Thesis submitted for the

Degree of Doctor of Philosophy at

The University of Newcastle

January 2013

Declaration

"This Thesis contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying subject to the provisions of the Copyright Act 1968."

(Signed) ______________________

Acknowledgement

It would not have been possible to finish this thesis without the help and support of the people around me, to only some of whom it is possible to give particular mention here.

First and foremost, my utmost gratitude to my supervisor, Prof. George Kuczera, whose sincerity and encouragement I will never forget. If it was not his patience and support it would not be possible for me to explore my wild ideas in this thesis. I will always remember the moments when I went to his office with worries and his patience and support made me calm.

I would like to show my gratitude to my invaluable network of supportive, forgiving, generous and caring friends without whom I could not have survived the process. Among them, I would like to thank my office mates, Andrew, Chaminda, Yousef, Mason, Nurul and Glen. Andrew, I really appreciate your effort to teach me Aussie words! If it was not because of your sense of humour I would feel more homesick here. Chaminda, I will always remember our inspiring discussions. Yousef, I am grateful for your help and support when you generously allowed me to use your computers to run my models.

I would like to thank the staff of our faculty for their help. Among them I would like to thank Lijie and Tom. Likewise, I would like to thanks Cherie, Ann and Josh for being accommodating to my queries and for all the help.

I would like to acknowledge eWater financial support for this research and their support for travelling to present my papers in conferences.

I owe my deepest gratitude to my family and especially to my parents. They always supported me during all years of my education. My mum always reminded me that it will be possible to finish this thesis when I was in down moods. Thank you for all of your help and support mum.

Table of Contents

List of Figures

List of Tables

Abstract

The provision of a water supply that is secure in the face of severe drought is a primary objective for urban water agencies – "running out of water" is not a viable option for a large city. However, there are other objectives that conflict with the primary one – these include minimizing costs and environmental impacts. A major challenge facing decision makers in the urban water sector is dealing with the trade-offs between these conflicting objectives. Multi-objective optimization methods have the potential to identify the optimal trade-offs between the competing objectives. The principal aim of this thesis is to address the shortcomings in existing multi-objective optimization applications to produce methods of greater practical relevance to urban water resource management.

Review of past studies identified three practically significant shortcomings. Focusing exclusively on either long-term (or infrastructure) options or on short-term options such as operation rules may lead to sub-optimal solutions. The use of short climate forcing data time series in simulation models to evaluate drought security can produce solutions that make the system highly vulnerable to severe drought. Finally, the setting of *a priori* environmental constraints may hide trade-offs between environmental, economic and security factors that are of considerable interest to decision makers. These shortcomings are addressed by a new multi-objective methodology that exploits the ability of evolutionary algorithms to handle complex objective functions and simulation models. The principal novelty is the explicit treatment of drought security. A case study based on the headworks system for Australia's largest city, Sydney, demonstrates the practical significance of these shortcomings and, importantly, the ability of the new approach to deal with these shortcomings in a practicable manner.

In the face of urban population growth and the accompanying growth in water demand, the performance of the urban water resource system is expected to deteriorate over time. This will result in the need to intervene and adapt the system to the changing conditions. The scheduling capacity expansion problem seeks to identify the optimal schedule for the changes to the system. In past studies, this problem has

been largely tackled by minimizing the total present worth of capital, operational and rationing costs. A significant drawback of minimizing the total present worth cost is that it is likely to produce solutions that lead to more severe and frequent rationing in the future. Such a solution is likely to be socially unacceptable. A multi-objective formulation for the scheduling capacity expansion problem is developed to overcome this shortcoming while addressing the need to explicitly deal with drought security and jointly optimize operating and infrastructure decisions. The formulation enables the trade-off between cost and equity (the equal sharing of the burden of restrictions over the planning horizon) to be explored. A case study based on the headworks system for Australia's capital city, Canberra, demonstrates the advantages of the new approach.

The optimization of urban water resource systems requires running simulation models tens of thousands of times. Given that simulation run times can range from less than a minute to thirty or more minutes, it is important to use a multi-objective optimization method which converges with the least number of evaluations (or simulations). To address this need, a detailed assessment is conducted of three benchmark multi-objective optimization methods and three newly developed methods based on ant colony optimization using case studies based on the Canberra and Sydney systems. No one method emerges as superior, although two of the six methods are identified as inferior.

Publications

One journal and seven conference papers have been published using work undertaken in this thesis:

1. Mortazavi, M., G. Kuczera and L. Cui (2012). "Multiobjective optimization of urban water resources: Moving toward more practical solutions." Water Resour. Res. 48(3): W03514.
2. Mortazavi N., S. M., G. Kuczera and L. Cui (2012). Application of multi objective optimization for managing urban drought security in the presence of population growth. 10th International Conference on Hydroinformatics (HIC 2012). Hamburg, Germany.
3. Mortazavi N., S. M., G. Kuczera and L. Cui (2010). New robust multiobjective ant colony optimization (MOACO) method. 9th International Conference on Hydroinformatics (HIC 2010). Tianjin, China: 2285-2292.
4. Kuczera, G., L. Cui, R. Gilmore, A. Graddon and S. M. Mortazavi N. (2010). Enhancing the robustness of water resource simulation models based on network linear programming. 9th International Conference on Hydroinformatics (HIC 2010). Tianjin, China: 2245-2252.
5. Mortazavi N., S. M., L. Cui and G. Kuczera (2009). Application of multiobjective optimization methods for urban water management: a case study for Canberra water supply system. H2009. Newcastle, Australia: 887-898.
6. Mortazavi N., S. M., L. Cui and G. Kuczera (2009). Comparison of Genetic Algorithm and Ant Colony optimization methods for optimization of short-term drought mitigation strategies. Hydroinformatics in Hydrology, Hydrogeology and Water Resources, Hyderabad, India, IAHS Publ. 331.
7. Kuczera, G., L. Cui, R. Grilmore, A. Graddon, S. M. Mortazavi N. and C. E. Jefferson (2009). Addressing the Shortcomings of Water Resource Simulation Models Based on Network Linear Programming. H2009 the 32nd Hydrology and Water Resources Symposium, Newcastle, Australia, Engineers Australia.
8. Cui, L., S. M. Mortazavi N. and G. Kuczera (2009). Comparison of multi-objective genetic algorithm with ant colony optimization: a case study for Canberra water supply system. 33rd IAHR Congress, 2009: Water Engineering for a Sustainable Environment, Vancouver, Canada.

Chapter 1

Introduction

1-1 Decision Support in Urban Water Management

The critical role of water in human society is incontestable. The development of many ancient civilizations close to water sources shows that humans have understood the importance of accessing water from early stages. The world's population has nearly tripled during the last century while exploitable water resources have remained largely unchanged. Increasing population has resulted in greater demand for water to support agriculture and urban water supply. Likewise the unprecedented growth in industry has further increased water demand in many developed and developing countries (UN-Water and FAO, 2007). This thesis considers the question of how to best manage urban water resources in the face of growing demand.

In an Australian industry position paper describing a framework for urban water resource planning, Erlanger and Neal (2005) state in the opening that: "A safe and reliable water supply system is of utmost importance to the community. It is expected and understood that water utilities manage their water resources so that communities never run out of water." Erlanger and Neal recognize that failure to supply minimum water needs for an extended period would most likely result in disastrous social and economic losses that could conceivably threaten the very existence of the urban community.

As an example of the challenges facing large urban centres, the water resource system supplying Sydney, Australian's largest city, is considered. The system, schematized in Figure 1-1, is complex with 11 major reservoirs serving Sydney's residential, commercial and industrial water demand. Its total capacity is more than 2,600 GL (gigalitres). The agencies responsible for Sydney's water supply have to deal with a number of future challenges including catering for a growing population, coping with high natural variability, future climate change and mitigating environmental impacts.

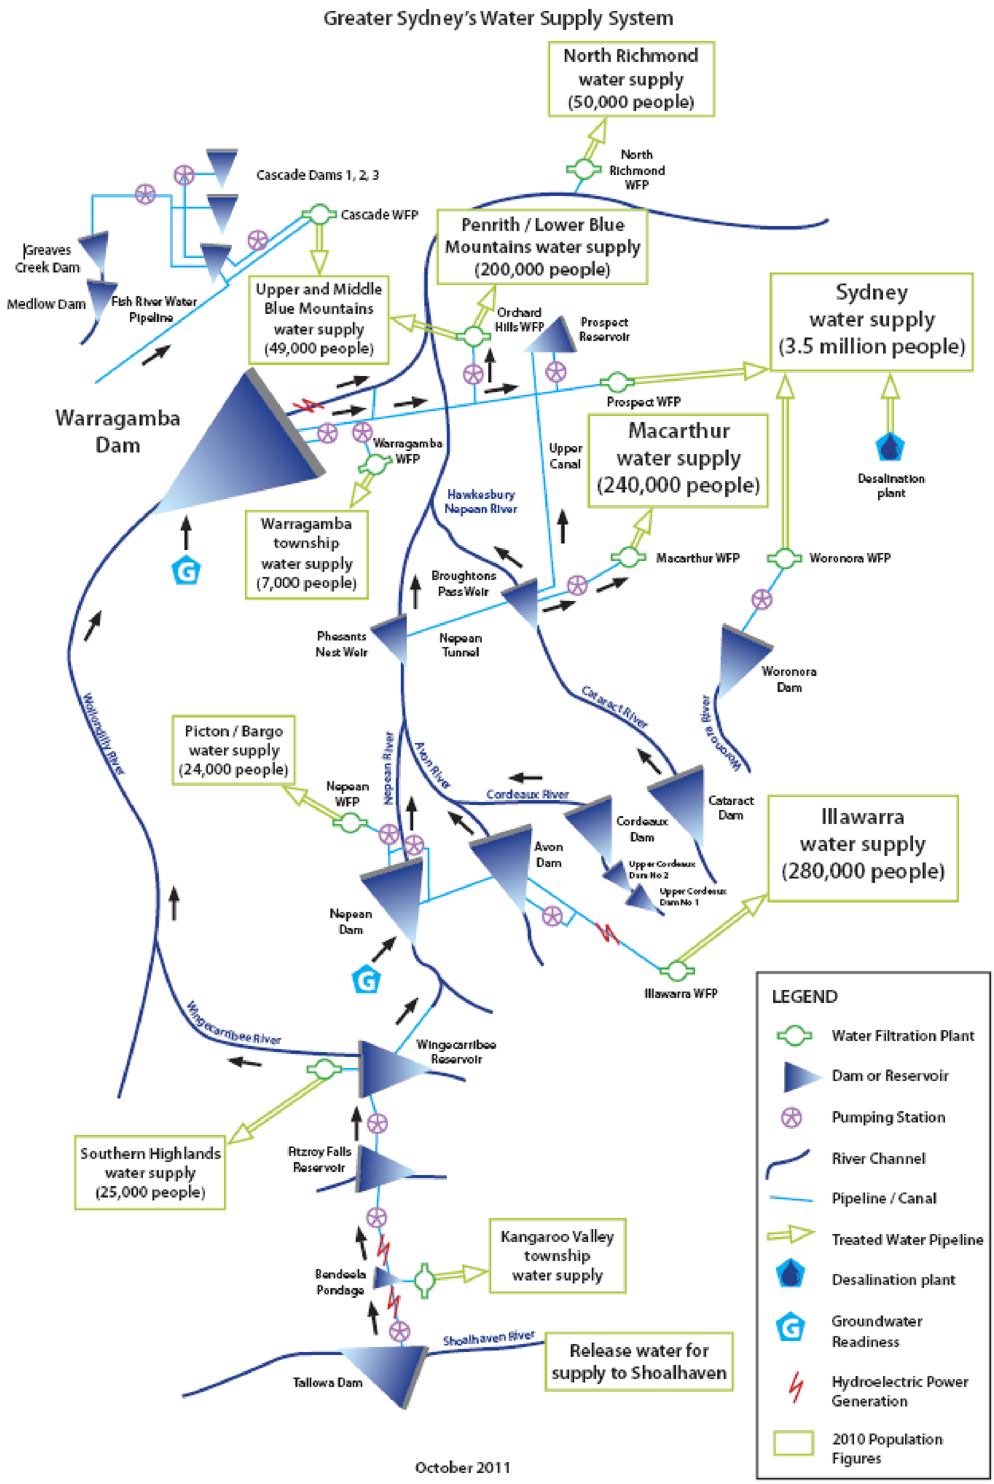

Figure 1-1 Sydney's water supply system

(Source : http://www.sca.nsw.gov.au/publications/publications/water-supply-diagram, last visit 05/05/2012)

Historical records indicate that Sydney's climate is highly variable and subject to prolonged drought, even without consideration of future climate change impacts. This is illustrated by Figure 1-2, which shows that Sydney's total storage dropped from 90% in 2001 to nearly 35% in 2007 during the severe drought that affected much of eastern Australia.

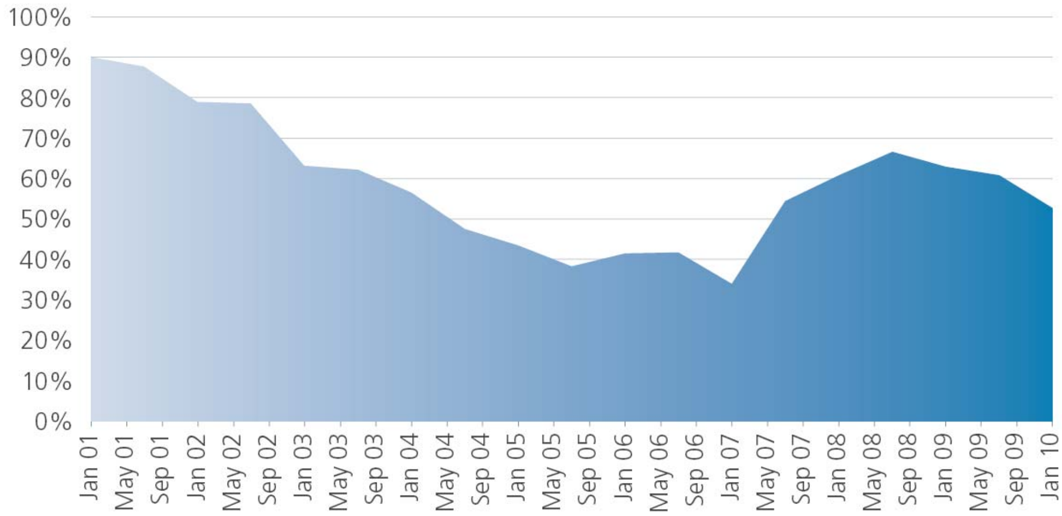

Figure 1-2 Sydney's total reservoir storage level (New South Wales Dept. of Environment and Water, 2010)

Sydney's population is expected to reach 5.7 million by 2035 which represents a 35% increase from 4.22 million in 2006. This population growth is likely to increase annual demand by 66 GL (New South Wales Dept. of Environment and Water, 2010). At the same time, the need to protect rivers and aquifers from environmental degradation has become a significant priority.

The agencies responsible for Sydney's water supply have available a range of options to tackle these challenges. The options involve either increasing supply by harvesting new sources of water or reducing demand by improving water use efficiency, pricing and rationing. As an example of the former, in 2007 when Sydney's total storage level was around 35%, the state government approved the construction of Sydney's first desalination plant as an emergency measure to reduce the risk of "running out" of water. With regard to demand reduction, annual demand has been reduced by over 100 GL since 1999 with the introduction of water efficiency programs. Moreover, several projects have been identified to increase recycled water from 33 GL/year in 2010 to 70 GL/year in 2015 (New South Wales Dept. of Environment and Water, 2010). Water sharing plans have been revised to ensure rivers and aquifers receive adequate environmental water and are not overused. In some cases, such as the upper Nepean River, major upgrades to a number of dams and weirs were implemented to improve natural fish passage in the Hawkesbury-Nepean River system (New South Wales Dept. of Environment and Water, 2010).

While the provision of a safe and secure water supply is a primary objective for urban water agencies, there are other objectives that conflict with the primary one – these include, inter alia, minimizing costs and minimizing environmental impacts. A major challenge facing decision makers in the urban water sector is dealing with the trade-offs between these conflicting objectives. Each solution involving a specific mix of options that increase supply and/or reduce demand will produce different outcomes with respect to supply security, cost and environmental impact. For example, construction of the 500 ML/day desalination plant would increase supply security at a cost of $1.9 billion. Alternatively, supply security could also be increased by imposing more frequent and severe restrictions at probably a lower economic cost but higher social cost than construction of the desalination plant. Yet again, supply security could be increased by relaxing environmental constraints in the Wollondilly River to allow greater transfers from the Shoalhaven River to Warragamba – this gain in security would be offset by a greater threat to the survival of water-dependent fauna in the Wollondilly River.

The challenges facing Sydney's water agencies are typical of those facing urban water agencies both in Australia and internationally. In broad terms, these agencies have to find solutions that maximize security of supply while minimizing cost and minimizing environmental impacts. The number of solutions involving different combinations of infrastructure investments and different operating rules can be astronomic. The challenge of finding the best set of solutions is considerable. The conventional approach of employing a trial-and-error search over a limited number of solutions runs the significant risk of missing good solutions and the consequent opportunity cost to the community.

Decision support systems are designed to assist decision makers to find good solutions. Of particular interest in this thesis is the use of multi-objective optimization methods to sift through all feasible solutions to identify those solutions that optimally trade-off two or more conflicting objectives. These methods have seen application to urban water resources planning and management over the last two decades. However, as will be shown, there remain significant shortcomings that limit their true potential.

1-2 Thesis Objectives

The primary goal of this thesis is to develop and demonstrate multi-objective optimization (MOO) approaches, which address the shortcomings of existing methods in order to move closer to providing practical and realistic optimization of urban water resources. To this end, several specific objectives are pursued in this thesis:

1. ***To address the shortcomings of existing multi-objective optimization in urban water resources planning and operation***

The practical value of MOO methods in urban water resources management depends on how well the optimization model represents the needs of the decision makers. Previous work in this area has fallen short in a number of practically important areas:

a) The performance of an urban water resource system is jointly dependent on the mix of infrastructure and operating rule options. Focussing exclusively on either infrastructure options or operating rules may lead to sub-optimal solutions.

b) Constraining objectives *a priori* may hide trade-offs of considerable interest to decision makers. For example, *a priori* specification of environmental flow rules can hide significant trade-offs between ecological, economic and security outcomes. Awareness of such trade-offs could result in communities willing to pay more in return for reduced environmental impacts.

c) Provision of adequate drought security is a key objective. Reliance on relatively short historic data to evaluate supply security runs a very significant risk of producing a system highly vulnerable to severe drought.

The first objective of this thesis has two components: 1) to formulate and solve the urban water resource optimization problem in a manner that better addresses the practical challenges of working with solutions that involve a mix of infrastructure and operating rule options, realistically account for drought risk and identify the trade-offs between economic, supply security and environmental factors; and 2) to demonstrate that failure to address these challenges comprehensively results in "optimal" solutions that can be far from optimal and of limited relevance to urban water resources management.

2. ***To extend application of multi-objective optimization to scheduling of options to cater for future changes***

The current world population of more than 7 billion is projected to reach 9.3 billion by the middle of this century (UN, 2011). Much of this growth will be in urban areas driving the demand for more water. This situation may be further exacerbated by future climate change (Palmer et al., 2008). The literature on optimal scheduling of future decisions is limited and tends to focus on minimizing present worth costs and on infrastructure investment. One consequence of focusing on minimizing present worth costs is that the temporal discounting of costs tends to result in lowered future levels of service. For example, the optimal minimum cost policy may lead to more severe restriction/rationing of consumption in future stages, an outcome unlikely to be acceptable to decision makers sensitive to temporal equity. The second objective of this thesis is to extend application of MOO to scheduling of decisions on infrastructure investment and operating rules to cater for future changes with consideration of equity over the planning period.

3. ***To identify and evaluate the most computationally efficient multi-objective optimization method for urban water resources application***

Water resources applications typically use computationally expensive methods for computing their objective functions (Pierro et al., 2009). For example, in one of the case studies presented in this thesis involving the Canberra water resource system, a 140-year simulation at monthly time steps takes approximately 2 CPU seconds. Hence, for an optimization involving 10,000 function evaluations, the turnaround time of nearly 6 hours is totally dominated by the simulation model rather than by the optimization algorithm. Urban water resource models typically use long stochastically generated records which can lead to simulation run times of the order of several minutes. These long run times are considered an impediment to the practical uptake of MOO. While parallel computing can reduce turnaround times (Cui and Kuczera, 2005), there remains a strong imperative to develop MOO methods which not only converge to the Pareto-optimal front with good diversity but do so with the fewest possible function evaluations. This is the final objective of this thesis.

There are two tasks to address the final objective of the thesis. The first is to identify and evaluate which of the existing MOO algorithms is best suited for urban water resource applications. The second is to explore the potential of a recently developed optimization method called ant colony optimization (ACO). The literature shows that ACO performs well for difficult combinatorial problems such as the travelling salesman problem. The challenge is to determine whether ACO can be successfully adapted to solve problems typical of urban water resources.

1-3 Thesis Outline

This thesis consists of six chapters. Chapter 2 presents an overview of simulation methods used in urban water resources and a review of existing MOO methods. It sets the scene for the following three chapters which represent the primary contribution of this thesis. Chapter 3 develops a MOO approach that addresses many of the shortcomings of existing applications and yet is computationally practicable. A case study involving the complex Sydney system demonstrates the benefits of the proposed approach. Chapter 4 considers the problem of optimally scheduling decisions over time from a MOO perspective. A case study involving the Canberra system demonstrates the advantages and insights that the MOO approach brings to this very difficult problem. Chapter 5 changes the focus from problem formulation to the algorithms that conduct the search for optimal solutions. It evaluates the performance of three benchmark algorithms using the Sydney and Canberra systems as exemplars of urban water resources applications and also investigates the potential of ACO. Chapter 6 concludes the thesis, summarizing its main findings and identifying future research directions.

Chapter 2
Simulation and Optimization in Urban Water Resource Management

2-1 Introduction

Decision support systems facilitate the process of decision making by providing insight to decision makers about the consequences of implementing various options. Two important components of a decision support system are the simulation model and tools that support optimizing outcomes important to the decision maker. Simulation models allow investigation of system behaviour under historical and future scenarios. In particular, they assist in answering “what-if” questions. Thus, by trial and error improved operating and planning strategies may be found. However, in most real-world problems, there is a huge number of technically feasible solutions. It is, therefore, problematic whether a trial-and-error search can identify near optimum solutions. The success of a trial-and-error search is likely to be very dependent on the skill of the analyst. There is a significant risk of missing good solutions and the consequent opportunity cost to the community. This challenge has inspired the development of a range of optimization methods in the last few decades to optimize operating rules and investment decisions in urban water resources systems. Indeed, this challenge is the primary interest of this thesis.

This chapter provides the necessary background for the ensuing chapters which present the main contributions of this thesis. In the first part, a review of simulation models used in urban water resources is conducted. This is followed by a more detailed description of the WATHNET5 simulation model that is used in the thesis case studies. In the second part, multi-objective optimization (MOO) is reviewed. Following a review of MOO concepts, the ε-dominance multi-objective evolutionary algorithm, εMOEA, is introduced and described prior to its use in the case studies presented in Chapters 3 and 4. In the third and final part, the software linkage between simulation and MOO models and the use of parallel computing is discussed to provide an understanding of how the case studies were implemented.

2-2 Water Resource Simulation Models

Simulation models are used widely to simulate the behaviour of the water resource systems for a given set of input conditions. These models can be generally categorized into two groups, namely reservoir-system-simulation models and system-analysis models based on a network-flow programming formulation (Wurbs, 1993; Labadie, 2004).

Reservoir-system-simulation models use operating rules to assign flows. Because most systems are operated using operating rules, these models are widely used by agencies responsible for planning and operation of water resource systems. While these models may be custom built, it is more common for a water agency to build a model of its system using a generalized model in which the system is represented by a network of nodes and arcs (which transfer water between nodes) and all system-specific data are stored in data files. The advantage of using a generalized model is that it is simpler to make changes to system configuration and operating rules.

There is a wide range of generalized reservoir-system-simulation models. The HEC-5 simulation flood control and conservation system (Hydrologic Engineering Center, 1998) has been used widely in studies of proposed new projects and operational modifications of existing systems. Other models include MITSIM and TAMUWARP (Wurbs, 1995), IRIS (Salewicz et al., 1991) and RiverWare (Zagona et al., 2001). MITSIM facilitates modelling alternative river basin development plans involving reservoirs, hydroelectric power plants, irrigation areas, and municipal and industrial water supply diversions. TAMUWARP is a simulation model developed for studies involving a priority-based allocation of water resources among many different water users. In a similar way IRIS, the interactive river system simulation, was developed with the aim of providing a useful tool for negotiating among stakeholders (Salewicz et al., 1991; Wurbs, 1995). A more recent development is RiverWare (Zagona et al., 2001). It can be used to simulate a wide range of river and reservoir configurations with diverse operational objectives and for applications ranging from small timescale scheduling to long-term planning.

System-analysis models are based on network-flow programming (NFP) which has been applied in a variety of operations research and systems engineering applications. System flows between nodes are not determined by operating rules but come from solution of the minimum cost capacitated network-flow problem (Wurbs, 1995). The NFP approach avoids much of the complexity of specifying rules for water transfers, which can be particularly challenging when multiple flow paths exist.

Nonetheless practical implementations of the NFP approach do use rules to determine the costs, capacities and requirements of the network flow program.

A range of algorithms have been developed to solve the NFP problem. They include the Out-of-Kilter (Fulkerson, 1961), RELAX (Bertsekas and Tseng, 1988) and simplex-on-a-graph NETFLO (Kennington and Helgason, 1980) algorithms. Kuczera (1993) compared RELAX and NETFLO and concluded that RELAX was the superior algorithm particularly when the NFPs were iterated many times at a given time step.

A significant limitation of the NFP approach arises when non-NFP constraints need to be imposed – for example, when the flow in one arc is related to the flow in another arc. The typical approach in such circumstances is to use fixed point iteration in which the solution from the previous iteration is used to change the non-NFP constraint into a NFP constraint. However, Ilich (2009) questioned the reliability of fixed point iteration. He presented examples in which the use of iteration to update non-network constraints may lead to convergence to the wrong solution. Kuczera et al. (2009) also highlight problems using fixed point iteration. In practice, modellers need to be careful when using the NFP approach and be aware that some systems cannot be robustly modelled using an NFP approach.

There exist a number of generalized models based on NFP: SIMYLD (Evenson and Moseley, 1970), ARSP (Sigvaldson, 1976), DWRSIM (Chung et al., 1989), CRAM (Brendecke et al., 1989), MODSIM (Labadie et al., 1986), KCOM (Andrews et al., 1992), WASP (Kuczera and Diment, 1988), REALM (Perera et al., 2005) and WATHNET5 (Kuczera et al., 2009). The REALM and WATHNET5 models are derivatives of the WASP model and are used by virtually all major urban water agencies in Australia.

In this thesis, the WATHNET5 simulation model is adapted. No claim is made about the superiority of WATHNET5 as a simulation model. Indeed the choice of simulation model is not central to this thesis. WATHNET5 was selected for three reasons: 1) the availability of the source code meant to software could be adapted to new and unplanned needs; 2) its architecture facilitates the implementation of multi-objective optimization which is the core focus of this thesis; and 3) a WATHNET5 model of the Sydney system, the most complex case study used in this thesis, was available.

2-3 The WATHNET5 Model

In an NFP model the water resource system is represented as a directed graph which is collection of nodes and a set of arcs. The nodes represent source, demand or transfer points on the network. The arcs represent flow paths from one node to another. In WATHNET5, two types of arcs are defined, namely stream arcs which represent rivers and conduit arcs which represent pipes. Six different nodes are defined in WATHNET5, namely stream, reservoir, demand, waste, harvest and junction nodes. Stream nodes represent a source of water to the system such as inflow to reservoirs or rainfall over a catchment. Reservoir nodes represent reservoirs and carryover storage from one time step to the next. Demand nodes represent sink points in the network. Junction nodes represent transfer points. Harvest nodes enable application of stochastic transfer functions such as in the modelling of domestic rainwater tank savings or run-of-river diversions at monthly time scales. Waste nodes act as a sink points to collect any water leaving the network.

In a network flow model, a transfer cost is assigned to all arcs. In order to force flow through an arc, for instance, an environmental flow arc, a high negative cost needs to be assigned. In WATHNET5, the NFP is formulated as follows:

$$\min_{Z} \quad c^T(x \mid Q_t, D_t, \theta) z \tag{2.1}$$

subject to

$$Az = b(x \mid Q_t, D_t, \theta) \tag{2.2}$$

$$0 \le z \le u(x \mid Q_t, D_t, \theta) \tag{2.3}$$

where Q_t and D_t are vectors of inflow and demand for time t respectively, θ is a vector of parameters assigned by the user, A is a node-arc incidence matrix and z is a vector of arc flows, x is a vector of decision variables (which can be optimized), $c(x \mid Q_t, D_t, \theta)$ is a vector of costs assigned to the arcs, $b(x \mid Q_t, D_t, \theta)$ is a vector of nodal requirements, either restricted demand or streamflow, and $u(x \mid Q_t, D_t, \theta)$ is a vector of maximum arc capacities. It is noted that $c(x|Q_t, D_t, \theta)$, $b(x|Q_t, D_t, \theta)$ and $u(x|Q_t, D_t, \theta)$ are vector functions of x, Q_t, D_t and θ whose algorithms are specified by the user using a FORTRAN-like script.

The formulation of the NFP in WATHNET5 is best described using an example based on the network shown in Figure 2-1. This network has two reservoirs and two demand nodes. Reservoir spill is collected by the waste node. The stream nodes provide stream inflow to the reservoirs. Figure 2-2 shows the full network including hidden arcs and the hidden balancing node. Without these hidden elements it would not be possible to simulate the system in Figure 2-1. The balance node ensures a mass balance for the network. The demand shortfall arc is assigned a very high cost and only conveys flow to the demand node if the demand cannot be satisfied by any other means. This ensures the NFP always returns a feasible solution even when demand cannot be satisfied by the real system. Waste nodes are connected to the balance node via waste arcs.

To simulate carryover of storage, one or more carryover arcs connect each reservoir to the balance node. By assigning sufficiently large gains (negative costs) to the carryover arcs, the NFP will assign flows to the carryover arcs in preference to assigning flows to a waste node.

WATHNET5 offers several options to assign carryover gains. These are illustrated in Figure 2-3 which shows the dialog box to assign carryover gains. All but one option involve some form of manual assignment of gains to individual carryover arcs. The remaining option, which is the one used in the thesis case studies, automates the assignment of gains using the following equation:

$$Gain(i) = BG + (i-1) * IG, i = 1, \dots, N \tag{2.4}$$

where *Gain(i)* is the gain assigned to the ith carryover arc, *BG* is the base gain, *IG* is incremental gain, and *N* is the number of carryover arcs. The capacity of each carryover arcs is set as follows:

$$u_i = \frac{\text{ResCap}}{N}, i = 1, ..., N \tag{2.5}$$

where u_i is the capacity of i[th] carryover arc for the reservoir and *ResCap* is the reservoir capacity.

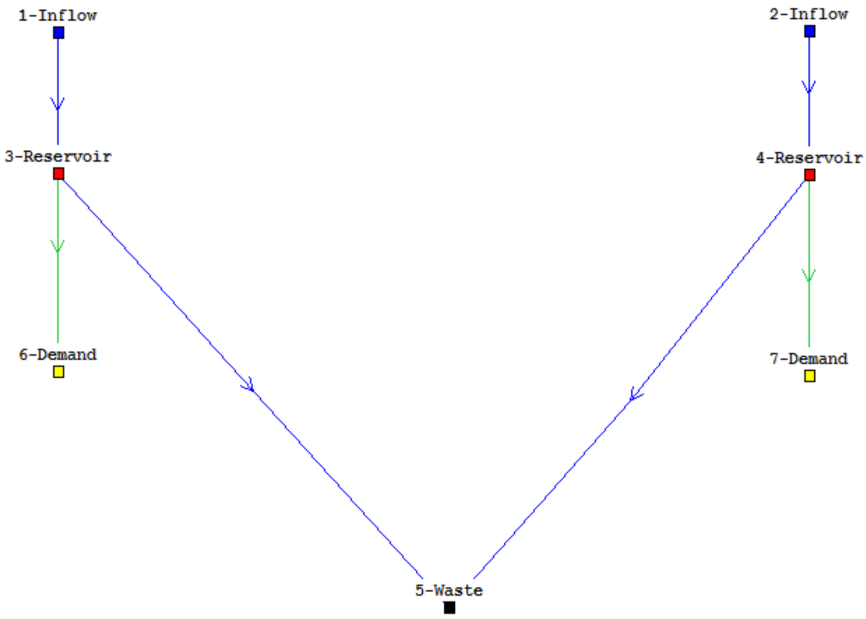

Figure 2-1 A simple network in WATHNET5 [adapted from Kuczera (1992)]

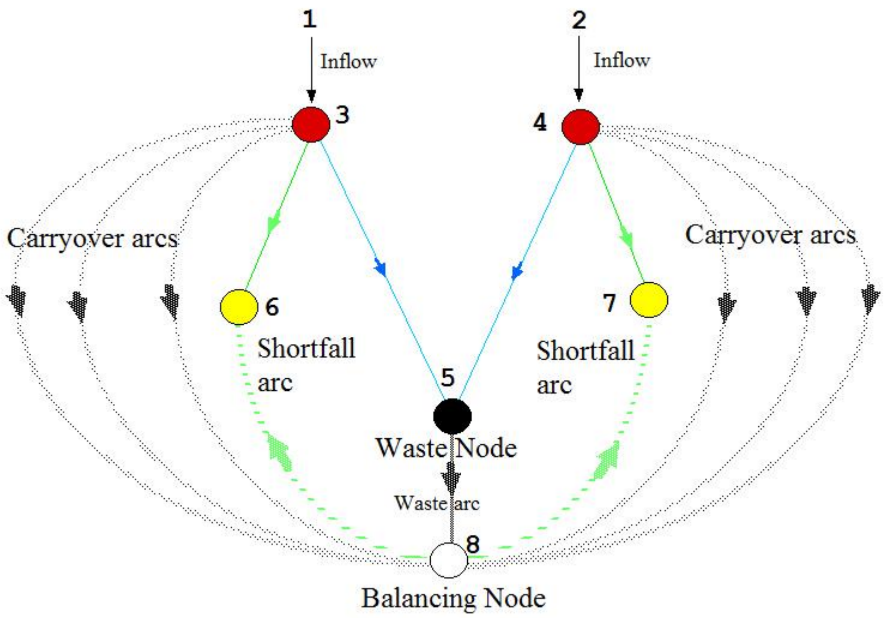

Figure 2-2 Full network including hidden arcs and nodes for network shown in Figure 2-1[modified from (Kuczera, 1992)]

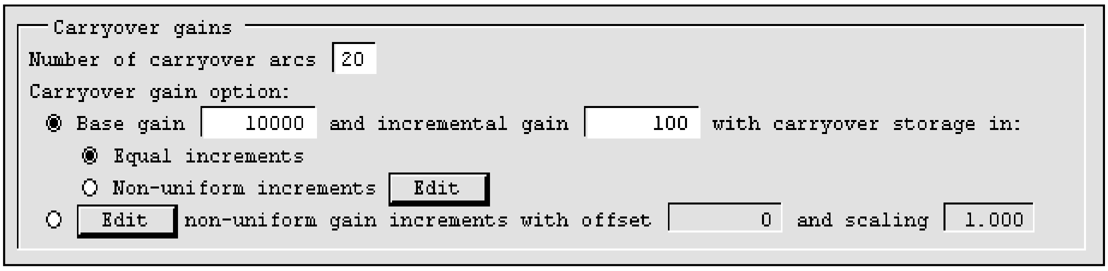

Figure 2-3 Carryover arcs input box

2-4 Multi-Objective Optimization

The main use of system operation models is to simulate system behaviour to answer “what-if” type of questions. However, the ultimate goal is to find the best solutions taking into account economic, social and environmental factors. Although it is theoretically possible to enumerate all possible solutions to find the optimum it is practically infeasible in most problems.

The purpose of optimization models is to facilitate finding optimum solutions. There is a vast literature on application of optimization methods in water resources planning with many studies focussing on reservoir operation. Yeh (1985), Wurbs (1993) and Labadie (2004) provide comprehensive reviews of optimization methods used in water resources.

Most of the reviewed studies have considered applications involving a single objective such as minimizing cost or demand shortages. However, inclusion of the environmental and social aspects into water planning naturally leads to multi-objective optimization in which there exist two or more objectives that conflict or cannot be optimized simultaneously.

Multi-objective optimization (MOO) has seen wide application in water resources management. Applications include reservoir operations (Ko et al., 1992; Liang et al., 1996; Kim et al., 2006; Reddy and Kumar, 2006; Chen et al., 2007; Reddy and Kumar, 2007a; Reddy and Kumar, 2007b; Consoli et al., 2008; Chang and Chang, 2009; Rani and Moreira, 2010), water distribution (Farmani et al., 2006; Mariano-Romero et al., 2007; Pierro et al., 2009), urban drainage (Barreto et al., 2007) and ground water (Kollat and Reed, 2006). More specifically, Chen et al. (2007) used MOO for optimizing a multi-purpose reservoir rule. Similarly, Consoli et al. (2008) optimized the operational rules for irrigation reservoirs employing two objectives. Yang et al. (2007) integrated a multi-objective genetic algorithm (MOGA) with constrained differential dynamic programming (CDDP); they applied CDDP to distribute optimal releases among reservoirs to satisfy water demand as much as possible and used multi-objective genetic algorithms to generate the various combinations of reservoir capacity. In a similar manner, Chang et al. (2009) hybridized the genetic algorithm (GA) and CDDP to optimize capacity expansion schedules for ground water supply; they used the GA to find the optimal capacity expansion options and the CDDP algorithm to find the optimal pumping policy associated with the selected expansion options.

The problems considered in this thesis involve optimization problems with K objectives, which are, without loss of generality, all to be minimized and all equally important. A solution is represented as a decision vector $x=(x_1, x_2, x_3, \ldots, x_n)$ in the decision space *X*. The quality of a specific solution is evaluated by a vector function $f(x) = \{f_1(x), \ldots, f_k(x)\}$ which assigns to a decision vector an objective vector $(f_1(x), \ldots, f_k(x))$ in the objective space F. The relation between the decision and objective spaces is illustrated in Figure 2-4.

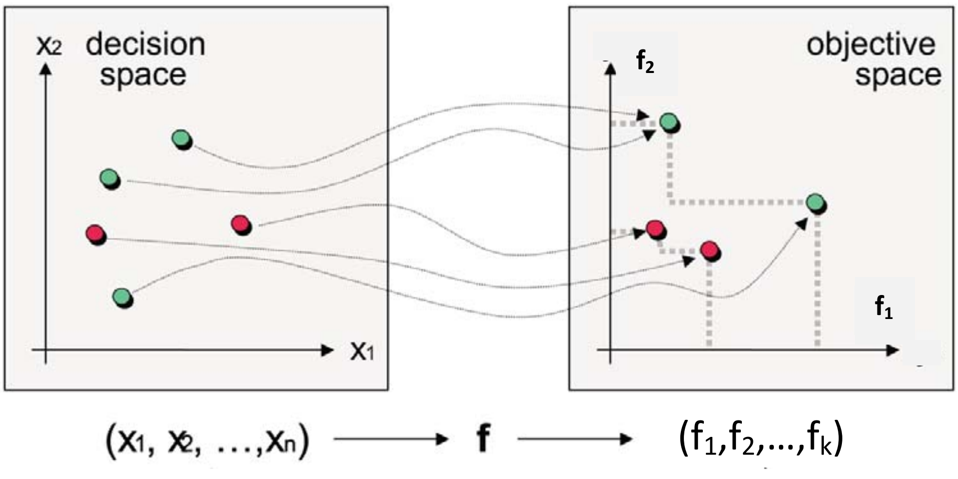

Figure 2-4 Illustration of decision and objective space of a multi-objective problem

In the case of a single-objective optimization problem, two solutions (x_1, x_2) can be compared easily based on their associated objective (or criterion) values $(f(x_1), f(x_2))$. However, in the case of multi-objective optimization it is necessary to introduce the concept of dominance.

Dominance definition: A solution[1] x_1 is said to dominate the solution x_2, if both of the following conditions are true (Deb 2001):

1. $f_j(x_1) \leq f_j(x_2)$ for all $j \in \{1,2,...,K\}$ - solution x_1 is no worse than x_2 in all objectives

2. $f_j(x_1) < f_j(x_2)$ for at least one $j \in \{1,2,...,K\}$ - solution x_1 is strictly better than solution x_2 in at least one objective

The set of non-dominated solutions for the whole search space X is called the Pareto-optimal set. The solutions belonging to the Pareto-optimal set are said to lie on the Pareto frontier or Pareto front.

The Pareto frontier is illustrated in Figure 2-5 for a two objective problem, minimizing cost and minimizing restriction frequency. Solution A dominates the solution B because it has a lower restriction frequency and a lower cost. However, A does not dominate C because A has a lower restriction frequency than C but a higher cost. For these reasons A and C are called non-dominated solutions. Indeed, it is not possible to find the optimum solution without any further information about criteria preferences.

Usually there is some higher-level information in every real optimization problem that depends on subjective assessment of social, political, economic and environmental factors not adequately captured in the formal optimization. This kind of information is usually non-technical, qualitative and experience-driven (Deb, 2001). Generally, there are two approaches to deal with this higher-level information. In Figures 2-6 and 2-7, schematics of these two approaches are shown. In the first approach called ideal multi-objective optimization, the Pareto-optimal solutions are found and then, using higher level information, one of the Pareto-optimal solutions is chosen as the preferred solution (Deb, 2001).

[1] The notation for decisions is context specific. Here x_1 refers to a decision vector with label "*1*" rather than the first component of the vector.

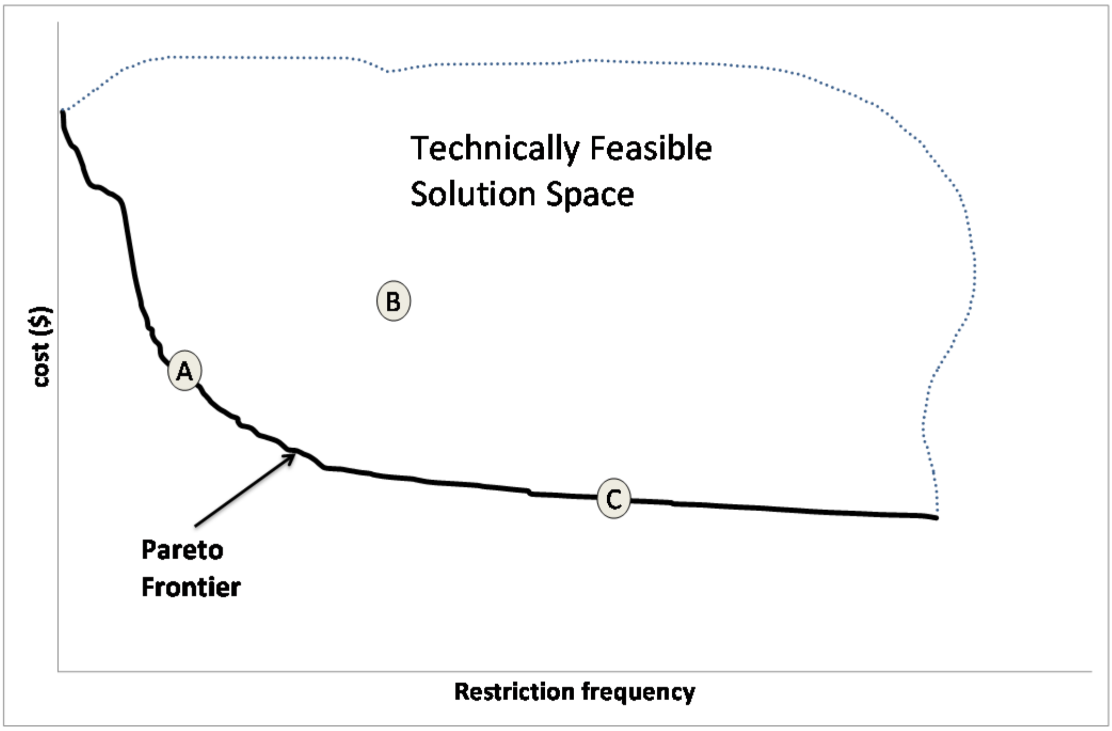

Figure 2-5 Concept of Pareto optimality

In complete contrast, the second approach, called the preference-based method, uses higher-level information at the start of optimization. This information is used to assign a relative importance vector which assigns a weight to each objective. Based on these weights a single objective function can be formulated. Let $w_1, w_2, ..., w_n$ be the weights assigned to the corresponding objectives. The single objective can thus be formulated as

$$F(x) = \sum_{i=1}^{K} w_i f_i(x) \tag{2.6}$$

It is important to note that the results obtained by using the preference-based method can be highly sensitive to the values assigned to the preference vector. Another problem with the preference-based method is that the weight vector needs to be supplied without any knowledge of the possible outcomes; it implicitly assumes that the weights are independent of outcomes. In many situations, decision makers would be reluctant to provide higher-level information or weights without knowledge of outcomes. It is thus concluded that the ideal method is the preferred approach. Therefore, methods for finding Pareto-optimal trade-off solutions will be the focus of the next section.

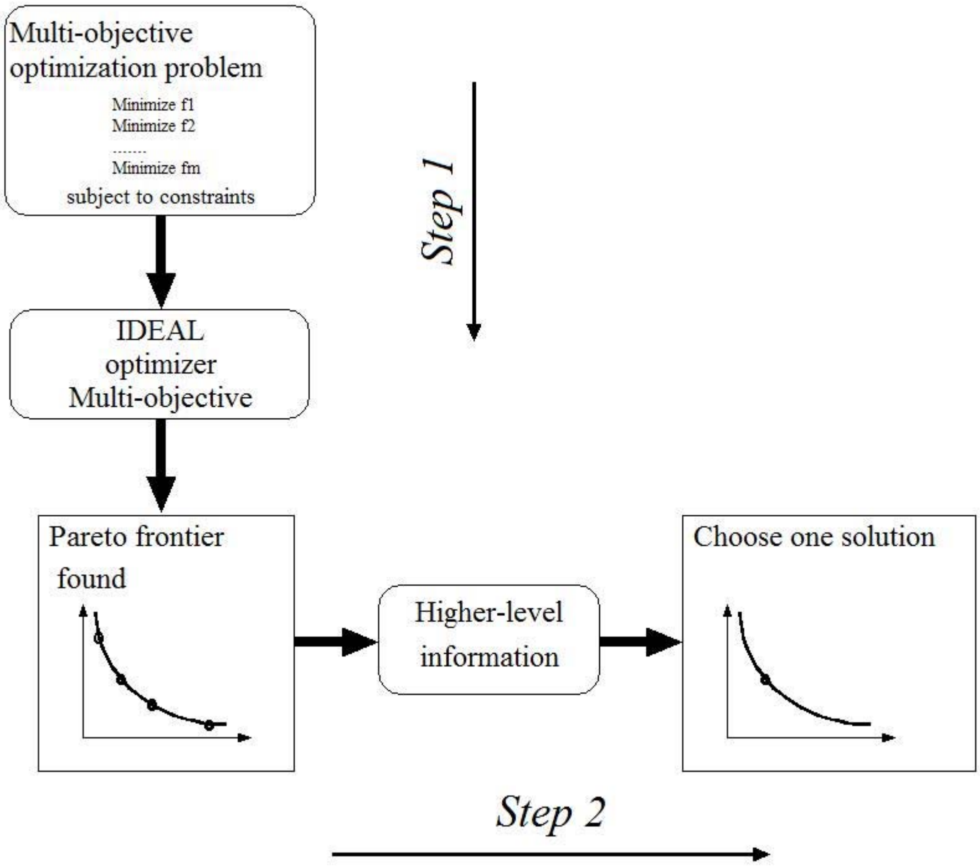

Figure 2-6 Schematic of ideal multi-objective optimization method (Deb, 2001)

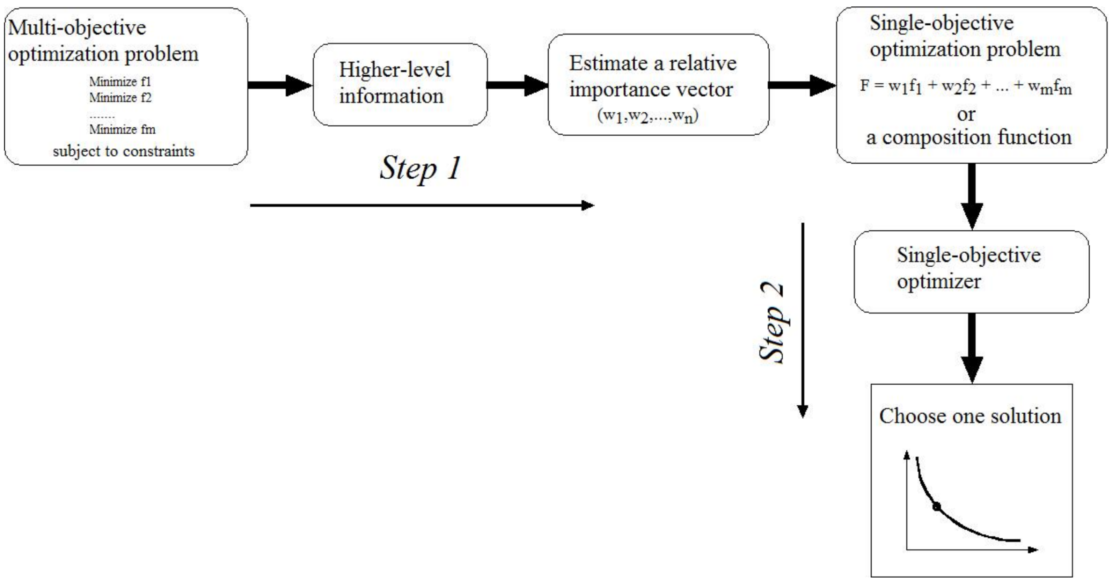

Figure 2-7 Schematic of preference-based multi-objective optimization method (Deb, 2001)

2-5 Methods for Identifying Pareto-Optimal Solutions

A good multi-objective optimization method should be able to converge to the Pareto-optimal front quickly as well as providing a good distribution of solutions along the front (Huang et al., 2006). There are several approaches described in the literature that seek the Pareto-optimal front. In the following sections, two broad classes are discussed, namely the classical and evolutionary methods.

2-5-1 Classical Optimization Methods

Classical optimization methods, which have been applied in the last four decades, are typically based on mathematical programming approaches that under certain conditions ensure convergence to a Pareto-optimal solution (Deb, 2001). Weighted sum, ε-constraint, weighted metric and goal programming approaches are some examples of classical methods. These methods convert the multi-objective optimization to a single-objective optimization problem to obtain one Pareto-optimal solution at a time. Therefore, they have to be applied many times, with the aim of finding a different Pareto solution at each iteration (Deb et al., 2002a). This can be grossly inefficient compared with heuristic methods that search for the Pareto-optimal set of solutions (Deb, 2001).

The classical methods have a number of drawbacks. First, classical methods suggest a way to convert a multi-objective problem to a single objective problem. In most cases the optimal solution to the single objective problem is expected to be a solution on the Pareto frontier. However, such a solution is subject to parameters used in the conversion approach. Thus to find N points on a Pareto front, at least N different sets of parameters should be used to form N single objective problems. Second, some of these methods will not be able to generate concave portions of the Pareto front. Finally, all methods require some problem knowledge to assign suitable weights or ε values (Deb, 2001). Martínez et al. (2007) note that the performance of these methods is sensitive to the choice of weights.

To avoid the above-mentioned significant shortcomings, many researchers have turned to heuristic methods such as evolutionary algorithms to solve multi-objective optimization problems.

2-5-2 Multi-Objective Optimization Evolutionary Algorithms

The term "evolutionary algorithm" (EA) represents a class of stochastic optimization methods that are based on the process of natural evolution. The origins of EAs were proposed in the late 1950s, and since the 1970s several classes of evolutionary methods such as genetic algorithms, evolutionary programming, and evolution strategies have been proposed. EAs have been employed in a variety of engineering applications and these algorithms have proven themselves as general, robust and powerful methods (Deb, 2001; Coello Coello et al., 2007). They have several characteristics that make them desirable for problems that have multiple objectives and large and highly complex search spaces.

Over the past decade, a number of multi-objective evolutionary algorithms (MOEAs) have been suggested (Fonseca and Fleming, 1993; Horn et al., 1994; Zitzler and Thiele, 1998; Deb, 2001; Coello Coello et al., 2007). Of these algorithms, NSGA-II and εMOEA were selected for use in this thesis. In Chapter 5, both of these algorithms are evaluated in a comparative assessment of performance involving a selection of different types of MOO algorithm.

For the case studies reported in Chapters 3 and 4, εMOEA was selected to perform the MOO search. Given that the focus of these chapters is on improved MOO problem formulation, the choice of MOO algorithm is not critical. It suffices to use a MOO algorithm with a good track record of providing a diverse set of approximately Pareto-optimal solutions. Although NSGA-II has been widely reported in the literature and used as a benchmark method in many studies, εMOEA was selected over NSGA-II for two reasons:

1) There was concern about NSGA-II's ability to provide a diverse set of Pareto-optimal solutions. This arises because NSGA-II limits the number of archived non-dominated solutions to the population size (Laumanns et al., 2002). εMOEA uses a different archiving strategy proposed by Laumanns et al. (2002) to overcome this limitation. Deb et al. (2003a) compared εMOEA against several evolutionary algorithms including NSGA-II and found εMOEA performed overall better in terms of convergence, diversity and computation time.

2) The comparative experiments reported in Chapter 5 confirmed this concern with εMOEA shown to be demonstrably superior to NSGA-II.

2-5-3 An Overview of εMOEA

The purpose of this section is to provide an overview of the ε-multi-objective optimization evolutionary algorithm (εMOEA). The distinguishing feature of εMOEA is the use of the ε-dominance concept which divides the objective space into hyperboxes of size ε and allows only one non-dominated solution to reside in each box (Laumanns et al., 2002). Inclusion of this concept in a genetic algorithm (GA) framework produces a method capable of maintaining a diverse and well-distributed set of solutions with a small algorithmic computational cost (Deb et al., 2003a).

As before, without loss of generality, it is assumed there are K objectives, all of which are to be minimized.

Definition of ε-Dominance: A solution x_1 is said to ε-dominate the solution x_2 for some $\varepsilon_j > 0$ if both of the following conditions are true (Coello Coello et al., 2007):

1. $f_j(x_1) \leq f_j(x_2) + \varepsilon_j$ for all

2. $f_j(x_1) < f_j(x_2) + \varepsilon_j$ for at least one $j \in \{1, 2, ..., K\}$

Figure 2–8 illustrates the ε-dominance concept geometrically. It shows two non-dominated solutions, P_1 and P_2. To check if P_1 ε-dominates P_2, P^*_2 is formed by adding ε_1 and ε_2 to the objective values of P_2. Since P^*_2 is dominated by P1 it follows that P_1 ε-dominates P_2. The box formed in Figure 2-8 leads to the idea of dividing the objective space into hyperboxes to facilitate checking whether solutions are ε-dominated. Figure 2-9 illustrates hyperboxes for a two-objective space. It shows that the solution P ε-dominates the entire region ABCDA while P only dominates the region PECFP. Indeed, any solution in the ABCDA region (except for the box in which P is located) would be ε-dominated by P because if ε_1 and ε_2 were added to objectives of such a solution it would lay in hatched area. However, all the solutions which share the same box with solution P ε-dominate each other. In this case the solution which has the shortest Euclidean distance to the bottom left corner of the box is deemed to dominate the other solutions. This situation is illustrated in Figure 2-9 for solutions 1 and 2. Since solution 1 is closer to the bottom left corner of the box, it is retained and solution 2 is eliminated.

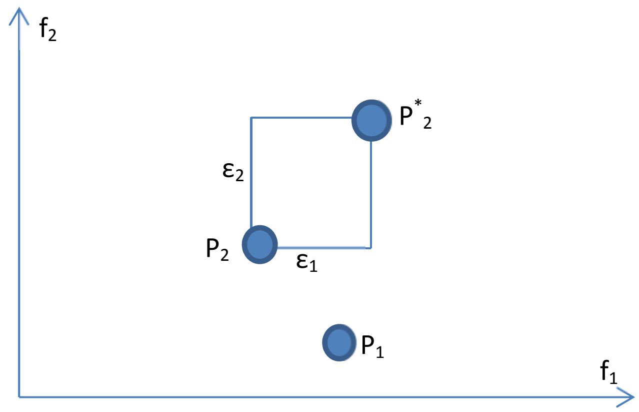

Figure 2-8 Schematic of ε-dominance concept

Figure 2-10 illustrates the application of ε-dominance in three steps. In the first step, for hyperboxes containing more than one solution, the solution which is closest to the bottom left corner of the hyperbox is retained (assuming minimization). For instance, in Figure 2-10, in two of the hyperboxes there are two solutions occupying the hyperbox; those marked by a red cross are eliminated. The next step applies the ε-dominance criterion to the remaining solutions. For example, the lower solution in the first column ε-dominates the higher solution in the first column. Eliminating the ε-dominated solutions produces the ε-dominance Pareto front in step 3.

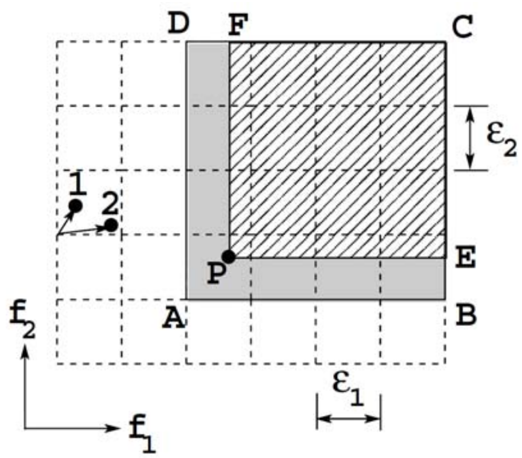

Figure 2-9 Illustration of ε-dominance concept for minimizing f_1 and f_2

(Deb et al., 2003a)

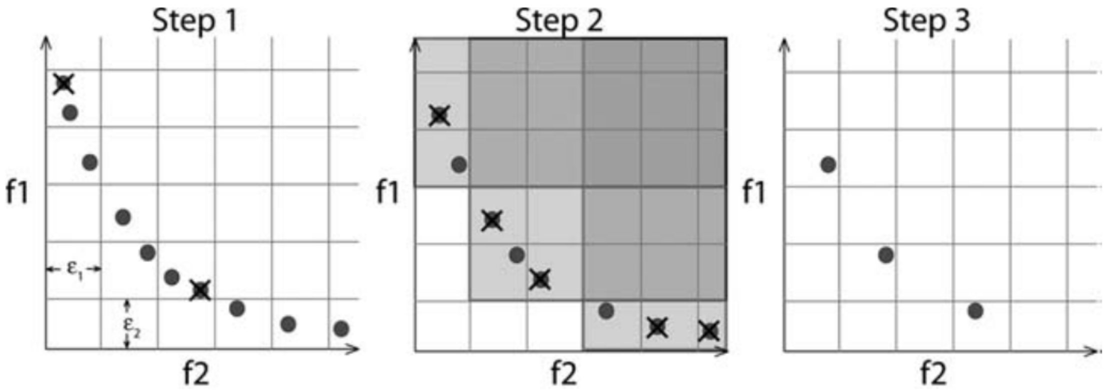

Figure 2-10 Illustration of Pareto frontier in conjunction with the ε-dominance concept (Kollat and Reed, 2006)

εMOEA uses two co-evolving populations: a current population archive P(t) and an archive of ε non-dominated solutions E(t), where t is the iteration counter. The initial population P(0) is selected randomly and the initial archive population is assigned the ε-non-dominated solutions of P(0). Thereafter, two solutions, referred as parents, one each from the current and the archive population are selected for mating. To select a parent from P(t), two solutions are chosen randomly. Then, if one of the solutions dominates the other one, that solution is chosen. Otherwise, the two solutions are non-dominated and one of the solutions is selected randomly. The parent from E(t) is simply chosen at random among the archive members. Applying crossover and mutation operations on the two parents produces two offspring solutions. This procedure is illustrated in Figure 2-11.

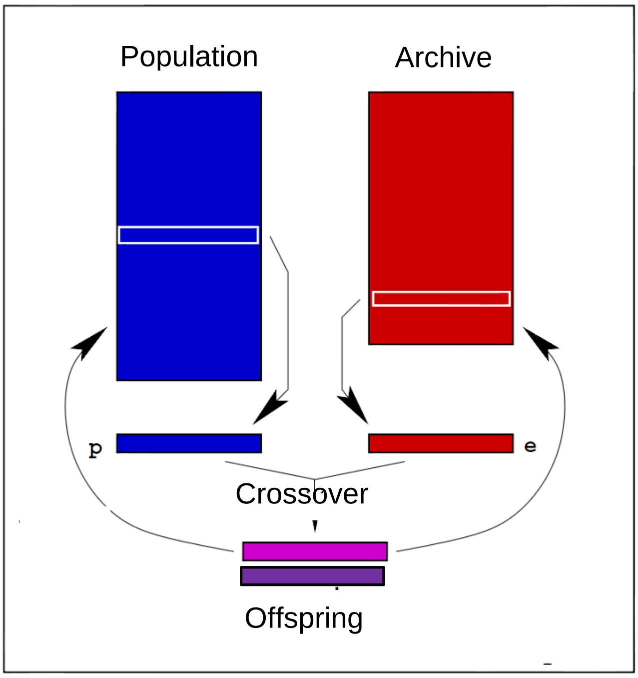

Figure 2-11 Schematic of εMOEA (Adapted from Deb et al., 2003a)

Each of the offspring solutions is evaluated and then compared with the current and archive populations for possible inclusion. First, tests are conducted to determine if an offspring should be accepted into the E(t) archive:

1. If the offspring solution is ε-dominated by any solution in E(t), it is rejected.

2. If the offspring ε-dominates any solution in E(t), that solution is deleted and the offspring added to E(t).

3. If both of the above cases fail, it indicates that the offspring solution is ε-non-dominated. In that case, the following tests apply:

    a. If the offspring solution does not share the same hyperbox with any solution in E(t), the offspring is added to E(t).

    b. If the offspring shares the same hyperbox with a solution, strict non-domination is applied. If the offspring solution strictly dominates the archive solution or it does not strictly dominate the archive solution but is closer to bottom left corner of the hyperbox (for minimization problems), then it is accepted into E(t) and the archived solution is rejected.

If an offspring is not accepted into E(t), then tests are conducted to determine if the offspring is to be accepted into P(t). To include the new offspring in P(t), three tests are conducted:

1. If the offspring solution is dominated by any existing member of the population, it is rejected.
2. If the offspring solution dominates one or more solutions in the current population, it replaces one at random.
3. If both of the above cases fail, it indicates the offspring solution is a non-dominated solution with respect to the current population. As a result, it replaces a random member of the population.

εMOEA and other heuristic search methods cannot guarantee finding Pareto-optimal solutions. For that reason it is a common practice to run these algorithms multiple times with different random seed numbers.

The search is terminated when certain conditions are satisfied. These may include the following:

1. No improvement in the non-dominated solution set for a prescribed number of iterations.
2. The number of iterations reaches a maximum value.
3. A prescribed value of convergence/diversity metric has been attained.

In the next two chapters, εMOEA is used in the case studies to conduct the search for Pareto-optimal solutions. However, Chapter 5 revisits the choice of MOO algorithm with the goal of identifying the algorithm best suited to urban water resources problems.

2-6 Optimization and Simulation Framework

A number of researchers have linked simulation and optimization methods. For instance, Cai et al. (2001) embedded a GA into a linear programming simulation model. Cui and Kuczera (2005) used a GA coupled to an earlier version of WATHNET5 to study single objective urban water resources problems. Shourian et al. (2008) coupled a single-objective particle swarm optimization method with the

MODSIM simulation model to allocate water optimally over time and space. They treated the capacities of reservoirs, transfer and pumping systems along with operational rules as decision variables.

The linking of the WATHNET5 simulation model with a MOO algorithm can be formulated mathematically using the notation of Eqs. (2.1) to (2.3) as follows:

$$\min_{x} \quad f_1[z(x)], f_2[z(x)], \ldots, f_K[z(x)] \tag{2.7}$$

subject to z(x) being the solution of the following minimization problem:

$$\min_{Z} \quad c^T(x \mid Q_t, D_t, \theta)\, z$$
$$\text{subject to } Az = b(x \mid Q_t, D_t, \theta), 0 \leq z \leq u(x \mid Q_t, D_t, \theta)$$

and

$$g(x,z) \leq 0$$

where $g(x,z)$ is a vector of constraints.

The implementation of the minimization problem given by (2.7) is schematized in Figure 2-12. The first step is to develop the simulation model and its input data Q_t, the decision space X and the objective functions $f_1, f_2, \ldots, f_K$. Then the optimization algorithm supervises the search for the Pareto-optimal solutions. At each iteration of the search, a set of decisions x is selected by the optimization method and passed to the simulation model. The simulation model simulates the system response to the inputs Q_t and parameters θ to produce outputs $z(x)$ which are used evaluate the objective function values $f_i[z(x)], i=1,..,K$ which, in turn, are passed to the optimization model. The optimization model then assigns a new set of decisions and passes them to the simulation model. This process continues until a termination condition is satisfied.

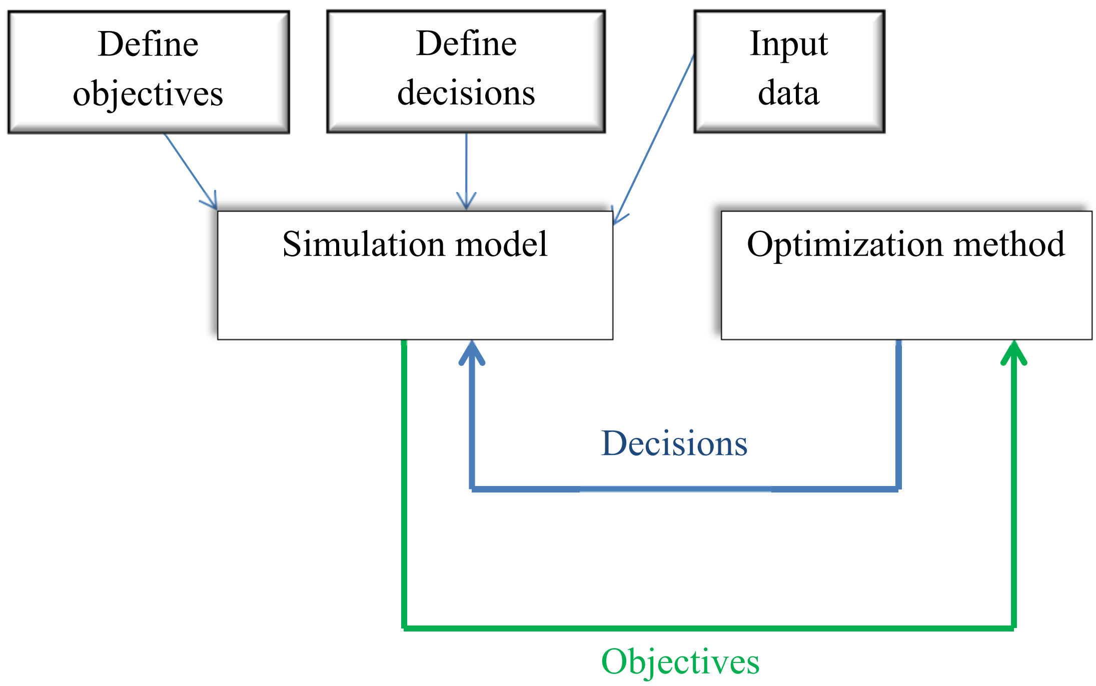

Figure 2-12 Schematic of communication between simulation and optimization models

2-7 Parallel Computing

In many real-world optimization problems, and particularly in water resource problems, the computation of the objective function is expensive. In addition, in some of these problems, it is necessary to evaluate a huge number of objective functions in order to find solutions close to the Pareto optimum (Jaimes and Coello, 2007). This means computation times may be days, weeks or even months.

Three strategies have been proposed in past studies to reduce computational time. Some researchers have developed optimization methods which converge to the optimal solution more efficiently (i.e. with fewer of evaluations) (Knowles, 2006; Pierro et al., 2009). Another strategy involves meta-modelling (Broad et al., 2010; Razavi et al., 2012). A meta-model is used to approximate the mapping between decisions and objective functions. If the mapping is sufficiently accurate, the meta-model can replace the computationally expensive simulation model. Finally, parallel computing has gained considerable attention since it can reduce the computation time very substantially. Civil engineering applications using parallel computing include applications in structural engineering (Kandil and El-Rayes, 2005), computational fluid mechanics and water engineering (Alonso et al., 2000; Cui, 2003; Cui and Kuczera, 2005).

One of the attractive features of evolutionary algorithms and other heuristic methods is their ability to support parallel computing. For instance, in EAs, the objective functions can be evaluated at each generation independently using the master-slave. Alternatively, Deb et al. (2003b) suggested a parallel MOEA approach based on NSGA-II which distributes the task of finding the whole Pareto-optimal front among participating processors with each processor dedicated to finding a particular part of the Pareto-optimal front.

There are several software protocols for implementing parallel computing, including Message Passing Interface (MPI) (Pacheco, 1997) and Parallel Virtual Machine (PVM) (Geist, 1994). Both MPI and PVM have been used widely. With respect to implementing the master-worker protocol, there is little difference between PVM and MPI. Accordingly, in this study PVM was adopted because of existing experience and support.

PVM (Geist, 1994) is a message-passing system which allows a user to create and access a parallel computing system consisting of multiple processors running on multiple hosts with possibly different operating systems. The PVM model accommodates a wide range of parallel computing models including the master-worker, node-only, tree computation, and hybrid computation models (Geist, 1994).

The master-worker protocol is the most natural model for parallelizing evolutionary algorithms (Cui, 2003). Following Cui and Kuczera (2005) the master-worker protocol is described by the pseudo-code presented in Figure 2-13. The master program hosts the MOEA. First, it spawns the PVM and determines the number of worker processes. The initial population is then generated and the corresponding decisions are sent to the workers for evaluation. When the initial population has been evaluated, the master enters the main iteration loop. The MOEA produces the next generation of decisions applying crossover and mutation operations on selected parents. These decisions are sent to the first available worker for evaluation. The master then waits until it receives a vector of objective function values from a worker.

That solution is then processed. It may be added to either the Pareto or population archive or discarded. The iterations continue until a termination criterion is met.

Master program

Spawn the PVM

Generate initial population

Send initial decision vectors to the workers for evaluation

Evaluate the initial objective functions

Do

  Produce a new decision vector using operations such as crossover and mutation

  Send decision vectors to an idle worker

  Wait and receive objective function values from any worker

  Process the new solution
  Stop if a termination criterion is met

End do

Display results

Terminate worker processes

End PVM program

Worker program

Do

  Wait for a message of decision vectors sent by the master

  Run simulation model and evaluate objective functions

  Send objective function values to master

End do

Figure 2-13 Pseudo code for master-worker protocol in PVM
(Adapted from (Cui, 2003))

In worker program, the worker waits until it receives a decision vector from the master. It then runs the simulation model, computes the objective functions values and send them to the master. It then waits for a new decision vector.

Because the message passing between the master and worker processors involves small data strings and because the time taken by a worker processor to conduct a simulation is, at least, several orders of magnitude longer than the time to pass a message, the speed-up is almost exactly proportional to the number of worker processors.

2-8 Summary

In this chapter a brief review of current simulation and optimization methods was presented with the objective of selecting a simulation model and an optimization method for use in subsequent chapters which constitute the main contribution of this thesis. Each selected model was then described in more detail to provide sufficient background for the case studies that appear in the subsequent chapters. The WATHNET5 model was selected for reasons of convenience – its source code was available, its software design facilitated linkage with MOO methods and a complex urban system had already been set up in WATHNET5. The εMOEA algorithm, an established method with good reported performance, was selected to conduct the multi-objective optimization case studies reported in the next two chapters. An overview of the communication protocols between the simulation and optimization models in a parallel computing environment concludes this chapter.

Chapter 3

Multi-Objective Optimization of Urban Water Resources: Moving Towards More Practical Solutions

3-1 Introduction

Recent Australian experience with arguably the severest drought on record and a potentially shifting climate has highlighted the vulnerability of urban water supplies to "running out of water". As storages dwindled in the major urban centres of Sydney, south-east Queensland, Perth, Melbourne and Adelaide, agencies responsible for urban water supply triggered drought contingency plans which started with the imposition of restrictions and, in most cases, led to the development of climate-independent sources of water such as desalination and wastewater reclamation. To secure Australian cities against drought, investments totalling tens of billions of dollars have been committed.

In an Australian industry position paper describing a framework for urban water resource planning, Erlanger and Neal (2005) state in their opening: "A safe and reliable water supply system is of utmost importance to the community. It is expected and understood that water utilities manage their water resources so that communities never run out of water." Erlanger and Neal recognize that failure to supply minimum water needs for an extended period would most likely result in disastrous social and economic losses that could conceivably threaten the very existence of the urban community.

Managing drought security in urban water supply is a complex and costly task, typically tackled using a two-pronged risk management approach, implementing short- and long-term options. The risk of exposure to severe drought is managed by application of long-term options such as policies that affect water use efficiency and provision of long lead-time infrastructure. Specifically, these long-term options control the probability of triggering short-term options or drought contingency plans, which may involve restrictions/rationing and short-lead time (and usually very expensive) source augmentation or substitution.

In view of the massive investments to secure Australian cities against drought, this chapter considers the question, what is the best mix of long- and short-term options in an urban headworks system? Here "headworks" refers to that part of the urban water supply infrastructure that harvests, stores and distributes water to major consumption zones. In seeking an answer to this question, several practical considerations deserve particular attention:

1. The maximization of drought security conflicts with the objectives of minimizing cost and environmental impacts. Recognizing the difficulty of quantifying environmental and social impacts solely in economic terms, multi-objective optimization (Deb, 2001) is needed to identify the trade-offs between conflicting objectives.

2. The consequences of an urban area "running out of water" are so severe that most systems are designed to have very high levels of security. This means that the probabilities of triggering drought contingency plans, particularly during extreme drought, are likely to be very small, while the probability of "running out of water" should be even lower. Because drought security criteria are often expressed in terms of probabilities of trigger events (Erlanger and Neal, 2005), it is vital that such probabilities be accurately estimated.

3. The performance of an urban headworks system is jointly dependent on the mix of short- and long-term options. Therefore, in a search for the best solution, it is essential that both short- and long-term options be evaluated jointly.

It is shown in the review of the water resource optimization literature in Section 3-2 that no previous work has adequately addressed all these practical considerations. The principal contribution of this chapter is twofold. First, the problem of optimizing the planning and management of urban water resources is formulated in a manner that addresses all these considerations. Specifically, the formulation addresses the practical challenges of identifying approximate Pareto-optimal solutions involving the full mix of short- and long-term options, while realistically accounting for drought risk and the trade-offs between economic, security and environmental factors. Second, a case study demonstrates the practical importance of addressing these challenges. It shows that failure to address these challenges can result in solutions that are significantly inferior and of limited practical value to headworks managers.

This chapter is organized as follows: Following a review of the literature, the shortcomings of existing methods are identified and motivate a new approach that more fully deals with the requirements of practical multi-objective urban water resource planning. An extensive hypothetical case study based on the headworks system for Sydney (Australia) demonstrates the practical importance of adopting this new approach and illustrates the challenges and insights identifying the approximate Pareto-optimal solutions that trade-off economic costs, environmental and drought-related social impacts.

3-2 Review of Urban Water Resources Optimization Literature

In the quest for securing urban water supplies against drought, water utilities use a mix of short- and long-term options to manage supply and demand. The short-term response to drought is embodied in a drought contingency plan (DCP). It is common practice to develop a staged DCP that progressively imposes severer restrictions on consumption while accessing emergency sources of water. The fundamental proposition is that the DCP reduces (and nowadays with the availability of climate–independent sources of water such as desalination, potentially eliminates) the risk of the system running out of water. A number of optimization studies have explored the benefit of imposing restrictions on demand to mitigate drought. For instance, Shih and ReVelle (1994, 1995) developed hedging rules for a single reservoir to reduce demand during drought. Tu et al. (2003, 2008) developed a mixed integer linear programming model that jointly considers reservoir release and hedging rules to minimize the shortages in current and future water supply. A limitation of these studies is that the social and economic cost of imposing restrictions was not addressed. Although imposing restrictions on demand reduces the risk of running out of water, frequent restrictions are not socially acceptable in major Australian cities (Erlanger and Neal, 2005).

In response to reducing the frequency of restrictions yet maintaining security, water utilities consider a range of long-term options to reduce demand and increase supply. However, each option imposes a cost on the community and environment. A number of studies have developed models to find the least-cost combination of short and long-term options. Lund (1987) evaluated the integration of water conservation measures with capacity expansion options showing that costs could be minimized by applying conservation measures to delay water treatment plant expansion. Rubinstein and Ortolano (1984) demonstrated the application of demand management in long-term water supply planning. In a similar vein, Dziegielewski et al. (1992) developed a framework to balance long-term water supply alternatives with short-term drought responses in order to identify the most cost-effective investments offering long-term drought protection. Subsequently, Wilchfort and Lund (1997) minimized the expected cost of a combination of long-term and short-term options. Jenkins and Lund (2000) integrated shortage management and yield models to identify operating rules that minimize operating and shortage costs. However, as Dziegielewski et al. (1992) emphasize, the usefulness of these approaches depends on the accuracy and validity of costs associated with short-term demand-reduction measures.

Due to difficulties in estimating costs associated with restrictions or shortages, a number of studies (Randall et al., 1990; Ko et al., 1992; Liang et al., 1996; Kim et al., 2006; Reddy and Kumar, 2006; Chen et al., 2007; Yang et al., 2007; Kim et al., 2008; Chang and Chang, 2009; Kasprzyk et al., 2009) have adopted a multi-objective optimization approach. All of these studies except Yang et al. (2007) and Kasprzyk et al. (2009) have focused on short-term decisions associated with reservoir releases and restriction rules. However, there is an interaction between short- and long-term options as demonstrated by Lund (1987). Yang et al. (2007) investigated the interaction between reservoir operating rules and reservoir capacity but did not incorporate any DCPs. Kasprzyk et al. (2009) focused on water marketing and portfolio-based management strategies in the context of a single reservoir system; they did not optimize infrastructure options nor DCPs.

The rationale for multi-objective optimization is strengthened when environmental impacts are considered. Rivers downstream of dams typically experience a hydrologic regime change which can adversely impact on the health of riverine ecosystems (Shiau and Wu, 2007). In recent years, in an effort to support sustainable ecosystems, releasing sufficient water to meet instream flow requirements – environmental flows – has received considerable attention from the water resources management community (Richter et al., 2006).

In many past studies, environmental flows have been considered as a constraint (Tu et al., 2003; Tu et al., 2008). However, this hides the trade-offs between cost, supply security and environmental impact. Suen and Eheart (2006) circumvented this shortcoming using multi-objective optimization to demonstrate the trade-off between human and ecosystem needs in which the ecosystem objective was to maximize the similarity between natural and flow released from the reservoir. Likewise Shiau and Wu (2007) applied multi-objective optimization to optimize weir operation to balance ecosystem and human needs. However, these studies ignored the cost dimension and only focused on operational rules. By explicitly presenting the trade-offs between cost, drought security and environmental impact, Erlanger and Neal (2005) suggest communities may be prepared to pay more in return for less environmental damage.

To evaluate the performance of an urban headworks system, a simulation model is typically constructed to model the behaviour of the system in response to a time series of hydro-climatic and demand inputs – see Labadie (2004) for a review. The length of the time series used as input is critical. Given that urban systems typically operate with high levels of reliability, the time series must be long enough to enable a meaningful assessment of drought risks. The significance of this issue is best illustrated by an example. The annual probability of triggering a DCP, p_{DCP}, can be estimated by counting the number of years the DCP is triggered in a simulation and dividing by the number of simulation years N. Assuming annual independence, the standard error of the estimate based on binomial probability model considerations is

$$stderr(\hat{p}_{DCP}) = \sqrt{\frac{1}{N}\hat{p}_{DCP}(1-\hat{p}_{DCP})} \tag{3.1}$$

Suppose in a 100-year simulation, the DCP was triggered once. Then $\hat{p}_{DCP} = 0.01$ and the standard error is 0.010. This large uncertainty can be presented more intuitively using return periods; it can be shown that the 95% confidence limits on the return period for the DCP trigger are 23 and 1580 years. This uncertainty arises solely because of the insufficient length of the simulation.

This example suggests that evaluating drought risks and associated drought security criteria using simulation models with insufficiently long input time series borders on being meaningless with the results being sensitive to the choice of the input time series. Indeed Ajami et al. (2008) suggest that use of historical data can lead to development of inefficient water management rules.

One way to reduce this sampling uncertainty is to increase the length of the input time series. This can be done by generating long stochastic input time series by sampling from probability models fitted to historical data (Salas et al., 2005). All but three of the reviewed multi-objective optimization applications to urban water resource systems used historical data. Though Kim et al. (2008) and Shiau (2009) used 100 and 40 years of synthetic data respectively, such record lengths are considered completely inadequate for use with high security urban systems. In their study of many-objective portfolio planning Kasprzyk et al. (2009) evaluated the performance of each proposed portfolio with 5,000 10-year Monte Carlo samples. However, their Monte Carlo strategy involved resampling 10-year samples from a 33-year historical record, which is statistically unlikely to include severe drought. That said, Kasprzyk et al. did consider solution robustness by investigating sensitivity to initial conditions and extreme drought/demand scenarios. As a result, all of the reviewed studies suffer from the potentially serious limitation that the Pareto solutions are not robust in the sense of the solutions being sensitive to the choice of input data used in the simulation.

3-3 A More Practical Multi-objective Optimization Methodology for Urban Water Supply

This section formulates a multi-objective optimization methodology for an urban headworks system which addresses all the shortcomings identified in previous work on this subject. In the following section a case study is used demonstrate the practical significance of addressing these shortcomings.

Generally, the urban headworks multi-objective optimization problem can be formulated as follows:

$$\min_{x} f_1[z(x)], f_2[z(x)], \ldots, f_k[z(x)] \tag{3.2}$$

subject to $z(x) = M[x, Q_{N,} D_N]$

$$g(x, z(x)) \leq 0$$

$$Sf_N = 0$$

where x is a vector of decision variables that are to be optimized.

The function $M[x, Q_N, D_N]$ represents the headworks simulation model which takes as input Q_N, a matrix of streamflow and climate values at multiple sites for an N-year period, and D_N, a matrix of unrestricted demand at multiple sites for the same N-year period, to produce simulation outputs $z(x)$. There are many simulation models in the literature (Labadie, 2004) capable of simulating urban headworks systems. All that matters is that the model satisfactorily simulates the actual operation of the headworks system using information that would be available to the operators. The simulation outputs are used to evaluate $f(Z_N)$, the vector of criterion (or objective function) values. The function $g(x, Z_N)$ is a vector of constraints.

The constraint $Sf_N = 0$ is essential to the urban headworks optimization problem. It requires that no unplanned demand shortfalls, denoted by Sf_N, occur during the simulation. Unplanned shortfalls occur when the demand, permitted by the DCP, cannot be supplied – such shortfalls typically would occur when reservoirs run dry or when limitations in transfer capacity result in demand zones being supplied less than the minimum permitted by the DCP.

The optimization problem (3.2) is largely intractable using classical optimization approaches which typically impose severe constraints on the form of the simulation model $M[x, Q_N, D_N]$ and constraints $g(x, Z_N)$ and therefore restrict the inclusion of variables in the decision space. However, the advent of evolutionary optimization algorithms – see (Deb, 2001) – has made the solution of (3.2) significantly more tractable. In the water resources field, many researchers have recognized and exploited this opportunity – see the recent review by Labadie (2004) and Nicklow et al. (2010). Of particular importance to this study is the greater freedom in specifying the decision vector. This enables optimization of the full mix of decision variables associated with short- and long-term options.

The formulation (3.2) differs from previous formulations in the way it deals with drought security. The specification of drought security in the sense of Erlanger and Neal (2005), namely urban "communities never run out of water" is problematic. Unless climate-independent sources of water (such as desalination) can guarantee a minimum supply, there will always be a finite probability that the system will run out of water. This is unavoidable. The best that one can do is manage the risk of running out of water.

The optimal solutions in (3.2) are conditioned on the input Q_N. A more useful interpretation is that the Pareto-optimal solutions (3.2) secure the system against droughts with return periods up to an expected value of N years. Seen this way, the expected return period N defines the drought security risk level for the system. As will be demonstrated, the explicit recognition of this risk level is vital to practical optimization outcomes.

3-4 Case Study

This section presents a case study to illustrate the application of the multi-objective optimization formulation (3.2) and to identify important insights arising from its application. It is motivated by the headworks system that supplies Sydney, Australia's largest city serving a current population of 4.5 million.

3-4-1 Optimization Implementation Issues

Chapter 2 provided an overview of the simulation and optimization models used to implement Eq. (3.2) in this study. This will be briefly reviewed here.

Similar to Cai et al. (2001), Cui and Kuczera (2005) and Yang et al. (2007), a two-level optimization approach is adopted. For the simulation model $M[x, Q_N, D_N]$, the WATHNET5 model (Kuczera, 1992; Kuczera et al., 2009) is adopted. The scripting language within WATHNET5 enables the user to specify quite complex run-time functions to assign arc capacities and costs and side constraints to the network linear program. The decision vector x is accessible to all scripts and, therefore, can fully control the specification of the network linear program.

A major implementation issue in this case study was the computational time to solve the optimization problem (3.2). The total computational time is proportional to N, the number of years of simulation. In this case study, a 10,000-year simulation using monthly time steps takes approximately 60 seconds on an Intel T7700 CPU running at 2.40 GHz. If the multi-objective optimization algorithm evaluates 20,000 different decision vectors, the total run time will be about 14 days. Extensive use was made of parallel computing in conjunction with εMOEA to make the multi-objective optimization tractable.

3-4-2 Description of Sydney Headworks System

The case study considers a simplified representation of the Sydney headworks system which, nonetheless, accounts for many of the interesting dynamics of the Sydney system. It considers several scenarios involving a hypothetical mix of short- and long-term options that cater for a future population of 7 million corresponding to a highly stressed system.

Figure 1-1 presented a graphical depiction of the Sydney headworks system. The representation of this system in WATHNET5 is presented in Figure 3-1 in which the nodes labelled “R” represent reservoirs, “S” stream nodes, “D” demand zones, and “W” waste/sink nodes. The network of reservoirs, pumping stations and water treatment plants supplies water to two demand zones labelled “Sydney” and “South” in Figure 3-1. The existing system has a total storage capacity of 3,343,487 ML (mega litres). Warragamba Reservoir is the largest reservoir in the system with a capacity of 2,031,000 ML. The Sydney demand zone, which serves approximately 90% of the population, is supplied by Warragamba Reservoir together with a number of smaller reservoirs, Avon, Woronora, Cataract, Nepean and Cordeaux. In contrast, the South demand zone, which serves the remaining 10% of the population, is only supplied by Nepean and Avon Reservoirs. An inter-basin transfer scheme augments the natural inflows into Warragamba and Nepean-Avon Reservoirs. The transfer scheme is located on the Shoalhaven River and involves a small pondage at Lake Yarrunga from which water is lifted over 500 m using two pumping stations to transfer water to Wingecarribee Reservoir from where it can be transferred to Warragamba or Nepean Reservoirs. The pump stations have a monthly transfer capacity of 46,600 ML.

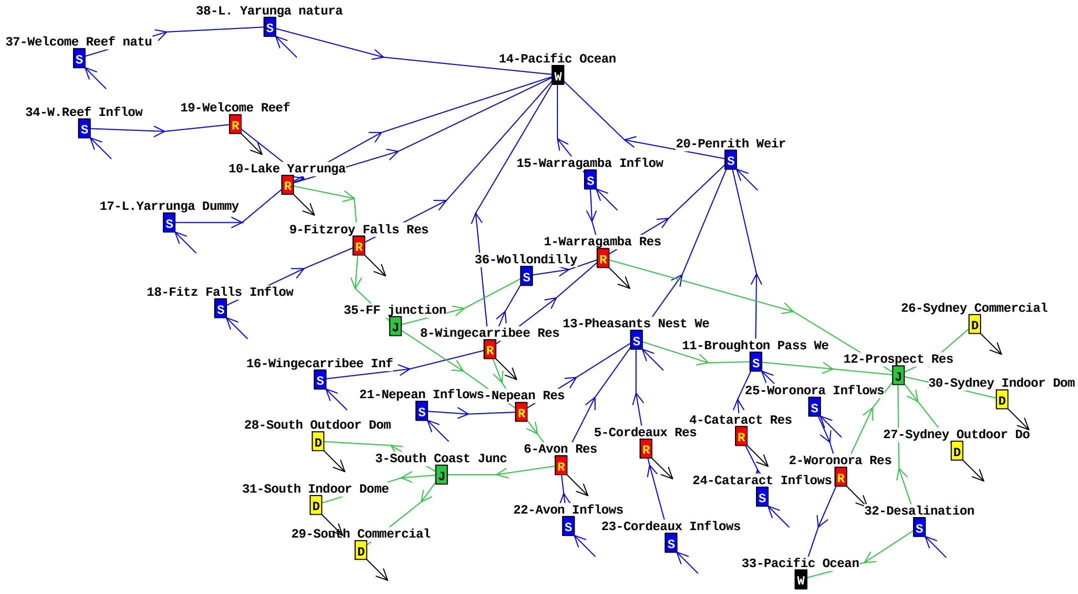

Figure 3-1 WATHNET5 schematic of Sydney water supply headworks system - the nodes labeled "R" represent reservoirs, "S" stream nodes, "D" demand zones, and "W" waste/sink nodes

For the purposes of this case study, environmental flow considerations are restricted to the Wollondilly River between Wingecarribee and Warragamba Reservoirs. The primary environmental issue is to limit high flows when pump transfers from the Shoalhaven are in progress, to avoid adverse impacts on riverine ecosystem function. Scott and Grant (1997) investigated the impacts of high flows on the riverine ecosystem and recommended maximum monthly regulated flows to avoid ecological impacts.

In this case study, two options for augmenting the supply are available. The first is a new dam at Welcome Reef on the Shoalhaven River, upstream of Lake Yarrunga. The second is a desalination plant serving the Sydney demand zone. This latter option is strategically different from Welcome Reef in that it provides a climate-independent supply of water.

The supply zones, Sydney and South, are each disaggregated into three demand nodes representing domestic indoor, outdoor watering and commercial/industrial consumption. In this case study, the DCP only restricts the outdoor watering usage. It is recognized that rationing during severe drought would be extended to the other usage categories.

3-4-3 Streamflow and Demand Data

The Sydney system experiences high natural climate variability. For instance, the annual coefficient of variation for inflows to Warragamba Reservoir is about 1.1. In view of this variability and the multi-year persistence of droughts, the reservoirs in the Sydney system have significant over-year carryover capacity. Therefore, when generating stochastic hydro-climate data for this system, it is important that the stochastic model produces sequences that are consistent with the multi-year observed statistics such as cumulative overlapping n-year sums (with n ranging from 1 to 5 years). The following two-step algorithm was used to generate stochastic streamflow and climate data: 1) annual values were generated using the Matalas (1967) lag-one multi-site model calibrated to non-contiguous historical streamflow and climate records up to 84 years long using the missing-data EM algorithm (Kuczera, 1987); and 2) monthly values were obtained by disaggregating the annual flows using the method of fragments. Extensive testing has revealed this model produces multi-year statistics consistent with the observed data. Indeed, Thyer et al. (2006) argue that more complex stochastic models describing decadal to multidecadal-scale variability are not identifiable using historical records of the length available in the case study.

To explore the sensitivity of the approximate optimal solutions to the choice of drought security return period, two sets of stochastic data were used: one 500 years long and the other 10,000 years long. It is noted that the 500-year series corresponds to the first 500 years of the 10,000-year series.

Demand for the 7-million population scenario was disaggregated into indoor domestic, outdoor domestic and commercial/industrial categories. Because outdoor domestic demand is correlated with rainfall, it was stochastically generated using the stochastically generated rainfall as input. This ensures that the higher outdoor water usage during droughts is preserved in the stochastic data Cui (2003).

The following steps were applied to determine water demand at each demand node:

1. Generate the monthly average indoor per capita water consumption (InWater(k), k=1,…,12) using the data in Table 3-1.

2. Generate concurrently with streamflow the number of rain days at Prospect reservoir denoted as $\mathrm{Rain}_t^i(k), k = 1, \ldots, 12$ which presents rain days for replicate i in year t.

3. Compute the total monthly demand for the following categories:

   a. Indoor residential demand

$$\mathrm{InHouse}_t^i(k) = \mathrm{InWater}(k) \times \mathrm{population}(t, k) \quad (3.3)$$

   where population(t,k) is the population for month t and year k

   b. Outdoor residential demand

$$\mathrm{ExHouse}_t^i(k) = \left[A(k) + B(k) \times \mathrm{Rain}_t^i(k)\right] \times \mathrm{population}(t, k) \quad (3.4)$$

   where A(k) and B(k) are parameters obtained from a monthly regression analysis between per capita outdoor water consumption and number of rain days (Cui, 2003).

   c. Commercial demand

$$\mathrm{Commercial}_t^i(k) = [\mathrm{InHouse}_t^i(k) + \mathrm{ExHouse}_t^i(k)] \times \frac{0.45}{0.55} \quad (3.5)$$

Table 3-1 Average monthly per capita indoor water consumption

Month	1	2	3	4	5	6	7	8	9	10	11	12
Indoor	3.98	3.89	4.15	4.10	4.31	4.05	4.22	4.20	4.01	4.23	4.17	4.29

3-4-4 Decision Variables

A large number of options is available to ensure a secure water supply for Sydney's 7-million population scenario. In this case study, eleven decision variables, listed in Table 3-2, were identified as being potentially important. They are classified as either infrastructure (which corresponds to a physical asset) or operational (which affects the way the system is operated).

Decisions x_1 and x_2 control the pump transfer of water from the Shoalhaven basin. x_1 is a pump mark that defines the Warragamba storage fraction which triggers transfer of water from Shoalhaven to Warragamba; if the storage fraction in Warragamba is below the pump mark x_1 at the start of a month, the maximum pump transfer capacity is activated. A separate pump mark x_2 is applied to Avon on account of it being the main supply to the South demand zone.

Table 3-2 List of decision variables x

Decision variable	Description	Lower limit	Upper limit	Category
1	Pump mark Warragamba	0.3	1	Operational
2	Pump mark Avon	0.3	1	Operational
3	Level 1 restriction trigger	0.05	0.95	Operational
4	Trigger increment	0.05	0.25	Operational
5	Desalination plant capacity (ML/day)	0	1,000	Infrastructure
6	Desalination plant trigger	0.05	0.95	Operational
7	Welcome Reef capacity (ML)	0	1,000,000	Infrastructure
8	Warragamba base gain	8,000	12,000	Operational
9	Warragamba incremental gain	10	200	Operational
10	Maximum Wollondilly flow during September to March (ML/month)	12,200	100,000	Operational
11	Maximum Wollondilly flow at other times (ML/month)	18,300	100,000	Operational

Decisions x_3 and x_4 define the first stage of the DCP. When the total storage fraction falls below the trigger x_3, the first level of restrictions is imposed on outdoor domestic water use with a target reduction of 33%. If the total storage fraction falls below (x_3 – x_4), then the second level of restrictions is imposed with outdoor domestic water use reduced by 67%. If the total storage fraction falls below (x_3 – $2x_4$), then the third level of restrictions is imposed with outdoor domestic water use totally banned.

Decisions x_5 and x_6 define the second stage of the DCP. When the total storage fraction falls below the trigger x_6, the already-constructed desalination plant with daily capacity of x_5 ML/day is activated.

Decision x_7 defines the capacity of Welcome Reef Reservoir.

Decisions x_8 and x_9 define the priority for storing water in Warragamba. All the reservoirs in the WATHNET5 network flow program were assigned 20 carryover arcs which "store" water for the next time step. Each carryover arc has a capacity equal to 1/20 of the reservoir capacity and a gain (i.e. negative cost) defined by

$$Gain(j) = BG + (j-1) * IG, j = 1, \dots, 20 \qquad (3.6)$$

where BG is the base gain and IG is the incremental gain. See Chapter 2 for a fuller discussion on the way carryover arcs are implemented in WATHNET5.

On account of Warragamba's dominant storage, all reservoirs except Warragamba were assigned a base gain of 10,000 and an incremental gain of 100. This implements the so-called space rule that seeks to keep each reservoir with the same storage fraction. Decisions x_8 and x_9 define the base and incremental gain for Warragamba respectively. Depending on the values assigned to x_8 and x_9, water may be preferentially stored in Warragamba or in the rest of the system.

Finally decisions x_{10} and x_{11} define the maximum monthly Wollondilly transfer capacity during September to March and at other times respectively. The lower limit on these decisions corresponds to that recommended by Scott and Grant (1997). These two decisions are active in the three-objective scenario and fixed in the other scenarios. These scenarios are discussed in Section 3-5.

3-4-5 Objectives and Constraints

Three objectives were judged to be relevant to the case study:

Minimize frequency of restrictions (%): Erlanger and Neal (2005) state that the supply system should be capable of maintaining an adequate level of supply most of the time. Accordingly, the frequency of restrictions describes the fraction of the time consumers will not have an adequate level of supply. Cui and Kuczera (2005) used willingness–to-pay concepts to estimate the economic cost of restrictions from which they estimated the economically optimal frequency of restrictions. Here, the restriction frequency is made an explicit criterion in recognition of the difficulty of accurately estimating the economic cost and the political/social sensitivity that is associated with imposition of restrictions.

Minimize the present worth cost ($): The present worth cost is the sum of capital and discounted expected operating costs and the costs of unplanned shortfalls. The capital cost represents the cost of building new infrastructure, which in this case study, is the Welcome Reef dam and/or the desalination plant. Table 3-3 summarizes the capital costs for Welcome Reef and the desalination plant. The capital cost model uses a binary function: if the asset is selected by the optimization, then the total cost is the sum of a fixed setup cost and a cost proportional to the size of the asset; however, if the asset is not selected, the capital cost is zero. The operating cost includes the costs for pumping transfers from the Shoalhaven and operation of the desalination plant. A 5% discount rate was used.

To ensure the DCP adequately copes with all droughts during the simulation period, solutions are constrained to avoid unplanned demand shortfalls. In this case study, an unplanned shortfall occurs when the system is unable to supply domestic indoor and commercial/industrial demand. This would occur when the highest restriction level, that bans all outdoor water use, is in force and the reservoirs become empty.

The constraint on unplanned shortfalls is imposed using a penalty function approach. Here, a penalty of \$100,000 per ML unplanned shortfall is added to the present worth cost. This penalty was selected to steer the optimization search away from solutions which allow reservoirs to “run dry” with consequent failure to supply minimum water needs.

Table 3-3 Cost summary for infrastructure decision variables

Decision Variable	Fixed and Unit Costs
Desalination plant capacity (ML/day)	\$1,250,000,000 + \$4,000,000 ML/day
Welcome Reef capacity (ML)	\$100,000,000 + \$1000/ML storage

Minimize environmental stress on the Wollondilly River: In this case study, the Wollondilly River between Wingecarribee and Warragamba Reservoirs has been identified as ecologically important. There is a vast literature that examines the ecological impacts of altered flow regimes. For example, Tharme (2003) documented over 200 individual environmental flow methodologies which have been utilized in 44 countries. Arthington et al. (2004) outlined the characteristics, strengths and limitations of the category of techniques termed holistic methodologies. In another study, Petts (2009) reviewed the advances in environmental flow science over the past 30 years. In a more targeted review, Dewson et al. (2007) reviewed literature on the consequences of natural low flows and artificially reduced flows on habitant conditions and on invertebrate community structure, behaviour and biotic interactions. These studies underscore the difficulty in characterizing ecological response.

As the purpose of this case study is illustrative, a notional response function is developed based on the field studies by Scott and Grant (1997) who identified potentially adverse impacts of altered flow regimes on platypus and water bird populations in the Wollondilly River. To avoid these impacts, they recommended that the maximum monthly regulated flow be limited to 18,300 ML during the winter months from April to August, and to 12,200 ML during the summer months. The ecological impact of exceeding these recommended maxima is not well understood (Grant and Temple-Smith, 2003). Nonetheless, it is known that during the summer months, high flows have the highest impacts on the breeding of platypus and water bird populations, while the impacts of high flows are significantly less severe during the winter months. Accordingly, the following environmental stress metric was adopted to penalize the adoption of maximum regulated flow limits, x_{10} and x_{11}, in excess of those recommended by Scott and Grant.

$$Stress(m) = \begin{cases} \max\left[0,\ 5\left(\dfrac{q_m - 12200}{12200}\right)\right] \text{if } m \in \{\text{Sept},..,\text{March}\} \\ \max\left[0,\ \left(\dfrac{q_m - 18300}{18300}\right)\right] \text{if } m \in \{\text{April},..,\text{August}\} \end{cases} \tag{3.7}$$

where q_m is the actual regulated release in the Wollondilly in month m and $Stress(m)$ is the penalty for exceeding the recommended flow limits in month m. The environmental stress criterion is the sum of the monthly stresses over the simulation. Unlike the first two criteria, the environmental stress criterion is based on limited field data and relies on subjective judgments such as the impact in summer months is 5 times that of winter months and that the impact is cumulative. Consequently, the trade-offs between environmental stress and the other criteria need to be interpreted with the understanding that there is considerable epistemic uncertainty about the environmental impacts.

Apart from the constraint on unplanned shortfalls, which was implemented using a penalty function approach, the only other constraints were the limits on the decision variables summarized in Table 3-2.

3-5 Case Study Scenarios

Seven case study scenarios are used to illustrate the importance of using an optimization formulation that deals with the shortcomings identified in the literature review. The first two scenarios demonstrate the importance of jointly optimizing the full mix of decisions, particularly when there are interactions between short and long-term and/or operational and infrastructure decision variables. Then, three different scenarios are used to demonstrate the influence of environmental constraints on system behaviour and the beneficial aspects of treating an environmental constraint as an objective. Finally, a comparison of two sets of scenarios with different input data length highlights the practically serious shortcoming arising from use historical or short-length synthetic data when there are expectations of high levels of drought security. A summary of these scenarios is presented in Table 3-4 as a reference.

Table 3-4 Summary of case study scenarios

Scenario number	Decisions	Optimization criteria	Record length (years)	Purpose
1	1 to 9; 10,11 set to upper limit	Present worth cost Restriction frequency	500	Study consequences of fixing operational decisions
2	5, 6, 7; 10,11 set to upper limit	Present worth cost Restriction frequency	500	
3	1 to 9; 10,11 set to lower limit	Present worth cost Restriction frequency	500	Contrast use of environmental constraints against environmental trade-offs
4	1 to 9; 10,11 set to upper limit	Present worth cost Restriction frequency	500	
5	1 to 11	Present worth cost Restriction frequency Environmental stress	500	
6	1 to 9; 10,11 set to upper limit	Present worth cost Restriction frequency	500	Contrast solutions with different levels of drought security
7	1 to 9; 10,11 set to upper limit	Present worth cost Restriction frequency	10,000	

For all seven scenarios, the εMOEA algorithm was run with 10 different initial random number seeds. Recognizing that an evolutionary algorithm cannot guarantee convergence to the Pareto-optimal solutions, the approximate Pareto-optimal solutions are taken here to be the non-dominated solutions from the 10 runs. The first four scenarios were run for 30,000 generations, while Scenario 5 was run for 100,000 generations and Scenarios 6 and 7 for 10,000 generations. εMOEA terminated its search if the maximum number of generations was reached or if the Pareto solutions did not change after 1000 generations. The εMOEA parameters were tuned to ensure good coverage and diversity along the Pareto front. The tuned parameters were the same as the parameters obtained in the tunning of εMOEA in Section 5-6-1: mutation rate = 0.01; crossover rate = 1.00; Inversion rate = 0.005; population size=100; and the hyper-box epsilons for the restriction criterion 0.005, for the present worth cost criterion \$1000 and for environmental stress criterion 0.001.

3-5-1 Joint Optimization of Operating and Infrastructure Decision Variables: Scenarios 1 and 2

This section compares the approximate Pareto-optimal solutions for two scenarios which differ in the mix of decisions to be optimized. In Scenario 1, all operational and infrastructure decisions were optimized except for decisions 10 and 11 which were fixed at their upper limit. In contrast, in Scenario 2, the optimization problem is akin to asking what is the best capacity expansion option with the rest of the system operated as normal. Accordingly, only two infrastructure decisions, the desalination plant and Welcome Reef reservoir capacities, and the desalination operational decision, the desalination plant trigger, were optimized; the remaining operational decisions were set to the following values guided by the desire to minimize operating costs and the frequency of restrictions: x_1=0.3; x_2=0.3; x_3=0.5; x_4=0.05; x_8=10,000; x_9=100. In each scenario, two objectives were considered, namely minimizing the present worth cost and restriction frequency.

Figure 3-2 shows the approximate Pareto-optimal fronts for the two scenarios. There is a considerable gap between the two Pareto fronts with the scenario 1 Pareto front dominating the Scenario 2 front. By optimizing all the operational decisions, considerably lower present worth costs can be achieved for the same restriction frequency. Clearly fixing some of the operational decisions severely limited the ability of the optimization to take full advantage of the desalination plant and the Welcome Reef Reservoir. In Scenario 2, because the restriction trigger x_3 was set to 0.50, it was impossible to produce outcomes with a restriction frequency greater than 20%. Likewise, the restriction frequency could not fall below 2.5%, because the desalination plant and Welcome Reef capacities were at their upper bounds.

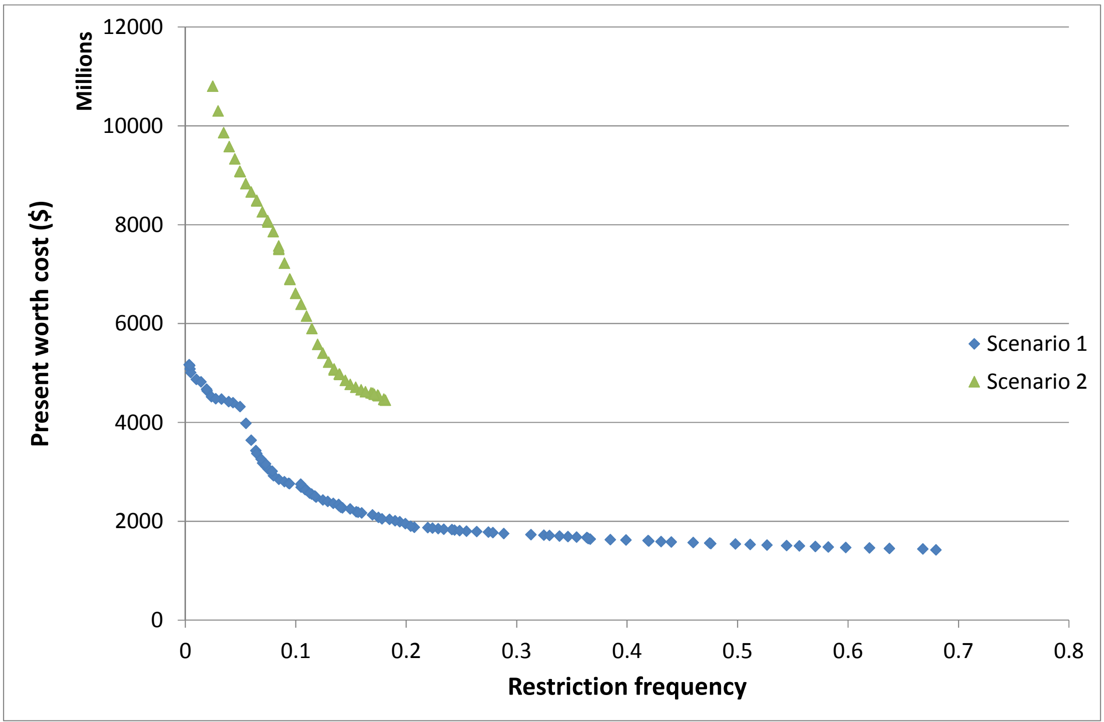

Figure 3-2 Approximate Pareto-optimal fronts for Scenario 1 (all decisions optimized) and Scenario 2 (two infrastructure decisions and one operational decision optimized)

While the gap between the Pareto fronts is dependent on the choice of values assigned to the decisions not optimized in Scenario 2, the important conclusion to be drawn is that when operational and infrastructure decisions interact, the failure to optimize all decisions can lead to inferior outcomes. Importantly, the ability to solve the optimization problem (3.2) makes it practically feasible to explore the whole decision space.

3-5-2 Moving From Environmental Constraints to Trade-Offs: Scenarios 3 to 5

This section investigates the insights and benefits that arise from considering environmental trade-offs rather than imposing environmental constraints. Three scenarios are considered. The first two, Scenarios 3 and 4, establish the sensitivity of the system to decisions x_{10} and x_{11}, which determine maximum regulated flows in the Wollondilly River.

3-5-2-1 Sensitivity to Environmental Flow Constraints – Scenarios 3 and 4

The sensitivity of the system to the specification of environmental flow constraints is explored using Scenarios 3 and 4. In Scenario 3, x_{10} and x_{11} are fixed at the values recommended by Scott and Grant (1997), while in Scenario 4, x_{10} and x_{11} are fixed at an arbitrarily large value that would not impose constraint on transfers from the Shoalhaven. Scenario 3 imposes nominally no environmental stress, while Scenario 4 would allow imposition of maximal environmental stress.

Figure 3-3 presents the approximate Pareto-optimal solutions for Scenarios 3 and 4. The imposition of the environmental flow constraint on the Wollondilly River substantially shifts the Pareto front outwards. For example, for a 10% restriction frequency, the imposition of the Wollondilly environmental flow constraint increases the present worth cost by ~$1,600 million. The reason for this sensitivity will be explained subsequently. Here the point to be made is that the imposition of environmental flow constraints can hide important trade-offs (Suen and Eheart, 2006) and consequently it may be more helpful to treat environmental needs as a criterion, albeit poorly defined, to better understand the trade-offs with other criteria.

The solutions presented by the filled symbols in Figure 3-3 are the only solutions in which the desalination plant has been selected. The steepening of the Pareto front just before a desalination plant is included in the solution set is attributed to the high fixed cost of constructing the desalination plant. What is particularly striking about the filled solutions is the sensitivity of the desalination plant to the Wollondilly environmental flow constraint. When no constraint is imposed (Scenario 4), the desalination plant is only selected if solutions produce restriction frequencies of less than 5%. In contrast, if the constraint is imposed (Scenario 3), a desalination plant is selected for all solutions with restriction frequencies less than 17%.

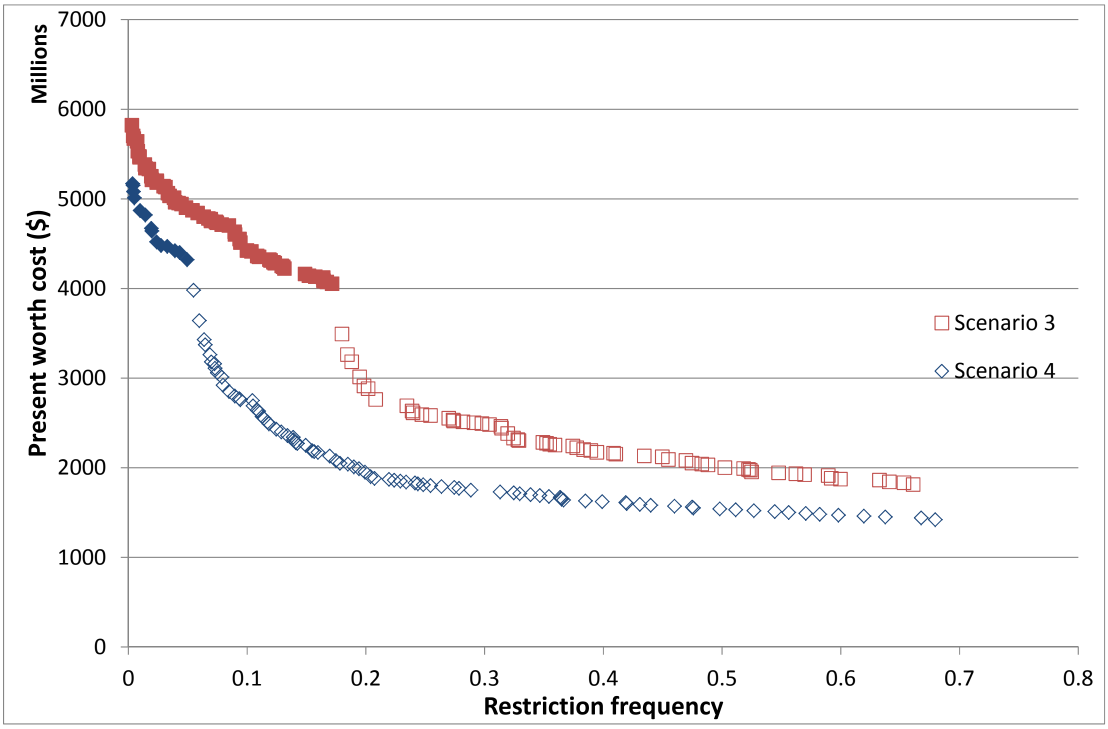

Figure 3-3 Approximate Pareto-optimal front for Scenario 3 (with environmental flow constraints) and Scenario 4 (without environmental flow constraints). The filled-in points represent solutions that include a desalination plant

Each solution on the Pareto front in Figure 3-3 corresponds to a particular decision vector. To gain a better understanding of the sensitivity of the solutions to the Wollondilly constraint, the relationships between subsets of the approximate Pareto-optimal decisions for Scenarios 3 and 4 are analysed. Figure 3-4(a) shows the relationship between decision x_5, the desalination plant capacity, and x_6, the desalination plant trigger, for the solutions that adopt desalination. Regardless of the desalination plant capacity, the trigger level lies between 0.5 and 0.75 for both scenarios. However, when there is no constraint on Wollondilly releases (Scenario 4), the desalination capacity lies in the range 200 to 300 ML/day. In contrast, for Scenario 3, the capacity ranges from 100 to 500 ML/day suggesting interaction with other variables. Figure 3-4(b) shows the relationship between decision x_8, the base gain for Warragamba, and x_9, its incremental gain. When the Wollondilly flow constraint is imposed (Scenario 3), all base gains except one, are greater than or very close to 10,000 and most of the incremental gains are greater than 100. This means the WATHNET5 simulation model assigns the highest preference to keeping water in Warragamba and thus will seek to supply the Sydney zone from other sources before accessing Warragamba. In contrast, when no constraint is imposed (Scenario 4), the situation is more complex with a negative linear relationship between base and incremental gain – increasing the base gain by 100 is offset by a reduction in incremental gain of about 50. This suggests there are complex interactions between the Warragamba gains and other decisions and, therefore, no simple interpretation can be made. Figure 3-4(c) displays the relationship between the capacity of Welcome Reef and restriction frequency. When the Wollondilly flow constraint is imposed, Welcome Reef has a consistently smaller capacity reflecting the fact that the Wollondilly constraint limits the utility of storage on the Shoalhaven River. Figure 3-4(d) shows the relationship between the decision x_4, the level-one restriction trigger, and restriction frequency. There is little difference between Scenarios 3 and 4, with a lower trigger associated with lower restriction frequencies. Furthermore, in virtually all cases, decision x_5 was at its lower limit of 0.05. This suggests that for the adopted criteria, the optimal strategy is to impose the severest restrictions as soon as possible – that said, such a strategy would be unlikely to be socially acceptable.

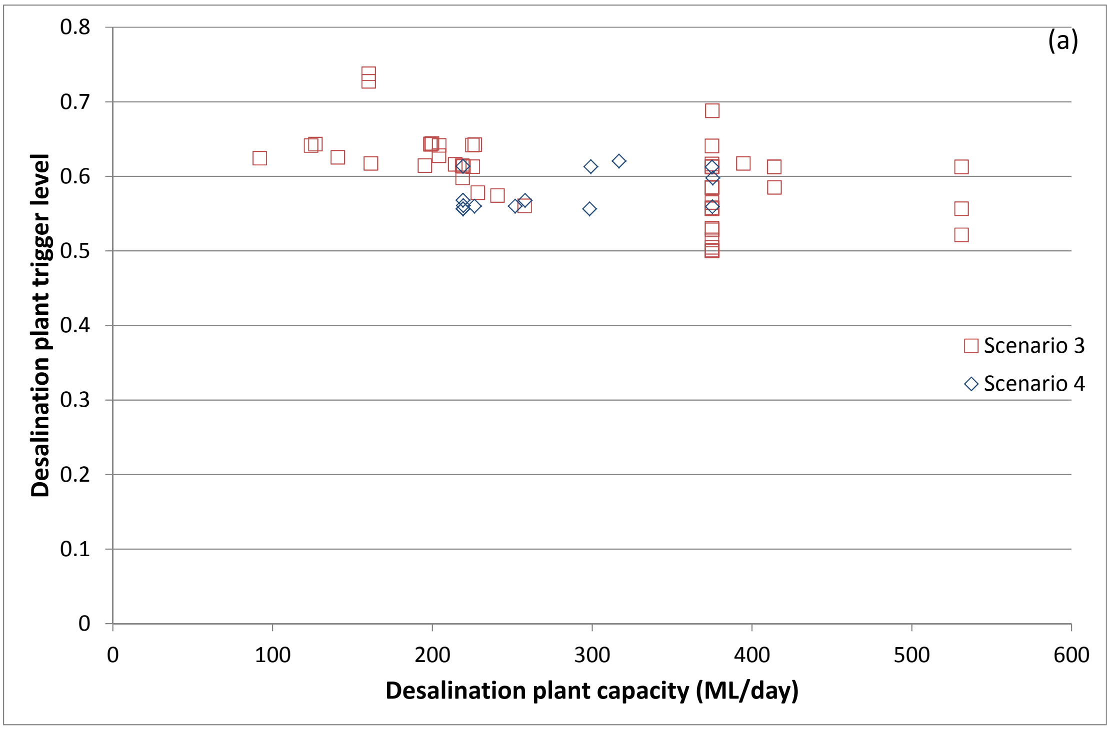
(a)
Desalination plant trigger level
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0
100
200
300
400
500
600
Desalination plant capacity (ML/day)
Scenario 3
Scenario 4

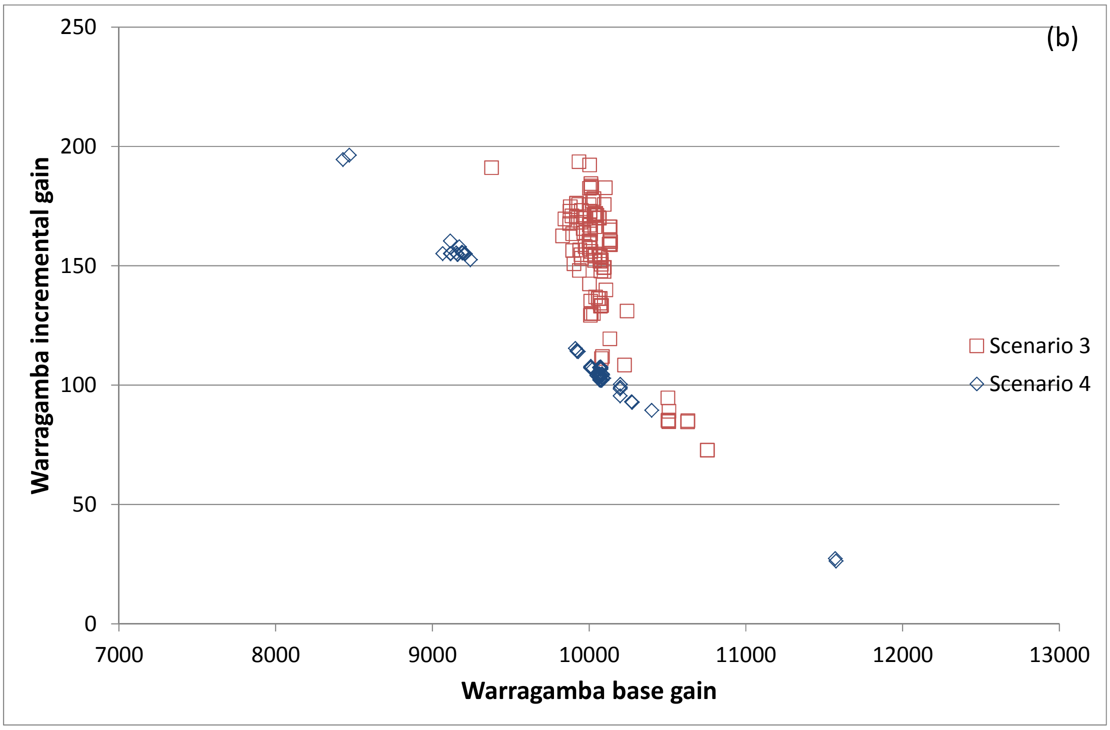
(b)
Warragamba incremental gain
250
200
150
100
50
0
7000
8000
9000
10000
11000
12000
13000
Warragamba base gain
Scenario 3
Scenario 4

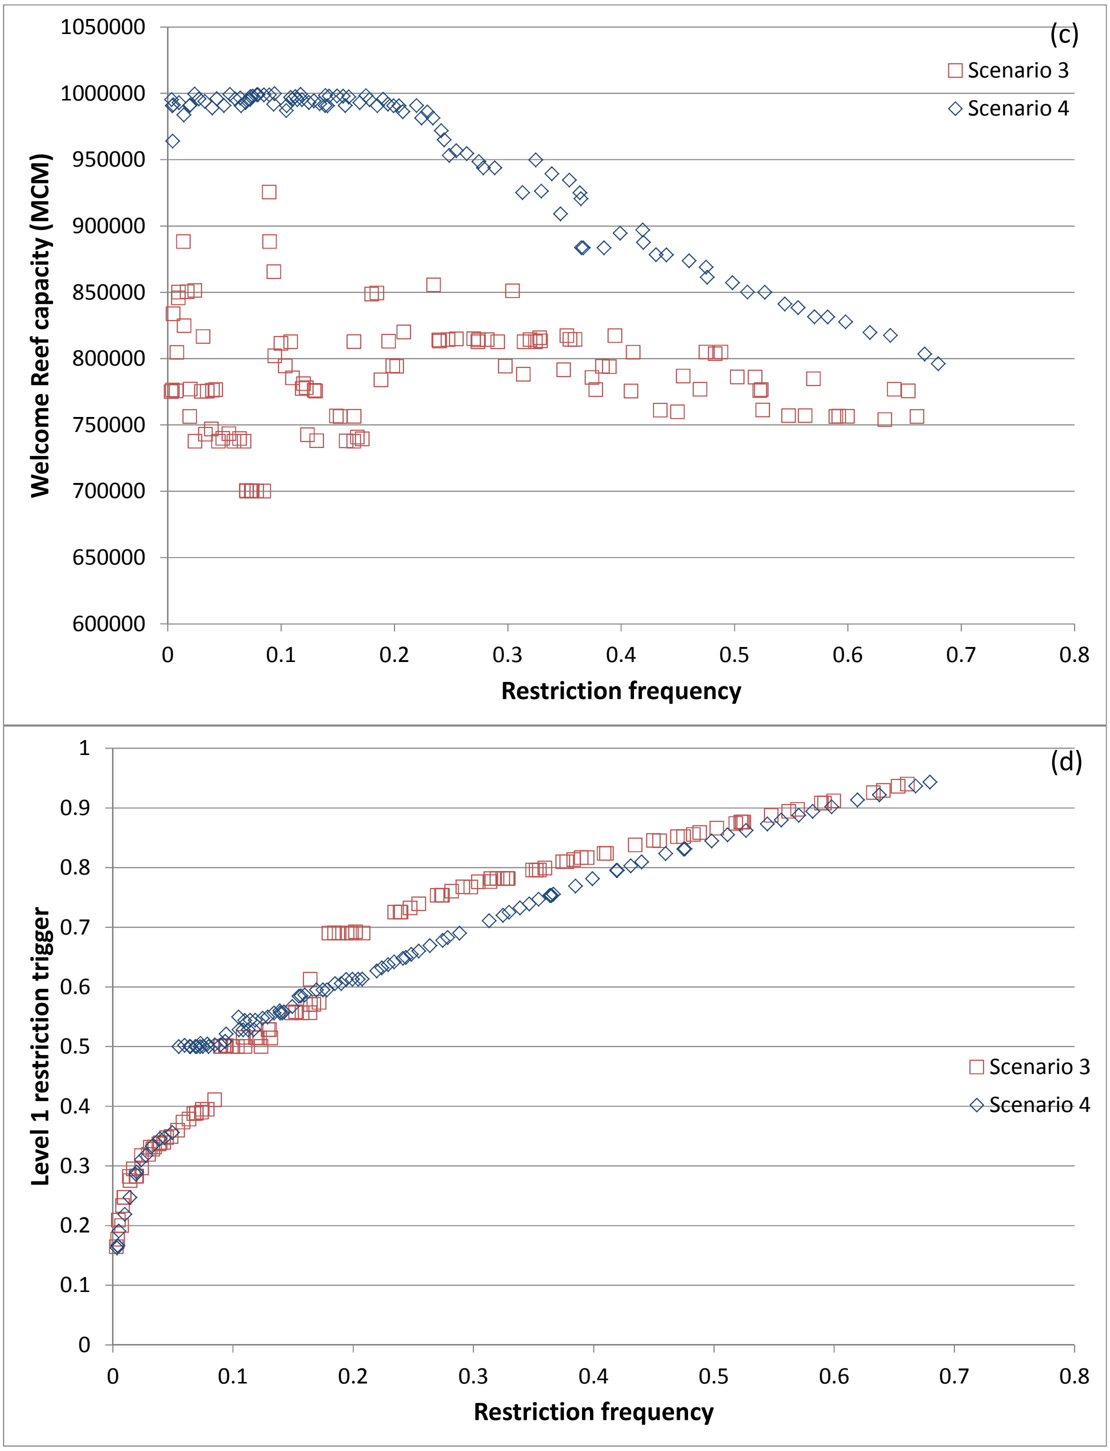

Figure 3-4 Comparison of approximate Pareto-optimal decisions for Scenario 3 (with environmental flow constraints) and Scenario 4 (without environmental flow constraints): (a) desalination plant capacity (ML/day) versus desalination plant trigger level; (b) Warragamba base and incremental gain; (c) Welcome Reef Capacity as a function of restriction frequency; and (d) restriction frequency versus level-one restriction trigger

Figure 3-5 displays the relationship between the two pump marks, x_1 and x_2, and the level-one restriction trigger x_4 against the restriction frequency for each scenario. For the scenario with no environmental constraint (Scenario 4), the Warragamba pump marks associated with the lowest restriction frequency are low because the presence of the desalination plant reduces the dependence of the system on transfers from the Shoalhaven. Without the desalination plant, however, the Warragamba pump mark jumps close to 1 and then declines to about 0.3 as the restriction frequency increases. With the exception of some interaction with Warragamba pump marks for restriction frequencies between 10 and 20%, the Avon pump mark largely lies in the range 0.3 to 0.4. In contrast, Scenario 3 (environmental constraint imposed) reveals a very different behaviour for the Warragamba pump mark which is mainly in excess of 0.7. This suggests the Wollondilly flow constraint forces transfers to start much earlier in the Warragamba drawdown. As a result, there is a higher chance that Warragamba will spill resulting in a wasted transfer and an overall higher pump cost than would occur with a lower pump mark.

The comparison of the Scenario 3 and 4 solutions highlights the complexity of the relationships between decisions. While Figure 3-3 displays a substantial cost trade-off between Scenarios 3 and 4, the analysis of Figures 3-4 and 3-5 suggests that it is not straightforward to interpret the difference in solutions. The interactions between decisions appear to involve, in many cases, more than two variables. This underscores the importance of conducting optimization using the full decision space.

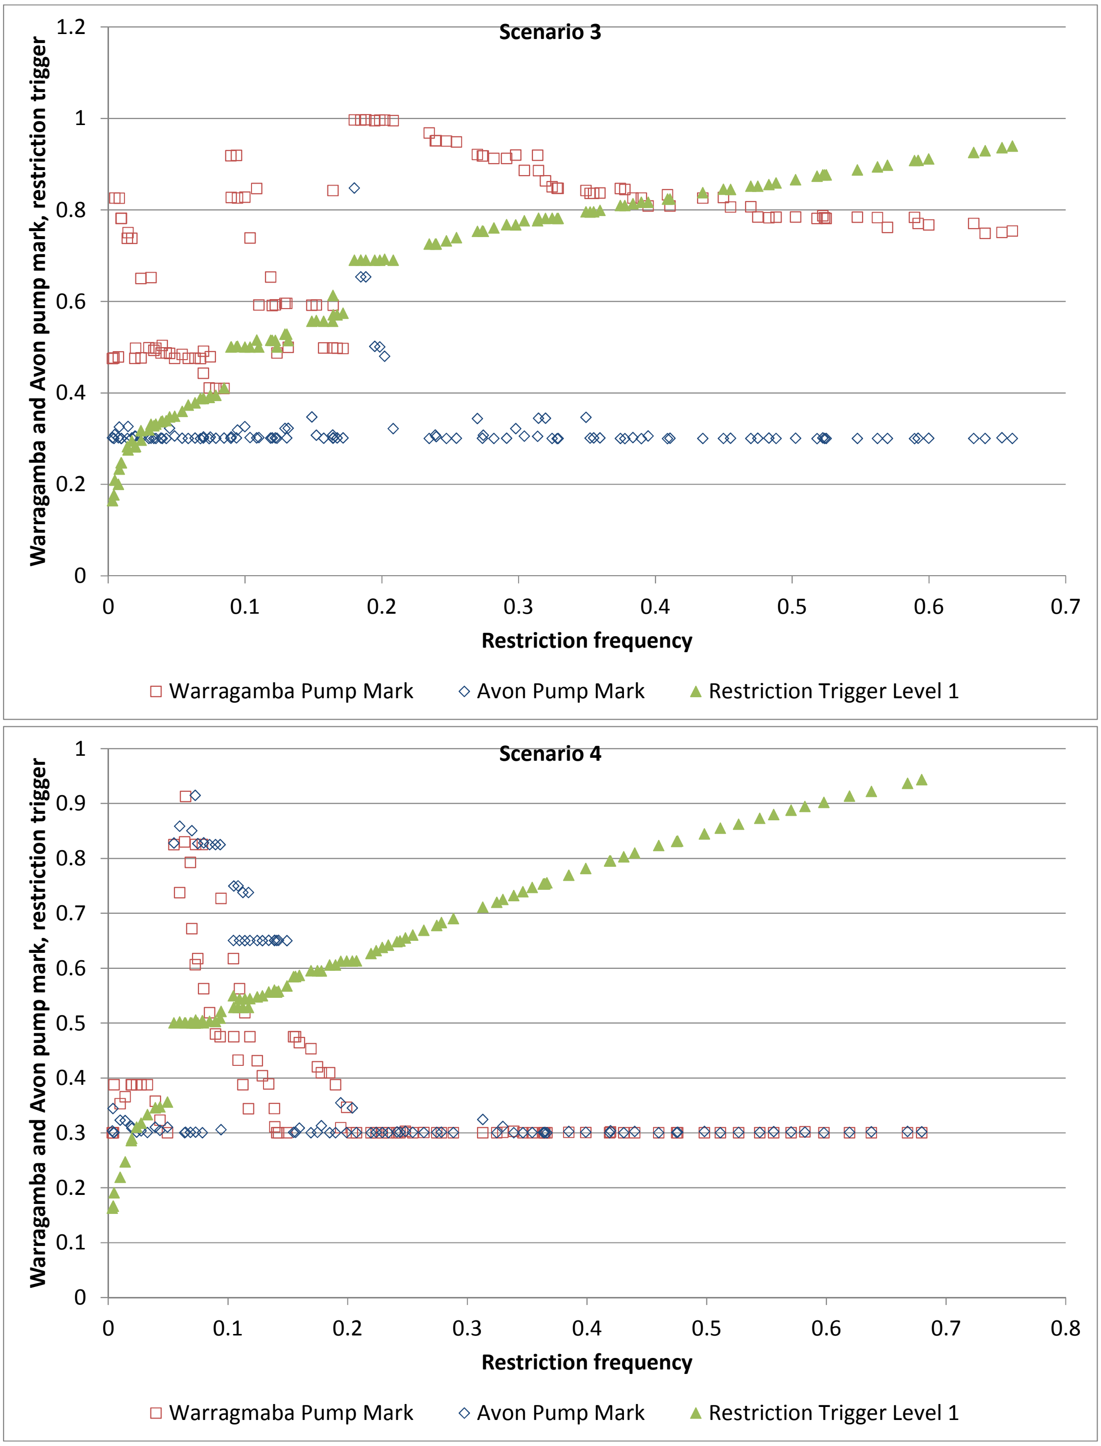

Figure 3-5 Plot of Warragamba and Avon pump marks and level-one restriction trigger against restriction frequency for Scenario 3 (with environmental flow constraints) and Scenario 4 (without environmental flow constraints)

3-5-2-2 Three-Objective Case Study – Scenario 5

Scenarios 3 and 4 represent the extremes in terms of environmental stress on the Wollondilly River and suggest there is a significant trade-off between the environmental stress and other objectives. For this reason, it is worth exploring the trade-offs more fully by undertaking an optimization using all three objectives, namely minimize restriction frequency, minimize present worth cost and minimize environmental stress on the Wollondilly River – this represents Scenario 5.

Figure 3-6 presents all the approximate Pareto-optimal solutions plotted against restriction frequency and present worth cost with a color-coded scale for environmental stress. For a given restriction frequency, reducing the environmental stress increases the present worth cost. However, what is of greater interest and practical significance is the variability in trade-offs between present worth cost and environmental stress as restriction frequency changes. For restriction frequencies less than 7%, the difference in present worth cost between the best and worst environmental outcomes ranges between $600 and $700 million. However, between 7% and 18%, the present worth cost difference increases by about a factor of two – this coincides with the transition to desalination. Beyond restrictions frequencies of 18%, no solution uses desalination and the cost gap rapidly closes.

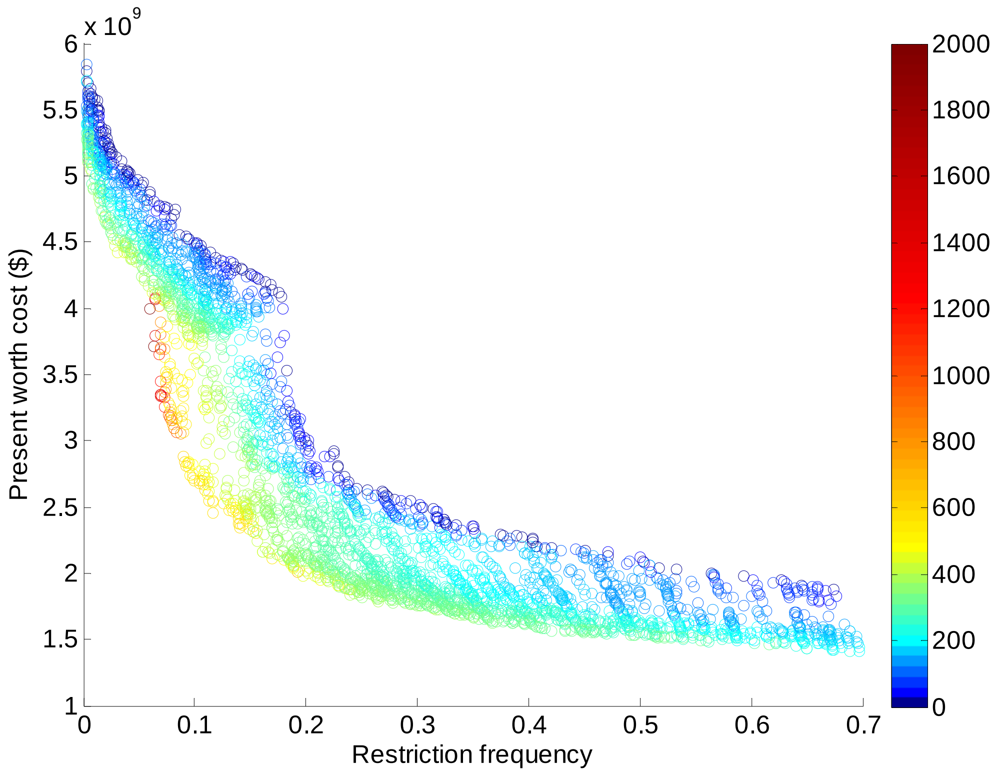

Figure 3-6 Approximate Pareto-optimal solutions for Scenario 5 are plotted against restriction frequency and present worth cost. The color code describes the environmental stress

3-5-3 The Drought Security Cost Trade-off: Scenarios 6 and 7

So far all the scenarios were based on an expected drought security return period of 500 years – that is, the approximate Pareto-optimal solutions ensure the system can cope with droughts having an expected return period up to 500 years without "running out of water". In this section, the sensitivity of the approximate Pareto-optimal solutions to the drought security return period is examined. Two scenarios, 6 and 7, which respectively use 500 and 10,000 year hydro-climate time series are considered. For both scenarios, restriction frequency and present worth cost are minimized with no environmental constraint on releases in the Wollondilly River.

As a prelude, one of the approximate Pareto-optimal solutions for the 500-year scenario was simulated using the 10,000-year input series – this solution had no desalination and a relatively high restriction frequency of 22%. Figure 3-7(a) shows the plot of total storage for the most severe drought in the first 500 years. It is observed that during this drought, the system ran dry but just avoided unplanned shortfalls. The fact that the optimized decisions just avoided unplanned shortfalls in the 500-year scenario would suggest the system becomes vulnerable when exposed to severer droughts. This is confirmed in Figure 3-7(b) which shows a plot of unplanned shortfall expressed as a percentage of total demand for the 10,000-year scenario. The limitations of the 500-year return period solution are abundantly clear. Unplanned shortfalls of up to 95% of demand, sustained for periods up to 6 months, would most likely lead to catastrophic outcomes. This vulnerability is unavoidable in systems totally reliant on climate-dependent sources of water.

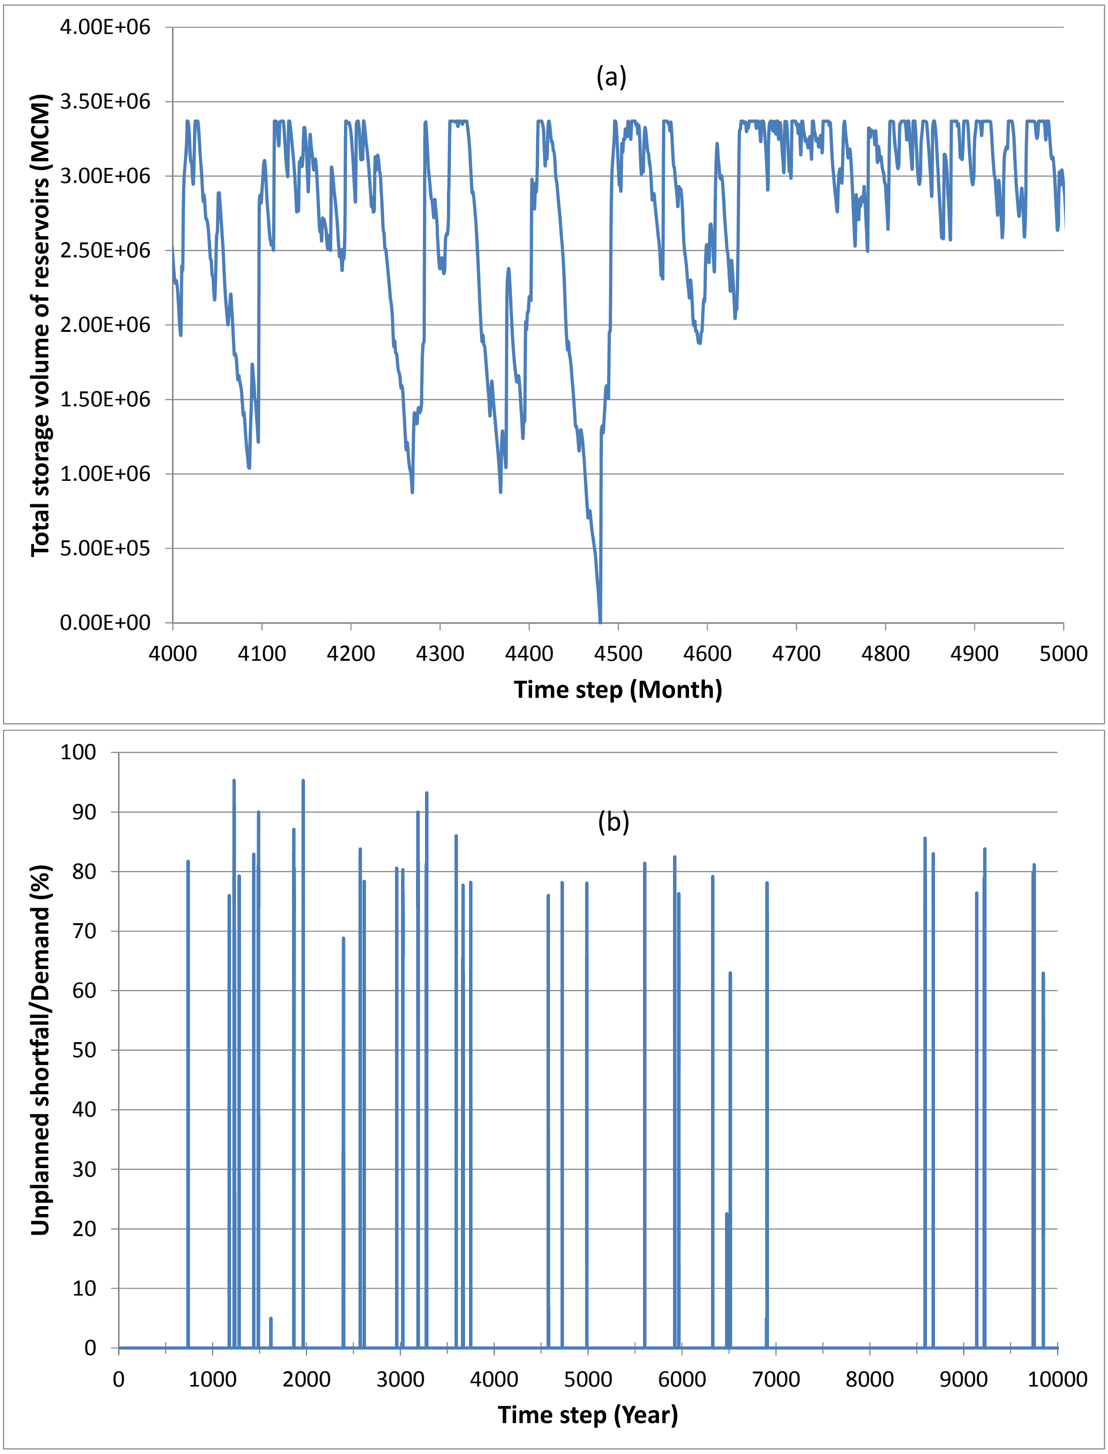

Figure 3-7 Pareto solution from Scenario 6: (a) Time series of total storage during the most severe drought in the first 500 years; and (b) Time series of unplanned shortfalls, expressed as a percentage of demand, for 10,000 years

Figure 3-8 presents the approximate Pareto-optimal fronts for the present worth cost and restriction frequency criteria for the 500 and 10,000-year scenarios. The shift in the Pareto front is striking. For a 10% restriction frequency, the present worth cost increases from $2,600 million to $8,300 million. This large jump in cost arises from the need to avoid unplanned shortfalls in droughts considerably more severe than experienced in the 500-year scenario. To better understand the impact of using the 10,000-year scenario, Table 3-5 presents three pairs of solutions on the Pareto fronts selected so that each pair has a similar restriction frequency. There are four key differences between the 500- and 10,000-year scenarios:

1. For the 500-year scenario no desalination plant was selected, while in the 10,000-year scenario, all solutions had the desalination plant capacity set close to the upper limit of 1000 ML/day.

2. The Warragamba pump mark jumps from 30% in the 500-year scenario to 68% in the 10,000-year scenario to commence transfers from the Shoalhaven much earlier in any drought.

3. For the 10,000-year scenario, the Warragamba base and incremental gains ensure that the Warragamba is preferentially drawn down. This strategy triggers transfers from the Shoalhaven earlier than if all reservoirs were balanced according to the space rule.

4. All three solutions for the 500-year scenario opt for Welcome Reef close to its maximum capacity of 1,000,000 ML. In contrast, for the 10,000-year scenario, the size of Welcome Reef decreases with increasing restriction frequency because the desalination plant capacity remains essentially unchanged close its maximum capacity.

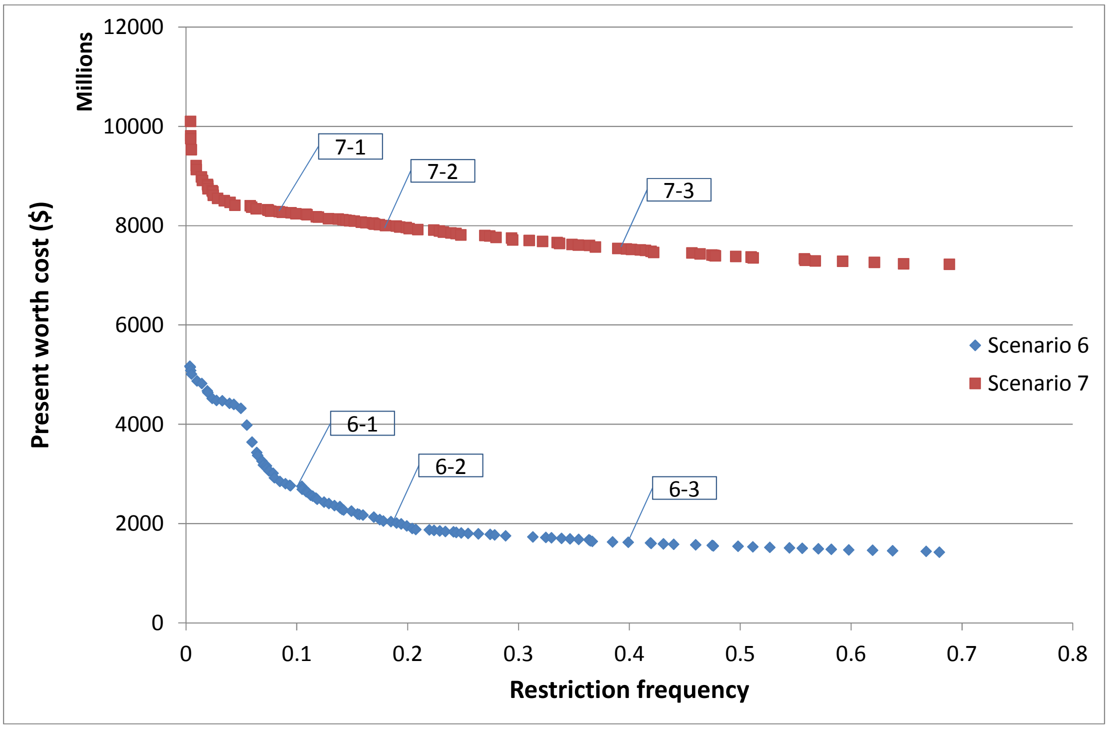

Figure 3-8 Comparison of approximate Pareto frontier for Scenario 6 (500-year record) and Scenario 7 (10,000-year record)

Table 3-5 Summary of labelled solutions on Pareto fronts in Figure 3-8

	Solution	6-1	7-1	6-2	7-2	6-3	7-3
	Drought return period, years	500	10,000	500	10,000	500	10,000
Objectives	Restriction frequency (%)	13	13	22	22	45	45
	Present worth cost ($m)	2400	8140	1860	7910	1570	7450
Decisions	Pump mark Warragamba	0.40	0.65	0.30	0.50	0.30	0.53
	Pump mark Avon	0.65	0.31	0.30	0.31	0.30	0.32
	Level 1 restriction trigger	0.55	0.56	0.63	0.61	0.82	0.80
	Trigger increment	0.05	0.05	0.05	0.05	0.05	0.05
	Desalination plant capacity (ML/day)	0	923	0	980	0	922
	Desalination plant trigger	---	0.52	---	0.50	---	0.54
	Welcome Reef capacity (ML)	994342	970629	981469	890867	873783	894823
	Warragamba base gain	10064	10267	10073	9984	10073	10141
	Warragamba incremental gain	104	11	107	24	102	13

3-6 Discussion

The seven case study scenarios have demonstrated the value of an optimization methodology that addresses the three shortcomings identified in the previous literature. The overarching conclusion from the case study is that, in the case of urban headworks systems, failure to optimize the full mix of operational and infrastructure decisions, failure to allow for droughts with high return periods and failure to explore trade-offs implicit in “soft constraints” can produce demonstrably inferior solutions.

The issue of drought security is of paramount importance for cities located in regions subject to severe prolonged droughts. The prospect of “running out of water” for an extended period would threaten the very existence of the city and its social and economic fabric. The case study highlighted the potentially serious shortcomings of solutions based on short historical or synthetic streamflow records. Very different approximately optimal solutions were found when securing against an expected 500- and a 10,000-year drought. For a 10% restriction frequency, the optimal solution for the 10,000-year record incurred a present worth cost over three times that for the 500-year record. While the optimal solution for the 500-year scenario just avoided unplanned shortfalls in the worst drought of the 500-year record, the more severe droughts in the 10,000 year record resulted in extended and unsustainable periods of unplanned shortfalls. It is therefore critically important that simulation record lengths, over which system performance is evaluated, are sufficiently long to match drought security expectations. It is not an uncommon industry practice to design for the worst historical drought and then add a reserve [for example, Cloke and Samra (2009)]. Such an approach does not communicate the risk of running out of water, which in the case of a large urban system could be potentially catastrophic. By rerunning the optimization problem (3.2) for different record lengths (as done in Scenarios 6 and 7), the trade-off between drought security and other criteria can be explicitly explored to enable an informed decision.

The issue of confidence in the drought return period deserves comment. The Pareto-optimal solutions given by (3.2) secure the system against droughts with return periods up to an expected value of N years. The actual return period may differ from the expected value. If one needs more confidence in the return period, the following pre-conditioning algorithm can be used to reduce uncertainty in the return period for the N-year record: Generate M replicates of length N years; rank the replicates using a suitable low-flow statistic such as the minimum k-year sum; select the replicate corresponding to the median rank.

Even if drought security is adequately accounted for, failure to optimize the full mix of infrastructure and operational decisions and explore trade-offs implicit in "soft" constraints can result in solutions that involve far greater economic cost than is necessary. The issue of soft constraints can be particularly challenging. To transform an environmental constraint into an objective requires the formulation of an environmental response function that maps decisions into a meaningful metric of environmental response. It is widely accepted that this is a difficult task constrained by limited data and difficulties in identifying causal mechanisms. It is acknowledged that the environmental stress function given by (3.7) is subjective and most likely incomplete. Accordingly, the main insight is not quantitative but an awareness that there are very significant trade-offs between environmental response and cost and that these trade-offs are a non-linear function of restriction frequency. In view of this, a strong case could be made to invest in studies to better inform the specification of the environmental response function and so better inform the trade-off process. Seen from this perspective, the optimization methodology advanced in this study is part of an iterative process involving progressive refinement of information and objectives.

3-7 Conclusions

This chapter has formulated and demonstrated a multi-objective optimization methodology for urban water supply headworks planning and management that produces solutions with demonstrably greater practical value. Its principal contribution is the identification of three practically significant shortcomings in the literature and a methodology to resolve these shortcomings. The case study, motivated by the headworks system for Sydney, Australia, demonstrated the significant manner in which these shortcomings can compromise the practical value of so-called Pareto-optimal solutions.

Urban headworks systems are typically planned and operated in a manner that ensures a very low risk of "running out of water" or catastrophic water shortages. The case study demonstrates the very considerable sensitivity of Pareto-optimal solutions to the return period of the worst drought. While this may seem self-evident, the literature has largely ignored this issue and repeatedly published optimal solutions conditioned on historical records or short stochastic records. Where high levels of drought security are required, such solutions are flawed and methodologies that produce such solutions should be avoided. Our approach addresses drought security explicitly. It identifies near-optimal solutions that are constrained so that the system does not "run dry" in severe droughts with expected return periods up to a specified value.

In many cases, the operating rules that control the operation of the headworks system are conditioned on the system infrastructure. It is therefore vital in optimization studies, in which new system infrastructure is to be added or existing infrastructure modified, that key operating rules are optimized jointly with the infrastructure options. While this may substantially increase the dimension of the decision space, it is not worth the risk of obtaining significantly inferior solutions. In a similar vein, the imposition of "soft" constraints, such as the environmental flow constraints in the case study, runs the risk of missing potentially good solutions. In the case of constraints to which system performance is sensitive, their reformulation as objectives within a multi-objective optimization framework can enable a more thorough and computationally efficient assessment of trade-offs, an outcome that would be difficult to achieve using conventional sensitivity analysis.

It is important that "good" solutions be found for the "right" problem. This study has made a significant contribution towards this goal by addressing identifiable shortcomings. However, in practice, planners have to deal with scenario uncertainty in which assumptions have to be made about model structure and exogenous factors such as system forcing and political and social constraints. There can be a considerable and difficult-to-quantify uncertainty about these scenarios. In such cases, one can argue that "good" solutions need to produce good outcomes across the range of plausible scenarios – in other words, "good" solutions need to be robust (Matalas and Fiering, 1977). There is a growing literature on robust optimization [see, for example, Deb and Gupta (2006)] whose concepts can be applied to the urban headworks problem.

In this chapter, a steady-state scenario was considered in which demand was assumed to be constant over the simulation interval. However, in the face of growth in urban populations and accompanying growth in demand for water, optimizing decision for a steady-state scenario is insufficient. While it may identify “good” solutions for a particular population, it provides no information on how to best schedule future infrastructure investments and future changes in operating rules to cope with the growing demand. This is the topic of the next chapter.

Chapter 4
Application of Multi-Objective Optimization to Scheduling Capacity Expansion of Urban Water Resource Systems

4-1 Introduction

With the worldwide trend of significant population growth in major cities, it is expected that most urban water resource systems will face a growing demand for water in addition to future climate change and changing expectations about level of service and acceptable impacts on environmental systems. In the face of such change, the performance of the water resource system is expected to change, most likely for the worse, resulting in the need to change the mix of infrastructure and operating rules. This chapter considers the scheduling capacity expansion problem from a multi-objective perspective. It generalizes the ideas developed in Chapter 3 to consider the question of when as well as how much should the infrastructure and operating rules be changed to serve the changing needs of a city.

Capacity expansion involves the provision of additional yield by increasing the capacity of existing infrastructure and the construction of new infrastructure harvesting new sources of water. In its simplest manifestation, capacity expansion deals with sizing reservoirs. For example, Khaliquzzaman and Subhash (1997) developed a model for sizing multiple reservoirs. Mousavi and Ramamurthy (2000) proposed an optimization method to determine the optimal multi-reservoir system design for water supply by converting two objectives, minimum cost and minimum water deficit, to a single objective function. Nainis and Haimes (1975) applied a multilevel approach for capacity expansion in water resource systems; they extended classical benefit-cost analysis, describing their approach as dynamic benefit-cost analysis. Yang et al. (2007) applied the concept of multi-objective optimization to reservoir capacity expansion trading off two objectives, minimizing capital costs and minimizing costs arising from water shortages.

Other studies have extended the concept of capacity expansion to include options other than those dealing with sizing reservoirs. For instance, Nakashima et al. (1986) developed a two-phase heuristic optimization technique to determine a water supply system layout and to size water production and transmission facilities. Hsu et al. (2008) developed a methodology to detect potential bottlenecks of a water distribution system with the aim of facilitating capacity expansion plans. Dziegielewski et al. (1992) incorporated drought management plans into their capacity expansion analysis; they assessed the trade-off between long-term and short-term options to manage drought by estimating the expected cost of coping with drought. Basagaoglu and Yazicigil (1994) considered capacity expansion in the context of a groundwater system.

All the aforementioned studies have focused on decisions at the start of the planning period. However, decisions to expand capacity can be implemented at different points of time over the planning period to take advantage of delaying a portion of investment outlays. Although the construction of large infrastructure at the start of the planning period exploits the economies of scale, the time discounting of costs and the dynamics of growth may nonetheless favour smaller projects staged over the planning period. To analyse this trade-off a number of studies have considered scheduling expansion (Grossman and Marks, 1977; Knudsen and Rosbjerg, 1977; Braga et al., 1985; Kim and Yeh, 1986; Lund, 1987; Watkins Jr and McKinney, 1998; Gillig et al., 2001; Voivontas et al., 2003; Mahmoud, 2006; Chang et al., 2009)

Scheduling expansion problems have typically been formulated to find the timing of predefined projects that minimizes the total present worth cost (PWC). Indeed, given this perspective, the main aim is to find the best sequence of projects (Luss, 1982). However, projects often can be implemented at different scales. Thus, the scheduling capacity problem can be generalized to find the optimum timing and scale of predefined projects – this is referred to as the scheduling capacity expansion problem.

Figure 4-1 illustrates the scheduling capacity expansion process. It plots demand and yield as a function of time. Given the initial yield of the system is Y_0, the system can meet demand up to time T_1. At time T_1, a decision is made to add extra yield ΔY_1. As a result, system yield will exceed demand until time T_2. In a similar manner, decisions are taken at later times to provide additional yield. Thus T_1, T_2 and so on represent change points at which decisions are made. The period between two consecutive change points is called a planning stage.

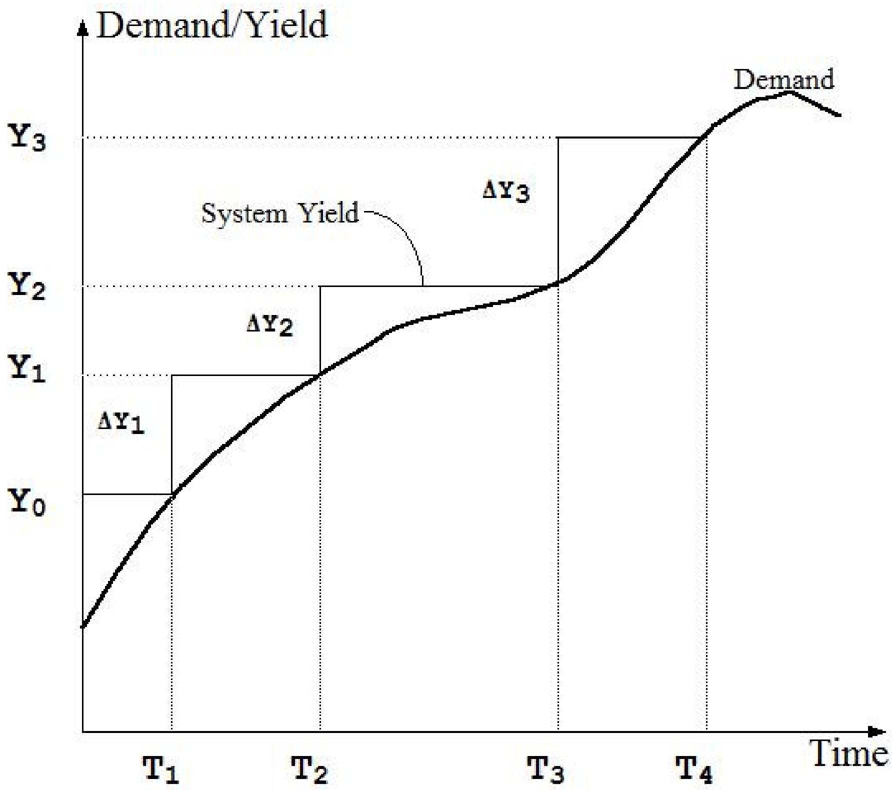

Figure 4-1 Schematic of scheduling capacity expansion over a planning horizon

A number of studies have investigated the scheduling capacity expansion problem in a water resources context. Knudsen and Rosbjerg (1977) developed a general dynamic programming algorithm to find the optimal scheduling of water supply projects. Kim and Yeh (1986) introduced a heuristic solution procedure to find an optimal sequence of capacity expansion projects. Connarty and Dandy (1996) used genetic algorithm optimization to find the optimum sequence involving nine reservoirs for a case study based on the southeast Queensland headworks system. Watkins Jr and McKinney (1998) developed a model involving capacity expansion of an integrated surface and groundwater system. In a similar way, Chang et al. (2009) applied an optimization model to determine the capacity expansion schedule for groundwater supply. They considered a variety of expansion options involving surface and groundwater sources such as increasing borehole, reservoir and desalination plant capacity. Mahmoud (2006) employed a high dimension dynamic programming model to determine the optimal expansion schedule of a desalination plant. In all these studies solely infrastructure options were considered as decisions. The interaction between infrastructure and operating rule options was not considered.

The high capital costs and environmental impacts associated with expanding or building new major urban water infrastructure warrant the investigation of scheduling system operating rules such as reservoir operating rules, demand reduction policies and drought contingency plans, as a way of delaying or avoiding the expansion of water supply infrastructure (Lund, 1987; Rosenberg et al., 2008). Lund (1987) incorporated conservation rules into the scheduling capacity expansion problem. He demonstrated the benefit of using conservation rules to defer water treatment plant expansion. In Lund's study the present worth of conservation cost and capacity expansion cost was minimized to find the optimum time to add new capacity to the system. However, a drawback of this approach is that discounting conservation costs can lead to higher levels of demand reduction in the future than in the present. This raises a socially-sensitive equity issue.

Decision makers usually set the level and frequency of demand restrictions based on a level of service acceptable to the community. To identify what is acceptable, it is important to identify the trade-off between conservation (or restriction) and infrastructure costs. Rubinstein and Ortolano (1984) applied a dynamic programming algorithm to demonstrate the trade-off between the present value of the cost of implementing projects and the expected value of the costs to cope with emergencies, i.e. imposing restrictions. Although the coping cost in emergency situations is separated from project capital cost, the fact that the coping cost is a discounted cost suggests that severer restrictions may be deferred to future planning stages.

All the reviewed studies suffer from one or more significant shortcomings:

1. Most of the studies considered only a single objective. The drawback of using a single objective is that it is not possible to identify the trade-off between capital and operating costs and the cost of restrictions. However, the main drawback is that discounting restriction costs can lead to higher levels of demand reduction in the future than in the present. The optimization process hides a socially-sensitive equity.

2. The studies optimized either decisions involving infrastructure alone or when operational decisions were included, they were not jointly optimized with infrastructure decisions. As shown in Chapter 3, changing the infrastructure

within a system without concomitant changes to operating rules can result in significantly inferior outcomes.

3. All the studies failed to address drought security adequately primarily because insufficient streamflow data was used to sample severe droughts. As this issue was treated extensively in Chapter 3, it was not explicitly reviewed here.

This chapter presents the application of a multi-objective optimization approach to scheduling capacity expansion in an urban water resource system that addresses the shortcomings identified in previous studies. The chapter is organized as follows: First, a new formulation of the multi-objective scheduling capacity expansion problem is presented. Using a case study based on the Canberra headworks system, twelve scenarios are investigated to demonstrate the significance of the identified shortcomings and how the proposed approach deals with them.

4-2 The Multi-Objective Scheduling Capacity Expansion Problem

This section presents a general formulation of the scheduling capacity expansion problem that addresses the shortcomings identified in previous applications. This involves generalizing the formulation presented in Section 3-3 to incorporate the staging of decisions and to allow for the stochastic nature of future inputs to the system. The section concludes with a review of optimization methods with the goal of identifying algorithms suited to the scheduling problem under consideration.

4-2-1 Formulation

Suppose the planning period of T years is subdivided into M planning stages with the i^{th} stage commencing at time T_i. To account for climate variability and other stochastic inputs, the inputs are replicated N times over the planning period by sampling from a suitably constructed probability model of the inputs. For each replicate r, q_{tr} is a vector of streamflow and climate values at multiple sites for year t, and d_{tr} is a vector of unrestricted demand at multiple sites for year t. The notation $Q^r_{u:v}$ denotes the time series of vectors $\{q_{tr}, t=u,...,v\}$.

Let $x_i = \{x_i^1,..,x_i^p\}$ denote a p-vector of decision variables that are implemented at the start of the i^{th} planning stage. The decision vector can represent a mix of infrastructure options and operating rules. A solution is defined as a sequence of decision vectors over M planning stages $x = \{x_1,\ldots, x_M\}$.

The simulation model produces N replicates of response denoted by $Z_{1:T}^r = M[x, Q_{1:T}^r, D_{1:T}^r], r = 1,..,N$ where $Q_{1:T}^r$ and $D_{1:T}^r$ represent the streamflow and demand for the r^th^ replicate of the T-year planning period. The performance of the system is evaluated using K objective function

$$f_i(x) = \sum_{t=1}^{T} \phi(t) E[f_i(Z_{1:t}(x_{1:t}))] \approx \frac{1}{N} \sum_{t=1}^{T} \phi(t) \sum_{r=1}^{N} f_i(Z_{1:t}^r(x_{1:t})),\ i = 1,..,K \tag{4.1}$$

where $x_{1:t} = \{x_1,..,x_j : T_j \leq \mathrm{t} < T_{j+1}\}$ is the sequence of projects or decision vectors implemented on or before year t and $\phi(t)$ is a temporal discounting factor. The term $E\left[f_i(Z_{1:t}(x_{1:t}))\right]$ is the expected value of the i^{th} objective function for year t and is evaluated by averaging over the N replicates – the notation emphasizes the fact that the objective function value depends on the response from the simulation model which in turn depends on the decision values.

The multi-objective optimization problem for the scheduling capacity expansion problem involves minimizing the K objective function over the decision space subject to constraints that include constraints on staging decisions, which are discussed further in Section4-3-3. This formulation addresses the shortcomings identified in previous applications in the following ways:

1. The use of multiple replicates of forcing data ensures that drought security can be adequately addressed. In Chapter 3, the issue of drought security was addressed by choosing an input record with sufficient length to ensure the system could cope with droughts up to a specified return period. In the case of scheduling this approach cannot be used because the planning period T is fixed and because the performance of the system changes over time. The use of multiple replicates of forcing data provides a solution to this problem. By selecting the appropriate number of replicates N, one can ensure the system will encounter droughts of appropriate severity.

2. The use of multiple replicates of forcing data ensures that the Pareto-optimal solutions are not dependent on any particular sequence of future climate and demand. This allows the use of a simulation model that can respond to changes in both infrastructure and operating rules. In turn, this enables both operating rules and infrastructure investments to be jointly optimized. The findings of Chapter 3 suggest that such capability is likely to produce significant benefits.

3. The potential equity issue arising from temporal discounting of costs can be addressed in a multi-objective context by exploring the trade-offs between economic and equity criteria.

In the following sections, the benefits of this formulation will be investigated using a case study.

4-2-2 Optimization Methods

The section briefly reviews the optimization methods that have been employed in capacity expansion problems and identifies those best suited for solving the problem described in Section 4-2-1. A review of the literature shows that a variety of optimization methods have been used in capacity expansion problems.

Approaches using some form of linear programming include Khaliquzzaman and Subhash (1997) who used network linear programming for sizing of reservoirs in a water resource system and Mousavi and Ramamurthy (2000) who integrated an optimal control theory approach with successive linear programming to determine the reservoir sizing. However, many capacity expansion problems are not amenable to linear programming approaches because of nonlinearities in objective functions and constraints. As a result, a number of studies have used nonlinear optimization methods. For instance, O'Laoghaire and Himmelblau (1974) applied the branch and bound method. Basagaoglu and Yazicigil (1994) developed three mixed-integer programming models to eliminate nonlinearity in the objective function. Watkins Jr and McKinney (1998) investigated application of two decomposition methods, namely generalized decomposition Benders and outer approximation, to solve problems involving cost functions with both discrete and nonlinear terms. Rosenberg et al. (2008) developed stochastic nonlinear programming to identify the benefit-maximizing options involving conservation and leak reduction programs, infrastructure expansions, and operational allocations under stochastic water availability.

Dynamic programming (DP) (Bellman, 1957) has been used in the sizing and sequencing water resources projects (Butcher et al., 1969; Morin and Esogbue, 1971; Erlenkotter, 1973; Morin, 1973; Erlenkotter and Trippi, 1976; Grossman and Marks, 1977; Knudsen and Rosbjerg, 1977). In a more recent study, Kim and Yeh (1986) presented a heuristic solution procedure that incorporates a shortest path DP method and a cyclic coordinate univariate direct search procedure. The main drawback of DP is that it can only be used for a relatively small number of projects because the number of possible states grows exponentially with the number of projects (Luss, 1982). This so-called curse of dimensionality limits the application of DP (Hsu et al., 2008). A variety of methods has been developed to overcome the curse of dimensionality. Some of these methods have been applied in capacity expansion problems. For instance, Mahmoud (2006) applied objective space dynamic programming (OSDP) in conjunction with mixed integer programming. OSDP is a variant of dynamic programming based on the use of the objective value function as the “state” variable to overcome the “curse of dimensionality” problem.

Evolutionary methods such as genetic algorithms (GAs) do not suffer from the curse of dimensionality or issues related to handling nonlinear equations. Dandy et al. (1985) applied a GA to a water supply system to find optimum water price and project sequences. In a similar way, Chang et al. (2009) hybridized a GA and constrained differential dynamic programming (CDDP) to optimize capacity expansion schedules for groundwater supply. They used GA to investigate capacity expansion alternatives and then applied the CDDP algorithm to compute the optimal pumping policy associated with the selected expansion options. It is worth noting that this hierarchical optimization approach is likely to produce a sub-optimal solution because the pumping policy and capacity expansion were not jointly optimized.

All of the above-mentioned studies have dealt only with a single objective. Rubinstein and Ortolano (1984) used DP in multi-objective capacity expansion. Because DP cannot optimize two objectives jointly, they weighted the multiple objectives to form a single objective. Yang et al. (2007) used a hierarchical approach to integrate a multi-objective genetic algorithm (MOGA) with CDDP; MOGA was used to generate various combinations of reservoir capacity and CDDP was used to distribute optimal releases among reservoirs to satisfy water demand to the extent possible.

Of the general approaches reviewed, those based on evolutionary methods appear best suited for the multi-objective problem described in the previous section. As discussed in Chapter 2, they can interface with complex non-linear simulation models, and handle multiple non-linear objectives and constraints. In view of the satisfactory performance of εMOEA in Chapter 3, it was decided to use εMOEA to solve the multi-objective scheduling capacity expansion problem associated with the case study presented in the remaining sections of this chapter. As in the case of the previous chapter, the main contribution in this chapter is with the improved problem formulation rather than the use of any particular optimization method.

4-3 Case Study: Description and Problem Formulation

This section introduces the case study for this chapter. It considers the water supply headworks system for Canberra, Australia's capital city. An overview of the Canberra system is presented followed by a detailed formulation of the multi-objective scheduling capacity expansion problem.

4-3-1 Description of Canberra System

The Canberra headworks system serves a current population of approximately 420,000. Figure 4-2 presents a schematic of the headworks system. Water is harvested from two catchments, Cotter and Googong, which flank the city to the west and east respectively. A network of pipelines, pumping stations and treatment plants connects four reservoirs to the Canberra demand zone. Releases from the reservoirs have to meet, not only the consumptive needs of the Canberra urban area, but also environmental flow requirements defined in the water authority's operating license.

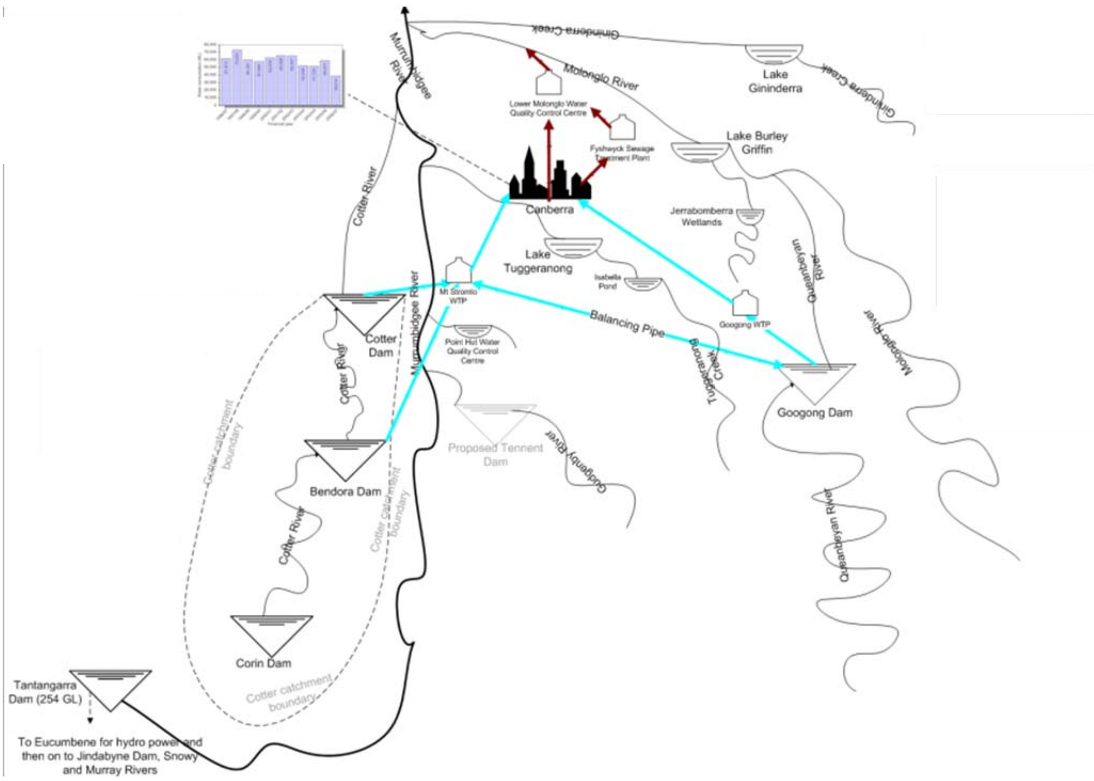

Figure 4-2 Schematic of Canberra headworks system (Adapted from http://www.actew.com.au/Water%20and%20Sewerage%20Systems/Water%20Supply%20System/ACT%20Water%20Supply%20Map.aspx, last visit 16/05/2012)

A WATHNET5 model of the Canberra system was constructed. Figure 4-3 presents the WATHNET5 schematic with the red nodes representing reservoirs, blue stream nodes, yellow demand zones, and black waste/sink nodes. The network of reservoirs, pumping stations and water treatment plants supplies water to the demand zone labelled "Canberra". The existing system includes four reservoirs, Corin, Bendora, Cotter and Googong. The reservoirs have a total storage capacity of 206,732 ML. Googong Reservoir is the largest reservoir in the system with a capacity of 121,084 ML. There are two water treatment plants, Googong and Stromlo WTP, serving the Canberra population.

In this case study, a hypothetical population scenario corresponding to a highly stressed system is presented. The base population is 175% of the current population and is assumed to grow at 1.2% per annum over the 30-year planning period. For simplicity, the same demand time series was used in all replicates; it is noted that this arrangement ignores the correlation between demand and climate and thus may underestimate the consequences of drought. Figure 4-4 presents the 30-year demand time series with and without population growth. What is evident is the high seasonality of water consumption with outdoor water usage during the hot, dry summer months more than doubling unrestricted consumption. Superimposed on this seasonality in consumption is a 43% increase in demand over the 30-year planning period. Multiple replicates of monthly future streamflow data from 2010 to 2040 were sampled from a stochastic model calibrated to an historical record from 1871 to 2009 – the stochastic model was the same as used in the Sydney case study in Chapter 3.

To cater for this increase in demand, three options are available for augmenting supply – these are highlighted in the WATHNET5 schematic by dashed ovals. The first is to increase the capacity of Cotter Reservoir by up to 100,000 ML. The second is to build a new pump station to divert up to 6,000 ML/month from the Murrumbidgee River into Googong Reservoir. The third option is to install domestic rainwater tanks in up to 15,000 houses.

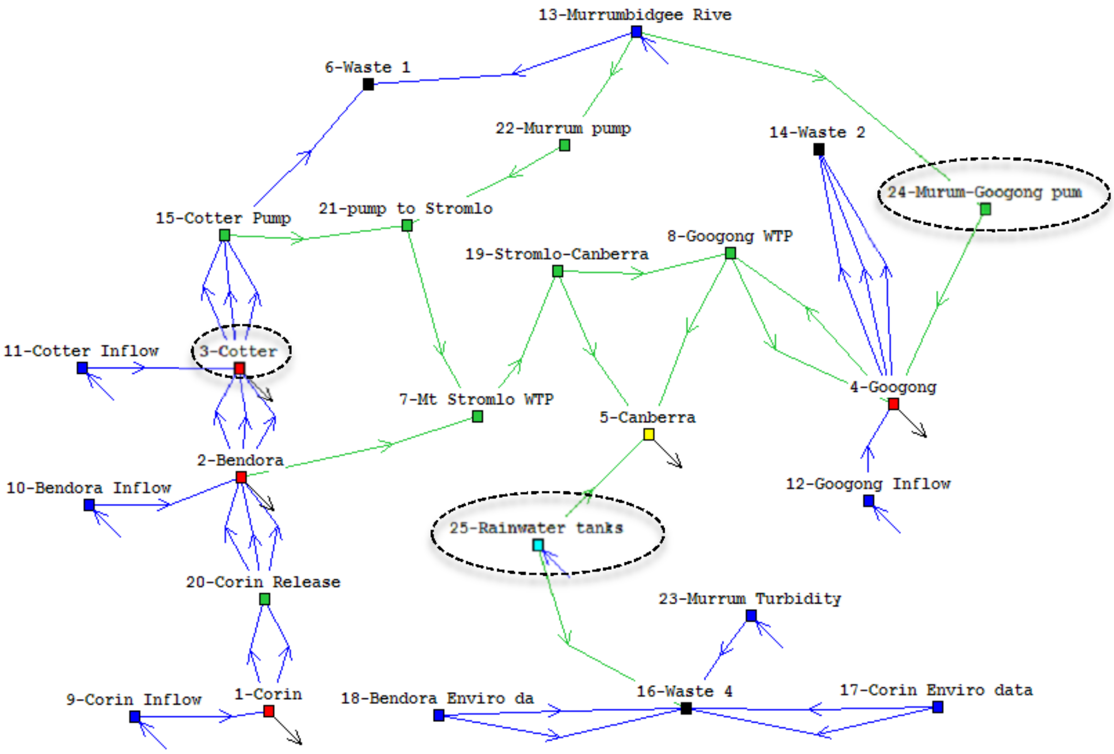

Figure 4-3 WATHNET5 schematic of Canberra headworks system

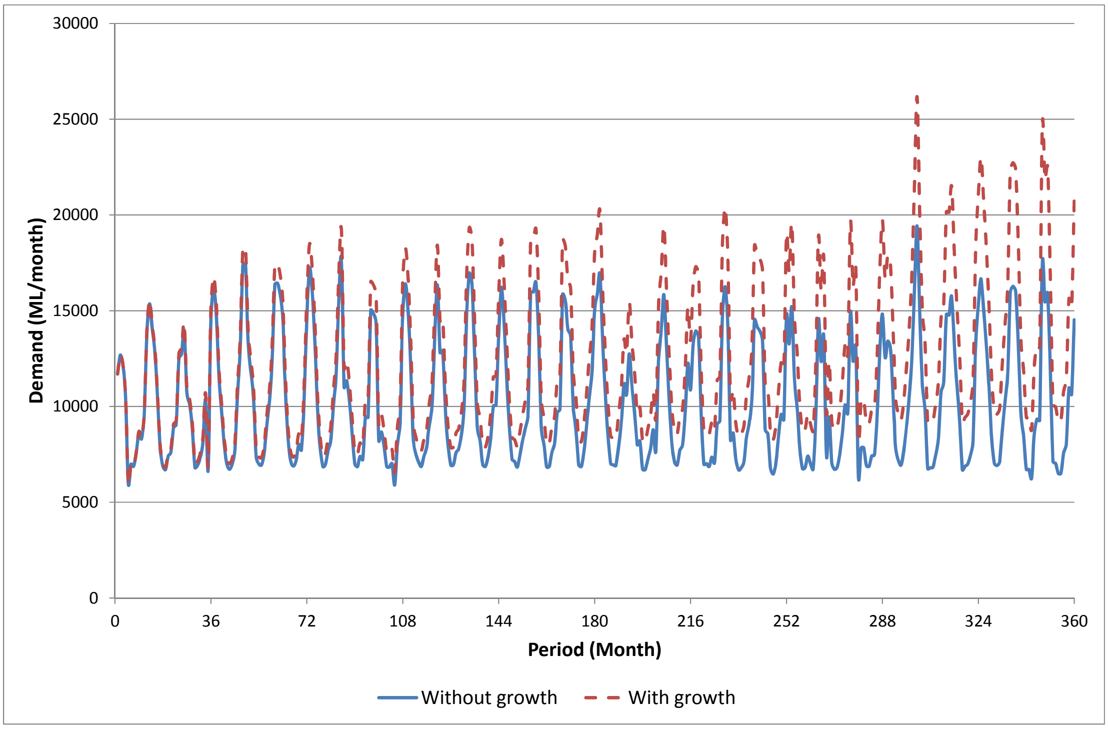

Figure 4-4 Comparison of Canberra unrestricted demand time series with and without growth

The high seasonality in water consumption due to summer outdoor water use indicates there is considerable scope for reducing demand by imposing restrictions on outdoor water use. In this study, four levels of restrictions are available with Table 4-1 presenting the ratio of restricted to unrestricted demand for each level.

Table 4-1 Demand fractions for each restriction level

Restriction level	Ratio of restricted to unrestricted demand
1	0.95
2	0.80
3	0.70
4	0.65

4-3-2 Decision Variables

The 30-year planning horizon, 2010 to 2040, was divided into three equal-length planning stages with change points occurring in 2010, 2020 and 2030. Six decisions associated with operational and capacity expansion options are considered at each change point. These decisions and their lower and upper limits are presented in Table 4-2.

Three decisions involve capacity expansion, namely Cotter Reservoir capacity, Murrumbidgee diversion capacity and the number of installed domestic rainwater tanks. The Murrumbidgee pump storage trigger controls the pumping of water from the Murrumbidgee River to Googong Reservoir after the Murrumbidgee diversion pump station is commissioned; when the storage fraction in Googong Reservoir level falls below the trigger level, pumping from the Murrumbidgee River up to the maximum capacity of the pump station is initiated. The level-one restriction trigger x^2 and increment x^3 are operational decisions that regulate the occurrence of restrictions on consumption during a drought drawdown. If the total storage fraction falls below x^2 then the first restriction level is imposed. If the total storage fraction falls below $x^2 + x^3$, then the second level of restrictions is imposed and so on.

Table 4-2 List of decision variables

Decision	Description	Lower limit	Upper limit	Category
1	Cotter capacity upgrade(ML)	0	100,000	“Zero-one” capacity expansion
2	Level-one restriction storage trigger	0	1	Operational
3	Restriction storage trigger increment	0.05	0.25	Operational
4	Murrumbidgee diversion (ML/month)	0	6,000	“Zero-one” capacity expansion
5	Murrumbidgee pump storage trigger	0	1	Operational
6	Number of houses with tanks	0	15,000	“Developing” capacity expansion

4-3-3 Constraints

A scheduling expansion problem is typically constrained. For example, decisions may belong to the “zero-one” category. If a non-zero value is assigned at a planning stage, then that value remains unchanged for all remaining planning stages. For example, if the capacity of Cotter is increased by 50,000 ML at the start of stage 2, then it will remain unchanged for the remainder of the planning period. Another decision category imposing a constraint is the “developing” category. In this case, the decision value cannot decrease at subsequent planning stages. For example, the number of installed domestic rainwater tanks can be increased but not decreased at each planning stage. The following equation formalizes these constraints:

$$x_{t+1}^i = x_t^i \text{ if } x_t^i > 0 \text{ and } x^i \in \text{"zero-one" decisions}$$
$$x_{t+1}^i \geq x_t^i \text{ if } x^i \in \text{"developing" decisions} \quad (4.2)$$

where x_t^i is the i^{th} decision at planning stage t.

4-3-4 Objective Functions

All reviewed studies dealing with capacity expansion have sought to minimize the present worth of capital, operating and other economic costs. In the context of the formulation described in Section 4-2-1, the total present worth cost can be expressed as

$$f(x) = \frac{1}{N}\sum_{t=1}^{T}\frac{1}{(1+r_o)^t}\sum_{r=1}^{N} C_t^r(x_{1:t}) + CR_t^r(x_{1:t}) + U_t^r(x_{1:t}) \quad (4.3)$$

where r_o is the discount rate and $C_t^r(x_{1:t})$ is the cost of infrastructure investments and operating costs for year t and replicate r, $CR_t^r(x_{1:t})$ is the economic cost of imposing restrictions on demand and $U_t^r(x_{1:t})$ is the cost of unplanned demand shortfalls. However, exclusive reliance on this objective can hide the trade-off between capital and operating costs and the social costs arising from restrictions and unplanned shortfalls.

To explore this trade-off, two multi-objective formulations are considered:

1. <u>Two-objective trade-off</u>

   The total present worth cost can be decomposed into its constituent costs to enable exploration of the trade-off between capital, operating and unplanned shortfall costs and costs due to restrictions. This yields the following two objective functions:

$$\min_x f_1(x) = \frac{1}{N}\sum_{t=1}^{T}\frac{1}{(1+r_o)^t}\sum_{r=1}^{N} C_t^r(x_{1:t}) + U_t^r(x_{1:t}) \quad (4.4)$$

$$\min_x f_2(x) = \frac{1}{N}\sum_{t=1}^{T}\frac{1}{(1+r_o)^t}\sum_{r=1}^{N} CR_t^r(x_{1:t}) \quad (4.5)$$

The second objective minimizes the discounted cost of imposing restrictions. However, minimizing discounted restriction costs can produce undesirable social outcome. Due to discounting, the same frequency and severity of restrictions in the future will be being costed less than if the same were to occur in the present. As a result, minimization of discounted restriction costs can lead to a higher frequency and severity of restrictions in the future, a situation which often would be deemed politically unacceptable on social equity grounds.

2. Three-objective trade-off

One way to overcome this practically significant shortcoming is to avoid discounting restriction costs. However, this in itself will not assure equity (or equal sharing of the burden of restrictions) over planning stages. To achieve this it is necessary to introduce a third objective which seeks to minimize the difference in undiscounted restriction costs over the planning stages. These considerations lead to the following three objective functions:

$$\min_x f_1(x) = \frac{1}{N}\sum_{t=1}^{T}\frac{1}{(1+r_o)^t}\sum_{r=1}^{N} C_t^r(x_{1:t}) + U_t^r(x_{1:t}) \tag{4.6}$$

$$\min_x f_2(x) = \frac{1}{M}\sum_{i=1}^{M}\frac{1}{N}\sum_{r=1}^{N}\sum_{t=T_i}^{t=T_{i+1}} CR_t^r(x_{1:t}) \tag{4.7}$$

$$\min_x f_3(x) = \sqrt{\frac{1}{M}\sum_{i=1}^{M}\left(\frac{1}{N}\sum_{r=1}^{N}\sum_{t=T_i}^{t=T_{i+1}} CR_t^r(x_{1:t}) - f_2(x)\right)^2} \tag{4.8}$$

The first objective seeks to minimize the present worth of capital, operating and unplanned shortfall costs. The second minimizes the expected cost of undiscounted restrictions in a planning stage. The third minimizes the standard deviation of undiscounted restriction costs between planning stages. This effectively seeks to ensure the burden of restrictions on the community is shared as fairly as is possible across all planning stages.

The capital cost of the infrastructure options is summarized in Table 4-3. These costs are indicative and therefore should not be taken literally. Two capital items involve a binary choice: if the item is selected by the optimizer, then there is a fixed setup cost along with a unit cost; if the item is not selected, there is zero capital cost. Operating costs include pumping and treatment costs for transfers from Cotter Reservoir ($250/ML), from the Murrumbidgee River ($285/ML), from Bendora Reservoir ($84/ML) to Stromlo water treatment plant, from the Murrumbidgee River to Googong Reservoir ($23/ML), and from Stromlo water treatment plant to Googong Reservoir ($36/ML). The unplanned shortfall cost was set to $\$1.0\text{x}10^9$/ML to ensure the optimizer steered away from solutions that resulted in "running out of water".

Table 4-3 Infrastructure cost of capacity expansion decisions for Canberra water headworks system

Decision Variable	Unit Cost
Cotter Reservoir capacity upgrade	$\$50\text{x}10^6 + \$1923$/ML
Murrumbidgee diversion	$\$20\ \text{x}10^6 + \$42623$/ML
Water tanks	$3000/ house

As a postscript to this section, a brief comment is made on how the economic cost of restrictions is estimated in this case study. Recognizing that restrictions in Australian urban areas are mainly targeted at outdoor water use (which in the case of Canberra is substantial; see Figure 4-4), the method developed by Dandy (1992) was adopted. Dandy assumed that:

i. All the households have the same price elasticity of demand for outdoor use

ii. The price elasticity for outdoor use is constant within the range considered.

iii. All households reduce their outdoor consumption in the same proportion in response to water restrictions.

Using a willingness-to-pay analysis, he showed that the economic cost of restrictions in a drought event could be approximated by

$$CR = \frac{\varepsilon}{1+\varepsilon} PQ[1-(1-R)^{\frac{1+\varepsilon}{\varepsilon}}] \tag{4.9}$$

where CR is the economic cost due to imposition of restrictions, P is the current price of water, Q is the unrestricted outdoor consumption, R is the fraction by which consumption is reduced and ε is the price elasticity of demand for outdoor water. In this study, following Cui (2003), ε and P were set equal to -0.25 and \$600/kL respectively.

4-4 Case Study Scenarios and Results

The section presents the main findings of the case study. Twelve scenarios are described with the intent of demonstrating in a structured manner the limitations of earlier applications and the performance of the formulation described in Section 4-2-1. The section then reports and discusses the results of each scenario.

4-4-1 Description of Scenarios

Table 4-4 summarizes 12 scenarios used to demonstrate the benefits of applying the multi-objective formulation of Section 4-2-1 to scheduling capacity expansion problems. The scenarios differ in the number of objectives, the demand growth rate, staging of infrastructure and operational decisions, the discount rate and the initial volume of the reservoirs. The first two scenarios are used to demonstrate the need for capacity expansion in the presence of demand growth. The next four scenarios, Scenarios 3 to 6, are used to demonstrate the benefit of scheduling operational decisions in addition to infrastructure decisions. The next three scenarios, Scenarios 7 to 9, investigate the sensitivity of results to choice of discount rate. Unlike the first 8 scenarios which minimize present worth cost, the remaining scenarios, Scenarios 10 to 12, use multiple objectives to demonstrate the advantages arising from of using multi-objective optimization particularly with regard to trading off equity against economic efficiency. Finally, the sensitivity of initial conditions is investigated. In all scenarios except Scenario 12, reservoirs are assumed to be full at the start of the first planning stage. In Scenario 12, the initial storage in the reservoirs is set to the historic 25th percentile volume.

For each scenario, the simulations were conducted using 50 replicates of stochastically generated streamflow. It is acknowledged that more replicates would be needed to ensure a high level of drought security. However, as this issue was already addressed in Chapter 3, a reduced number of replicates was adopted to make the computation manageable for the twelve scenarios. Because εMOEA is a probabilistic method, it is unable to guarantee convergence to the Pareto front. Accordingly, to reduce the chance of premature convergence affecting the results, each scenario was optimized 10 times with different random number seeds. The results presented in the subsequent sections are the best out of 10 runs. As in the Chapter 3 case study, the εMOEA parameters were: probability of crossover = 1, probability of mutation = 0.01 and probability of inversion = 0.005. The maximum number of iterations for the single objective scenarios, 1 to 9, was set equal to 10,000, while for the multiple-objective scenarios, 10 to 12, it was set to 30,000. The εMOEA epsilon was set to 100,000 for the single objective cases and to 10,000 for the first objective and 1000 for the second and third objectives in the multi-objective optimization.

Table 4-4 List of scenarios

Scenario	Number of objectives	Growth rate	Timing of decision		Discount rate	Initial reservoir volume	Purpose
			Infrastructural	Operational			
1	1	0	N/A	Any stage	5%	Full	Impact of demand growth
2	1	1.2%	N/A	Any stage	5%	Full	
3	1	1.2%	Stage 1	Stage 1	5%	Full	Consequence of different timing of infrastructural and operational decisions
4	1	1.2%	Stage 1	Any stage	5%	Full	
5	1	1.2%	Any stage	Stage 1	5%	Full	
6	1	1.2%	Any stage	Any stage	5%	Full	
7	1	1.2%	Any stage	Any stage	1%	Full	Sensitivity to choice of discount rate
8	1	1.2%	Any stage	Any stage	5%	Full	
9	1	1.2%	Any stage	Any stage	10%	Full	
10	2	1.2%	Any stage	Any stage	5%	Full	Use multiple objectives to deal with equity issues
11	3	1.2%	Any stage	Any stage	5%	Full	
12	3	1.2%	Any stage	Any stage	5%	25^{th}-percentile	Sensitivity to initial reservoir volumes

4-4-2 Scenarios 1 and 2: Impact of Demand Growth

Scenarios 1 and 2 are used to justify the need for capacity expansion in the Canberra case study. Scenario 1 has the annual demand growth rate set to zero, while Scenario 2 has the rate set to 1.2%. In both scenarios, no capacity expansion is allowed; only operational decisions namely level one restriction trigger and trigger increment can be changed at each planning stage.

Figure 4-5 shows the demand, unplanned shortfalls and restricted demand time series for the first replicate of Scenario 1 using the optimized decisions. Unplanned shortfalls occur when the demand, permitted by the DCP, cannot be supplied – such shortfalls typically would occur when reservoirs run dry or when limitations in transfer capacity result in demand zones being supplied less than the minimum permitted by the DCP. Restricted demand represents amount of water supplied to the demand node after restrictions imposed. Because there is no demand growth, the system could avoid unplanned shortfalls by imposing frequent restrictions. However, in the presence of demand growth in Scenario 2, Figure 4-6 shows that unplanned shortfalls could not be avoided even though severe and frequent restrictions were imposed. This highlights the need to augment the capacity of the system to cater for the demand growth, as optimizing operational decisions alone cannot prevent the occurrence of unplanned shortfalls. Table 4-5 presents the total present worth cost and associated decisions for the two scenarios. Because unplanned shortfalls attract a punitive cost, the total present worth cost for Scenario 2 is an order of magnitude higher than for Scenario 1. The decisions controlling the imposition of restrictions are at or close to their most severe values for Scenario 2. When the level-one trigger is 1.0 and trigger increment is 0.05, the highest frequency and maximum severity of restrictions are imposed on the system. It is noted that the restriction decisions in Scenario 2 did not assume the most severe values in the third stage. This is because shortfalls were unavoidable even when the restriction decisions were set at their most severe values. Figure 4-7 shows the total storage time series for the first replicate of Scenario 2. There are two periods during which the system was empty and, consequently, unplanned shortfalls occurred. These shortfalls occurred even if the stage-three level-one trigger was 1.00 and the trigger increment was 0.05.

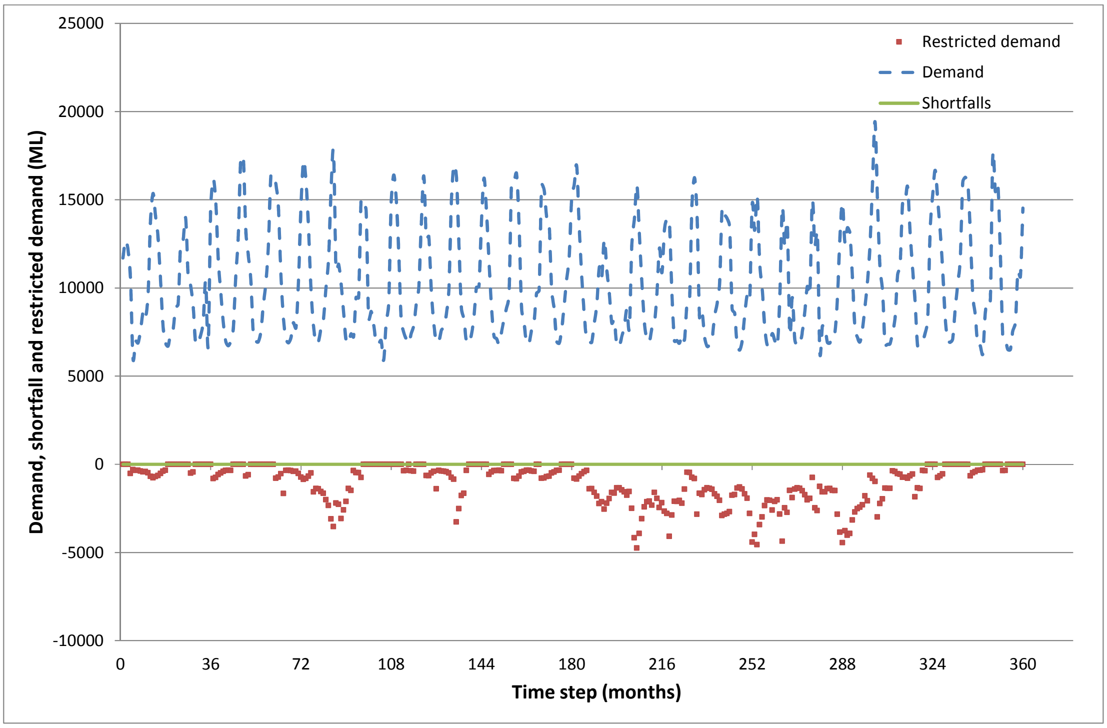

Figure 4-5 Demand, unplanned shortfalls and restricted demand for the first replicate of Scenario 1

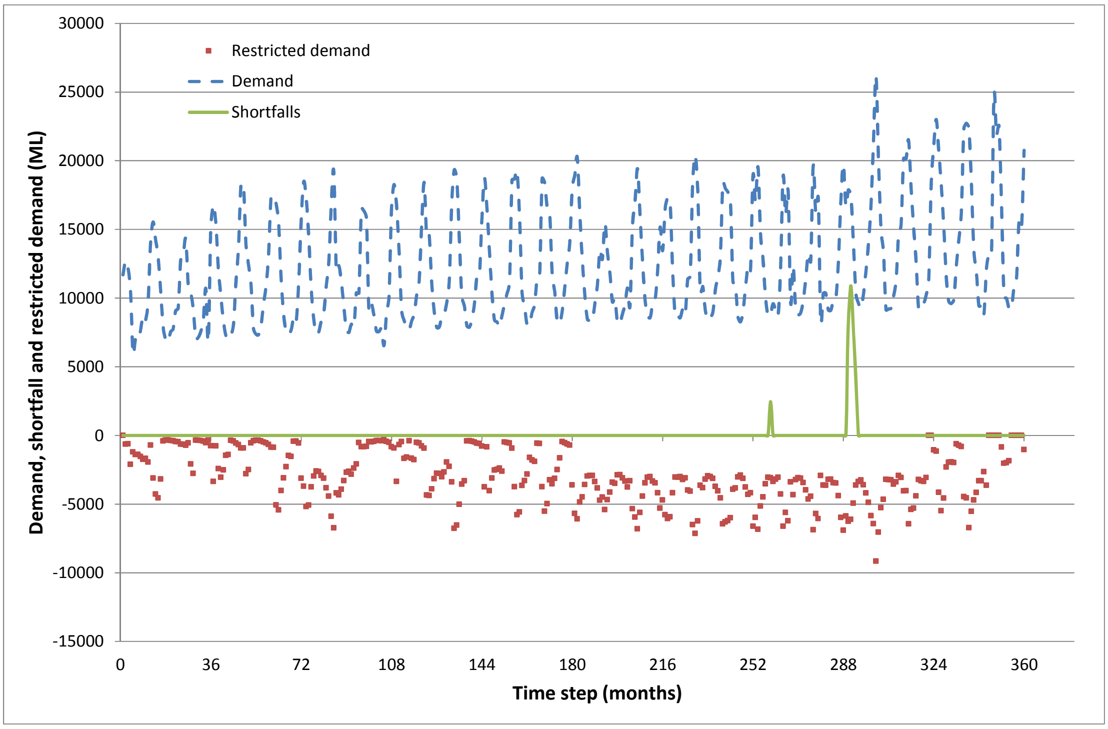

Figure 4-6 Demand, unplanned shortfalls and restricted demand for the first replicate of Scenario 2

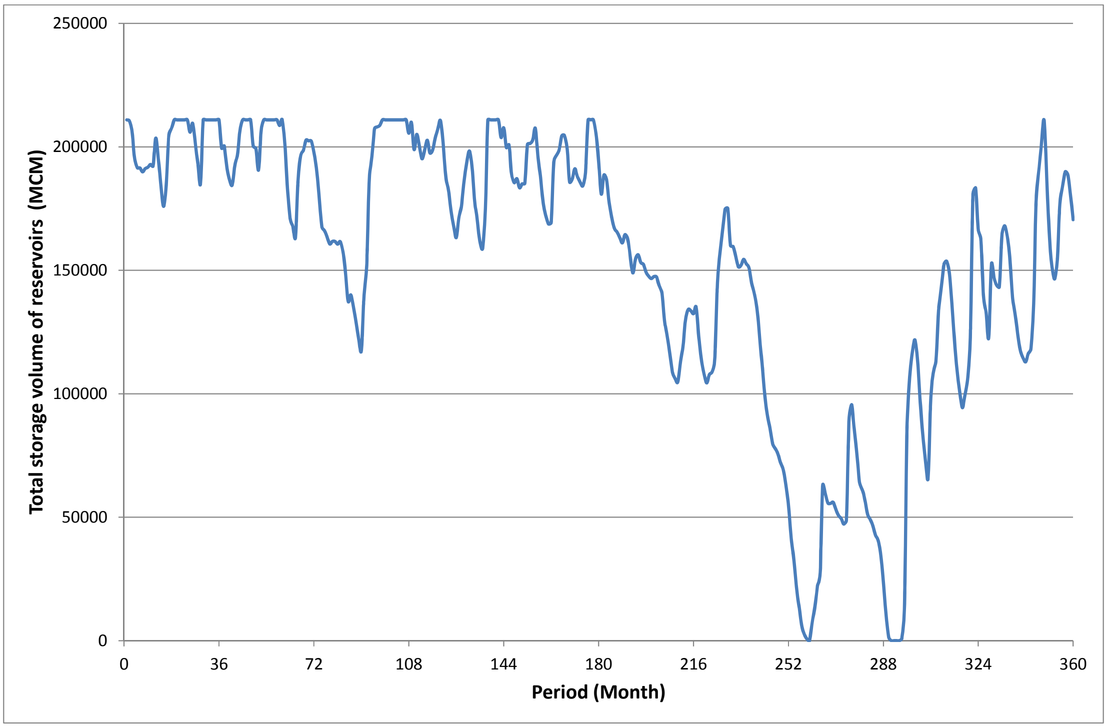

Figure 4-7 Time series of total storage for the first replicate of Scenario 2

Table 4-5 Comparison of present worth cost and decisions for Scenarios 1 and 2

Scenario	**Total present worth cost ($million)**	**Planning stage 1**		**Planning stage 2**		**Planning stage 3**	
		Level-one restriction trigger	Trigger increment	Level-one restriction trigger	Trigger increment	Level-one restriction trigger	Trigger increment
1	362	0.949	0.221	0.933	0.180	0.8	0.248
2	3690	1.000	0.0524	1.000	0.0516	0.827	0.0892

4-4-3 Scenarios 3 to 6: Impact of Different Timing of Infrastructure and Operational Decisions

Four scenarios are used to investigate the relative impacts of scheduling infrastructure (or capital) and operational (or operating rule) decisions over the planning period. These scenarios differ in the timing of their infrastructure and operational decisions. In Scenario 3, all decisions are made at start of the first planning stage. In Scenario 4, all infrastructure decisions are made at start of the first planning stage, while operational decisions are flexible in the sense they can be changed at any of the planning stages. In contrast, in Scenario 5 all operational decisions are made at start of the first planning stage, while infrastructure decisions are flexible. Finally, in Scenario 6, all decisions can be made at any planning stage subject to constraints on the infrastructure decisions.

Table 4-6 and 4-7 present respectively the costs and decisions for the four scenarios. To provide a better understanding of how these scenarios deal with restrictions, the undiscounted restriction cost is presented for each planning stage in Table 4-6. Scenario 5 has effectively the same total PWC as Scenario 3. One would expect Scenario 5 to produce a smaller total PWC than Scenario 3; the fact that Scenario 5 produced a marginally higher cost reflects premature convergence by the optimization algorithm. Bearing this in mind, the near equal costs for Scenarios 3 and 5 suggests that for the Canberra system, scheduling infrastructure decisions while fixing operational decisions offers no significant benefit over making all decisions at the start of the planning horizon. Scenario 5 has the highest restriction cost of the four scenarios indicating a heavy reliance on imposing restrictions. The level-one restriction trigger of 0.815 and the trigger increment of 0.144 confirm the selection of a severe restriction policy. Of interest is the finding that the Cotter upgrade was not selected and the Murrumbidgee diversion was delayed to stage two.

Table 4-6 Results for Scenarios 3 to 6

Scenario	**Total present worth cost ($million)**	**Capital and operational present worth cost ($million)**	**Total present worth cost of restrictions ($million)**	**Undiscounted restriction cost ($million)**			**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of undiscounted restriction costs over three stages ($million)**
				Stage 1	Stage 2	Stage 3		
3	462	393	69	0.060	50.6	99.7	57.3	54.1
4	445	391	54	0	59.4	89.6	49.6	50.5
5	464	362	102	64.94	55.3	94.3	71.5	47.9
6	444	396	48	0.056	34.4	53.2	35.9	36.9

Table 4-7 Optimum decisions for Scenarios 3 to 6

	Scenario 3			**Scenario 4**			**Scenario 5**			**Scenario 6**		
	Planning stage			Planning stage			Planning stage			Planning stage		
Decisions	One	Two	Three	One	Two	Three	One	Two	Three	One	Two	Three
Cotter capacity upgrade(ML)	0	Same as stage one	Same as stage one	0	Same as stage one	Same as stage one	0	0	0	0	0	0
Level-one restriction storage trigger	0.77	Same as stage one	Same as stage one	0.019	0.831	0.627	0.815	Same as stage one	Same as stage one	0.4	0.627	0.752
Restriction storage trigger increment	0.111	Same as stage one	Same as stage one	0.224	0.149	0.063	0.144	Same as stage one	Same as stage one	0.096	0.055	0.149
Murrumbidgee diversion (ML/month)	2460	Same as stage one	Same as stage one	2414	Same as stage one	Same as stage one	0	3091	3091	0	4221	4221
Murrumbidgee pump storage trigger	1	Same as stage one	Same as stage one	1	1	1	0.989	Same as stage one	Same as stage one	-----	1	1
Number of houses with tanks	0	Same as stage one	Same as stage one	0	Same as stage one	Same as stage one	0	0	0	0	0	0

The capital and operational present worth costs for Scenarios 3, 4 and 6 are almost the same but their restriction present worth costs vary significantly. Out of these scenarios, Scenario 3 had the highest restriction cost. This is because all decisions had to be made at the start of the first planning stage. In contrast, Scenarios 4 and 6 exploited the flexibility of adjusting restriction triggers over the planning period. In stage one, a low level-one trigger could be adopted because that stage experienced the lowest demand and benefited the most from the full state of the storages at the start of the planning period – indeed virtually no restrictions were experienced in the first stage. In the subsequent stages, the trigger was increased to cope with the growing demand.

Scenario 6 demonstrates the benefits of having all decisions flexible. It has the lowest average restriction cost among all scenarios. Since this scenario could schedule capacity expansion decisions, it defers the capacity expansion of the Murrumbidgee diversion to the second planning stage. This choice takes advantage of the discounted construction cost and the fact that the system is initially full. As a result, even though Scenario 6 has a substantially larger Murrumbidgee diversion capacity than Scenarios 3 and 4, its discounted capital and operational costs are virtually the same.

Not surprisingly, the optimal strategy is to provide flexibility in timing and sizing for both infrastructure and operational decisions. This is clearly demonstrated in Table 4-6 where Scenario 6 has the lowest total PWC and also the lowest restriction present worth cost of restriction. However, the more significant finding is that virtually all of the benefit from scheduling comes from allowing the operational decisions to change over time. Indeed, having flexible operational rules reduces the incidence of severe restrictions particularly in the last planning stage. It appears this allows the optimizer to better adapt to the fact that in the first stage the initially full system and the lowest demand impose the least stress on the system, while in the third stage, the benefit of the initially full system is no longer "felt" and the demand is at its highest.

4-4-4 Scenarios 7 to 9: Sensitivity to Discount Rate

In all of above mentioned scenarios, the discount rate was set equal to 5%. However, as Luss (1982) has observed, the estimation of discount rate is subjective. The reality is that typically the discount rate used by the private sector is different from that used by the public sector. To investigate the effect of discount rate on the optimum solution, three scenarios with different discount rates, i.e. 1%, 5% and 10% are compared. The total present worth costs are presented in Table 4-8. There are large differences in PWC across the scenarios due to the spread in discount rates. The PWC of Scenario 7 is about three times greater than for Scenario 9. However, the important point here is that the discount rate exerts considerable influence on the severity and frequency of restrictions over the three stages. As shown in Table 4-8, Scenarios 8 and 9 have very similar total discounted restriction costs but their average discounted restriction costs over the three stages are vastly different. It is also evident that the higher the discount rate, the higher the undiscounted restriction costs in later planning stages.

The optimum decisions for the three scenarios are presented in Table 4-9. It is noted that only in Scenario 7 is the Cotter upgrade option invoked with an upgrade capacity of 52,000 ML. This occurs because the use of the low discount rate of 1% would result in a blowout of restriction costs if additional storage were not available to reduce the frequency of restrictions.

The overall conclusion is that the discount rate determines how much reliance the optimizer places on the imposition of restrictions to avoid unplanned shortfalls and on how restrictions are distributed over the planning stages. Comparison of Scenarios 7 to 9 clearly shows that as the discount rate increases, the investment in infrastructure decreases at the expense of more restrictions imposed in future stages.

Table 4-8 Comparison of three scenarios with different discount rates

Scenario	**Discount rate %**	**Total present worth cost ($million)**	**Capital and Operational cost ($million)**	**Total present worth cost of restrictions ($million))**	**Undiscounted restriction cost ($million)**			**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of restriction costs over three stages ($million)**
					Stage 1	Stage 2	Stage 3		
7	1	775	708	67	0	33.3	48.8	27.4	28.6
8	5	444	396	48	0.056	34.4	53.2	35.9	36.9
9	10	267	221	46	45.66	50.3	80.5	58.8	53.4

Table 4-9 Optimum decisions for Scenarios 7 to 9

	Scenario 7 (r=1%)			**Scenario 8 (r=5%)**			**Scenario 9 (r=10%)**		
	Planning stage			**Planning stage**			**Planning stage**		
Decisions	**One**	**Two**	**Three**	**One**	**Two**	**Three**	**One**	**Two**	**Three**
Cotter capacity upgrade(ML)	0	0	14352	0	0	0	0	0	0
Level-one restriction storage trigger	0.004	0.815	0.752	0.4	0.627	0.752	0.8	0.68	0.647
Restriction storage trigger increment	0.18	0.162	0.162	0.096	0.055	0.149	0.211	0.061	0.061
Murrumbidgee diversion (ML/month)	3995	3995	3995	0	4221	4221	0	3091	3091
Murrumbidgee pump storage trigger	1	1	1	----	1	1	----	1	1
Number of houses with tanks	0	0	0	0	0	0	0	0	0

4-4-5 Scenarios 10 and 11: Revealing Equity Tradeoffs

In all the scenarios considered so far, only one objective, namely minimization of the total present worth cost, was considered. This cost includes capital, operating and restriction costs – unplanned shortfall costs were always zero because of their punitive unit value. However, there is a trade-off between capital, operating and unplanned shortfall costs and restriction costs. Indeed, more investment in infrastructure results in less need to impose restrictions and vice versa. To demonstrate this trade-off, Scenario 10 considers a multi-objective optimization jointly minimizing capital, operating and unplanned shortfall costs and minimizing restriction costs. The two objectives are described by Eqs. (4.4) and (4.5). In Figure 4-8 the Pareto frontier for Scenario 10 is presented. The results for Scenario 8, which is a special case of Scenario 10, are also shown in this figure. As expected, the Scenario 8 result is located on the Pareto frontier, which confirms that Scenario 8 represents only one of the possible solutions for Scenario 10.

Figure 4-8 shows there is a distinct trade-off between capital, operating and unplanned shortfall costs and the cost of imposing restrictions. Indeed, the restriction cost can be very large in the absence of sufficient infrastructure investment. The figure shows there is initially a very favourable trade-off between higher capital investment and reduced restriction cost (see labeled points 1 and 2) followed by a progressively worsening trade-off culminating with virtually zero restriction costs when the present worth of capital, operating and unplanned shortfall costs exceeds $750 million. Up to $750 million, there are no unplanned shortfall costs. However, beyond that, unplanned shortfall costs grow rapidly to produce minute reductions in restriction costs. This segment of the Pareto frontier would be of no interest to a decision maker. It is presented here to document the full Pareto frontier.

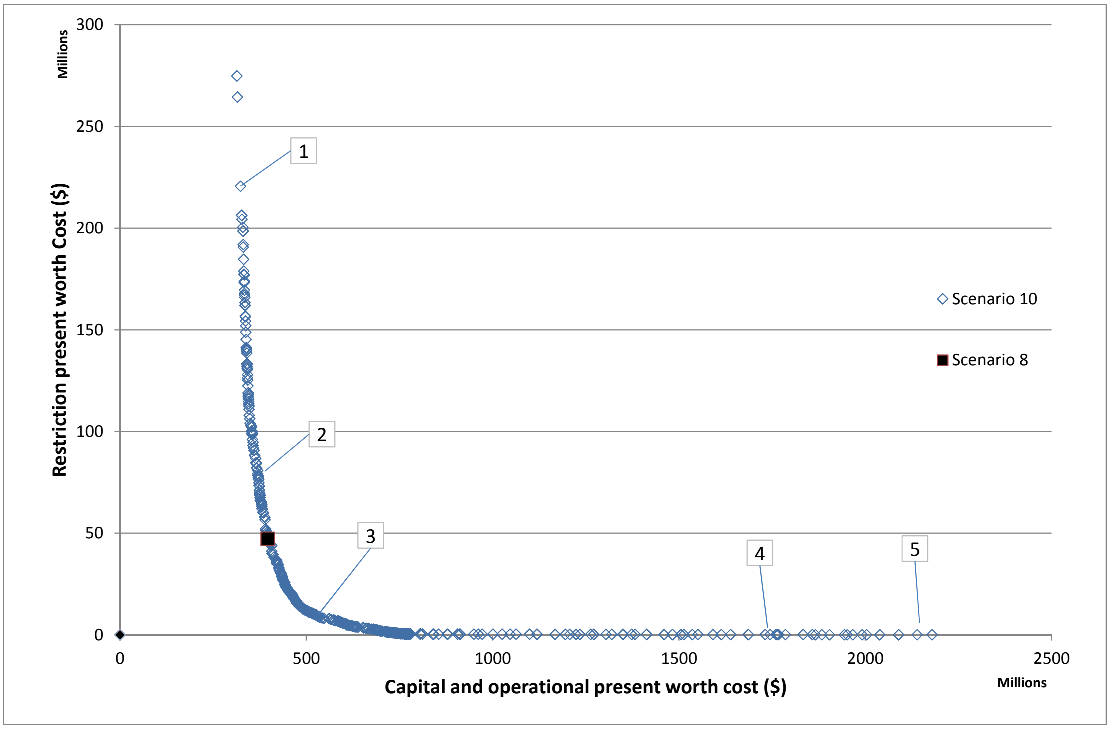

Figure 4-8 Pareto frontier for Scenario 10

Discounting can hide the significance of the impact of restrictions on the community. To highlight this, five solutions on the Pareto frontier in Figure 4–8 were selected with Table 4–10 providing a summary of these solutions. The results show that for all the solutions, progressively more severe restrictions are imposed in future planning stages highlighting the implicit inequity associated with discounting.

To deal explicitly with this equity issue and to offer the opportunity to moderate differences across planning stages, the three-objective formulation described by Eqs. (4.6) to (4.8) is considered in Scenario 11. The first objective minimizes total present worth of capital, operating and unplanned shortfall costs, while the remaining two objectives introduce equity considerations. The second objective seeks to minimize the magnitude of restriction costs across the stages while the third objective seeks to minimize the difference in restriction costs between stages.

Figure 4-9 presents the Pareto frontier for Scenario 11. What is striking is the absence of a surface. The trade-offs essentially lie on a one-dimensional thread. Once significant restriction costs are encountered, there is a strong almost linear dependence between objectives two and three, namely the average cost of undiscounted restrictions and the variability of cost across stages. To offer more insight into this trade-off, Figure 4-10 to Figure 4-12 present projections of the threedimensional Pareto front onto three two-dimensional objective planes. Figure 4-10 shows that the average of undiscounted restriction costs decreases substantially as capital, operating and unplanned shortfall cost increases. A similar trend can be seen in Figure 4-11 for the standard deviation of undiscounted restriction costs. Figure 4-12 presents the trade-off between the average and standard deviation of undiscounted restriction costs. It shows that, unless there is sufficient investment to eliminate restrictions, it is not possible to share equally the burden of restrictions across stages; moreover, as the average level of restriction costs in a stage grows there will be greater variability across the stages.

To highlight the difference between Scenarios 10 and 11, five solutions were selected for each scenario in order to have equal capital and operating costs for each pair of solutions. Table 4-11 presents the results. What is striking is the fact that Scenario 11 produces solutions with lower average (undiscounted) restriction costs across the planning stages and less variability in restriction cost between stages. This significantly improved equity outcome arises solely from the choice of objective functions. The use of three objectives enabled a more thorough exploration of cost and equity with the consequent identification of solutions with more equitable outcomes for the same capital and operating present worth cost.

Table 4-10 Comparison of five solutions marked on Figure 4-8 of Pareto frontier for Scenario 10

Solution label	**Capital and operational present worth cost ($million)**	**Unplanned shortfall cost ($million)**	**Restrictions present worth cost ($million)**	**Undiscounted Restriction cost ($million)**			**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of undiscounted restriction costs over three stages ($million)**
				Stage 1	Stage 2	Stage 3		
1	323	0	221	77.6	195.8	240.6	171	94.7
2	373	0	79.9	4.0	80.7	132.9	72.5	67.2
3	525	0	9.86	0	7.9	21.1	9.68	12.7
4	779.47	910.54	0.8067	0	0	0.269	0.0897	0.1268
5	779.39	1360.61	0.0085	0	0	0.028	0.0096	0.0136

Table 4-11 Comparison of five marked solutions on the Pareto frontiers for Scenarios 10 and 11

	Scenario 10				Scenario 11			
Solution label	**Capital and operational present worth cost ($million)**	**Unplanned shortfall cost ($million)**	**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of undiscounted restriction costs over three stages ($million)**	**Capital and operational present worth cost ($million)**	**Unplanned shortfall cost ($million)**	**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of undiscounted restriction costs over three stages ($million)**
1	323	0	171	94.7	323	0	116	61
2	373	0	72.5	67.2	373	0	51.4	40.7
3	525	0	9.68	12.7	525	0	9.47	12.2
4	779.47	910.54	0.0897	0.1268	780.94	907.67	0.0897	0.1268
5	779.39	1360.61	0.0096	0.0136	779.49	1362.12	0.0096	0.0136

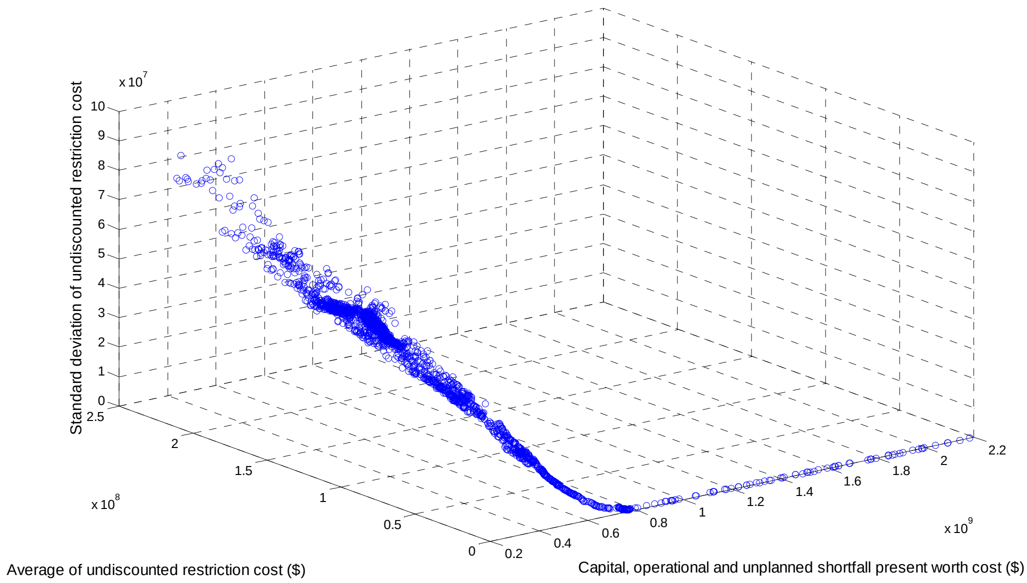

Figure 4-9 Pareto frontier for Scenario 11

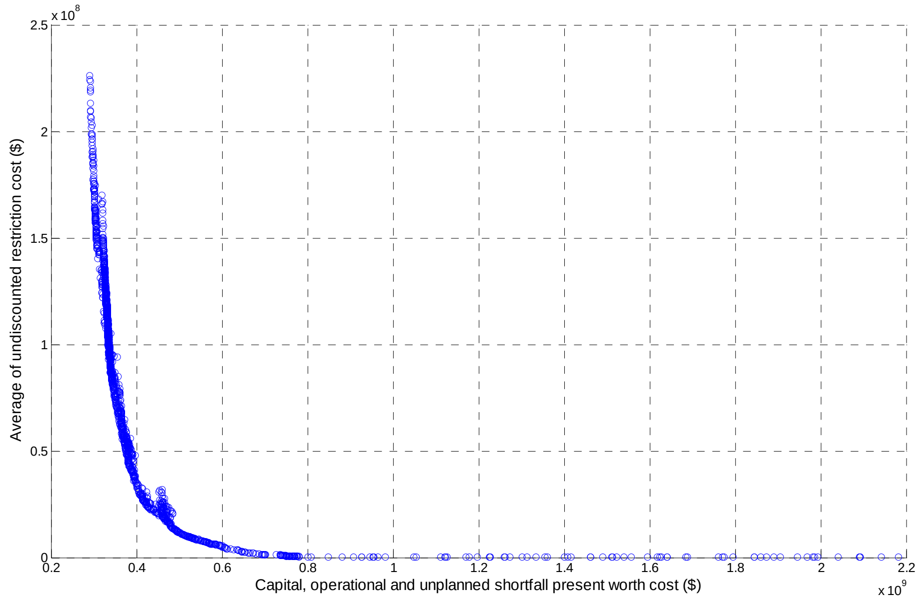

Figure 4-10 Pareto trade-off between present worth of capital, operational and unplanned shortfall costs and average of undiscounted restriction costs over three planning stages for Scenario 11

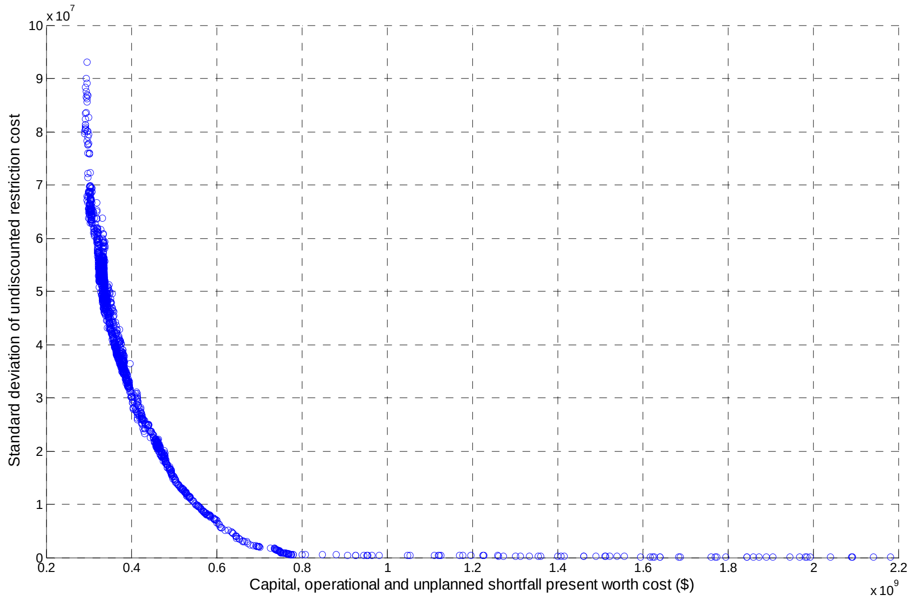

Figure 4-11 Pareto trade-off between present worth capital, operational and unplanned shortfall cost and standard deviation of undiscounted restriction costs over three planning stages for Scenario 11

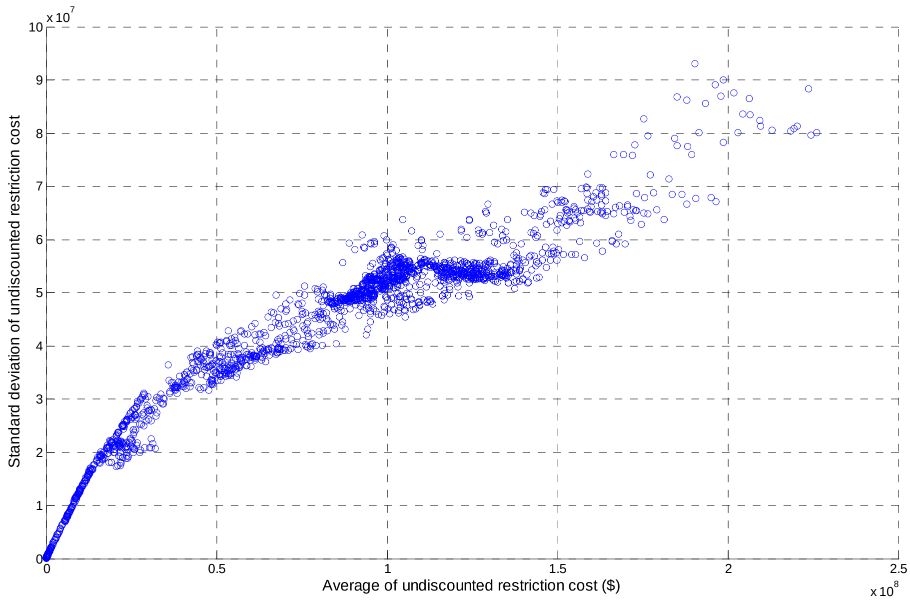

Figure 4-12 Pareto trade-off between the average and standard deviation of undiscounted restriction costs over three planning stages for Scenario 11

Table 4-12 presents the decisions associated with the five Scenario 11 solutions in Table 4-11. The table ranks the solutions from smallest to highest capital and operating cost. The first solution has no capacity expansion except for the Murrumbidgee diversion in the second planning stage. The first and second stage level-one restriction triggers are very high indicating a high frequency of restrictions. In solutions 2 and 3, the size of the Murrumbidgee diversion increases. For solution 3, the Murrumbidgee diversion is brought forward to stage one and a rollout of rainwater tanks over the three stages is adopted with the number of tanks hitting the upper bound in stage two. Offsetting this increased capital investment are lower level-one restriction triggers leading to a lower frequency of restrictions. Solutions 4 and 5 are the most costly with the Cotter upgrade and Murrumbidgee diversion maximized in stage one and rainwater tank installations maximized in stage two. The level-one restriction triggers are low resulting in virtual elimination of restrictions. It is noted that a huge increase in unplanned shortfall cost is required to bring about a minute reduction in restriction costs. As already noted, this is due to the punitive cost assigned to unplanned shortfalls. Of interest, all solutions opted for the Murrumbidgee diversion and set the pump trigger to one. This maximizes the yield from what is the most cost effective capital option.

Table 4-12 Decisions associated with the five solutions presented in Table 4-11 for Scenario 11

Solution	Capital, operational and unplanned shortfall present worth cost ($m)	Planning stage 1						Planning stage 2						Planning stage 3					
		Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks
1	323	0	0.996	0.140	0	----	0	0	0.929	0.143	1962	1	0	0	0.760	0.129	1962	1	0
2	373	0	0.827	0.239	0	----	0	0	0.752	0.0837	3402	1	0	0	0.564	0.0618	3402	1	0
3	525	0	0	0.195	5774	1	1256	0	0.564	0.097	5774	1	5932	0	0.568	0.119	5774	1	14686
4	1690	100000	0.215	0.179	6000	1	1699	100000	0.039	0.182	6000	1	14941	100000	0.125	0.074	6000	1	15000
5	2140	100000	0.219	0.235	6000	1	175	100000	0	0.186	6000	1	14882	100000	0.019	0.122	6000	1	15000

4-4-6 Scenario 12: Sensitivity to Initial Conditions

In the previous scenarios, it was observed that in the first planning stage restrictions were typically low. This is attributed to the fact that the system was full at the start of stage one and that stage one had the lowest demand. To separate the contributions of these two factors, this section investigates the sensitivity of the Pareto-optimal solutions to the initial reservoir storage. In Scenario 11, the reservoirs are full at the start of the planning period, while in Scenario 12 the initial volume of all reservoirs was set equal to the 25th percentile storage volumes obtained from a 130-year simulation using historical flows and demand corresponding to the start of the planning period. Figure 4-13 shows the Pareto frontiers for Scenarios 11 and 12 together with five selected solutions on each front. These solutions were selected to produce five pairs where each member of a pair was located on a different Pareto front but had a near equal average undiscounted restriction cost. Tables 4-13 and 4-14 present the three objective function values for each solution as well as the undiscounted restriction costs for each stage for Scenarios 11 and 12 respectively.

There is a striking shift in the Pareto frontier with all five Scenario 12 solutions experiencing unplanned shortfalls during the first stage of the planning horizon. In some of the replicates, there was a significant drought during the first planning stage. The optimizer was unable to find a solution that could compensate for the low starting storage in those replicates. There was no solution that could avoid "running out of water".

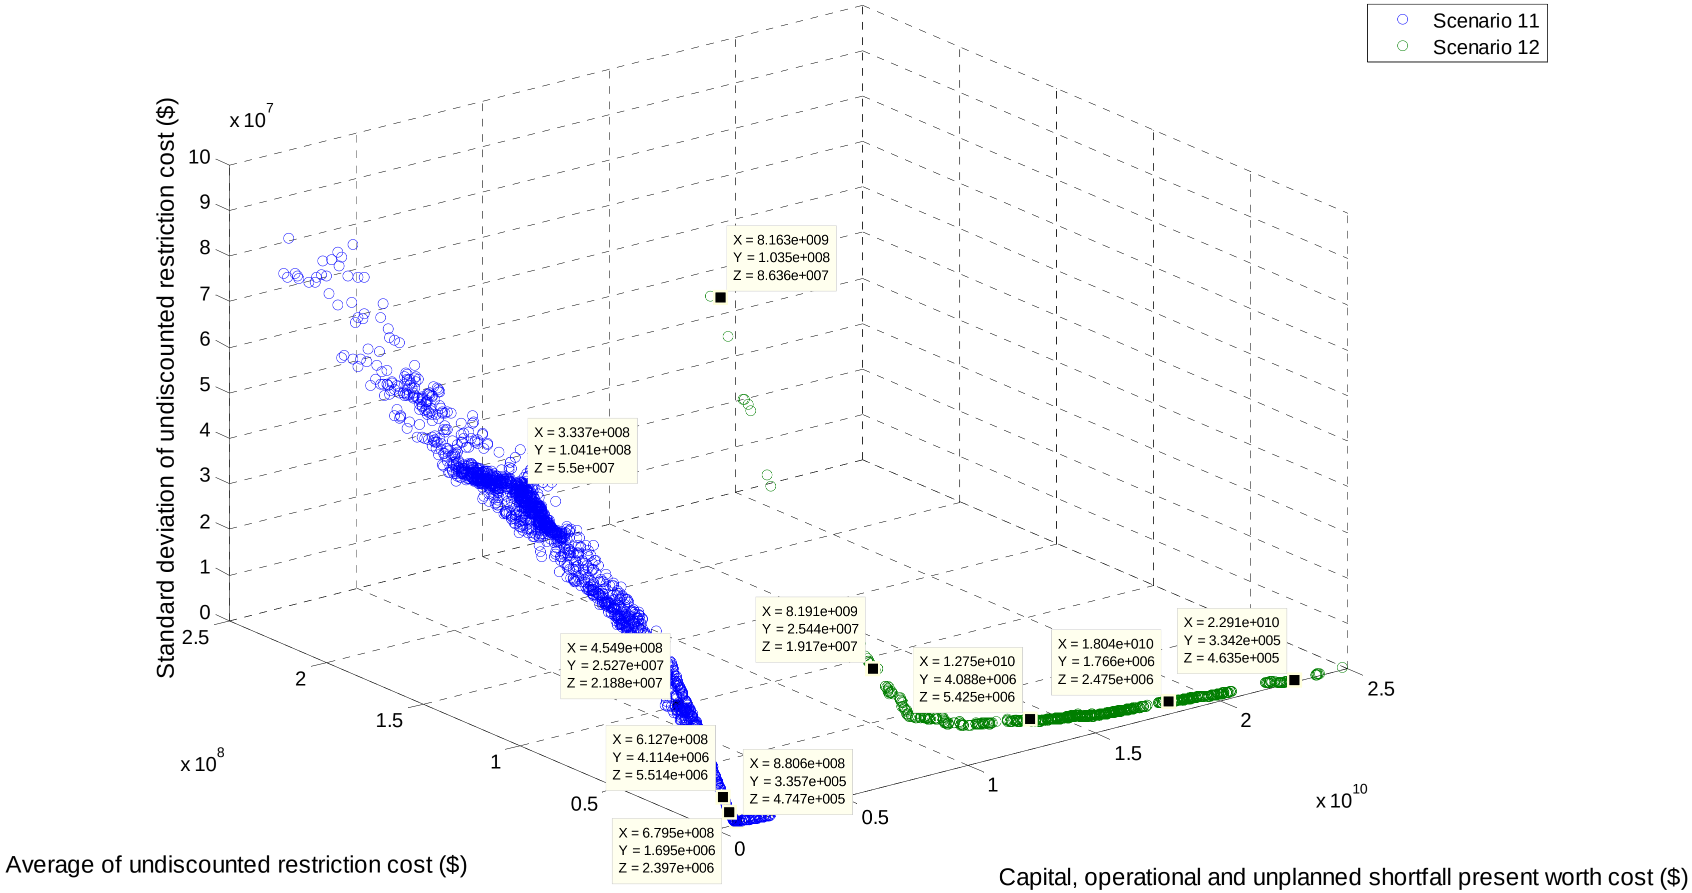

Figure 4-13 Comparison of Pareto frontiers for Scenarios 11 and 12 showing location of 5 selected solutions on each front

More insight about the impact of initial conditions can be obtained by examining Tables 4-15 and 4-16, which present the decisions associated with each marked solution in Figure 4–13 for Scenarios 11 and 12 respectively. In stage one of Scenario 11, there is limited uptake of rainwater tanks and no expansion of Cotter capacity (except for Solution 5). In contrast, for stage one of Scenario 12, all solutions opt for maximum rainwater tank uptake and maximum Murrumbidgee diversion capacity. This is because these options can immediately provide additional yield to the system. Upgrading Cotter in stage one is not as effective because there is no storage in the upgraded Cotter at the start of stage one. Consequently, if in a replicate a drought occurs at the start of stage one, the upgraded Cotter remains effectively empty unable to moderate the impact of the drought. In contrast, rainwater tanks will harvest any available roof runoff, while the Murrumbidgee diversion will be able to divert any river flow to Googong. In Scenario 12, solutions 3, 4 and 5 upgrade Cotter to maximum capacity in stage one in order to reduce the burden of restrictions.

In this case study, the optimal scheduling policy is profoundly affected by the initial state of the storages. The low initial storage in Scenario 12 makes the system much more vulnerable to drought in the first planning stage. The stage-one decisions reflect this vulnerability. They bring forward to stage one the maximum capital investments that were deferred to latter stages in Scenario 11. Despite this, it was not possible to avoid “running out of water” in some of the replicates. The short-term vulnerability in Scenario 12 could not be adequately managed given the constraints on the capital investment mix.

Table 4-13 Comparison of five points marked on Figure 4-13 of Pareto frontier for Scenario 11

Solution	**Capital and operational present worth cost ($million)**	**Unplanned shortfall cost ($million)**	**Undiscounted restriction cost ($million)**			**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of undiscounted restriction costs over three stages ($million)**
			Stage 1	Stage 2	Stage 3		
1	333.46	0	130.85	96.56	85.01	104.14	54.98
2	454.93	0	10.33	31.83	33.65	25.26	21.87
3	612.70	0	0.03	5.41	6.90	4.11	5.51
4	679.29	0	0.00	3.75	1.34	1.69	2.39
5	779.67	100.91	0.00	0.00	1.01	0.335	0.474

Table 4-14 Comparison of five points marked on Figure 4-13 of Pareto frontier for Scenario 12

Solution	**Capital and operational present worth cost ($million)**	**Unplanned shortfall cost ($million)**	**Undiscounted restriction cost ($million)**			**Average of undiscounted restriction cost over three stages ($million)**	**Standard deviation of undiscounted restriction costs over three stages ($million)**
			Stage 1	Stage 2	Stage 3		
1	647.85	7514.72	39.98	53.12	217.48	103.52	86.36
2	673.77	7517.62	27.32	29.76	19.24	25.43	19.16
3	801.62	11947.28	10.88	0.08	1.31	4.09	5.42
4	802.39	17234.71	5.27	0.00	0.03	1.76	2.48
5	802.81	22108.49	0.99	0.00	0.01	0.334	0.463

Table 4-15 Optimum decisions for five solutions presented in Table 4-13 for Scenario 11

Solution	Planning stage 1						Planning stage 2						Planning stage 3					
	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks
1	0	1.000	0.169	0	0.216	0	0	0.937	0.150	2414	1	0	0	0.761	0.145	2414	1.000	0
2	0	0.820	0.216	4391	1.000	0	0	0.882	0.234	4391	1	0	0	0.627	0.134	4391	1.000	0
3	0	0.263	0.197	6000	1.000	486	52000	0.522	0.123	6000	1	885	52000	0.502	0.148	6000	1.000	4804
4	0	0.004	0.223	6000	1.000	372	100000	0.337	0.051	6000	1	896	100000	0.298	0.099	6000	1.000	13431
5	100000	0.176	0.183	6000	1.000	409	100000	0.004	0.249	6000	1	413	100000	0.341	0.173	6000	1.000	15000

Table 4-16 Optimum decisions for five solutions presented in Table 4-14 for Scenario 12

Solution	Planning stage 1						Planning stage 2						Planning stage 3					
	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks	Cotter Capacity expansion	First restriction Trigger	Trigger intervals	Murrumbidgee diversion capacity	Murrumbidgee pump trigger	Number of houses with tanks
1	0	0.471	0.096	6000	1.000	15000	52000	0.741	0.150	6000	0.498	15000	52000	0.980	0.059	6000	0.996	15000
2	0	0.318	0.095	6000	1.000	14941	47765	0.643	0.149	6000	0.498	15000	47765	0.576	0.097	6000	1.000	15000
3	100000	0.165	0.155	6000	0.996	14941	100000	0.259	0.193	6000	1	15000	100000	0.247	0.062	6000	1.000	15000
4	100000	0.157	0.155	6000	0.988	14941	100000	0.016	0.190	6000	1	15000	100000	0.020	0.134	6000	1.000	15000
5	100000	0.000	0.150	6000	0.996	14941	100000	0.008	0.197	6000	1	15000	100000	0.004	0.137	6000	1.000	15000

4-5 Summary

Various options are available to water agencies responsible for meeting the growing demand for water arising from urban population growth. These options include operational decisions such as imposing restrictions, rules controlling water transfers and allocations, policies promoting more efficient water use and infrastructure investments such as harvesting new sources of water. Because the performance of the urban water resource system will change over time, the challenge is to find the best combination of these options over both time and scale (or magnitude).

Many studies have investigated methods to find the optimum size and timing of capacity expansion of projects with the aim of minimizing the total present worth cost. However, review of these studies has identified a number of shortcomings. These include the following:

1. Minimizing a single objective based on present worth cost hides a socially-sensitive equity issue related to the sharing of the burden of restrictions across planning stages.

2. Failure to optimize jointly infrastructure and operational decisions.

3. Failure to address drought security adequately due to inadequate sampling of severe droughts.

This chapter presented a multi-objective formulation that addresses these shortcomings in a practicable manner. The formulation uses a multi-replicate approach in which multiple realizations of future inputs are simulated. It permits use of a full simulation model that enables the tracking of system performance over time and enables the optimization algorithm to search for the best mix of both infrastructure and operational decisions.

A case study based on the Canberra headworks system demonstrated the ability of this formulation to address in a practical manner the shortcomings identified in earlier studies. The following conclusions based on the case study are considered to have applicability beyond the case study itself:

1. The joint scheduling of operational and infrastructure decisions can produce significantly better outcomes than just scheduling infrastructure decisions. Indeed, in the Canberra case study, virtually all of the benefit of scheduling over time was attributed to scheduling operational decisions associated with the imposition of restrictions. This arose because the first stage of the planning period benefited from the storages being full at the start of the stage.

2. Minimizing total present worth cost can lead to more severe and frequent restrictions in future planning stages. This is potentially an unacceptable social outcome. The magnitude of this inequity is dependent on the discount rate with higher discount rates leading to greater inequity in restriction outcomes.

3. The use of a multi-objective formulation, which minimizes the present cost of capital, operating and unplanned shortfall costs together with the level and variability of restriction costs across planning stages, makes the equity issue visible to a decision maker.

4. The optimal scheduling solution can be sensitive to the initial state of the system. In the Canberra case study, a low initial storage elevated the short-term vulnerability of the system to drought. This is by no means an undesirable finding. Indeed, by being able to schedule both infrastructure and operational decisions across multiple planning stages, it is possible to adapt to changing circumstances. This capability is arguably the most important feature of the formulation developed in this chapter.

5. The current generation of multi-objective evolutionary algorithms makes the multi-objective scheduling capacity expansion formulation developed in this chapter practicable for urban systems with complexity similar to the Canberra case study. The case study was conducted on a four-core desktop computer with typical run times of 16 hours. With access to large computer clusters, more complex systems can be studied with many more replicates than the 50 considered in this chapter.

Chapter 5

Efficient Multi-Objective Optimization Methods for Computationally-Intensive Urban Water Resources Models

5-1 Introduction

The preceding chapters have focused on the formulation of multi-objective optimization (MOO) problems in the context of urban water management. In those chapters the εMOEA method was employed to search for Pareto-optimal solutions. This chapter considers the question whether there are MOO methods superior to εMOEA for urban water resource applications. A number of heuristic algorithms, including evolutionary, particle swarm and ant colony optimization methods, have been developed for solving multi-objective problems (Czyz et al., 1998; Deb, 2001; Deb et al., 2002a; Coello Coello, 2006; Huang et al., 2006; Martı´nez et al., 2007). The performance of these algorithms has been investigated mostly using well-known benchmark problems (Zitzler et al., 2000; Deb et al., 2002b) with their results being compared using a range of indicators that measure the convergence and diversity of the solutions after a relatively large pre-defined number of function evaluations.

Unlike the benchmark problems, water resource applications typically use computationally expensive methods for computing their objective functions (Pierro et al., 2009). For example, in the case study presented in the previous chapter involving the Canberra headworks system, a 30-year simulation with 50 replicates at monthly time steps takes approximately 6 CPU seconds, which is several orders of magnitude longer than the standard benchmark problems. Hence for an optimization involving 10,000 function evaluations, the turnaround time of nearly 17 hours is totally dominated by the simulation model rather than by the optimization algorithm. Our experience with urban water supply headworks models using long stochastically generated streamflow at monthly time steps is that simulation run times of the order of several minutes are typical. For instance, in the case study presented in Chapter 3 involving the Sydney headworks system, a 10,000-year simulation at monthly time steps takes about 40 seconds. These long simulation run times are considered an impediment to the practical uptake of MOO. While parallel computing can reduce turnaround times (Cui and Kuczera, 2005), there is also a strong imperative to identify or develop MOO methods which not only converge to the Pareto-optimal front with good diversity but do so with the fewest possible function evaluations. This is the subject of this chapter.

In recent years a considerable number of studies have sought to address this issue; for example Eskandari et al. (2007), Pierro et al. (2009), Santana-Quintero et al. (2006) and Toscano-Pulido et al. (2007). The focus of these studies was identifying which algorithm can produce a better Pareto front after a relatively small number of evaluations. Durillo et al. (2010) took a different perspective comparing the number of evaluations to reach a certain convergence criterion threshold for seven multi-objective methods, namely NSGA-II, SPEA2, PAES, SMPSO, GDE3, AbYSS and MOCell. They used three convergence criteria, the number of Pareto-optimal solutions and the convergence and hypervolume metrics and concluded that SMPSO was the best of the seven algorithms.

In the context of urban water resource optimization, the question as to which MOO method is the best choice for a given function evaluations budget remains unexplored. There is no study comparing the efficacy of MOO methods constrained by a limited number of function evaluations nor is there any study evaluating the number of evaluations to reach convergence thresholds. The primary objective of this chapter is to address this question.

A recent development in probabilistic optimization, called ant colony optimization (ACO), was proposed by Dorigo et al. (1996). ACO emulates the foraging behaviour exhibited by ant colonies in their search for food. ACO algorithms have been successfully applied to a number of benchmark combinatorial optimization problems, such as the travelling salesman and quadratic assignment problems (Dorigo and Stützle, 2004; Stützle et al., 2010). The good performance of ACO in single objective optimization motivated researchers to apply ACO to multi-objective problems (Iredi, 2001; Shelokar et al., 2002; García-Martínez, 2004; Alaya et al., 2007; Angus, 2007b; Bui et al., 2008; Angus and Woodward, 2009). However, these studies have focussed on combinatorial problems, while many engineering problems include decision variables that have a continuous, real-valued domain. A limited number of studies have applied multi-objective ACO methods to problems with continuous real-valued search spaces (Shelokar et al., 2002; Angus, 2007a; Afshar et al., 2009).

Computationally expensive problems such as encountered in urban water resources provide a strong motivation to develop new optimization methods that require fewer evaluations to converge. The secondary objective of this chapter is to explore whether the multi-objective ant colony optimization approach can be successfully adapted to solve computationally-intensive problems typical of urban water resources.

This chapter is organized as follows: First, a review of existing MOO methods is presented from which three benchmark methods are selected. This is followed by a discussion of the performance metrics to be used in the case studies and a brief description of the two urban water resource case studies and the benchmark problems. Then the principles of ant colony optimization are described after which a new MOACO algorithm, MOACO-state, incorporating the best features of the existing MOACO algorithms, is proposed and further enhanced. Finally, using two urban water resource case studies, the performance of the new MOACO methods is compared against three benchmark methods.

5-2 Review of Existing MOO Methods

The last decade has seen considerable effort towards developing efficient MOO methods for computationally-intensive problems. Eskandari et al. (2007) proposed a new algorithm called fast Pareto genetic algorithm (FastPGA) which uses a new fitness assignment and ranking strategy. They compared their method against NSGA-II using four benchmark problems known as the Ziztler-Deb-Thiele (ZDT) test suite and found that FastPGA outperformed NSGA-II in terms of convergence and diversity after completion of a relatively small number of evaluations (6500 and 10000). Pierro et al. (2009) applied two hybrid algorithms, ParEGO and LEMMO, to optimize cost and pressure deficit in water distribution network systems and compared their performance against an evolutionary algorithm called PESA-II. They found for a medium sized network that LEMMO generated solutions after 10,000 evaluations that were comparable with those produced by PESA-II results after 100,000 evaluations. However, for a large network involving 600 decisions LEMMO did not perform well. Santana-Quintero et al. (2006) developed a new particle swarm optimization in conjunction of a local search method. They tested their method for two sets of benchmark problems, ZDT and DTLZ. After 4,000 evaluations, the method was shown to outperform the well-established benchmark NSGA-II algorithm for the ZDT problems but it did not perform well for DTLZ problems. In a similar study, Toscano-Pulido et al. (2007) presented an efficient multi-objective particle swarm optimization method EMOPOS, which, after 2000 evaluations, produced a solution closer to the optimal Pareto front than did NSGA-II for the same number of evaluations.

The focus of above-mentioned studies was evaluating the performance of algorithms after a pre-defined number of evaluations. This type of analysis can be approached from a different perspective. For instance, Nebro et al. (2008) ranked six multi-objective optimization methods, namely NSGA-II, SPEA2, PAES, OMOPSO, AbYSS and MOCell, according to the number of evaluations required to produce Pareto front with a certain accuracy. They found MOCell, OMOPSO, and AbYSS the most competitive algorithms. In a similar study, Durillo et al. (2010) analysed the performance of similar multi-objective methods, except SMPSO replaced OMOPSO and GDE3 was added, using three criteria, the number of Pareto-optimal solutions, the convergence metric and the hypervolume metric. They concluded that SMPSO performed the best of the seven algorithms.

In this study three methods, namely NSGA-II, εMOEA and SMPSO, were selected for comparison based on their usage and performance reported in the literature and on the availability of computer codes. NSGA-II has been widely applied in the MOO literature, often being used as a benchmark for new developed methods in computationally-intensive problems (Santana-Quintero et al., 2006; Eskandari et al., 2007; Toscano-Pulido et al., 2007; Nebro et al., 2008; Durillo et al., 2010). In Chapter 2 it was argued that εMOEA may perform better than NSGA-II and thus was selected for use in the case studies reported in Chapters 3 and 4. Accordingly, εMOEA was selected to test this; moreover, there is no study evaluating the performance of εMOEA when the number of evaluations is constrained. Finally, SMPSO was chosen because of its superior performance among the seven state-of-the-art MOO methods evaluated by Durillo et al. (2010).

In following sections SMPSO and NSGA-II algorithms are described briefly. The εMOEA algorithm was described in Chapter 2.

5-2-1 SMPSO

Particle swarm optimization (PSO) is a population-based metaheuristic method mimicking the social behavior of bird flocking. The initial ideas on particle swarms were proposed by Kennedy and Eberhart (1995). Since then many studies have been carried out to develop and improve PSO; see Poli et al. (2007). PSO has been shown to produce good results at a very low computational cost (Kennedy et al., 2001; Engelbrecht, 2006; Reyes-Sierra and Coello, 2006).

The good performance of PSO in single objective optimization applications motivated researchers to extend it to multi-objective problems. Moore and Chapman developed the first multi-objective implementation of PSO in 1999 and since then more than twenty different methods have been reported (Reyes-Sierra and Coello, 2006).

Each particle in PSO is composed of three vectors, its current position $\vec{x}_i$, the best solution that particle *i* has viewed $\vec{p}_i$, and its current velocity $\vec{v}_i$. The position $\vec{x}_i$ represents a set of coordinates in the search space. Each particle is influenced by the best point found by any member of its topological neighbourhood. This best particle is denoted as a leader ($\vec{p}_g$). In multi-objective optimization problems all non-dominated solutions are considered to be leaders. During the optimization search the velocity of each particle is iteratively adjusted so that the particle stochastically fluctuates around $\vec{p}_i$ and $\vec{p}_g$. In speed-constrained multi-objective PSO (SMPSO), the leader particle $\vec{p}_g$ is selected by sampling two solutions from the external archive and selecting the one which has the largest crowding distance to its neighbor in the archive (Nebro et al., 2009).

Figure 5-1 presents pseudo code for SMPSO. The first step initializes the swarm by assigning a population of particles with random positions and velocities. The second step initializes the leaders in the external archive with non-dominated solutions in swarm. Thereafter the main loop of algorithm is executed until termination criteria met. Termination may be defined by a maximum number of evaluations, the attainment of a prescribed accuracy or the maximum number of evaluations during which no improvement occurs.

The first step in the main loop is to calculate velocity of each particle as follows:

$$\vec{v}_i(t) = X \times (w.\vec{v}_i(t-1) + C_1.r_1.(\vec{p}_i - \vec{x}_i) + C_2.r_2.(\vec{p}_g - \vec{x}_i)) \tag{5.1}$$

where w is the inertia weight of the particle which controls the trade-off between global and local experiences, r_1 and r_2 represent random numbers uniformly distributed in [0,1], C_1 and C_2 are specific parameters which control the effect of the local and global best particles and X is defined by

$$X = \frac{2}{\phi - 2 + \sqrt{\phi^2 - 4\phi}} \tag{5.2}$$

where

$$\phi = \begin{cases} C_1 + C_2 & if \quad C_1 + C_2 > 4 \\ 0 & if \quad C_1 + C_2 \leq 4 \end{cases} \tag{5.3}$$

Nebro et al. (2009) introduced a mechanism to bind further the speed of each variable j in particle i as follows:

$$v_{i,j}(t) = \begin{cases} delta_j & if\ v_{i,j}(t) > delta_j \\ -delta_j & if\ v_{i,j}(t) \leq -delta_j \\ v_{i,j}(t) & otherwise \end{cases} \tag{5.4}$$

where

$$delta_j = \frac{(upper\,limit_j - lower\,limit_j)}{2} \tag{5.5}$$

The position of each particle is updated based on:

$$\vec{x}_i(t) = \vec{x}_i(t-1) + \vec{v}_i(t) \tag{5.6}$$

After updating the particle's position, the polynomial mutation (Deb, 2001) on particle's velocity is performed with a given probability. The objective functions values associated with the new particle are evaluated. If these values dominate the objective values at the previous position, the position of the particle is updated and the new objective values are compared with the leaders archive. If the new solution dominates a leader in the archive then that leader will be replaced by the new solution. On termination, the leader archive represents the approximate Pareto-optimal solutions.

In SMPSO, as in NSGA-II, the leader archive size is fixed. As a result, the number of leaders can exceed the archive size. Thus, if the leader archive is full and a new solution does not dominate any archive solution, a crowding distance approach is used to decide which particle may be retained in the leader archive (Nebro et al., 2009).

This study used the jMetal code, which is a Java implementation of SMPSO by Durillo and Nebro (2011).

```
Initialize Swarm
Initialize LeadersArchive
while (termination criteria are not met) do
    ComputeVelocity()/ Eqs. (5.1) to (5.5)
    UpdatePosition() / Eq.  (5.6)
    Mutation()
    Evaluation()
    UpdateParticlesMemory()
    UpdateLeadersArchive()
end while
ReturnLeadersArchive()
```

Figure 5-1 Pseudo code for the SMPSO algorithm (Adapted from Nebro et al. (2009))

5-2-2 NSGA-II

The non-dominated sorting genetic algorithm (NSGA) proposed by Srinivas and Deb (1994) was one of the first EAs in context of multi-objective application. Criticisms of the NSGA approach included the high computational complexity of non-dominated sorting, lack of elitism and the need for specifying sharing parameters (Deb et al., 2002a). Deb et al. (2002a) addressed these issues by introducing the fast non-dominated sorting approach and crowding distance feature in an improved version of NSGA called NSGA-II.

In the fast non-dominated sorting approach, two entities are calculated for each solution, p, in the population: 1) the domination count, $n_{p,}$ which is the number of solutions that dominate solution p; and 2) S_p which is a set of solutions that the solution p dominates. The domination count of all solutions in the first non-dominated front is zero. For each solution p with $n_p = 0$, each member (q) of its S_p set, is visited and the associated n_p of that member is reduced by one. If n_p of any member becomes zero it means this member belongs to the next non-dominated front. When all the members of the current front are visited, the procedure is repeated for the next non-dominated front. This process terminates when all fronts have been identified.

The crowding-distance approach is the second key enhancement to NSGA. The first step is to sort the population according to each normalized objective function value in ascending order. Then, for each objective function, solutions with the smallest and largest function values are assigned an infinite distance value. For all other intermediate solutions, a distance equal to the absolute normalized difference in each function value of the two adjacent solutions is calculated. This procedure is repeated for all objectives. The crowding-distance is then calculated as the sum of individual distance values corresponding to each objective. In Figure 5-2, which depicts the minimization of two objectives, the crowding-distance of the i^{th} solution is the average side length of the cuboid.

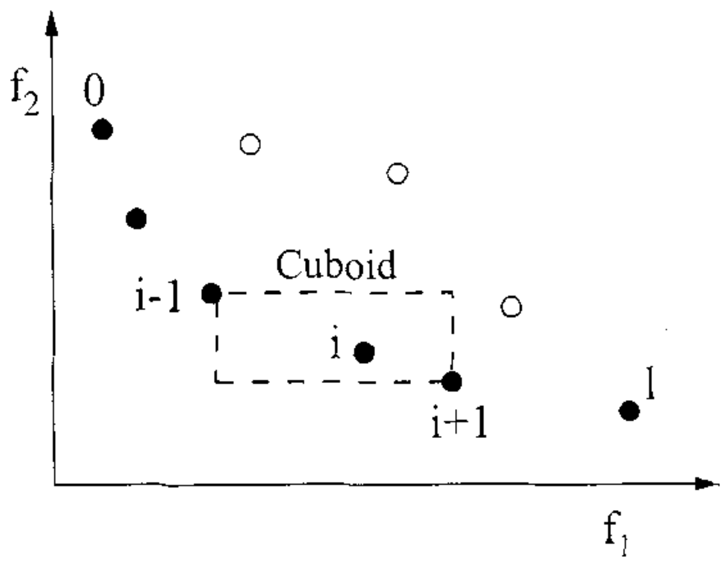

Figure 5-2 Illustration of crowding-distance algorithm. The points marked as filled circles are solutions of the same non-dominated front (Deb et al., 2002a)

The NSGA-II algorithm is straightforward to apply. Initially, a random population P_0 is created. Then the population is sorted into a number of non-dominated fronts using a fast non-dominated sorting approach, after which each solution is assigned a fitness value equal to its non-domination front. Applying selection, recombination and mutation operators creates an offspring population Q_0 of population size (N). As presented in Figure 5-3, the two populations are combined to form R of size 2N. Combining all previous and current population members ensures elitism. Then, all solutions are sorted using non-dominated sorting approach. The solutions belonging to the first non-dominated front are the best solutions in the combined population. If the number of solutions in the first front is fewer than N then all the solutions will be selected. The remaining members of the new population are selected from the subsequent non-dominated fronts in the order of their rankings. This process continues until all slots in the new population filled. During this process, it is possible to have more number of solutions in a front compared with the available slots in the new population. In this case, application of the crowding-distance algorithm ensures solutions within the less crowded regions will be selected (Deb et al., 2002a). This improves the diversity of population.

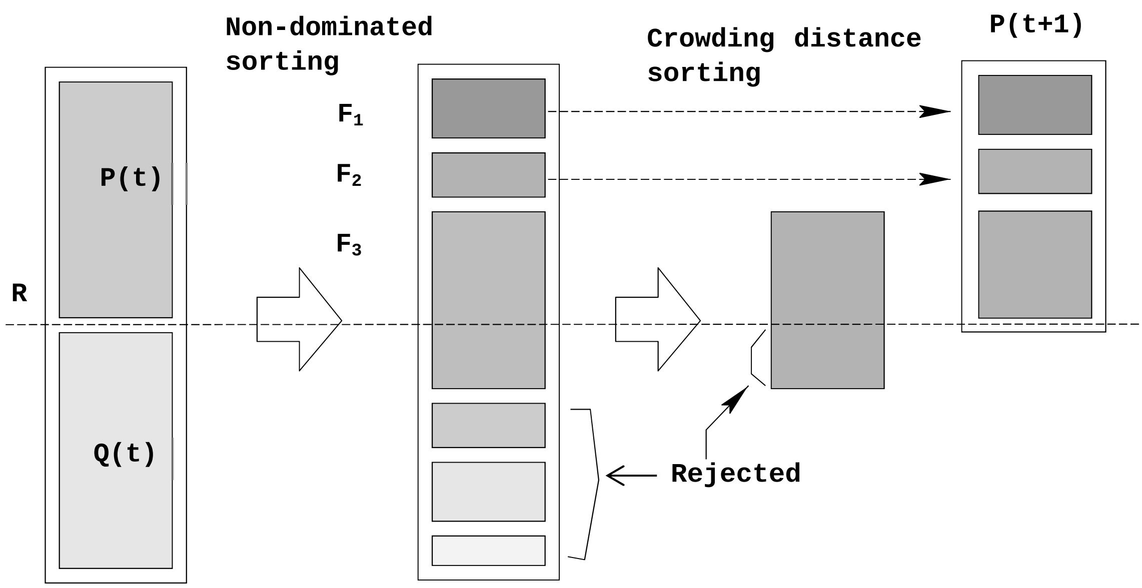

Figure 5-3 NSGA-II procedure (Deb et al., 2002a)

The NSGA-II code used in this study was obtained from the Kanpur Genetic Algorithms Laboratory web site (http://www.iitk.ac.in/kangal/codes.shtml, last visit 14/05/2012).

5-3 Evaluation of Multi-Objective Performance

Various performance metrics for measuring the quality of a Pareto-optimal set have been proposed to compare the performance of different multi-objective algorithms (Deb, 2001). However, there is no clear consensus in the literature on how the performance of multi-objective methods should be evaluated or compared. Deb and Jain (2002) suggested the use of metrics to characterize the two main functional objectives of MOO methods, proximity and diversity. Hadka and Reed (2011) investigated a broad range of performance metrics including hypervolume, generational distance (GD), inverse generational distance, additive epsilon indicator (ε_+-indicator) and spread. They recommended three metrics to characterize the three main functional objectives of MOO methods associated with proximity, diversity and consistency. The GD and hypervolume metrics are used to assess proximity and diversity respectively, while the ε_+-indicator is used to assess the consistency of the proximity of solutions. In this chapter the three measures recommended by Hadka and Reed (2011) were used to compares the competing MOO algorithms. In all of these measures, normalized objective values are used. The following sections discuss each of these measures in more detail.

5-3-1 Convergence (Generational Distance) Metric

The convergence metric is a proximity or distance measure describing how close a set of non-dominated solutions is to the Pareto-optimal front (Van Veldhuizen and Lamont, 2000). The minimum normalized Euclidean distance from each point i in the non-dominated solution set (Q) to the reference solution set (P*) is calculated using the following equation (Deb and Jain, 2002)

$$d_i = \min_{j \in P^*} \sqrt{\sum_{n=1}^{K} \left(\frac{f_n(i) - f_n(j)}{f_n^{\max} - f_n^{\min}} \right)^2} \tag{5.7}$$

where $f_n^{\max}$ and $f_n^{\min}$ are the maximum and minimum function values of the n^{th} objective function in P*. $f_n(i)$ is the n^{th} function value of point i in the set Q and $f_n(j)$ is the n^{th} function value of point j in the set P*. K is the number of objectives. The average of d_i is taken to be the convergence metric. The smaller the value of this metric, the closer the solutions are to the reference solution set.

The principal shortcoming of the convergence metric is that it contains no information about diversity. For instance, consider Figure 5-4, which illustrates two sets of non-dominated solutions along with the reference set. Set 1 has a smaller convergence metric than set 2. However, it is clear that set 2 has a superior coverage of the reference solution set. For this reason, it is necessary to use other measures that monitor the diversity of the non-dominated solution set.

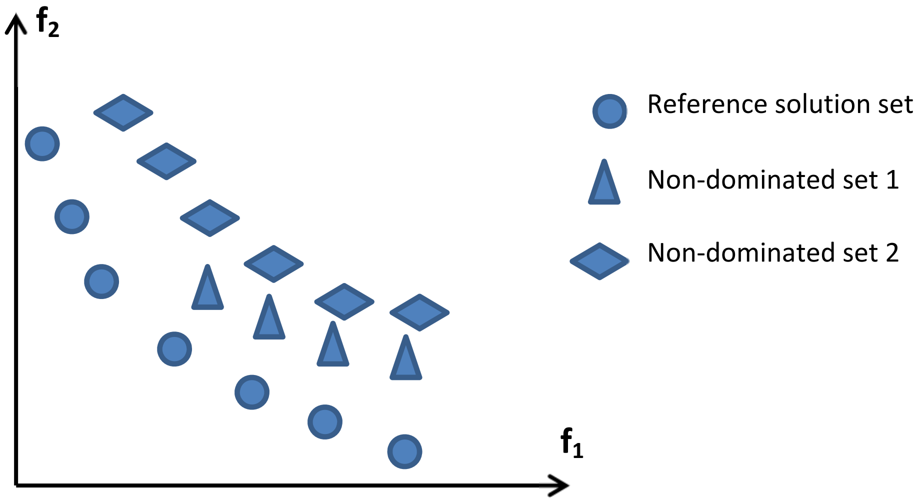

Figure 5-4 Schematic showing three non-dominated fronts to illustrate shortcoming of convergence metric

5-3-2 Hypervolume Ratio (HVR)

The hypervolume (HV) metric is defined as the volume (in objective space) enclosed by a reference point and the non-dominated solution set. The reference point can be defined using the worst objective function values. To illustrate this concept consider Figure 5-5 which shows a Pareto-optimal front, a non-dominated solution set (A, B and C) and a reference point denoted by W. The dashed lines define the hypervolume enclosed by non-dominated solutions. When the same reference point is used for multiple non-dominated fronts, the front with the larger HV is considered to be superior. In this study, the method developed by Fonseca et al. (2006) is used. This method is coded in an R package called “emoa” which can be downloaded from “http://www.statistik.tu-dortmund.de/~olafm/software/” (last visit 24/05/12).

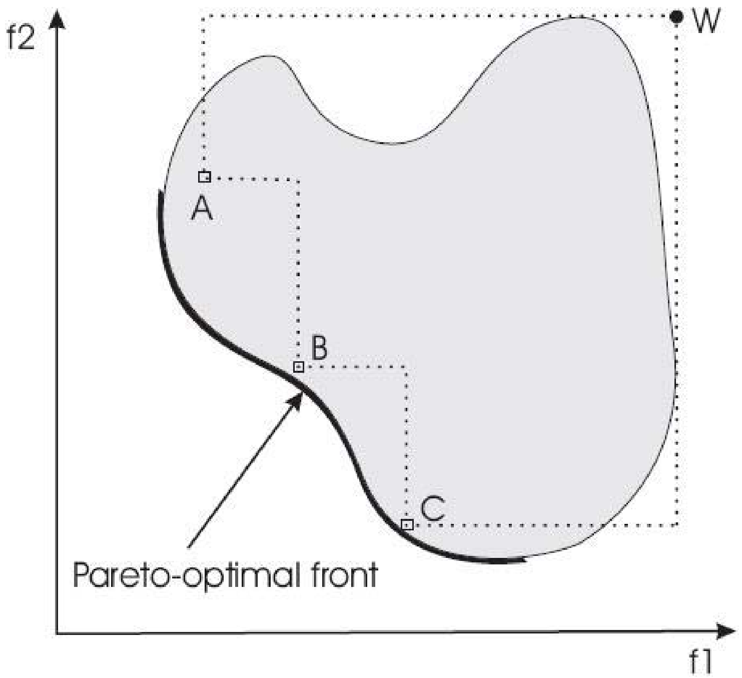

Figure 5-5 Hypervolume defined by the non-dominated solutions A, B and C (Durillo et al., 2010)

The hypervolume ratio (HVR) normalizes the HV to facilitate comparisons (Deb, 2001). It is defined as the ratio of the HV for a non-dominated solution set Q and the HV of a reference solution set P* which is taken to be the approximate Pareto-optimal solution set:

$$HVR = \frac{HV(Q)}{HV(P^*)} \tag{5.8}$$

5-3-3 Additive Epsilon Indicator ($I_{\varepsilon+}$)

The convergence and hypervolume metrics measure the proximity and diversity of a non-dominated solution set. However, these measures fail to identify a non-dominated solution set which contains one or more solutions with poor proximity. To deal with this another measure called $I_{\varepsilon+}$ is introduced. It is defined as the smallest distance one would need to translate every point in the non-dominated solution set, Q, so that it dominates a reference solution set, P* (Zitzler et al., 2002). Formally, if x_1 is an element of Q, x_2 is an element of P* and K is the number of objectives, the $I_{\varepsilon+}$ metric is (Durillo and Nebro, 2011):

$$I_{\varepsilon+}(Q) = \inf_{\varepsilon \in \mathbb{R}} \{\forall x_2 \in P^* \exists x_1 \in Q : x_1 < \varepsilon x_2\} \quad (5.9)$$

where, $x_1 < \varepsilon x_2$ if and only if $\forall 1 \le i \le K : f_i(x_1) < \varepsilon + f_i(x_2)$

Figure 5-6 illustrates the importance of this measure as a way dealing with the shortcoming of convergence and hypervolume metrics. Figure 5-6(a) shows a good approximation set, indicated by filled circles, and the reference set, indicated by the dashed line. In Figure 5-6(b) a new approximation set with a gap is illustrated. This new set is the same as the set in Figure 5-6(a) except the missing points are shaded light grey. The convergence measure fails to identify this gap because it averages the distance between the approximation set and reference set thereby reducing the impact of large gaps. The hypervolume measure also fails to identify the gap since the change in hypervolume due to a gap is small relative to the entire hypervolume - this is illustrated in Figure 5-6(c). However, the $I_{\varepsilon+}$ measure readily identifies the gap because it will be 2-3 times worse for the set with missing points as shown in Figure 5-6(d) (Hadka and Reed 2011). The $I_{\varepsilon+}$ measure can be interpreted as a measure of the quality or consistency of the coverage of the reference set. Therefore, the smaller the value of this metric, the smaller the gap in the solution set.

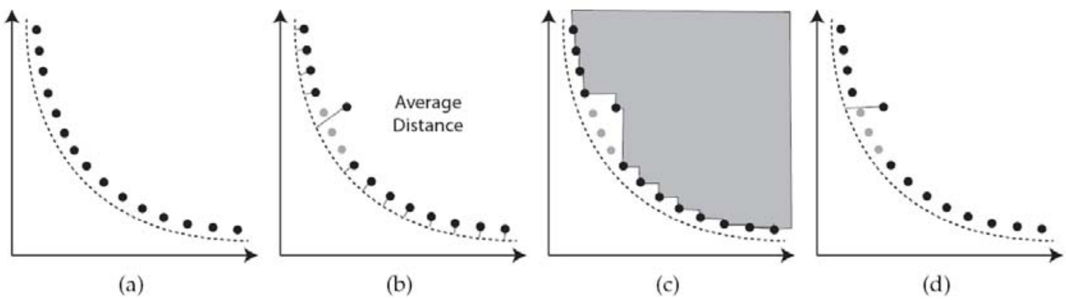

Figure 5-6 Illustration of $I_{\varepsilon+}$ as a measure of consistency (Hadka and Reed, 2011)

It is stressed that the convergence, HVR and $I_{\varepsilon+}$ metrics require knowledge of the reference solution set which in this chapter is referred to as the approximate Pareto-optimal solutions set.

5-4 Overview of Case Studies to Evaluate MOO Methods

In this chapter two urban water resource case studies based on the Canberra and Sydney headworks supply systems, are used to assess the performance of different

MOO methods. Detailed descriptions of these case studies can be found in Chapters 3 and 4. In addition, six benchmark problems and two water management case studies are used to test performance of MOACO variants. The purpose of this section is to summarize briefly the decisions and objectives used in these case studies.

5-4-1 Canberra Headworks System

The Canberra headworks system has four reservoirs supplying water to the city of Canberra. The layout of the system is presented in Figure 4-3. Details of the system can be found in Section 4-3-1. The main difference between the simulations conducted in this chapter and Chapter 4 is that there is no population growth and only one streamflow replicate is used based on the historical data for the period 1871 to 2009 during which several major droughts were experienced.

Thirteen decision variables are considered which categorized as either operational in that they control the running of the system or as infrastructure in that they define the physical characteristics of the system. The decisions are summarized in Table 5-1. Operational decisions include storage triggers for imposing restrictions, a pump mark for turning on the Murrumbidgee-Googong diversion, and parameters that determine the balance of storage between the Googong and Corin catchments. Infrastructural decisions involve major capital works to upgrade the capacity of Cotter Dam, a pumping station, the Stromlo water treatment plant and the construction of the Murrumbidgee diversion. In addition, the installation of rainwater tanks on individual allotments is supported to harvest roofwater and use it for non-potable indoor and outdoor uses.

Table 5-1 List of decision variables used in Canberra case study

Decision	Decision variable	Lower limit	Upper limit	Category
1	Murrumbidgee pump trigger	0	1	Operational
2	Murrumbidgee diversion capacity	0	6000	Infrastructure
3	Cotter pump capacity	3050	10000	Infrastructure
4	Googong incremental gain	20	500	Operational
5	Googong base gain	9000	11000	Operational
6	First restriction trigger level 1	0	0.999	Operational
7	Trigger intervals	0.05	0.25	Operational
8	Stromlo water treatment plant capacity	7625	15000	Infrastructure
9	Cotter incremental gain	20	500	Operational
10	Cotter base gain	9000	11000	Operational
11	Googong water treatment plant capacity	8235	15000	Infrastructure
12	Cotter Reservoir capacity upgrade	0	100000	Infrastructure
13	Number of houses used water tank	0	15000	Infrastructure

Up to three objectives were used:

1) **Minimize the frequency of restrictions** expressed as the percentage of months during which restrictions on water consumption are imposed. The restriction time fraction criterion is an important level-of-service measure.

2) **Minimize the expected present worth cost ($)** defined as the sum of capital and discounted expected operating costs and the costs of unplanned shortfalls. The capital cost represents the cost of building new infrastructure such as dams or water treatment plant capacity upgrades. Table 5-2 summarizes the capital cost items. It is noted the costs are hypothetical and thus should not be taken literally. Two capital items involve a binary choice: if the item is selected by the optimization, then there is fixed setup cost along with a unit cost; if the item is not selected, there is zero capital cost. The operating cost includes the costs for pumping from Cotter Reservoir and the Murrumbidgee River to the Stromlo water treatment plant and pumping from the Murrumbidgee River to Googong Reservoir, and the transfer and treatment costs associated with Stromlo and Googong water treatment plants. A 5% discount rate was used.

   An unplanned shortfall arises when the system is unable to supply, demand that may be restricted; in most cases, an unplanned shortfall arises when reservoirs empty and there is insufficient streamflow. To steer the optimization away from solutions that result in unplanned shortfalls, a penalty of $1,000,000 per ML unplanned shortfalls is added to the present worth cost.

3) **Minimize the fraction of time that total reservoir storage falls below 20%**: This objective measures the vulnerability of the supply system to drought condition.

Apart from the constraint on unplanned shortfalls, which was implemented using a penalty function approach, the only other constraints were the limits on the decision variables summarized in Table 5-1.

Table 5-2 Capital decision variables for the Canberra water headworks system

Decision Variable	Unit Cost
Cotter pump capacity upgrade	\$1000/ML
Cotter Reservoir capacity upgrade	\$50x10^6+\$1923/ML
Stromlo and Googong WTP capacity upgrade	\$9180/ML
Murrumbidgee to Googong pump diversion	\$20 x10^6+\$42623/ML
Water tanks	\$3000/ house

5-4-2 Sydney Headworks System

A full description of the Sydney headworks system can be found in Section 3-4-2. In this chapter, 150 years of stochastically generated streamflow data were used.

A large number of options are available to ensure a secure water supply for the 7-million population scenario. In this case study, eleven decision variables, listed in Table 5-3, were selected. Decisions 1 and 2 control the pump transfer of water from the Shoalhaven basin. Decisions 3 and 4 define the first stage of the drought contingency plan (DCP) to determine restriction levels. Decisions 5 and 6 define the second stage of the DCP. When the total storage fraction falls below the trigger given by decision 6, the already-constructed desalination plant with capacity given by decision 5 is activated. Decision 7 defines the capacity of Welcome Reef Reservoir. Decisions 8 and 9 define the priority for storing water in Warragamba. Depending on the values assigned to decisions 8 and 9, water may be preferentially stored in Warragamba or in the rest of the system. Decisions 10 and 11 define the maximum monthly Wollondilly transfer capacity during September to March and at other times respectively. The lower limit on these decisions corresponds to that recommended by Scott and Grant (1997). These two decisions are active in the three-objective scenario and fixed in the other scenarios.

Table 5-3 List of decision variables used in Sydney case study

Decision variable	Description	Lower limit	Upper limit	Category
1	Pump mark Warragamba	0.3	1	Operational
2	Pump mark Avon	0.3	1	Operational
3	Level 1 restriction trigger	0.05	0.95	Operational
4	Trigger increment	0.05	0.25	Operational
5	Desalination plant capacity (ML/day)	0	1,000	Infrastructure
6	Desalination plant trigger	0.05	0.95	Operational
7	Welcome Reef capacity (ML)	0	100,000	Infrastructure
8	Warragamba base gain	8,000	12,000	Operational
9	Warragamba incremental gain	10	200	Operational
10	Maximum Wollondilly flow during September to March (ML/month)	12,200	100,000	Operational
11	Maximum Wollondilly flow at other times (ML/month)	18,300	100,000	Operational

Up to three objectives were considered:

1) **Minimize frequency of restrictions (%)** defined as the percentage of months during which restrictions on water consumption are imposed.

2) **Minimize the present worth cost ($)** defined as the sum of capital and discounted expected operating costs and the costs of unplanned shortfalls. The capital cost represents the cost of building new infrastructure, which in this case study, is the Welcome Reef dam and/or the desalination plant. Table 5-4 summarizes the capital costs for Welcome Reef and the desalination plant. The capital cost model uses a binary function: if the asset is selected by the optimization, then the total cost is the sum of a fixed setup cost and a cost proportional to the size of the asset; however, if the asset is not selected, the capital cost is zero. The operating cost includes the costs for pumping transfers from the Shoalhaven and operation of the desalination plant. A 5% discount rate was used.

   The constraint on unplanned shortfalls is imposed using a penalty function approach. Here, a penalty of $100,000 per ML unplanned shortfall is added to the present worth cost. This penalty was selected to steer the optimization search away from solutions which allow reservoirs to “run dry” with consequent failure to supply minimum water needs.

3) **Minimize environmental stress on the Wollondilly River:** The following environmental stress metric was adopted to penalize the adoption of maximum regulated flow limits, defined by decisions 10 and 11, in excess of those recommended by Scott and Grant.

$$Stress(m) = \begin{cases} \max\left[0,\ 5\left(\dfrac{q_m - 12200}{12200}\right)\right] \text{if } m \in \{\text{Sept},..,\text{March}\} \\ \max\left[0,\ \left(\dfrac{q_m - 18300}{18300}\right)\right] \text{if } m \in \{\text{April},..,\text{August}\} \end{cases} \tag{5.10}$$

where q_m is the actual regulated release in the Wollondilly in month m and $Stress(m)$ is the penalty for exceeding the recommended flow limits in month m. The environmental stress criterion is the sum of the monthly stresses over the simulation.

Apart from the constraint on unplanned shortfalls, which was implemented using a penalty function approach, the only other constraints were the limits on the decision variables summarized in Table 5-3.

Table 5-4 Cost summary for infrastructure decision variables in Sydney case study

Decision Variable	Fixed and Unit Costs
Desalination plant capacity (ML/day)	\$1,250,000,000 + \$4,000,000 ML/day
Welcome Reef capacity (ML)	\$100,000,000 + \$1000/ML storage

5-4-3 Benchmark Problems

The performance of MOO algorithms is usually assessed using well-known benchmarks such as the Ziztler-Deb-Thiele (ZDT) test suite (Zitzler et al., 2000) and the DTLZ problems (Deb et al., 2002b). The eight-benchmark problems provide a sample of different types of Pareto fronts and different numbers of decision variables, thereby improving the chance of identifying efficient and robust MOO methods. Table 5-5 summarizes the benchmark problems that were used to evaluate the performance of different MOACO algorithms.

Table 5-5 Summary of benchmark problem characteristics

Name	Number of objectives	Number of variables	Type of Pareto Front
ZDT1	2	30	Convex
ZDT3	2	30	Convex, disconnected
ZDT4	2	10	Convex, multimodal
ZDT6	2	10	Concave, non-uniformly spaced
DTLZ1	3	7	Linear, multimodal
DTLZ2	3	12	Concave
DTLZ3	3	12	Concave, multimodal
DTLZ6	3	22	Degenerate

5-5 Ant Colony Optimization

In the preceding sections, three benchmark MOO methods, namely EMOEA, NSGA-II and SMPSO, were introduced. The focus of this section is the investigation of the potential of ant colony optimization (ACO) for applications involving urban water management. If the reader is mainly interested in the evaluation of MOO algorithm performance in the urban water resource context, then the reader may skip this section and proceed directly to Section 5-6. This section is organized as follows: First, the ACO method is briefly explained. Then existing multi-objective ant colony optimization (MOACO) algorithms are critically reviewed. Three MOACO methods, called MOACO-State, EMOACO and EMOACO-I, are then proposed to overcome the shortcomings identified in existing methods. These ACO methods are included in the evaluation of MOO algorithms in Section 5-6.

5-5-1 Overview of Ant Colony Optimization

Ant colony optimization is a recently developed heuristic optimization method. It was inspired by the fact that some species of ants are blind but nonetheless can find the minimum path between their nest and food. This is because of a chemical substances called pheromone that ants deposit when they travel on a route (Dorigo and Stützle, 2004). Based on the behaviour of real ants, Dorigo et al. (1991) and Dorigo et al. (1996) developed the first ant colony optimization method called Ant System (AS) to solve the travelling salesman problem (TSP) and job-shop scheduling problem (JSP).

The first step in ACO is to represent the search space as a graph. In the literature, two approaches have been used to represent the search space. The nodal method was applied by Abbaspour et al. (2001) and Kumar and Reddy (2006) using a graph similar to Figure 5-7 while the link method was used by Maier et al. (2003) using a graph similar to Figure 5-8. In both approaches, the first step is to split the range of each decision variable into a specific number of segments with a representative value assigned for each segment. In the nodal method, each variable value is represented by a node, and in the link method by a route. Ants travel between nodes corresponding to different variables to define a route as illustrated in Figures 5-7 and 5-8. In the nodal method, the route is denoted as (k, i, j) which means an ant travels from node i of variable k to node j of variable k+1, while in the second approach, the route is denoted as (k,i) which means an ant travels along route i from node k which corresponds to the i^{th} segment of variable k. To handle constraints in ACO, a tabu list can be defined to prevent ants travelling on infeasible routes.

One of the drawbacks of the nodal method is the potentially huge number of route combinations. To illustrate this, Figure 5-9 shows the possible routes for the nodal and link method when there are three segments. It shows there are nine possible routes for the nodal method and six possible routes for the link method. This difference rapidly grows as the number of variables and segments is increased. The large number of possible routes in the nodal method limits its ability to explore the search space. Mortazavi N. et al. (2009) compared these two methods and concluded the link method is the better choice. Accordingly, in this study, the link method is used.

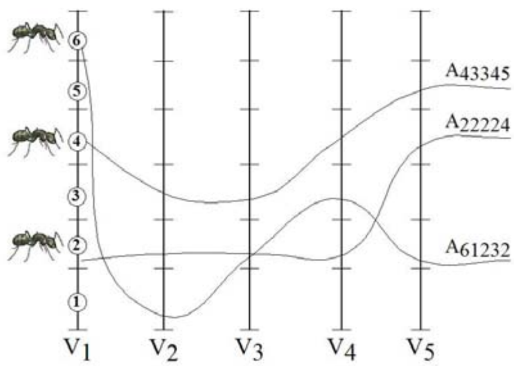

Figure 5-7 Schematic for nodal method showing ant routes between five discrete variables (Abbaspour et al., 2001)

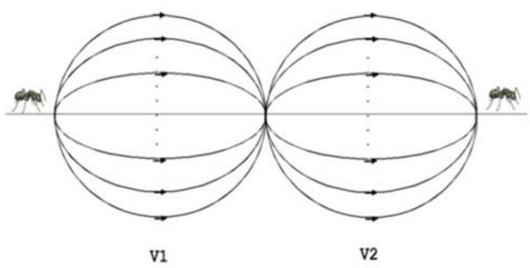

Figure 5-8 Schematic for link method showing ant routes between two discrete variables

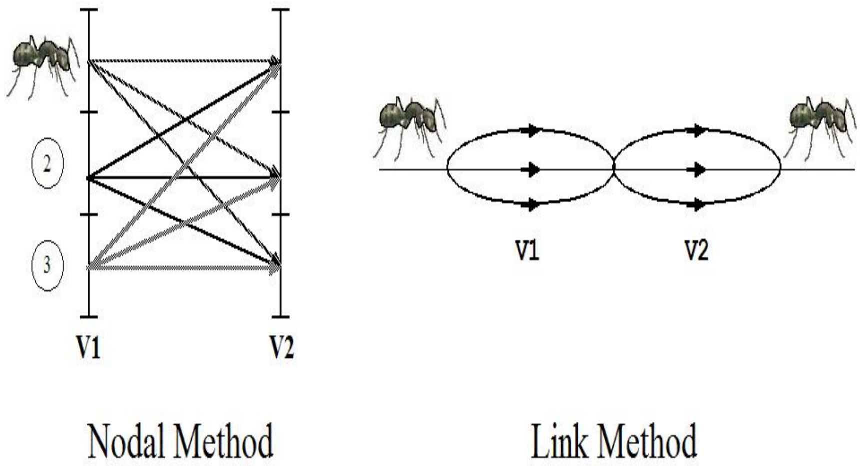

Figure 5-9 Depiction of possible routes for nodal and link methods

The next step is to define a transition rule that describes how ants select their route. The transition rule for the i^{th} variable in the link model is (Dorigo and Stützle 2004):

$$P_{ij} = \frac{[\tau_{ij}]^{\alpha}[\eta_{ij}]^{\beta}}{\sum_{i=1}^{N}[\tau_{ij}]^{\alpha}[\eta_{ij}]^{\beta}} \tag{5.11}$$

where P_{ij} is the probability the ant at node i will travel on link j, τ_{ij} is the pheromone trail strength and η_{ij} is heuristic information. The parameters α and β are introduced to control the relative importance of the pheromone and heuristic information respectively.

The pheromone trail strength encodes a long-term memory about the entire ant search process, and is updated by the ants themselves. In contrast, the heuristic information represents a priori information about the problem or run-time information provided by a source different from ants. Review of the literature shows there is no general rule for defining heuristic information. For instance, Dorigo et al. (1996) defined the inverse of distance between cities as heuristic information in the TSP while Maier et al. (2003) applied the inverse of cost of individual pipes.

When all ants complete their tours, the pheromone information is updated in two ways according to the following equation:

$$\tau_{ij} = (1-\rho)*\tau_{ij} + \Delta\tau \tag{5.12}$$

First, pheromone is reduced in strength by evaporation where ρ is the fraction of pheromone that evaporates. Second, ants deposit an amount of pheromone on the paths they visited with the magnitude of $\Delta\tau$ based on the quality of the solutions they found.

Various ACO methods have been developed to improve AS performance. The first improvement, based on the concept of elitism, was introduced by Dorigo et al. (1991) and Dorigo et al. (1996). The main idea is to add significant additional pheromone onto the arcs belonging to the best tour found since the start of the search.

The MAX-MIN AS introduced four modifications with respect to AS (Stützle and Hoos, 1996; Stützle and Hoos, 1997; Stützle and Hoos, 2000). The first modification strongly exploits the best tour found. Only either the iteration-best route (that is, the best route in the current iteration) or the best-so-far route is allowed to receive pheromone. The drawback of this strategy is that it can lead to stagnation where all ants follow the same, although suboptimal, tour. To overcome this shortcoming, the second modification limits the pheromone values on all routes to the interval $[\tau_{\min}, \tau_{\max}]$. The third modification initialized the pheromone on all routes to τ_{max} which in conjunction with a small evaporation rate increases the exploration at the start of the search process. Finally, pheromone values were reinitialized when the system reaches stagnation (Dorigo and Stützle, 2004).

Stützle and Hoos (1997) defined the maximum amount of pheromone for a TSP as follows

$$\tau_{\max} = \frac{1}{\rho}.\frac{1}{L_{opt}} \tag{5.13}$$

where ρ is evaporation rate and L_{opt} is the shortest route found in an iteration. The minimum amount of pheromone is defined by (Stützle and Hoos 1996):

$$\tau_{\min} = \tau_{\max}\frac{1-\sqrt[n]{p_{best}}}{(avg-1)\sqrt[n]{p_{best}}} \tag{5.14}$$

where *avg* is the average number of different choices available to an ant at each step and *n* is number of segments. It is assumed that a run of MAX-MIN AS has converged if the best found route is constructed with a probability significantly higher than 0 – this probability is assigned a specific value p_{best} (Stützle and Hoos 2000).

The Ant Colony System (ACS) (Dorigo and Gambardella, 1997) is another improved method based on AS which differs in three aspects, namely transition rule, global and local updating. ACS explores the search space more strongly due to its more aggressive transition rule. The transition rule in ACS is as follows:

$$j = \begin{cases} \arg\max\limits_{l \in N_i^k} \left\{ \tau_{il}[\eta_{il}]^{\beta} \right\} & if \;\; q \le q_0 \\ J & otherwise \end{cases} \tag{5.15}$$

where q is a random number uniformly distributed in [0,1], q_0 is a parameter and J is a random variable selected according to Eq. (5.11) with α=1.

In ACS pheromone is only updated on the best-so-far routes. The pheromone level is updated by applying the global updating rule using Eq. (5.12) where τ_{ij} denotes route (i, j) that belongs to the best-so-far solution. The local updating rule is applied to all routes in a tour. In this rule pheromone is evaporated on any route traversed by ants making that route slightly less desirable. The local rule is:

$$\tau_{ij} = (1-\zeta)*\tau_{ij} + \zeta\tau_0 \tag{5.16}$$

where τ_0 is the initial amount of pheromone on routes and ζ is local evaporation rate.

All these improved AS methods have been successfully applied to a number of benchmark combinatorial optimization problems (Dorigo and Di Caro, 1999; Stützle et al., 2010).

5-5-2 Review of Existing Multi-objective ACO Methods

Several issues need to be addressed when adapting ACO to multi-objective optimization including the number of pheromone and heuristic matrices and the pheromone updating procedure. The fact that there is considerable choice has resulted in a wide range of published multi-objective ant colony optimization (MOACO) methods (Martı´nez et al., 2007; Angus and Woodward, 2009). In this section, the main features of existing methods are described and their potential drawbacks highlighted.

Most MOACO approaches are extensions of well-known single objective ACO methods. For example, Baran and Schaerer (2003) and Doerner et al. (2003) adapted the Ant Colony System (ACS) while Bui et al. (2008) adapted the Ant System (AS). Although MOACO methods differ in detail, all share the following common steps:

Step 1: Initialize parameters
Step 2: Construct solutions
Step 3: Find and archive non-dominated solutions
Step 4: Update pheromone
Step 5: Go to step 2 if the termination condition is not satisfied

Generally MOACO methods can be categorized according to their number of pheromone and heuristic matrices (Martı´nez et al. 2007; Angus and Woodward 2009). Iredi (2001) proposed an approach for bi-criterion optimization problems which uses cooperative ant colonies and multiple pheromone and heuristic matrices. Doerner and Gutjahr (2004) and Cardoso et al. (2003) developed the Pareto ant colony optimization (P-ACO) and multi-objective network optimization based on ACO (MONACO) respectively with a single heuristic matrix and several pheromone matrices. In contrast, crowding population-based ant colony optimization (CPACO) and multiple ant colony system (MACS) methods were applied with multiple heuristic matrices and a single pheromone matrix (Baran and Schaerer, 2003; Angus, 2007b). Indeed, in these methods, diversity is achieved across the Pareto front through the use of heuristic rather than pheromone information. McMullen (2001), Gravel et al. (2002) and T'Kindt et al. (2002) developed methods which used a single pheromone and a single heuristic matrix.

In some of above methods, a single ant colony was used (Doerner and Gutjahr, 2004; Alaya et al., 2007), while in the other methods, multiple colonies were used (Iredi, 2001; Baran and Schaerer, 2003; Doerner et al., 2003). The main reason for having multiple colonies is to treat objectives independently. Doerner et al. (2003) introduced COMPETants with multiple colonies. Each colony corresponds to an objective. One drawback of this approach is that by allowing ants to explore individual objectives independently, they are more likely to explore the extremes of the Pareto front and neglect the compromise trade-off points. For this reason Doerner et al. (2003) introduced the *spies* idea to facilitate sharing and exchanging information between colonies. In a similar way, Alaya et al. (2007) used multiple ant colonies with each colony dedicated to a single different objective using its own pheromone and heuristic information to build solutions. To avoid exploring extremes of Pareto front they introduced an extra colony that aims at optimizing all objectives. They compared four MOACO methods with different numbers of colonies and pheromone matrices and found the method using a single colony and multiple pheromone matrices performed best. Other researchers used multiple colonies for other purposes. Iredi (2001) used multiple colonies with the aim of forcing ants to find good solutions along the whole the Pareto front. The Iredi (2001) approach is different from Alaya et al. and Doerner et al. in that he used multiple pheromone matrices in each colony. This approach is conceptually similar to having a single ant colony with multiple start points but it is different because altered weights have been applied in each colony to weight the pheromone and heuristic information.

Except in the case of multiple colonies where each colony has its own pheromone and heuristic information, it is necessary to integrate multiple pheromone or heuristic matrices in the transition rule. There are two methods for integration, namely weighted product (Iredi 2001; Baran and Schaerer 2003; Cardoso et al. 2003; Angus 2007b) and weighted sum (Doerner et al. 2003; Doerner and Gutjahr 2004).

In the weighted product method the transition rule is defined as:

$$p_{ij} = \frac{\prod_{l=1}^{L} (\tau_{ij}^{l})^{\alpha w_l} \prod_{m=1}^{M} (\eta_{ij}^{m})^{\beta w_m}}{\sum_{i=1}^{N} \prod_{l=1}^{L} (\tau_{ij}^{l})^{\alpha w_l} \prod_{m=1}^{M} (\eta_{ij}^{m})^{\beta w_m}} \tag{5.17}$$

where L is the number of pheromone matrices, M is the number of heuristic information matrices and the w_l and w_m are pheromone and heuristic information weights respectively.

In the weighted sum method the transition rule is defined as:

$$p_{ij} = \frac{\sum_{l=1}^{L} (w_l \tau_{ij}^{l})^{\alpha} \sum_{m=1}^{M} (w_m \eta_{ij}^{m})^{\beta}}{\sum_{i=1}^{N} (\sum_{l=1}^{L} (w_l \tau_{ij}^{l})^{\alpha} \sum_{m=1}^{M} (w_m \eta_{ij}^{m})^{\beta})} \tag{5.18}$$

When applying either the weighted sum or weighted product approach the main challenge is to maintain diversity along the Pareto front. For that reason in the literature different approaches have been proposed to define weights. Iredi (2001) developed an approach that changed weights dynamically according to the following equation:

$$w = \frac{k-1}{m-1} \tag{5.19}$$

where w is the weight for ant k and m is the number of ants.

Doerner and Gutjahr (2004) in P-ACO assigned a set of weights randomly at each iteration for each ant. Alaya et al. (2007) did not apply weights. However, since at each iteration a randomly selected objective was optimized, they implicitly applied binary weights (0 or 1). Angus (2007b) used the average-rank-rate method in which higher scoring objectives were assigned a greater weighting.

The use of different approaches to set weights can result in different search behaviours (López-Ibáñez et al., 2004). The increase in required memory associated with multiple pheromone matrices can be of concern if the actual problem size is sufficiently large (Angus and Woodward 2009). Most of above mentioned approaches were applied to bi-objective cases. Thus, there is little guidance on application of these approaches to problems with three or more objectives.

One of the challenging aspects of MOACO is the definition of a heuristic information matrix. This is problem-specific and not necessarily easy to establish (Doerner et al., 2003; Coello Coello et al., 2007). Developing a proper heuristic information matrix is likely to be even more challenging in a multi-objective setting because of the need to define heuristic information for every objective.

One of the most important steps of MOACO is the pheromone update. This involves two considerations that affect performance. The first is the selection of the routes to be updated and the second is the amount of pheromone to be deposited on the selected routes.

Several approaches have been proposed in the literature for selecting routes. The selected routes may be based on the non-dominated solutions within an iteration (Iredi, 2001), the non-dominated solutions found so-far (Baran and Schaerer, 2003; Alaya et al., 2007; Afshar et al., 2009) or the best (and the second-best) solutions according to each objective (Doerner and Gutjahr, 2004). Bui et al. (2008) compared several pheromone updating methods including updating based on all solutions from the current iteration, non-dominated solutions in the current iteration, and non-dominated solutions of all iterations. Their conclusion was that updating the non-dominated solutions of all iterations outperformed other updating methods. This finding is consistent with the good performance of the elitism strategy in single objective optimization.

However, a drawback of updating non-dominated solutions found in all iterations is that adding pheromone to these routes continuously may induce premature convergence and thus prevent the algorithm from generating an even coverage of the Pareto front (Angus and Woodward 2009). To abate loss of diversity, several methods have been suggested in the literature. Angus (2007b) used dominance ranking according to a non-dominated sorting technique to produce an even coverage of the Pareto front. Alaya et al. (2007) updated the pheromone value of a route only once despite how many solutions contain it. Bui et al. (2008) introduced an aging factor to deal with this issue. The main idea is to deposit more pheromone on routes associated with more recent solutions in the archive of non-dominated solutions. They defined the age factor (AF) as:

$$AF = \frac{1}{Current\,iteration\,number - iteration\,number\,solution\,added\,to\,archive + 1} \tag{5.20}$$

A variety of pheromone updating methods has been developed. In a single objective ACO minimization problem, pheromone is updated according to the inverse of objective function value - that is, depositing more pheromone on routes with the smaller objective function values encourages ants to follow those routes. Several researchers have tried to extend this idea to multi-objective problems. Baran and Schaerer (2003) used the inverse of the product of two objective function values for updating pheromone. A drawback of this method is that the amount of pheromone can be very sensitive to the objective value. This can cause premature convergence. Alaya et al. (2007) used the following equation to update pheromone.

$$\Delta\tau^i(c) = \frac{1}{1 + f_i(S^i) - f_i(S^i_{best})} \tag{5.21}$$

where $\Delta\tau^i(c)$ is the quantity of pheromone deposited on a route (c) at the i^{th} iteration, $f_i(S^i)$ is the value of an objective function for the current iteration and $f_i(S^i_{best})$ is the value of the best solution found so-far. Indeed, in Eq. (5.21) the value of $\Delta\tau^i(c)$ is scaled between 0 and 1.

To avoid the drawbacks associated with objective-dependent updating pheromone approaches, several researchers introduced updating methods independent of the objectives. Iredi (2001) suggested an updating rule where every ant is allowed to update the amount of pheromone equal to 1/L where L is the number of ants that are allowed to update in the current generation. Doerner et al. (2003) updated pheromone for only a number of the best ants ranked according to solution quality. They deposit pheromone based on the ant's rank. In a similar way, Angus (2007b) updated pheromone based on the ant's ranking. He used dominance ranking according to a non-dominated sorting approach such as that of the NSGA-II algorithm. López-Ibáñez et al. (2004) suggested all ants deposit a constant amount of pheromone. Similarly, Doerner and Gutjahr (2004) used a constant amount of 10 and 5 for the best and the second best ants respectively.

5-5-3 Towards An Improved MOACO Algorithm

There are several studies (Martı´nez et al., 2007; Angus and Woodward, 2009; López-Ibáñez and Stützle, 2010) which have sought to investigate all existing MOACO methods. However, these studies have mainly considered a particular combinatorial problem such as the travelling salesman problem (López-Ibáñez and Stützle 2010). As a result, it is difficult to these findings in urban water management applications.

The review of existing MOACO methods in the preceding section showed that there are several important aspects that need to be addressed in MOACO algorithms. These include the number of heuristic and pheromone matrices, the transition rule, the pheromone updating procedure and the specification of heuristic information. To investigate the importance of each of these aspects, fifteen MOACO variants are constructed from the review of existing MOACO methods. These variants are compared against each other, using the benchmark problems and the Canberra case study with two and three objectives.

In the interest of brevity, the details of these variants and the findings are presented in Appendix A. Here the focus will be on the best of these variants and on further enhancement. Out of the variants investigated, the variant called MOACO-State appeared to best explore the search space regardless of number of objectives and objective scales. The main features of MOACO-State are as follows:

1. Use a single colony with a single pheromone matrix regardless of the number of objectives. This avoids the complexity of colonies communicating in the search process and also avoids the need to assign weights.

2. Do not use problem-specific heuristic information by setting β to zero in Eq. (5.11). This ensures MOACO-State is problem-independent.

3. Apply a constant amount of pheromone (C) when updating routes corresponding to non-dominated solutions to maximize diversity in the Pareto front. This is

motivated by the fact that all the points on the Pareto front should be treated equally.

4. Use the MAX-MIN AS method with $\tau_{\max}$ defined by:

$$\tau_{\max} = \frac{C}{\rho} \tag{5.22}$$

5. Introduce a pheromone aging factor to reduce the chance of premature convergence. This is accomplished using the following pheromone update :

$$\Delta\tau = \frac{C}{AF} \tag{5.23}$$

where AF is the number of iterations since the current non-dominated solution was added to the archive as described in Eq. (5.20).

Figure 5-10 presents pseudo code for the MOACO-State algorithm. It provides a reference point for the next section which explores ways of enhancing this algorithm.

Initialization

Assign parameters /*Number of decisions, Number of objectives, Number of segments (N), Number of ants, Initial pheromone (τ_0), pheromone deposit (C), evaporation rate (ρ), α, β, *avg*, P_{best} */

Calculate /*τ_{max}, τ_{min} */

For i= 1 **to** number of decisions **do**

  For j = 1 **to** number of segments **do**

    Create a route from decision i to decision i+1 ($R_{i,j}$) (Link method)

  End for

End for

Initialize pheromone trail /*put initial pheromone (τ_0) on all routes*/

Initialize archive of non-dominated solutions found so-far (ND) = {ϕ}

While stopping criteria is not met **do**

/*a sufficient good fitness or a maximum number of iterations or no results improvements after a specific number of iterations*/

Solution construction

  For all ants **do**

   For all decisions **do**

    Select the route based on probabilistic transition rule Eq. (5.11)

   End for

   Assign decision values to solutions based on the route traversed by the ant

  End for

Evaluation

  For all solutions **do**

   Evaluate the objective function values corresponding to the solution

   If (new solution dominates any in ND) **then**

    Delete dominated solutions

    Add the new solution to ND

   Else if (none of ND dominate the new solution) **then**

    Add the new solution to ND

   End if

  End for

Updating pheromone

  Evaporate pheromone of all routes

  For all ND **do**

   Calculate aging factor (Eq. (5.20))

   Calculate $\Delta\tau$ (Eq. (5.23))

   Update pheromone trail for the solution

  End for

  For all decisions **do**

   If (pheromone trail< τ_{min}) **then**

    pheromone trail = τ_{min}

   Else if (pheromone trail> τ_{max}) **then**

    pheromone trail = τ_{max}

   End if

  End for

End while

Output

Print set ND

Figure 5-10 Pseudo code for MOACO-State algorithm

5-5-4 Towards a More Efficient MOACO Algorithm

Comparison of MOACO-State against other MOO methods, reported in Section 5-6, revealed that MOACO-State was outperformed by the other methods. This prompted a careful assessment of the factors affecting the rate at which MOACO-State converged to the Pareto-optimal front. The ensuing insights led to the development of superior methods called EMOACO and EMOACO-I, which are described in this section.

The way pheromone is updated in MOACO algorithms has a major effect on performance. In the absence of heuristic information, the pheromone update assumes an even more important role balancing exploration and exploitation. In MOACO-State it was found that the amount of pheromone on all routes was decreased to a small value, except for the limited number of routes belonging to the current non-dominated solution set. To illustrate this, the amount of pheromone on all segments for the first decision for a range of evaluation numbers is shown in Figure 5-11. This figure shows that only a limited number of segments, which belong to non-dominated solutions, have high pheromone. As a consequence, ants mostly explore routes belonging to the non-dominated solutions resulting in a very narrow exploration of the search space. While this behaviour impedes the convergence rate, it does offer an opportunity for improvement.

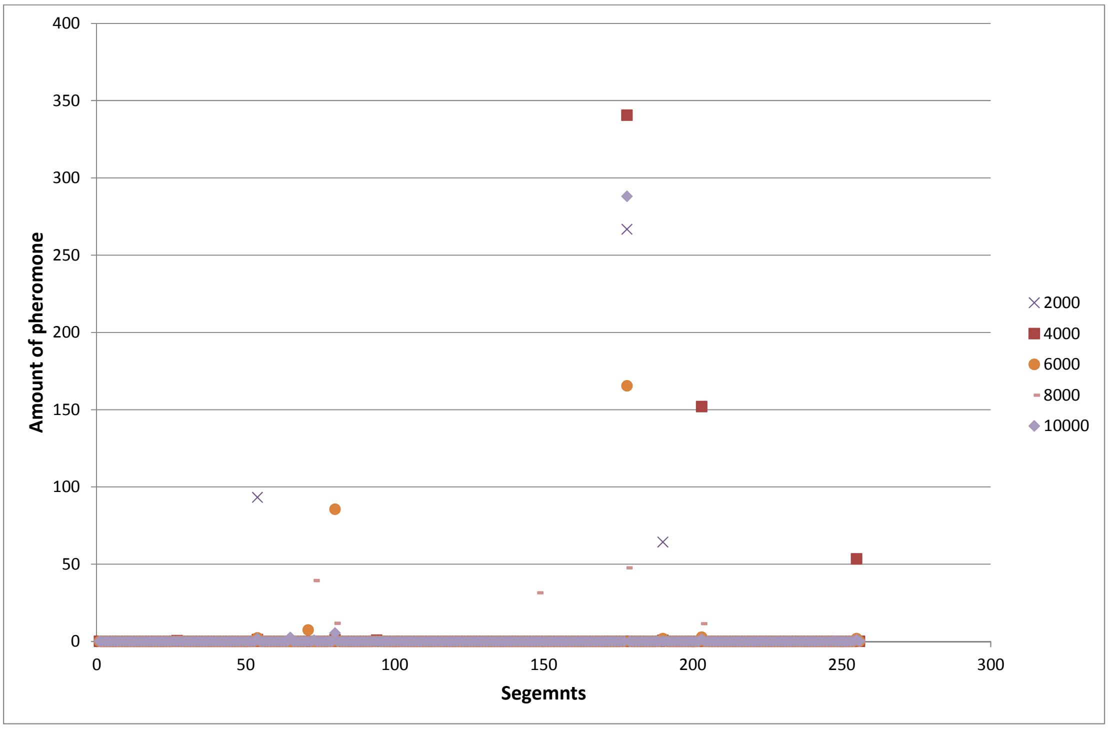

Figure 5-11 Amount of pheromone on segments for the first decision after 2000, 4000, 6000, 80000 and 10000 evaluations (for Canberra case study with two objectives)

To improve MOACO-State performance several ideas are adapted from evolutionary algorithms. It is worth stressing that the methods developed in this section are not hybrid GA-ACO methods since they do not use any GA or other evolutionary algorithm steps in any part of their algorithm.

It is widely accepted that the mutation operator in evolutionary algorithms helps avoid the algorithm being trapped at local minima and fosters diversity (Srinivas and Deb, 1994; Deb et al., 2002a). To facilitate this feature in MOACO, the following route selection process is proposed. For ant k and decision i, the route r_{ki} is selected using:

$$r_{ki} = \begin{cases} Randomly\ select\ one\ of\ the\ N_i\ routes & if\ \ q \le q_0 \\ Select\ route\ applying\ Eq.(5.11) & otherwise \end{cases} \tag{5.24}$$

where N_i is number of routes available for decision i, q_0 is a parameter in the interval (0,1) and q is a random sample from a uniform distribution over the interval (0,1).

A key feature of genetic algorithms is the crossover operator which allows, for much of the time, exploration of the search space in the neighbourhood of the parent decisions. This motivated the introduction of the adjacency concept in pheromone updating as a way of mimicking the ability of the crossover operator to explore the search space. The adjacency concept is implemented by depositing pheromone on routes adjacent to the current non-dominated solution routes. For decision i, the proximity of route (i, j) to the nearest non-dominated route is given by its adjacency score k, defined as |j-j*| where (i, j*) is the closest non-dominated route to route (i, j). The pheromone deposit on route (i, j) then becomes

$$\Delta\tau_{ij} = \begin{cases} \dfrac{c}{(AF)} & if\ (i,j)\ has\ an\ adjacency\ score\ k = 0 \\ \dfrac{c}{(AF)(1+k)} & if\ (i,j)\ has\ an\ adjacency\ score\ 0 < k \le k_{\max}\ and\ u < P_{adj} \\ 0 & otherwise \end{cases} \tag{5.25}$$

where u is a random uniformly distributed number over the interval (0,1), P_{adj} is the adjacency probability which determines the probability of depositing pheromone on an adjacent route and k_{max} is the maximum number of adjacent routes.

It was found that the adjacency pheromone update improved convergence to the Pareto-optimal frontier as long as new solutions were being added to the set of non-dominated solutions. However, the longer it took to find a new non-dominated solution, the greater was the likelihood of stagnation. To overcome this problem, the strategy of revisiting and mutating one of the current non-dominated solutions was introduced. The strategy commences when the number of iterations during which no new non-dominated solution is added, exceeds a predefined value, NoImprovement. One of the current non-dominated solutions is then selected randomly and one of its decisions is changed randomly. This procedure continues until a new non-dominated solution is found. The NoImprovement value is defined as the number of iterations required to reduce pheromone by evaporation from $\tau_{\max}$ to τ_{min}. It can be shown to be:

$$NoImprovement = \frac{-\log\left(\frac{(avg-1)\sqrt[n]{p_{best}}}{1-\sqrt[n]{p_{best}}}\right)}{\log(1-\rho)} \tag{5.26}$$

Adding these enhancements to the MOACO-State method results in a more efficient multi-objective ant colony optimization method hereinafter referred to as the EMOACO algorithm. To summarize these changes more formally Figure 5–12 presents the pseudo code for EMOACO.

EMOACO, like all other MOACO methods, randomly selects the initial routes traversed by each ant. The rate of convergence is affected by how close these initial routes are located to the Pareto-optimal routes. With this in mind, the following simple heuristic was adopted: EMOACO starts with the decision space being split into 8 rather than 256 segments – this reduces the number of decision combinations and thus improves the chance of EMOACO finding routes in the neighbourhood of Pareto-optimal solutions. Once a predetermined number of evaluations (500 in this study) is completed, the routes of the current non-dominated solutions are mapped, as initial routes, to the final search space where the number of segments for each decision is substantially increased (256 segments in this study). This enhancement to EMOACO is referred to as EMOACO-I.

Initialization

Assign parameters /* adjacency probability(P_{adj}), Number of decisions, Number of objectives, Number of segments (n), Number of ants, Initial pheromone (τ_0), pheromone deposit (C), evaporation rate (ρ), Number of adjacency routes (K_{max}), mutation rate (q_0), *avg*, P_{best} /*
Calculate /* τ_{max}, τ_{min}, NoImprovement */
For i= 1 **to** number of decisions **do**
  For j = 1 **to** number of segments **do**
    Create a route from decision i to decision i+1 ($R_{i,j}$)
  End for
End for
Initialize pheromone trail /*put initial pheromone (τ_0) on all routes*/
Initialize archive of non-dominated solutions found so-far (ND) = {ϕ}
Initialize Count = 0 /*count number of iterations without any changes in ND */
While stopping criteria is not met **do**
/*a sufficient good fitness or a maximum number of iterations or no results improvements after a specific number of iterations*/

Solution construction

  N = -1 /*flag changes in ND */
  For all ants **do**
    If (Count > NoImprovement) **then** /*revisiting solutions */
      Chose randomly one of the solutions in ND
      Change randomly one of its routes
      Save the changed solution as a new solution
    Else
      For all decisions **do**
        Select a random number between 0 and 1 (r)
        If (r < q_0) **then** /* mutation */
          Select route randomly
        Else
          Select route based on probabilistic transition rule Eq. (5.24)
        End if
      End for
      Assign decision values to solutions based on the route traversed by the ant
    End if
  End for

Evaluation

  For all solutions **do**
    Evaluate the objective function values corresponding to the solution
    If (new solution dominates any of ND) **then**
      Delete dominated solutions
      Add the new solution to ND
      N=0 /*flag changes in ND */
    Else if (none of ND dominate the new solution) **then**
      Add the new solution to ND
      N=0 /*flag changes in ND */
    End if
  End for
  If (N <> 0) **then** /* No changes in ND */
    Count = Count + 1
  Else
    Count = 0
  End if

```
Updating pheromone
        Evaporate pheromone of all routes
        For all ND do
                Calculate aging factor (Eq.(5.20))
                Calculate Δτ (Eq.(5.25))
                Update pheromone trail for the solution
                For all decisions do
                        Select a random number between 0 and 1 (r)
                        If (r < P_adj) then /* adjacency */
                                Add pheromone to adjacency routes (Eq.(5.25))
                        End if
                End for
        End for
        For all decisions do
                If (pheromone trail< τ_min) then
                        pheromone trail = τ_min
                Else if (pheromone trail> τ_max) then
                        pheromone trail = τ_max
                End if
        End for
    End while
Output
    Print set ND
```

Figure 5-12 Pseudo code for EMOACO algorithm

5-6 Evaluation of MOO Algorithm Performance

This section evaluates the performance of three benchmark MOO algorithms, NSGA-II, εMOEA and SMPSO, and three MOACO methods, MOACO-State, EMOACO and EMOACO-I. The primary objective is to identify the most efficient algorithm for urban water resources problems for a relatively small number of evaluations – this is motivated by the fact that function evaluations for urban water resource problems are computationally expensive. This section is organized as follows: First, the parameters of all six MOO methods are tuned. Then the performance of these methods is compared using the three metrics described in Section 5-3 for the Canberra and Sydney case studies.

5-6-1 Tuning

In this section the six MOO methods are tuned to obtain "good" parameters to ensure a fair comparison. To ensure consistency across methods, binary coding with 8 bits (equivalent of 256 segments) and the same number of evaluations, i.e. 10000, were used. All methods were run 10 times with different initial random number seeds.

p_{best} in Eq. (5.14) was set to 0.05 (Stützle and Hoos, 2000). Polynomial mutation and uniform crossover operators were applied to SMPSO and NSGA-II respectively. One-point crossover with bitwise mutation was used in εMOEA with an initial population of 100.

The parameters of the six MOO methods are listed in Table 5-6. To ensure all six methods were compared in a fair manner, a structured search was used to optimize the performance metrics using a related problem, namely the two-objective Canberra system simulated between 1970 and 1990. For each method, a set of default parameters based on values recommended in the literature was adopted. Then a range of values for each parameter was generated by perturbing the default values. A combination of tuning parameters was formed by selecting a value for one parameter from the available range while keeping the other parameters at their default values. These combinations are described in Table 5-6 . The notation “Id” and “n” are used in this table to express the combinations. “Id” represents the name of a method, eg “E” denotes EMOACO method, and “n” shows the combination number. Finally, the performance metrics for up to 2000 evaluations were evaluated for different combinations of tuning parameters to identify the best set of tuning parameters for each method. The results of the HVR and convergence measures of all combinations for all methods are presented in Figure 5-13 to Figure 5-17. In the case where no combination of tuning parameters was superior over all evaluations, the combination which had the best HVR was selected on the grounds that HVR assesses both proximity and diversity while the convergence measure only assesses proximity. Both the HVR and convergence metrics require knowledge of the approximate Pareto-optimal front. For the purposes of this study, the reference or “true” Pareto-optimal set was obtained from the Pareto-optimal set extracted from 60 runs, obtained from 10 runs for each the six MOO methods with each run involving 10,000 evaluations. The reference set will be referred to as the “true Pareto front” or TPF.

Table 5–6 summarizes the adopted parameters for each method. As only a limited number of combinations for each method were explored, there is a distinct possibility that the best set was not identified. Because EMOACO and EMOACO-I had 7 parameters for which only 20 combinations were tested, it is more likely that a better set of parameters was found for the non-EMOACO methods. Therefore, the tuning procedure intrinsically favoured better outcomes for the non-EMOACO methods. It is acknowledged that the structured search is premised on the assumption that there is little interaction between parameters. As most of the tuned parameters were at the default literature values, this issue is considered to be of secondary importance.

Table 5-6 Summary of parameters used in the six MOO methods

Method	Parameter	Default Value	Range of values used in tuning(presented number in Figures)	Number of parameter combinations tested	Tuned value
EMOACO and EMOACO-I	Number of ants	1	1(E1); 2(E2); 5(E3); 10(E4)	20	1
	ρ	0.02	0.05(E5);0.1(E6);0.5(E7)		0.02
	C	10	5 (E8); 20(E9)		10
	τ_0	20	10(E10); 30(E11); τ_{max}(E12)		20
	K_{max}	5	2(E13); 10(E14)		5
	P_{adj}	0.05	0.01(E15); 0.1(E16)		0.05
	q_0	0.005	0(E17); 0.01(E18); 0.05(E19); 0.1(E20)		0.1
MOACO-State	Number of ants	1	1(M1); 2(M2); 5(M3); 10(M4)	12	1
	ρ	0.05	0.02(M5); 0.1(M6); 0.5(M7)		0.02
	C	10	5(M8); 20(M9)		10
	τ_0	20	10(M10); 30(M11); τ_{max}(M12)		30
SMPSO	Swarm size	100	100(S1); 50(S2); 200(S3)	8	100
	Archive size	100	50(S4); 200(S5)		100
	$P_{Mutation}$	1/number of decisions	0.005(S6); 0.05(S7); 0.1(S8)		0.005
NSGA-II	$P_{Crossover}$	0.9	0.9(N1); 0.95(N2); 1(N3)	8	0.9
	$P_{Mutation}$	0.005	1/length of string (N4); 0.015(N5); 0.05(N6)		0.005
	Population	100	50(N7); 200(N8)		50
εMOEA	$P_{Crossover}$	1.0	1.0(T1); 0.95(T2); 0.9(T3)	9	1.0
	$P_{Mutation}$	0.005	0.01(T4); 0.015(T5); 0.05(T6)		0.01
	$P_{Inversion}$	0.005	0.01(T7); 0.015(T8); 0.025(T9)		0.005

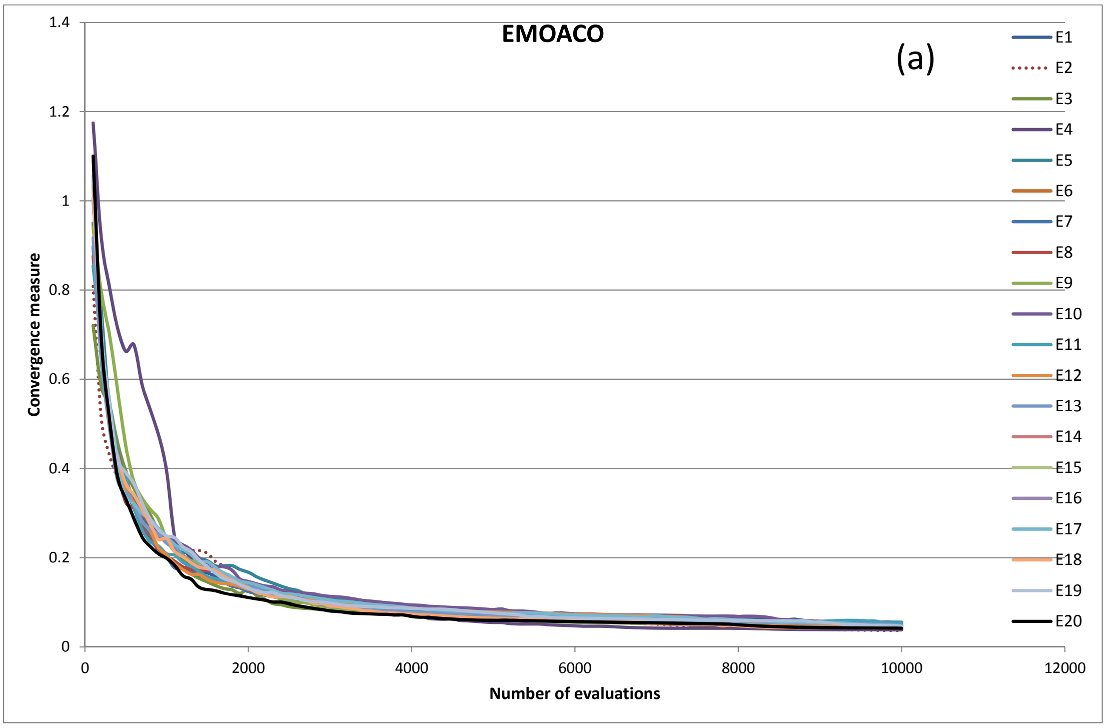

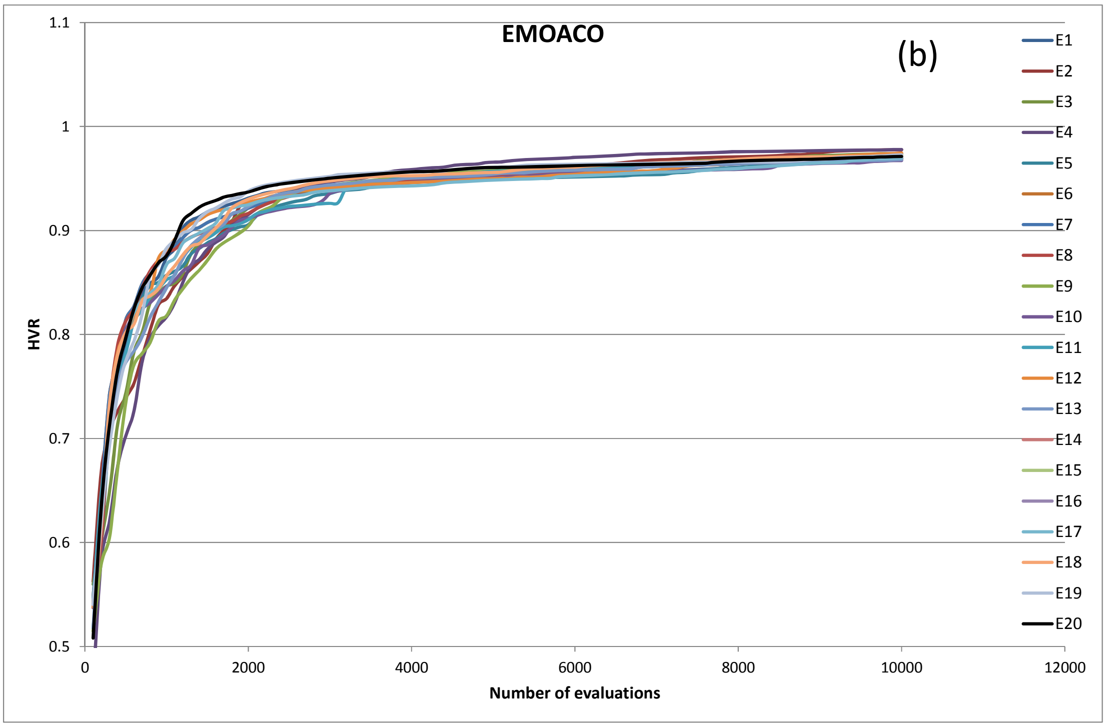

Figure 5-13 Results of EMOACO tuning as a function of number of evaluations for Canberra case study minimizing two objectives: present worth cost and restriction frequency (a) Convergence measure (b) HVR measure

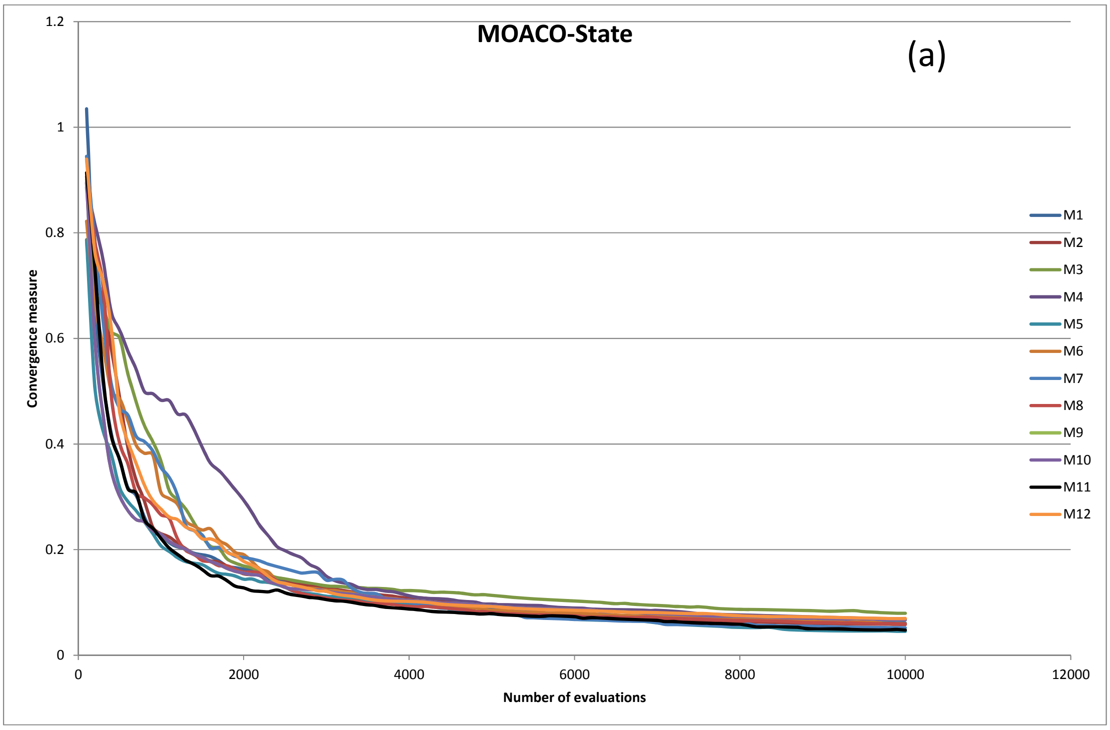

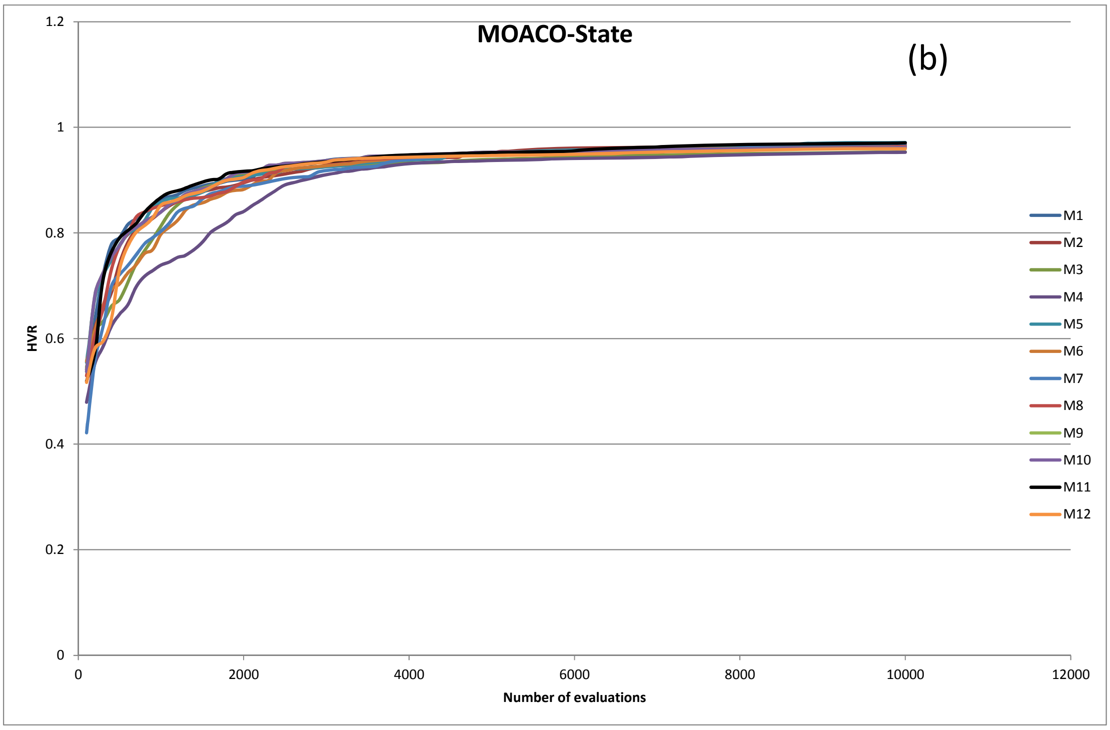

Figure 5-14 Results of MOACO-State tuning as a function of number of evaluations for Canberra case study minimizing two objectives: present worth cost and restriction frequency (a) Convergence measure (b) HVR measure

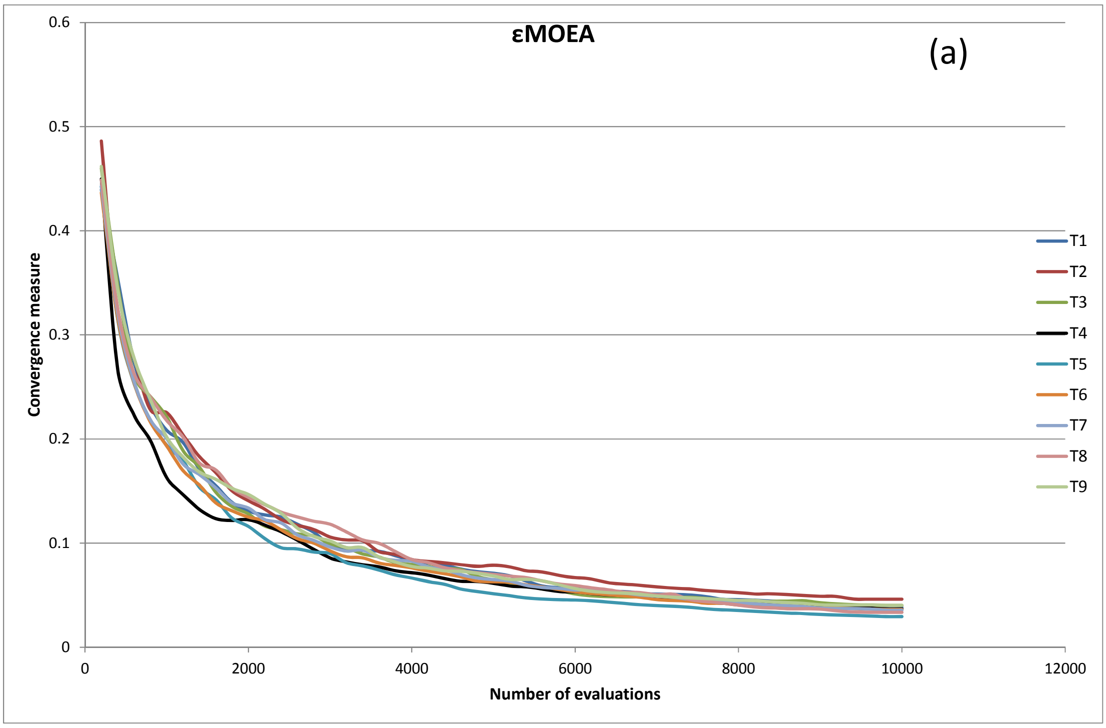

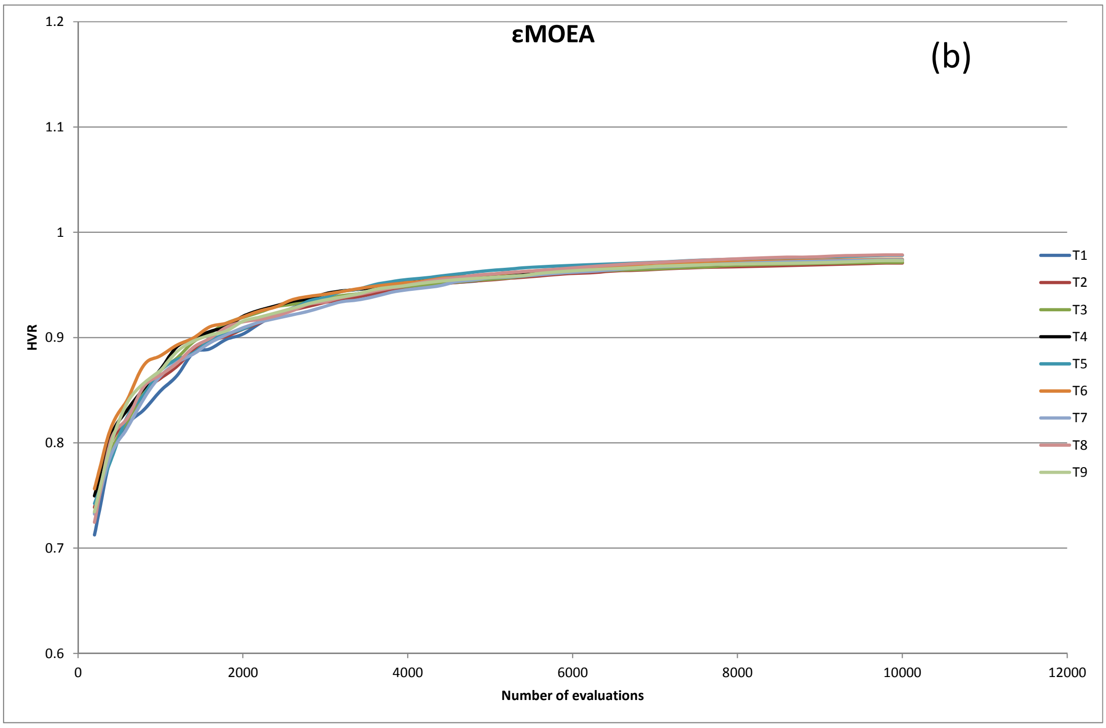

Figure 5-15 Results of εMOEA tuning as a function of number of evaluations for Canberra case study minimizing two objectives: present worth cost and restriction frequency (a) Convergence measure (b) HVR measure

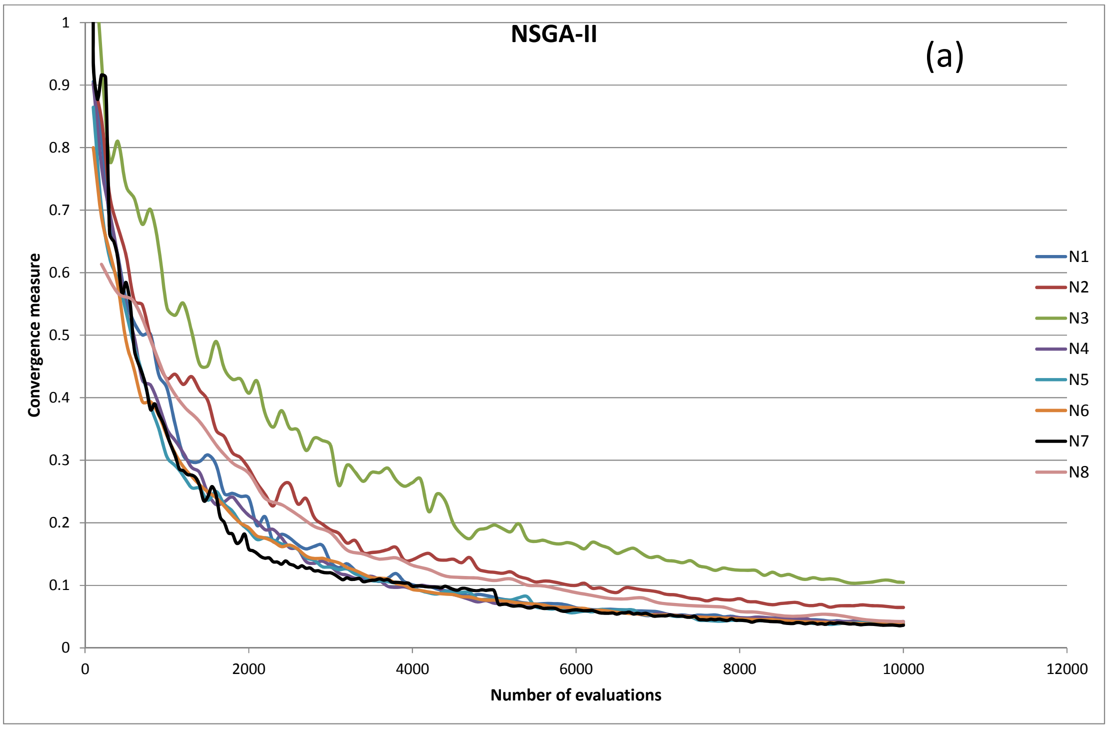

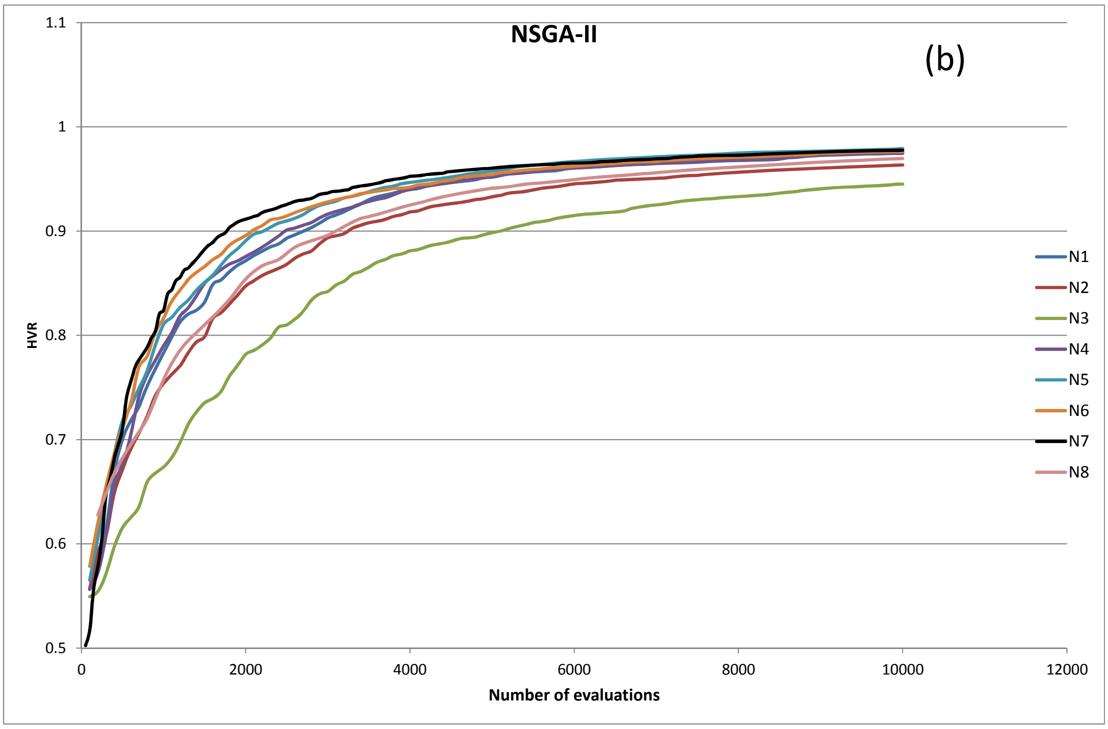

Figure 5-16 Results of NSGA-II tuning as a function of number of evaluations for Canberra case study minimizing two objectives: present worth cost and restriction frequency (a) Convergence measure (b) HVR measure

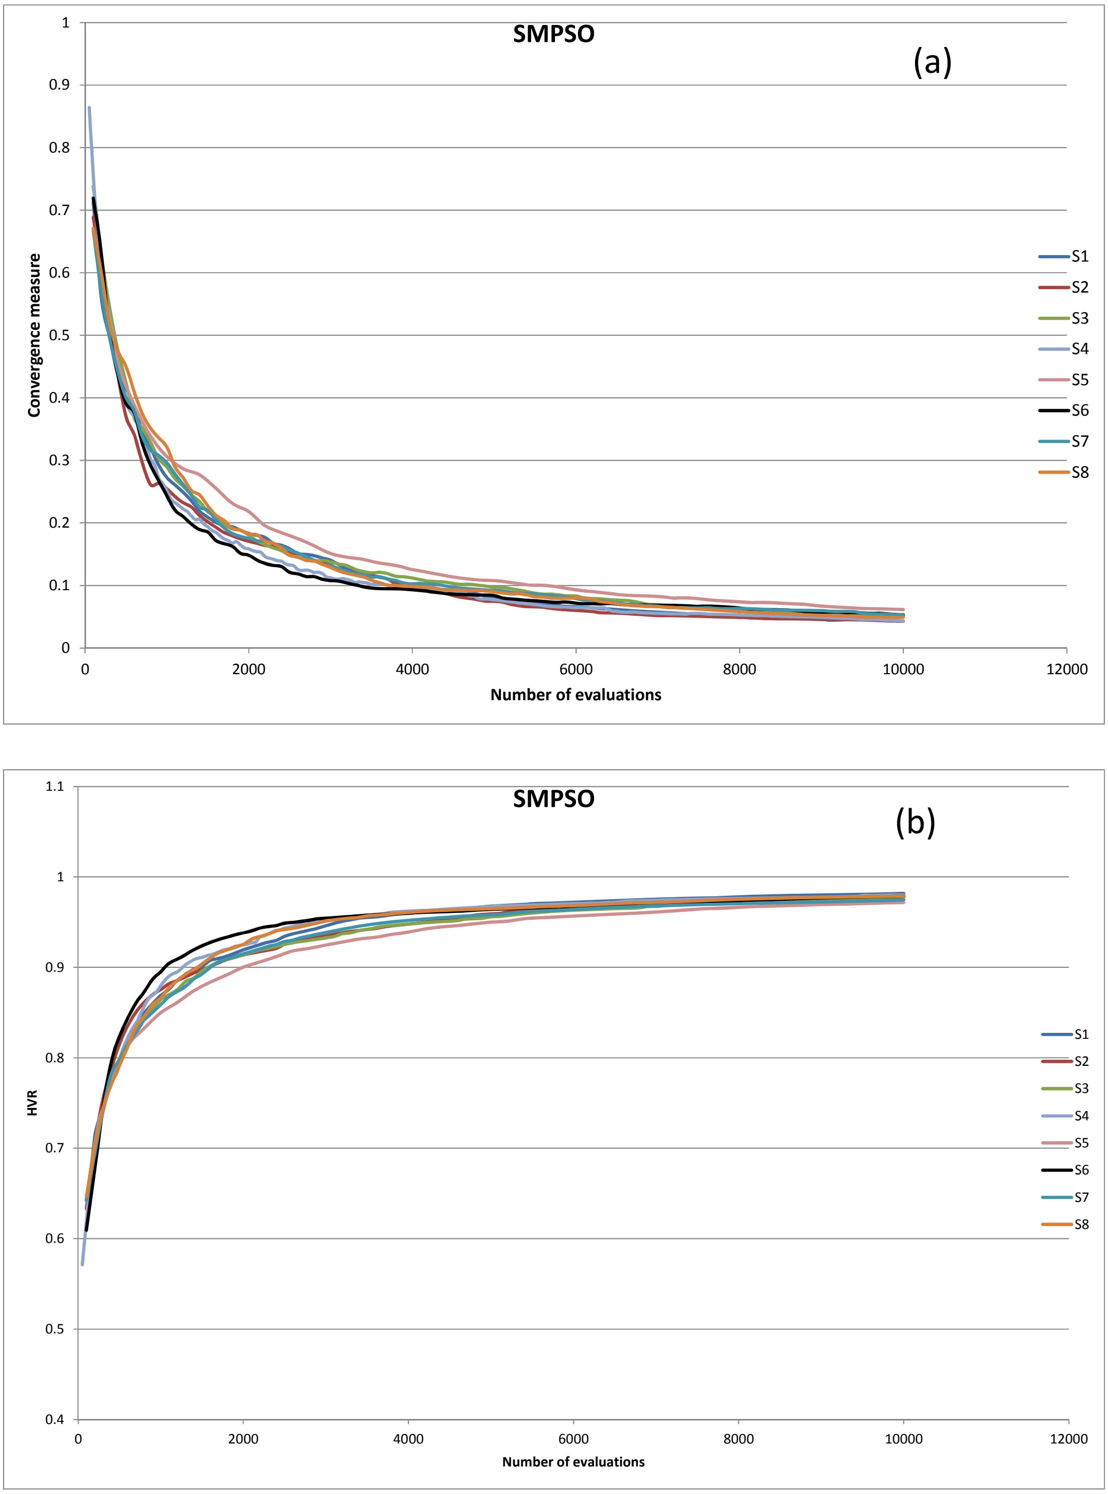

Figure 5-17 Results of SMPSO tuning as a function of number of evaluations for Canberra case study minimizing two objectives: present worth cost and restriction frequency (a) Convergence measure (b) HVR measure

5-6-2 Results and Discussion

In this section, results of six MOO methods are evaluated. All methods applied with tuned parameters were obtained in the last section. The first employed case study is Canberra water supply system. To investigate the performance of these six MOO methods for a more complex system another case study namely Sydney water supply system also is employed. Moreover, since MOO methods are tuned to the Canberra case study, applying these methods in the Sydney case study will test capability of these methods more. Presented measures in this section are the average of obtained measures from 10 runs.

5-6-2-1 Case Study – Canberra Water Supply System: Two objectives

Figure 5-18 shows a plot of the convergence measure for the six MOO methods against a range of function evaluations for the Canberra system minimizing restriction frequency and present worth cost. EMOACO-I unequivocally outperforms the other methods demonstrating very rapid convergence. εMOEA is the best method among non-ACO methods. This graph clearly shows that the ranking of the methods varies as the number of function evaluations changes; for instance, NSGA-II outranks SMPSO and MOACO-State after 2000 evaluations. Figure 5-19 presents a similar plot for the HVR metric. Once again, EMOACO-I outperforms the other methods except when the number of evaluations is 5000. EMOACO and εMOEA are ranked as second best. Again this graph shows changes in rankings as the number of evaluations progress; for example, the rank of SMPSO is four after 5000 evaluations but climbs to two after 10000 evaluations. Figure 5-20 shows a plot of the $I_{\varepsilon+}$ measure for the six MOO methods. This plot shows even more variation in the ranking highlighting the sensitivity of this measure. Except for 1000 evaluations, εMOEA is the best method. Interestingly, all non-ACO methods continue to improve as the number of evaluations increase.

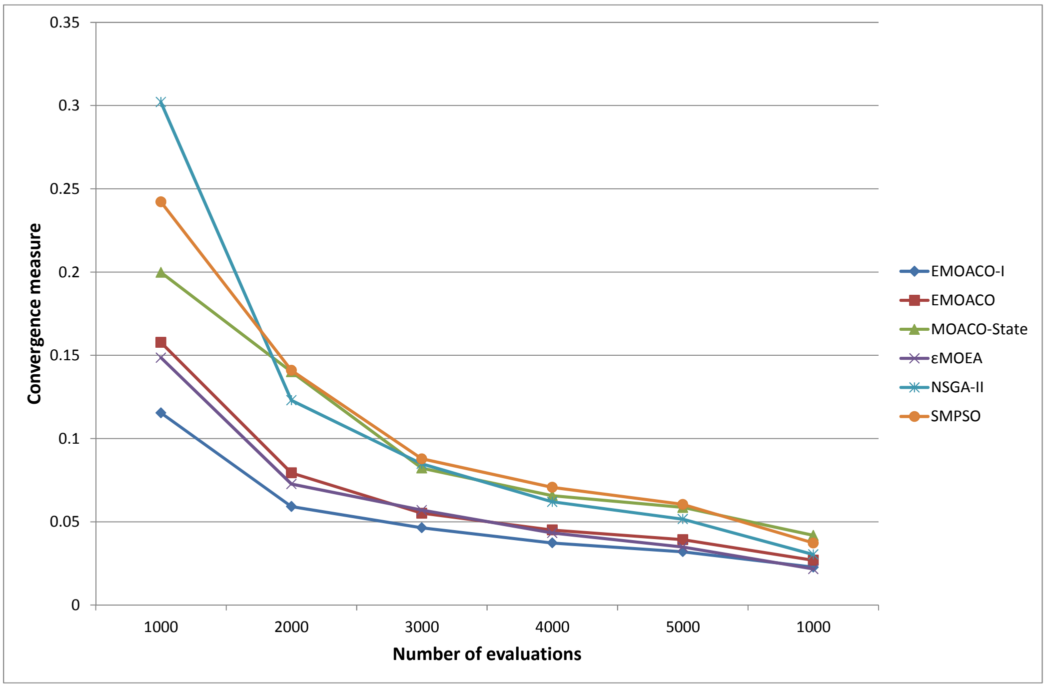

Figure 5-18 Convergence measure for six MOO methods as a function of the number of evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

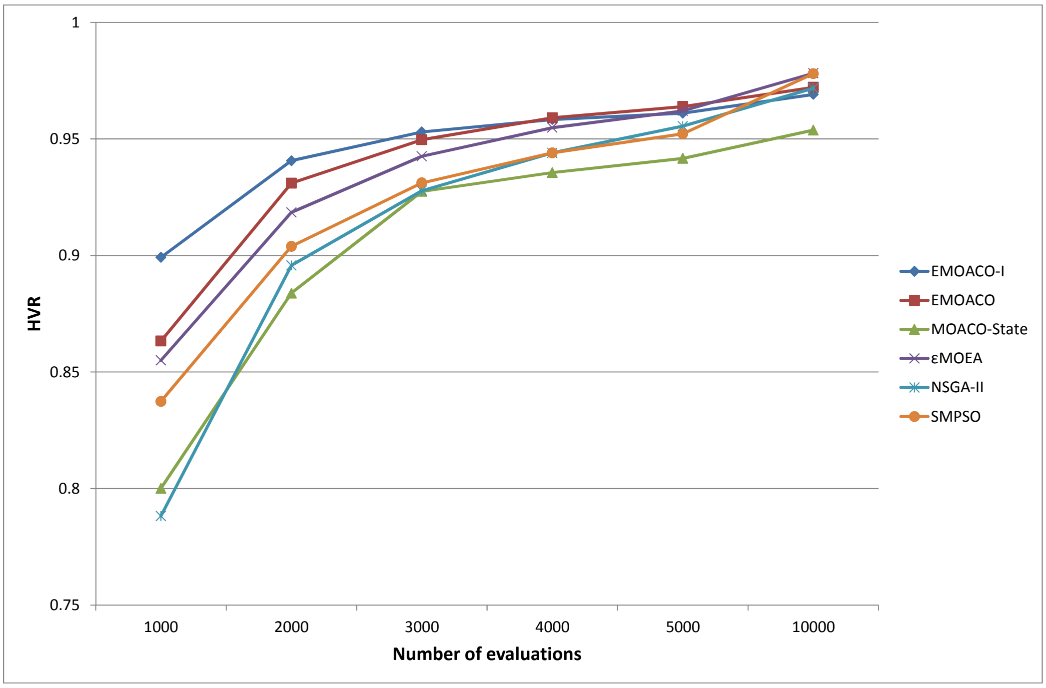

Figure 5-19 HVR measure for six MOO methods as a function of the number of evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

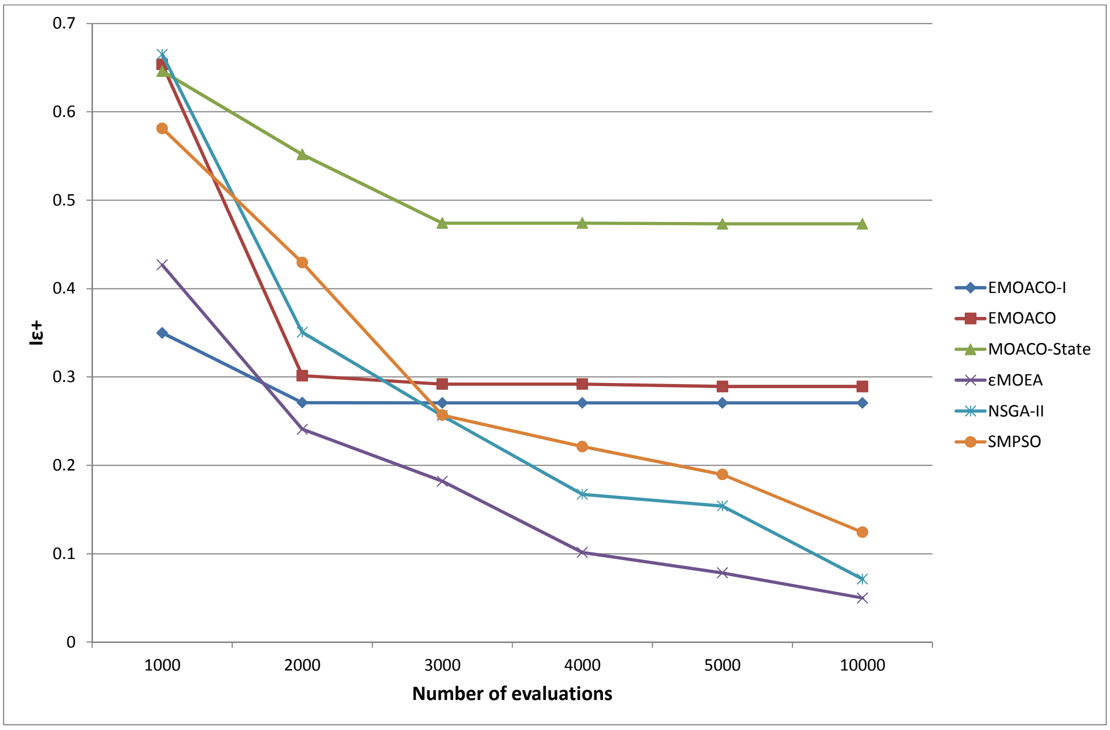

Figure 5-20 $I_{\varepsilon+}$ measure for six MOO methods as a function of the number of evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

To elaborate more on the difference between non-dominated solutions of each method and how the three measures reflect these differences, the results of the six methods are compared in Figure 5-21 for the first run after 1000 evaluations. This figure shows that EMOACO-I is the best method and MOACO-State is the worst in terms of all three measures. The non-dominated solutions obtained for the best and worse methods in terms of convergence measure are presented in Figure 5-22. This plot shows clearly that the EMOACO-I solutions are closer to the TPF. The convergence measures for EMOACO-I and MOACO-State are 0.087 and 0.358 respectively. It is noted that although the difference between the two convergence measures is 0.271, the gap between the two non-dominated solution sets shown in Figure 5-22 is considerable.

In Figure 5-23, the non-dominated solutions for the best and worse methods in terms of HVR measure, EMOACO-I and MOACO-State methods are plotted. In this figure the reference point is also shown. The area enclosed by green solid lines represents the hypervolume of MOACO-State and the area enclosed by blue dashed lines denotes the hypervolume of EMOACO-I. The figure clearly indicates that EMOACO-I has a larger hypervolume, which associates with better convergence and diversity.

In Figure 5-24 two methods that have similar convergence measure value, EMOACO and εMOEA, are presented to demonstrate the difference between $I_{\varepsilon+}$ measures for these methods. As discussed in Section 5-3, the $I_{\varepsilon+}$ measure is sensitive to outlier solutions. In Figure 5-24 the solution for each method that has maximum distance from the TPF is marked by dashed circle. The $I_{\varepsilon+}$ value for εMOEA is larger than for EMOACO because the outlined solution of this method is further from the TPF.

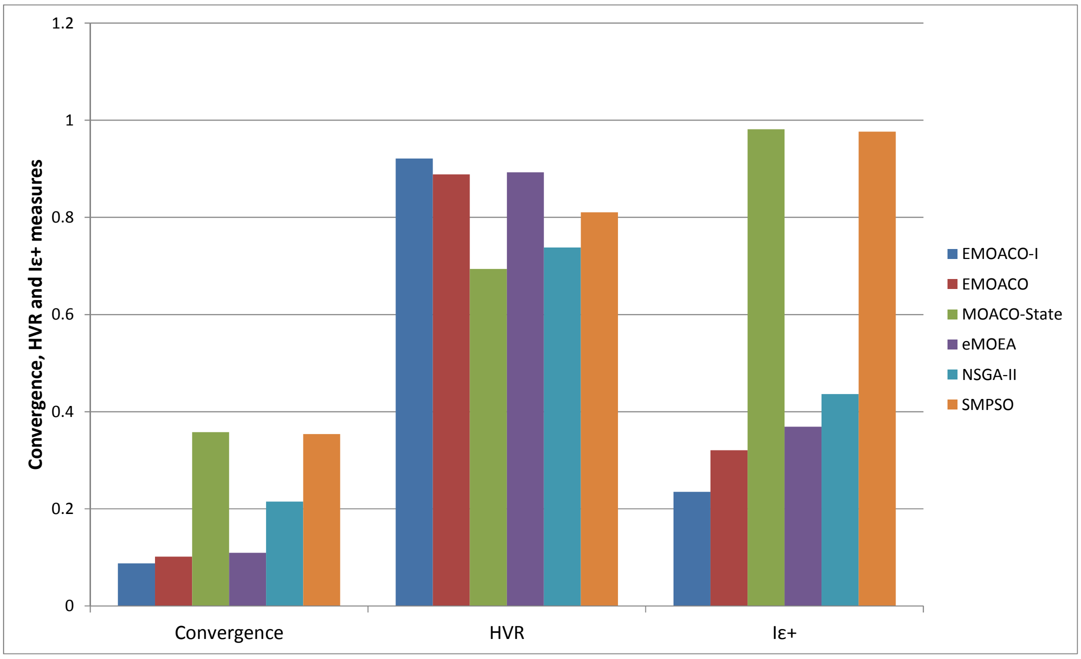

Figure 5-21 Three performance measures for the first run of the six MOO methods after 1000 evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

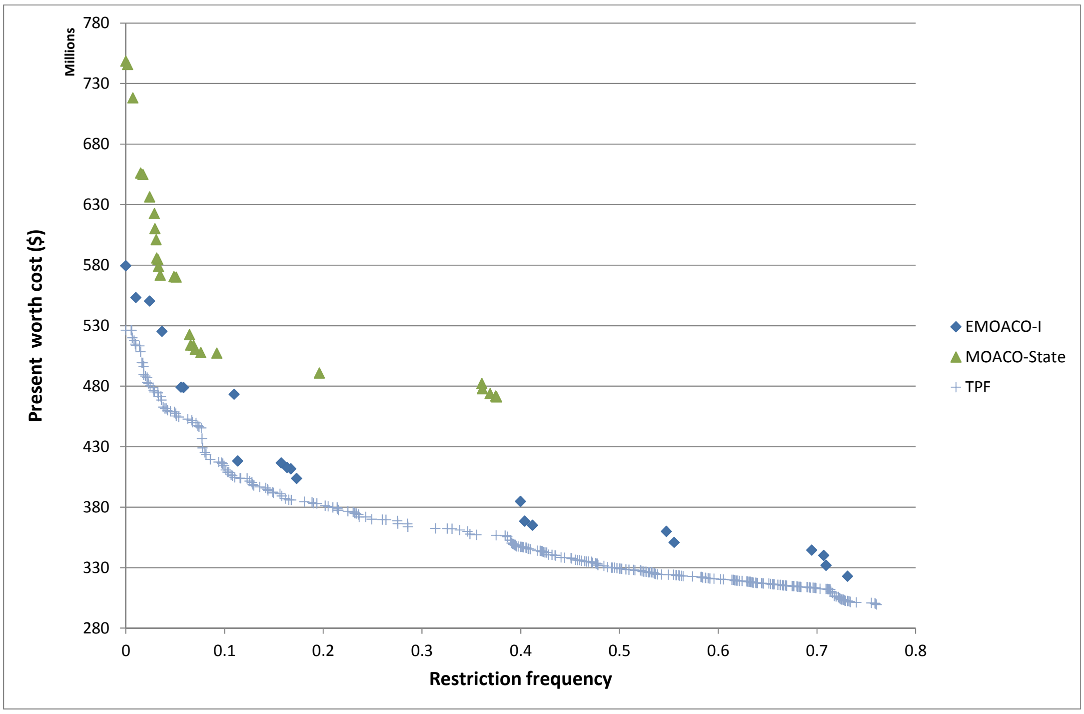

Figure 5-22 Non-dominated solutions of the best and worst methods in terms of convergence measure obtained at the first run of six MOO methods after 1000 evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

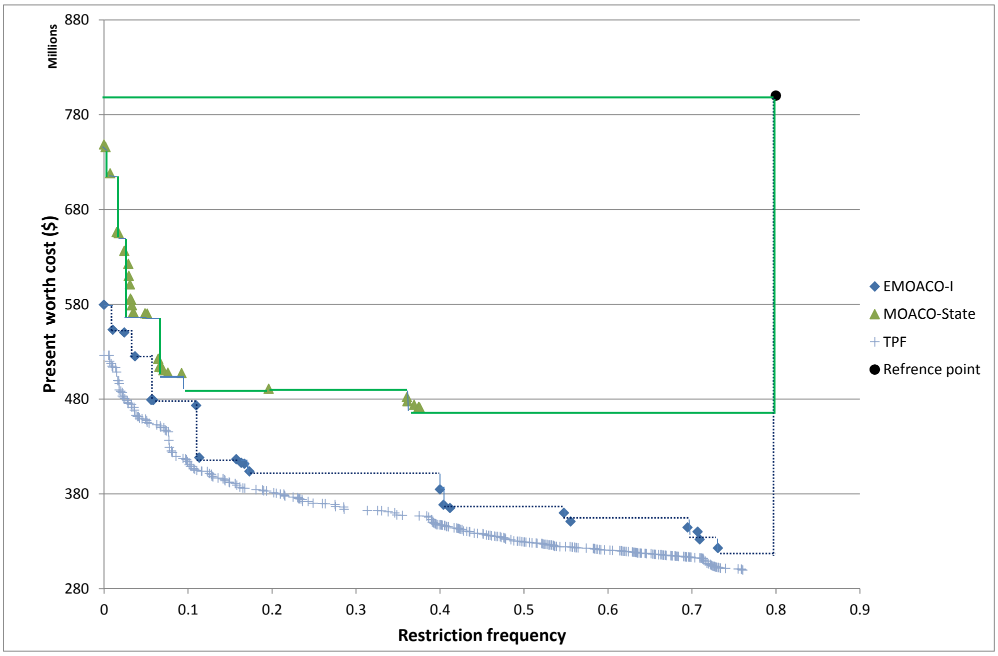

Figure 5-23 Non-dominated solutions of the best and worst methods in terms of HVR measure obtained at the first run of six MOO methods after 1000 evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

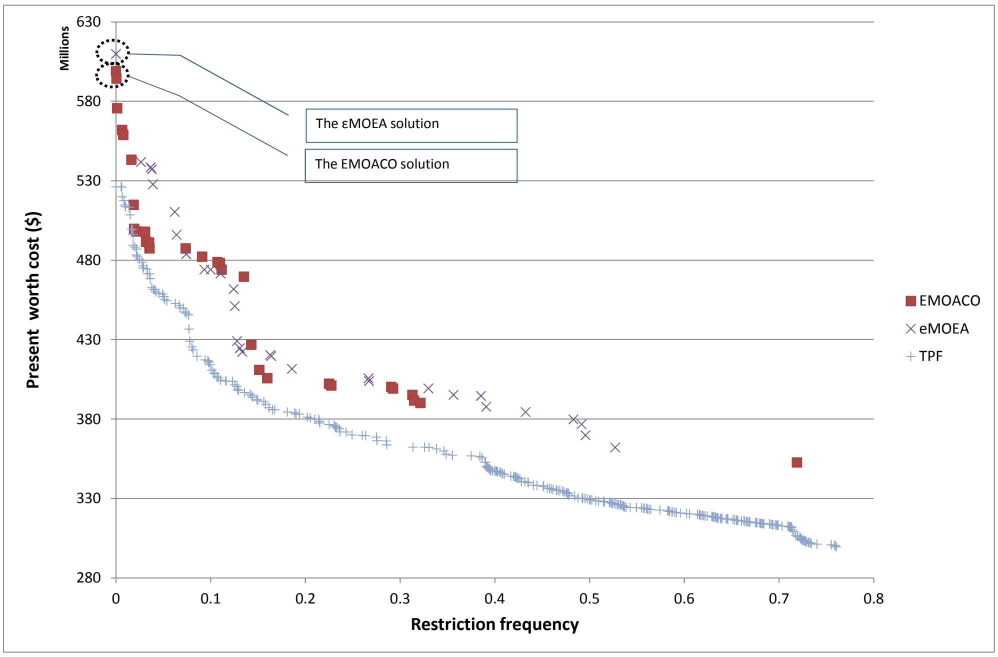

Figure 5-24 Non-dominated solutions of two methods with similar convergence measure value obtained at the first run of six MOO methods after 1000 evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

Figure 5-20 indicated failure of ACO-based methods to improve their performance in terms of $I_{\varepsilon+}$ after 2000 evaluations. To elaborate more on this issue, Figure 5-25 illustrates the ants' exploration of the search space. In this figure the number of evaluations is plotted against the value of the first objective obtained in each evaluation. The pattern in Figure 5-25 shows ants explore the search space locally; that is, they only explore only part of Pareto frontier at any one time. The shortcoming of this is that the ants rarely explore part of the Pareto frontier that was found in the early stages of exploration. In this particular case study, the outlined point in Figure 5-24 was found in the early stages of exploration. This provides insight as to why the $I_{\varepsilon+}$ measure of EMOACO and EMOACO-I fails to improve for evaluations beyond 2000.

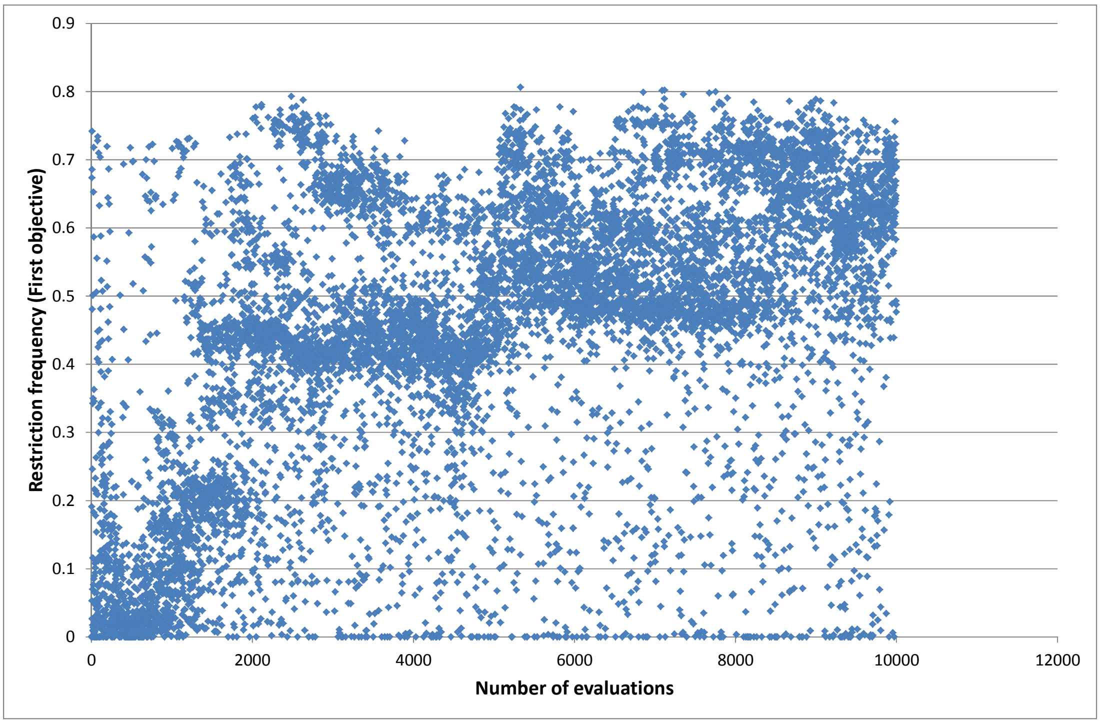

Figure 5-25 Pattern of EMOACO search space exploration for the first run

To demonstrate the variability of the MOO methods, Figure 5-26 presents the best of the 10 Pareto fronts for each method at 1000 and 2000 evaluations. Figure 5-26(a) shows the Pareto fronts after 1000 evaluations. There is a gap, considerable in places, between the NSGA-II, SMPSO and εMOEA Pareto sets and the MOACO-based sets. In Figure 5-26(b) the Pareto fronts are presented for 2,000 evaluations in which it is observed that the Pareto sets have moved closer to the “true” set. However, the NSGA-II and SMPSO fronts appear to be dominated over much of the front by the other methods. To highlight the differences between the methods, the present worth cost at a restriction frequency of 0.1 for 1000 and 2000 evaluations is presented in Table 5-7. There is considerable variation between the methods. For the 1000-evaluations case, the difference between the best and the worst cost is $73.4 million, which is about 17% of the lowest cost of $428.4 million.

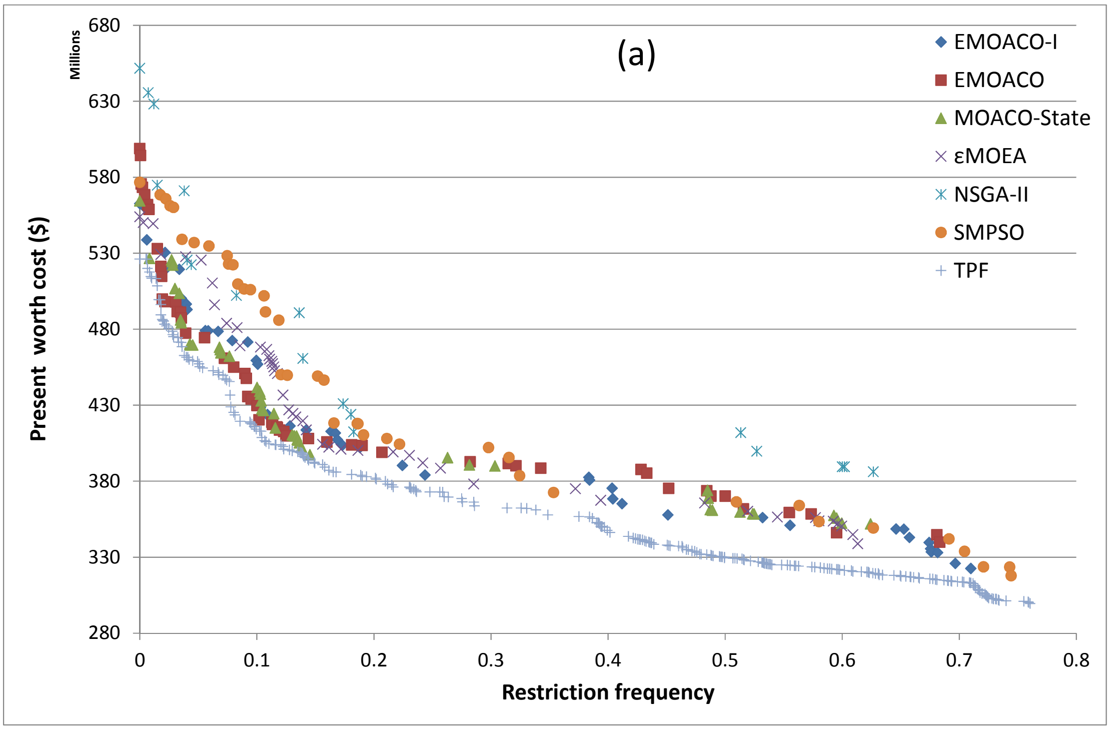

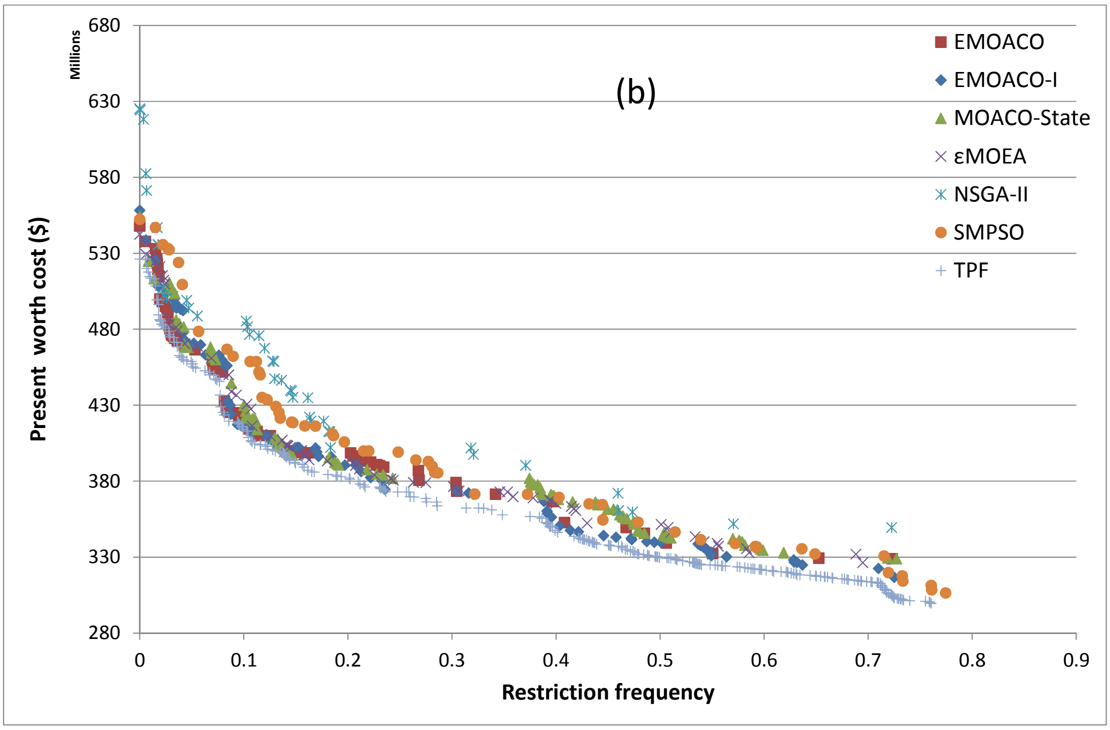

Figure 5-26 Comparison of the "true" Pareto-optimal front (TPF) against the best Pareto sets (out of 10 runs) produced by the six MOO methods for the Canberra case study for (a) 1000 and (b) 2000 evaluations

To demonstrate the value of the initial phase in EMOACO-I, the EMOACO and EMOACO-I Pareto fronts are compared in Figure 5-27 for a range of evaluations for a randomly selected run. Recall that EMOACO uses 256 segments for each decision, while EMOACO-I uses 8 segments for the first 500 evaluations and thereafter 256 segments. It can be seen that EMOACO required at least 1000 evaluations before its Pareto front was in the neighbourhood of the Pareto front produced by EMOACO-I after 500 evaluations. Although these results are for one run, they provide insight into the performance of EMOACO-I. For computationally expensive evaluations, this saving is particularly valuable.

Table 5-7 Best (out of 10 runs) present worth cost ($ million) for a restriction frequency of 0.1 for the six methods

Evaluations	EMOACO-I	EMOACO	MOACO State	εMOEA	NSGA-II	SMPSO
1000	428.4	429.7	441.6	468.2	498.5	501.8
2000	416.7	416.2	426.1	430.7	485.3	459.0

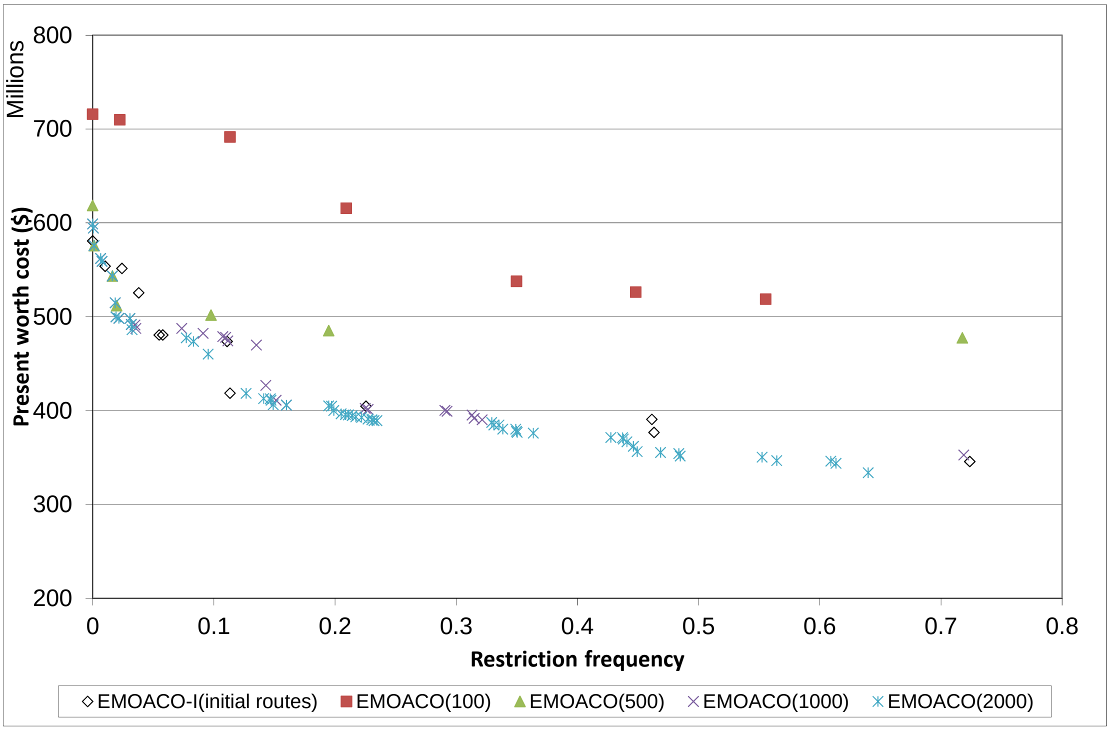

Figure 5-27 Pareto fronts produced by EMOACO and EMOACO-I for a single run for different number of evaluations

The convergence, HVR and $I_{\varepsilon+}$ metrics are based on the average obtained from 10 runs with different random seed numbers. This ensures that comparisons are not significantly affected by sampling variability. However, it is insightful to compare the inter-run performance variability of the six MOO methods. Table 5-8 summarizes the standard deviation for the convergence, HVR and $I_{\varepsilon+}$ measures over 10 runs after completing 1000, 2000 and 3000 evaluations. MOACO-State exhibits the highest variability for all metrics after 1000 and 2000 evaluations. The results confirm the overall superior performance of EMOACO-I. Its convergence, HVR and $I_{\varepsilon+}$ standard deviations are the lowest for 1000 and 2000 evaluations. For 3000 evaluations, EMOACO-I is no longer superior in terms of convergence and HVR measures but remains competitive. These results suggest EMOACO-I is more robust than the other methods in the sense of its performance being less affected by choice of random number seed.

Table 5-8 Convergence, HVR and $I_{\varepsilon+}$standard deviations for 1000, 2000 and 3000 evaluations (the best is shown as bold italics and the worst is underlined)

	MOO methods / Number of evaluations	EMOACO-I	EMOACO	MOACO-State	εMOEA	NSGA-II	SMPSO
Convergence standard deviation	1000	***0.023***	0.057	<u>0.095</u>	0.058	0.116	0.089
	2000	***0.009***	0.024	<u>0.071</u>	0.024	0.037	0.031
	3000	0.011	0.011	0.023	***0.010***	<u>0.031</u>	0.016
HVR standard deviation	1000	***0.011***	0.046	<u>0.085</u>	0.029	0.054	0.029
	2000	***0.012***	0.012	<u>0.044</u>	0.014	0.024	0.017
	3000	0.011	***0.006***	0.012	0.009	<u>0.023</u>	0.011
$I_{\varepsilon+}$ standard deviation	1000	***0.136***	0.228	<u>0.286</u>	0.203	0.215	0.239
	2000	***0.093***	0.140	<u>0.296</u>	0.153	0.174	0.201
	3000	***0.093***	0.145	<u>0.278</u>	0.119	0.188	0.114

5-6-2-2 Case Study – Canberra Water Supply System: Three Objectives

The convergence, HVR and $I_{\varepsilon+}$ metrics for the six MOO methods optimizing the Canberra headworks system with three objectives are presented in Figure 5-28 to 5-30. EMOACO-I has the best convergence metrics and εMOEA has the best convergence among benchmark MOO methods. However, in contrast to Figure 5-18, Figure 5-28 shows a greater variability between methods – even after 10,000 evaluations, there remain significant differences in the convergence measure. The addition of the third objective, to which the MOO parameters were not tuned, has made the optimization more challenging. With regard to the HVR metric, Figure 5-29 shows that EMOACO-I is superior up to 3000 evaluations but is then marginally overtaken by all methods except MOACO-State. The poor HVR but satisfactory convergence performance of MOACO-State is suggestive of its inability to fully explore the decision space and the consequent loss of diversity. Indeed this confirms earlier experience with MOACO-State and the motivation for EMOACO and EMOACO-I. It is noted in Figure 5-30 εMOEA is clearly the best of all methods in terms of $I_{\varepsilon+}$ with EMOACO-I unambiguously ranked second.

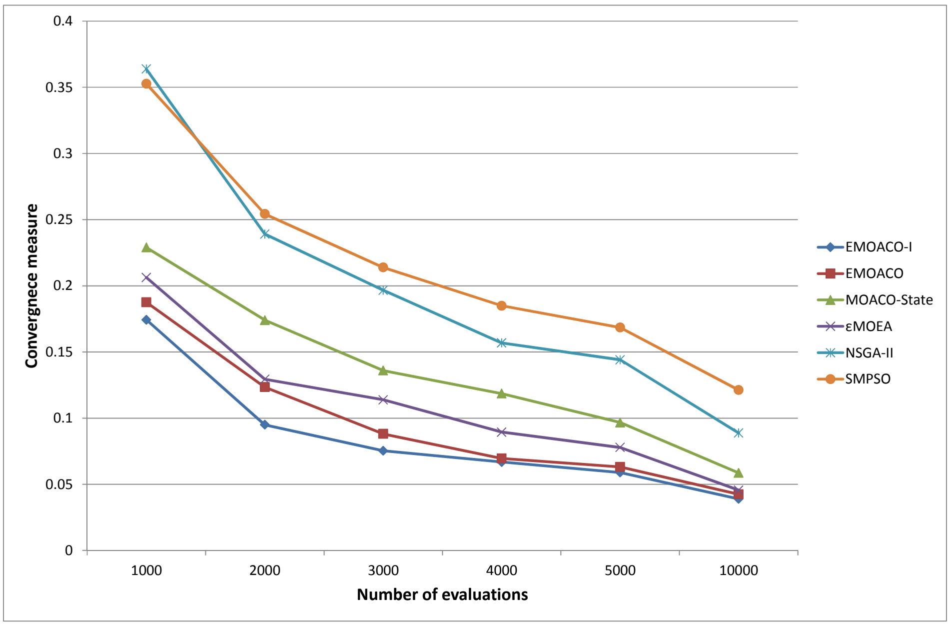

Figure 5-28 Convergence measure for six MOO methods as a function of the number of evaluations for the Canberra case study minimizing three objectives: present worth cost, restriction frequency and time storage less than 20%

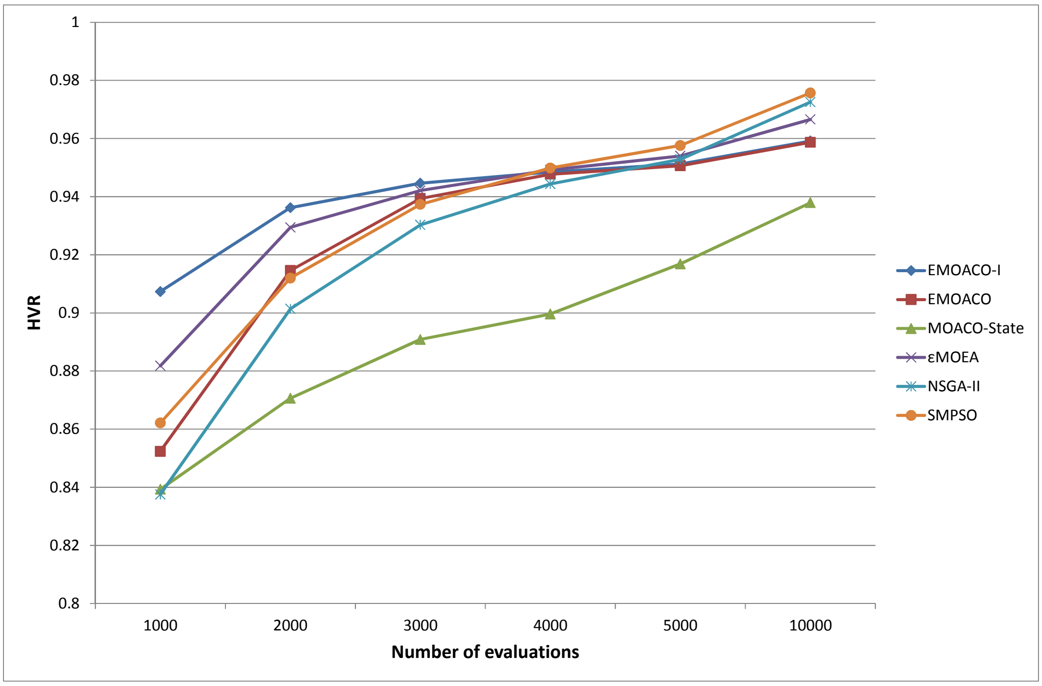

Figure 5-29 Measure for six MOO methods as a function of the number of evaluations for the Canberra case study minimizing three objectives: present worth cost, restriction frequency and time storage less than 20%

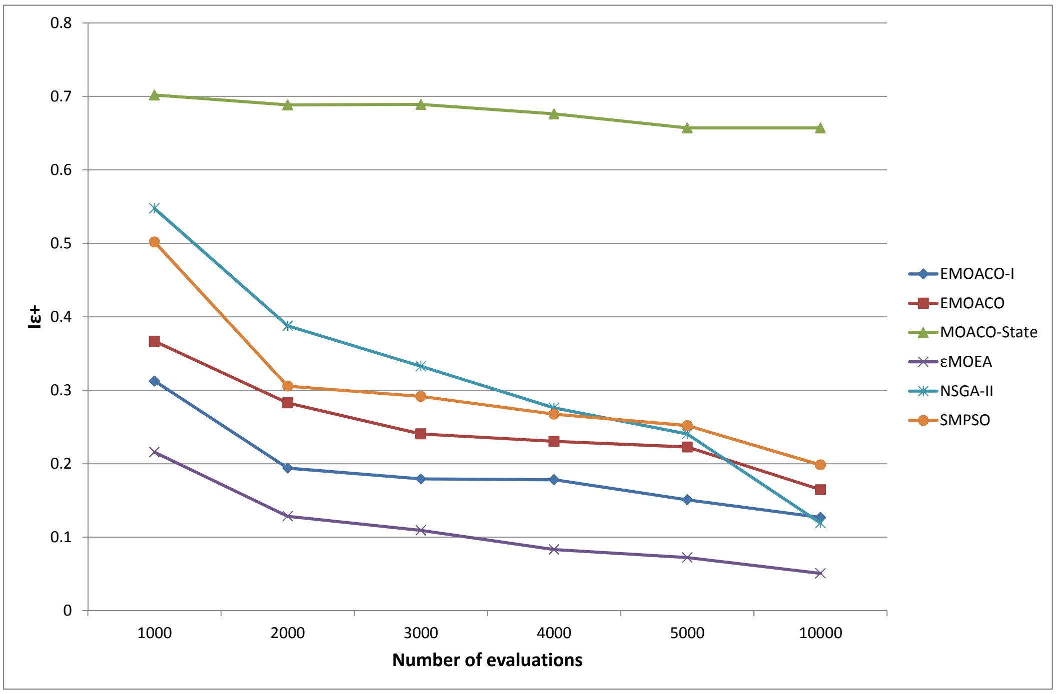

Figure 5-30 $I_{\varepsilon+}$ measure for six MOO methods as a function of the number of evaluations for the Canberra case study minimizing three objectives: present worth cost, restriction frequency and time storage less than 20%

5-6-2-3 Case Study – Sydney Water Supply System: Two and Three Objectives

The MOO parameters, presented in Table 5-6, were tuned to the Canberra system. In the Sydney case study, the performance of the MOO methods is evaluated without any further tuning. All six MOO methods were run 10 times for 10,000 evaluations with different initial random number seeds.

Figure 5-31 through 5-33 present the convergence, HVR and $I_{\varepsilon+}$ measures for the six MOO methods for the Sydney case study minimizing two objectives, present worth cost and restriction frequency. EMOACO-I has the best convergence measure except for 1000 evaluations. With the exception of 1000 evaluations, SMPSO is the best method among benchmark methods in terms of convergence. EMOACO-I outperforms other methods in terms of HVR measure. SMPSO is the best method among benchmark methods except for 1000 evaluations. Although Figure 5-33 indicates good $I_{\varepsilon+}$ performance of εMOEA for 1000 and 2000 evaluations, SMPSO overtakes εMOEA after 2000 evaluations. It also shows EMOACO-I is competitive with these methods with respect to $I_{\varepsilon+}$.

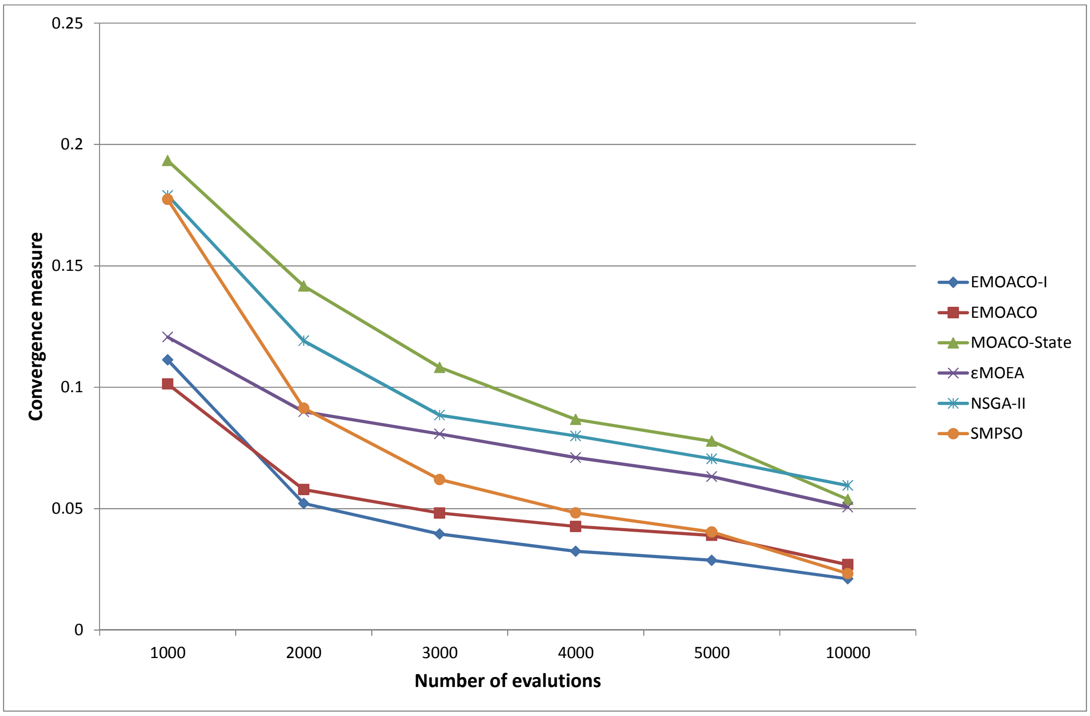

Figure 5-31 Convergence measure for six MOO methods as a function of the number of evaluations for the Sydney case study minimizing two objectives: present worth cost and restriction frequency

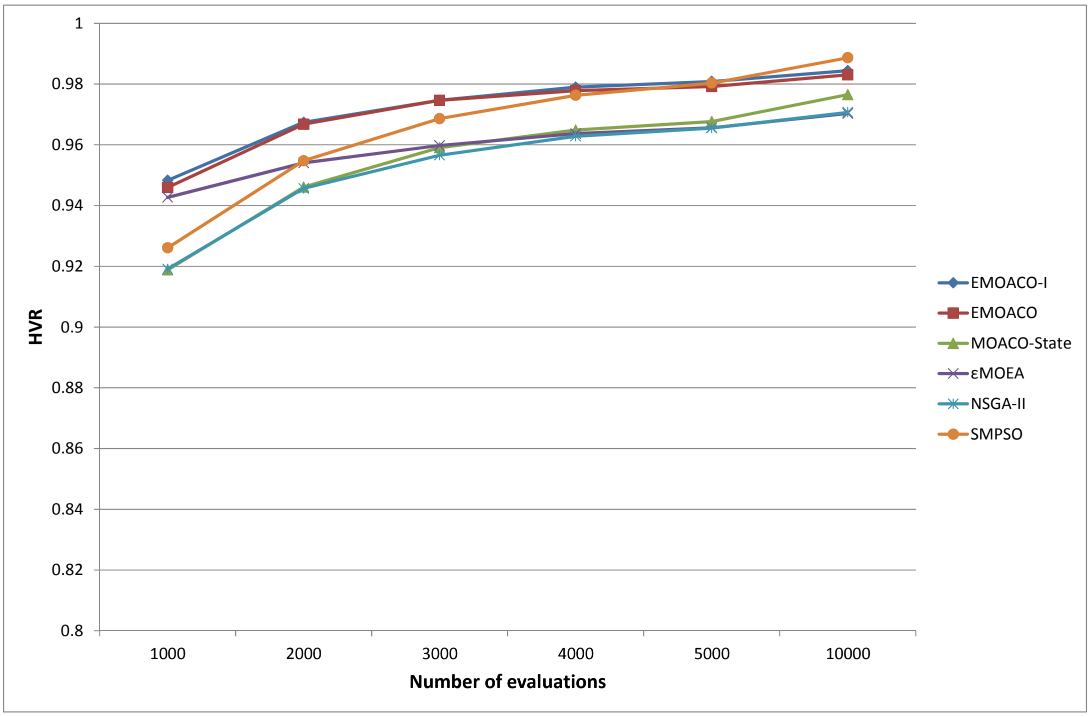

Figure 5-32 HVR measure for six MOO methods as a function of the number of evaluations for the Sydney case study minimizing two objectives: present worth cost and restriction frequency

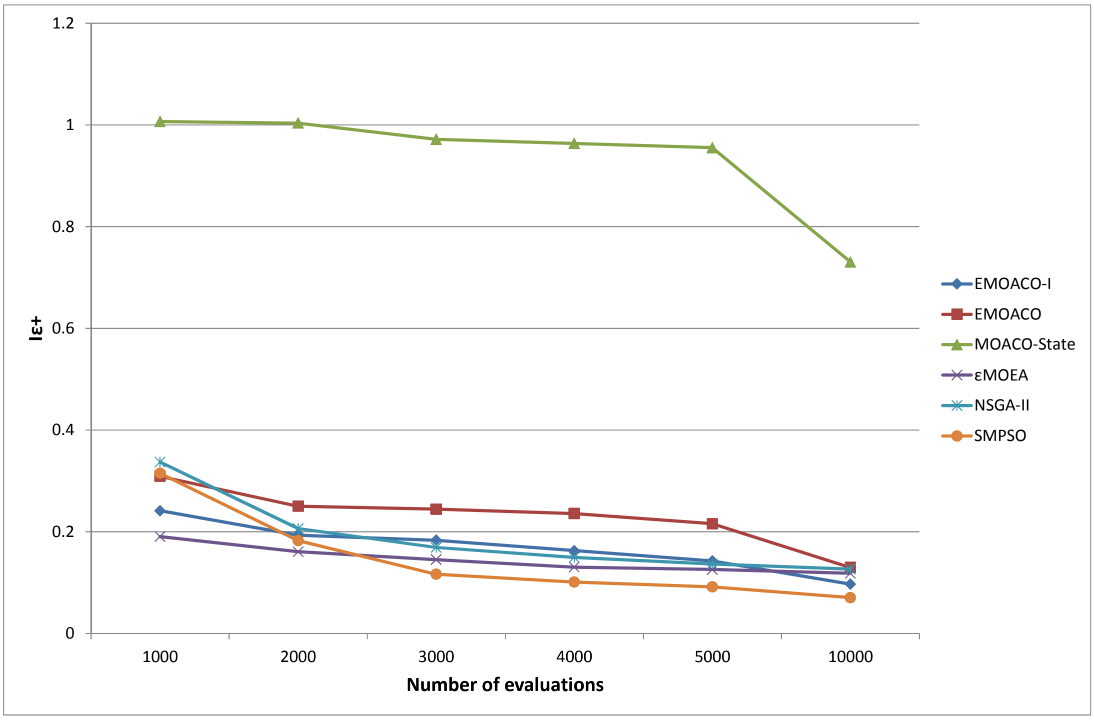

Figure 5-33 $I_{\varepsilon+}$ measure for six MOO methods as a function of the number of evaluations for the Sydney case study minimizing two objectives: present worth cost and restriction frequency

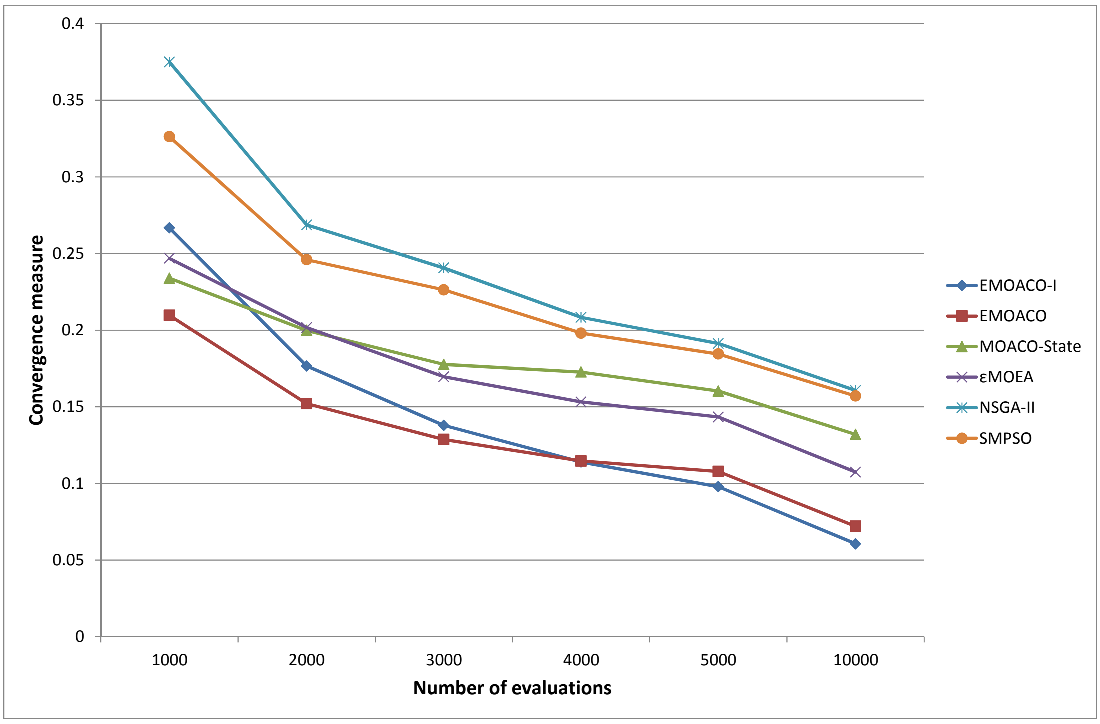

Figure 5-34 Convergence measure for six MOO methods as a function of the number of evaluations for the Sydney case study minimizing three objectives: present worth cost, restriction frequency and environmental stress

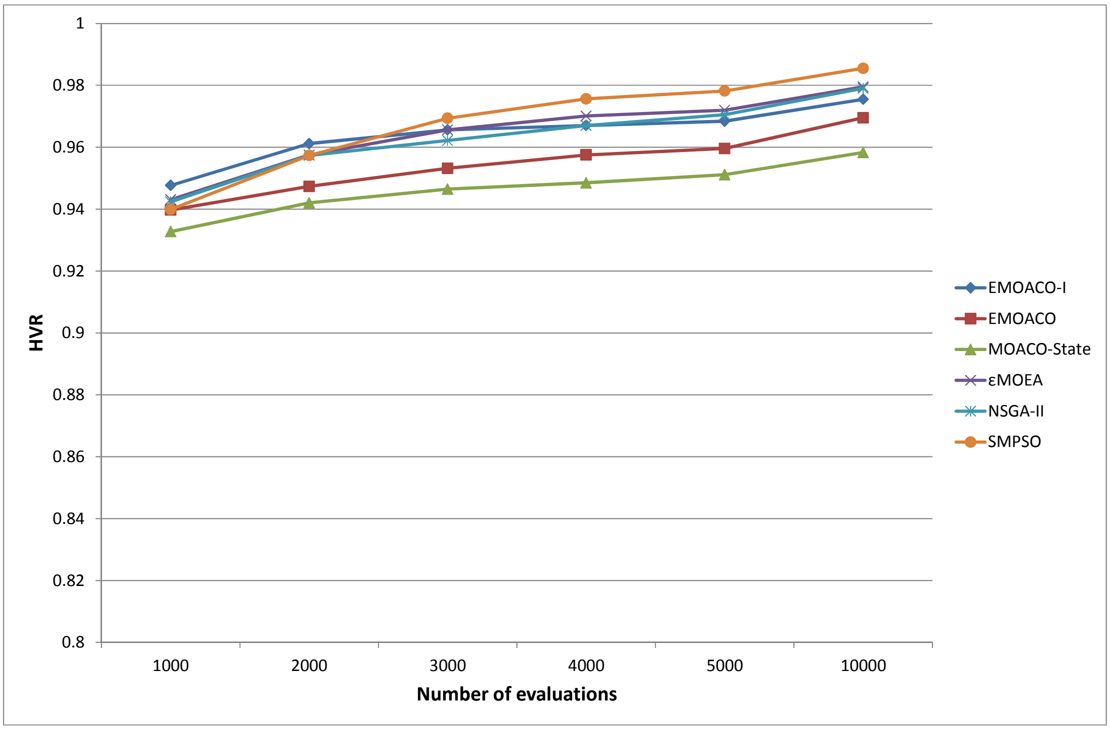

Figure 5-35 HVR measure for six MOO methods as a function of the number of evaluations for the Sydney case study minimizing three objectives: present worth cost, restriction frequency and environmental stress

The Sydney three-objective case study reveals a significantly greater divergence in performance among the methods. Figure 5-34 through 5–36 present the convergence, HVR and $I_{\varepsilon+}$ measures for the six MOO methods for the Sydney case study minimizing three objectives, present worth cost, restriction frequency and environmental stress. EMOACO and EMOACO-I are significantly superior to the other methods for the convergence measure, but EMOACO-I did not perform well after 1000 evaluations. εMOEA is the best of the benchmark methods. For the HVR measure, EMOACO-I is clearly superior up to 2000 evaluations, after which SMPSO and εMOEA marginally overtake it. Figure 5-36 illustrates the $I_{\varepsilon+}$ measure for the six methods. What is striking is the very poor performance of MOACO-State up to 4000 evaluations. The dramatic decrease of MOACO-State $I_{\varepsilon+}$ value after 4000 evaluations, not observed in the other metrics, highlights sensitivity of the $I_{\varepsilon+}$ measure to outlier solutions. As shown in the enlargement, εMOEA is the overall best method in terms of $I_{\varepsilon+}$ with EMOACO-I ranked overall as second at 2000 or more evaluations.

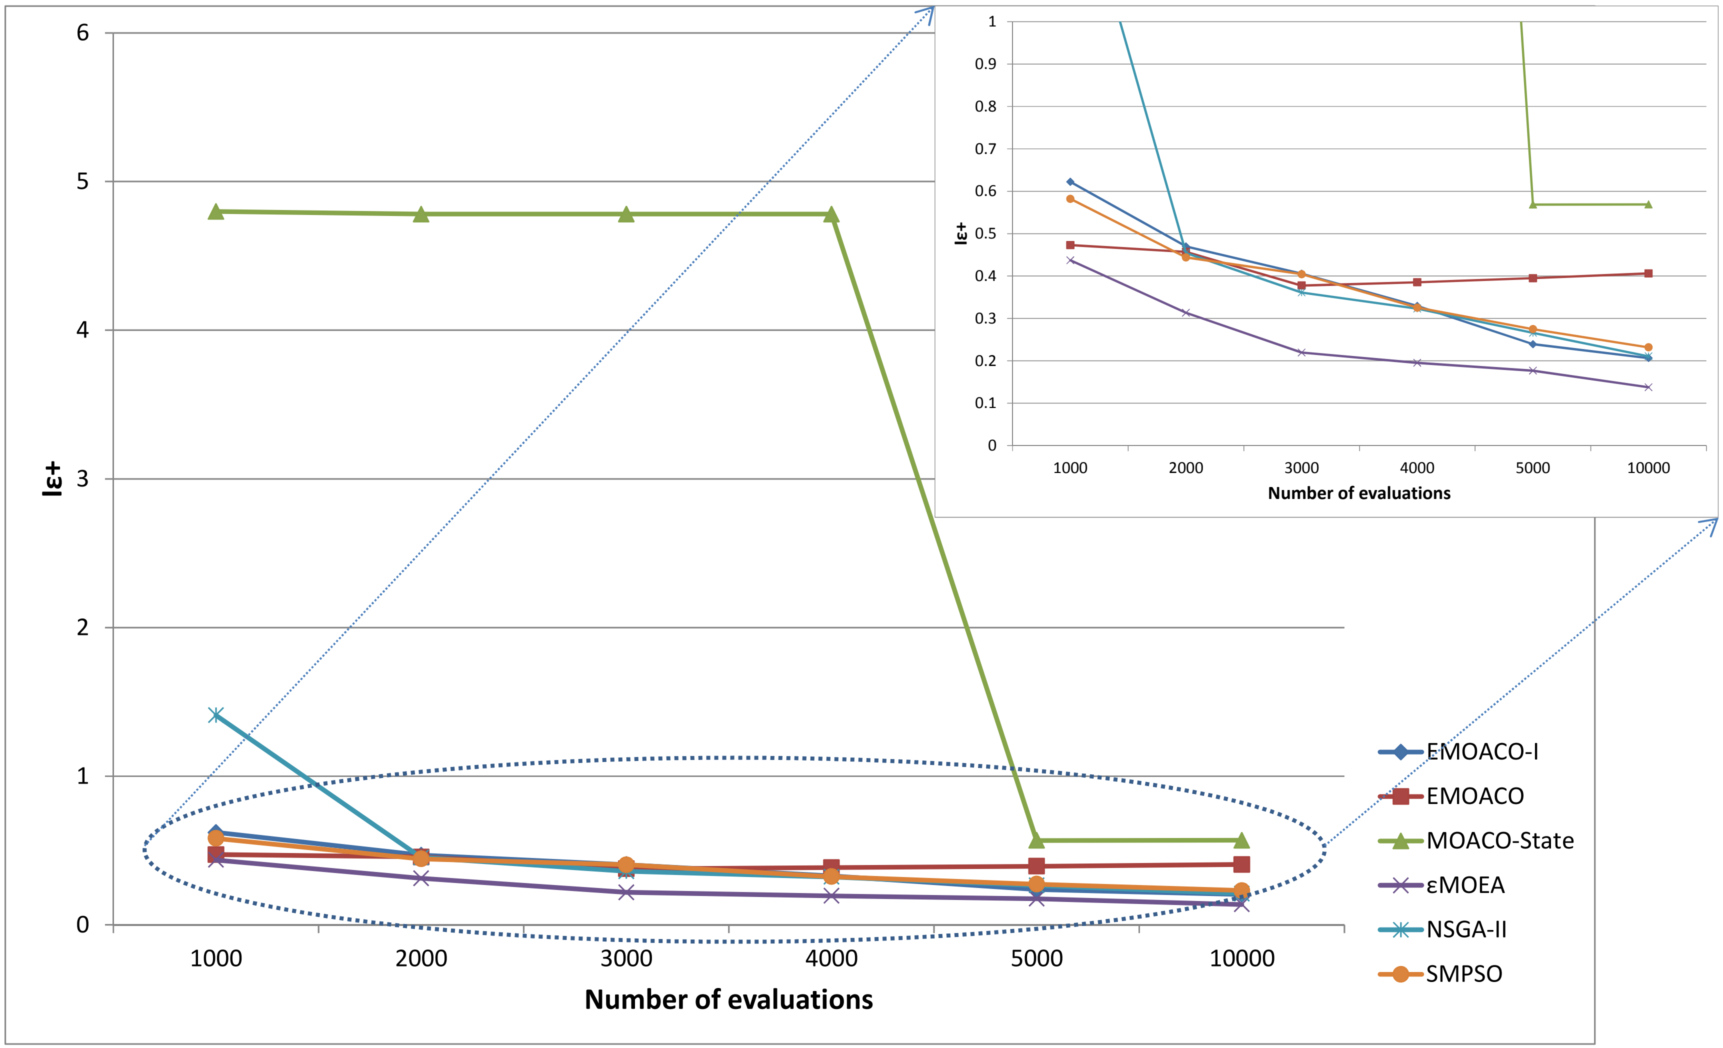

Figure 5-36 $I_{\varepsilon+}$ measure for six MOO methods as a function of the number of evaluations for the Sydney case study minimizing three objectives: present worth cost, restriction frequency and environmental stress

Comparison of all six MOO methods in terms of the three measures for the four case study combinations reveals that no one method unambiguously outperforms the other methods. For instance, while EMOACO-I performed well in terms of convergence and HVR measure for the two-objective Canberra case it did not perform as well in terms of $I_{\varepsilon+}$. Similarly, εMOEA is the best method among benchmark methods for three cases while SMPSO has better convergence performance for the two-objective Sydney case.

The results clearly demonstrate the improvement of EMOACO and EMOACO-I over MOACO-State. Indeed MOACO-State is among the worse methods in almost all cases. It is also noted that NSGA-II was ranked last among the benchmark methods in most of the cases and was particularly poor in the three-objective cases with respect to the convergence measure.

5-7 Conclusions

The optimization of water resource systems in the presence of conflicting objectives necessitates the use of multi-objective optimization methods. Modern day MOO methods based on probabilistic methods require many thousands of objective function evaluations. Unfortunately, these evaluations typically require running simulation models, which for complex water resource systems, can be computationally very expensive. Therefore there is a strong practical need for MOO methods that converge quickly while maintaining diversity along the Pareto front. Recently, a number of studies, mainly using evolutionary algorithms and particle swarm optimization methods, have focussed on MOO methods that converge more quickly. There is a strong practical motivation to identify which of these methods is best suited to urban water management applications.

This chapter approached the task of identifying the best-suited MOO methods for urban water resource applications in three steps:

1) A review of the literature was conducted to shortlist a number of existing MOO methods for detailed evaluation. The criteria for selecting the methods were evidence of good performance, uptake and availability of codes. The review identified three benchmark MOO methods, NSGA-II, εMOEA and SMPSO.

2) The good performance of ACO in finding optimal solutions in single objective problems motivated investigation of its potential in multi-objective optimization. As most multi-objective implementations of ACO have focussed on combinatorial problems, the approach taken was to adapt existing ACO methods to develop a MOACO algorithm suitable for urban water resource applications.

   A review of the MOACO literature identified a number of shortcomings, the main one being their problem-specific implementation. Based on this review, an algorithm called MOACO-State method was developed incorporating the best features of existing ACO methods whilst avoiding their shortcomings. Important features of MOACO-State include the use of a single ant colony with one pheromone matrix, a pheromone updating process independent of the number of objectives and the scale of objective function values, and the elimination of heuristic problem-specific information. However, it was found that MOACO-State did not perform better than existing benchmark methods and was prone to stagnation or premature convergence.

   To improve MOACO-State's performance, two concepts borrowed from evolutionary search methods, namely adjacency and random selection, were implemented in the ACO framework. Adjacency exploits the proposition that potentially good solutions lie in the neighbourhood of current non-dominated solutions. Random selection allows ants to visit routes with low pheromone. The inclusion of these features in MOACO-State led to a new method called EMOACO. Furthermore, the use of a simple heuristic to reduce the number of decision combinations in the initial phase of EMOACO was added to accelerate initial convergence. This method was called EMOACO-I.

3) To identify the best existing MOO methods for urban water resource applications and to assess the performance of newly developed MOACO methods, a performance comparison was conducted using two case studies based on the urban headworks systems serving the Australian cities of Canberra and Sydney. Each case study considered a two- and three-objective optimization problem with about a dozen decision variables affecting infrastructure investment and system operation. Three performance metrics were used to evaluate performance: i) the convergence metric to assess proximity; ii) the hypervolume ration (HVR) metric

to assess proximity and diversity; and iii) the $I_{\varepsilon+}$ measure to assess consistency. The comparison was conducted for function evaluations in the range 1,000 to 10,000.

For the non-MOACO methods, εMOEA and SMPSO had comparable performance with NSGA-II ranked behind them. It was found that in most cases EMOACO-I was the best performing method in terms of convergence and HVR with EMOACO ranked second. With respect to the $I_{\varepsilon+}$ metric EMOACO-I was the best in one case and competitive in the other cases. It was observed that the $I_{\varepsilon+}$ metric was the most sensitive of the metrics, primarily because it focuses on outliers on the non-dominated solution set.

Overall none of six MOO methods was superior in terms of all measures and for all case study problems. However, it was clear that MOACO-State was the worst performing method, a finding which vindicated the enhancements leading to the EMOACO-I algorithm. Of particular interest was the greater variability in the performance of the MOO methods when moving from two to three objective problems and from the Canberra case study, for which MOO parameters were tuned, to the more complex Sydney case study, for which the MOO parameters were not tuned.

Appendix A
Fusing the Best Features of Existing MOACO Methods

A-1 Introduction

This appendix evaluates the performance of sixteen MOACO algorithms (or variants) which are based on what appear to be the best features identified in the review of existing MOACO methods. These variants enable systematic investigation of the benefits and shortcomings of using single or multiple pheromone matrices in transition rules and of using the different approaches in updating pheromone.

Table A-1 summarizes the fifteen MOACO variants (M1 to M15) plus the best performing variant called MOACO-State. These variants are used to systematically trial the algorithms that control the transition rule and pheromone updating. Three methods for integrating multiple pheromone matrices into the pheromone transition rule are considered: 1) weighted sum 2) weighted product and 3) random. Pheromone updating involves two steps, determining the amount of pheromone to be deposited and selection of the routes to be updated. Two methods are considered that determine the amount of pheromone to be deposited: 1) scaled objective function value; and 2) inverse of objective function values. Three methods are considered that determine which routes are to be updated with pheromone: 1) select the non-dominated solutions in the current iteration; 2) select the non-dominated solutions found in all iterations; and 3) select the best routes according to each objective.

The transition rules based on weighted sum and weighted product are given by Eqs. (5.17) and (5.18). However, since no heuristic information is used, these two equations can be simplified as follows:

$$p_{ij} = \frac{\sum_{l=1}^{L} (w_l \tau_{ij}^l)^\alpha}{\sum_{i=1}^{N} \sum_{l=1}^{L} (w_l \tau_{ij}^l)^\alpha} \tag{A-1}$$

$$p_{ij} = \frac{\prod_{l=1}^{L} (\tau_{ij}^l)^{\alpha w_l}}{\sum_{i=1}^{N} \prod_{l=1}^{L} (\tau_{ij}^l)^{\alpha w_l}} \tag{A-2}$$

The random transition rule is based on Eq. (A-1) but w_l must be either 0 or 1 and $\sum_{l=1}^{l} w_l = 1$.

A number of MOACO parameters were common to all the variants. The initial pheromone (τ_0) should bet set at a moderate value. If τ_0 is too large it will take a long time before ants explore good solutions, and if it is too small, search performance is sensitive to early search outcomes (Dorigo and Stützle, 2004). Based on literature values, τ_0 is set to 20. ρ and α are set to 0.02 and 1, respectively, which are the values recommended by (Dorigo and Stützle, 2004). P_{best} is set to 0.05 (Stützle and Hoos, 1996). The number of decision segments and number of ants were set to 256 and 10 respectively. The pheromone deposition constant C used in Eq. (5-23) was set to 10.

Table A-1 Summary of MOACO variants

Variant	Transition rule	Routes on which pheromone is deposited	Amount of pheromone deposited
M1	Random	Non-dominated solutions found so far	Scaled objectives
M2	Random	Non-dominated solutions within an iteration	Scaled objectives
M3	Random	Best of objectives	Scaled objectives
M4	Weighted product	Non-dominated solutions found so far	Scaled objectives
M5	Weighted product	Non-dominated solutions within an iteration	Scaled objectives
M6	Weighted product	Best of objectives	Scaled objectives
M7	Weighted sum	Non-dominated solutions found so far	Scaled objectives
M8	Weighted sum	Non-dominated solutions within an iteration	Scaled objectives
M9	Weighted sum	Best of objectives	Scaled objectives
M10	Random	Best of objectives	Inverse of objectives
M11	Weighted product	Best of objectives	Inverse of objectives
M12	Weighted sum	Best of objectives	Inverse of objectives
M13	Random	Non-dominated solutions found so far	Inverse of objectives
M14	Weighted product	Non-dominated solutions found so far	Inverse of objectives
M15	Weighted sum	Non-dominated solutions found so far	Inverse of objectives
MOACO-State	*Single Pheromone using Eq. (5.11)*	*Non-dominated solutions found so far*	*Constant with aging factor (Eq. (5.23))*

A-2 Results

In this section the performance of the fifteen variants M1 to M15 is summarized for the six benchmark problems and for the Canberra case study with two and three objectives. For each variant, 25,000 and 10,000 evaluations were performed on the benchmark problems and Canberra case studies respectively. The primary goal of any MOO method is to produce a diverse non-dominated solution set close to the true Pareto front. Given that perspective and given the considerable number of variants to be compared, it was decided in the interest of clarity to adopt a single performance metric. HVR was chosen because, of the three metrics considered in this study, it is the only one that evaluates both proximity and diversity. The goal of this approach is to screen out the poorly performing variants and to identify the most promising one for inclusion in the more rigorous assessment in Section 5-6.

As already noted, the variants are constructed to enable investigation of three key aspects of MOACO methods. These deal with the selection of routes on which pheromone is to be deposited, the amount of pheromone deposited on routes and the number of pheromone matrices. In what follows, the variants M1 to M12 are first compared to assess the impact on performance of route selection and the amount of deposited. Then M1, M4 and M7 are compared against M13, M14 and M15 to investigate further the role of the amount of pheromone deposited on routes. Finally to assess the value of using single and multiple pheromone matrices, the M1, M4 and M7 variants are compared with MOACO-State.

In Figures A–1 to A–10 the hypervolume measure for M1 to M12 is presented as a function of the number of evaluations for the eight benchmark problems and the two- and three-objective Canberra case studies. In all cases, M1, M4 and M7 are found to be the best of the 12 variants. The feature common to these three variants is that pheromone is updated on the non-dominated solutions found so-far. This finding is in agreement with Bui et al. (2008). Indeed, although M1, M4 and M7 have different transition rules, there is little difference in their hypervolume measures. This suggests that the pheromone updating algorithm may be more important than the transition rule.

In some problems, the M2, M5 and M8 variants were ranked consistently as second best. Common to these variants is that pheromone is updated on the non-dominated solutions within an iteration. One of the drawbacks of this approach is that non-dominated solutions within an iteration can be potentially far from the non-dominated solutions found so-far. Moreover, the non-dominated solutions found in an iteration can vary considerable between iterations. This may result in a more diffuse, less structured exploration of the search space.

All of the variants that used the best-of-objective option to update pheromone performed very poorly in the Canberra and ZDT4, ZDT6 and DTLZ6 problems. For instance, for ZDT4, these variants reached stagnation after 5,000 evaluations, while for the two- and three-objective Canberra problem, they stagnated after 2,000 evaluations. The main drawback of this updating approach is that because only the routes associated with the best-of-objective routes are updated at each iteration, only a limited number of routes is updated. For instance, in the case of two objectives only two routes are updated. It is likely that the best-of-objective values and their associated routes do not change over a number of iterations. This can quickly lead to stagnation. In cases where the best-of-objective values vary among iterations this approach performs better as in the case of ZDT1. When there are more than two objectives it is likely there will be more than one route identified as having the best-of-objective result for a particular objective. In such cases it is necessary to decide which routes should be updated. This issue has not been addressed in the literature. In this study, only one of the routes was randomly selected and updated.

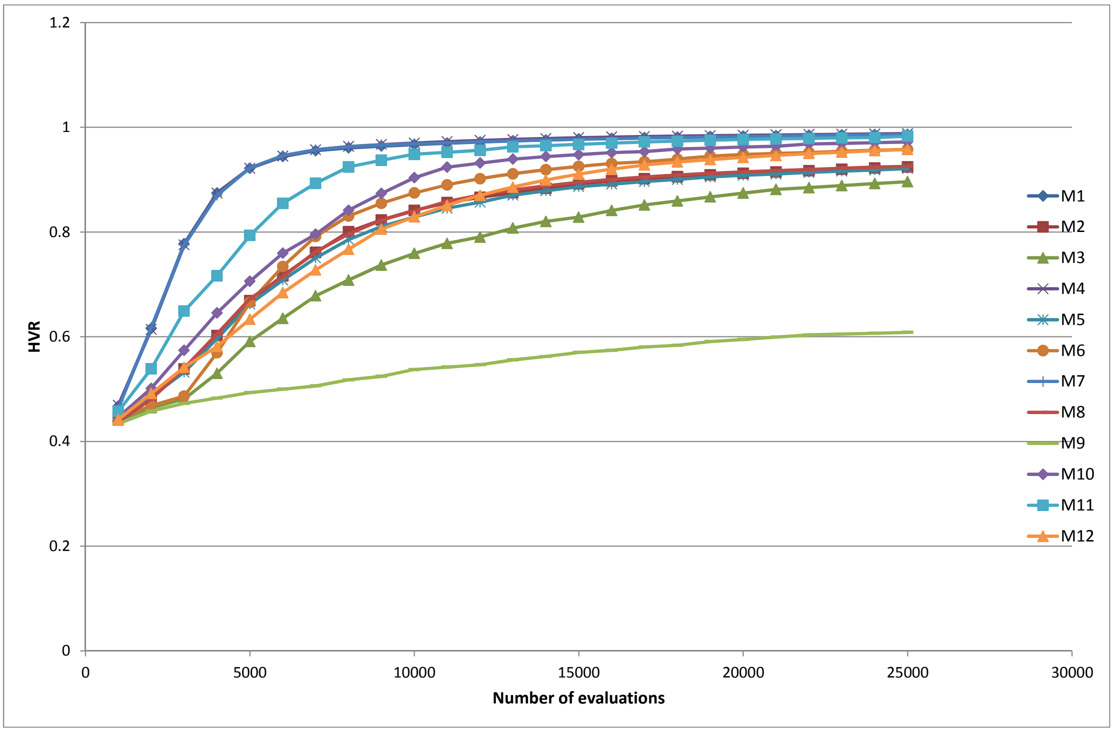

Figure A-1 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem ZDT1

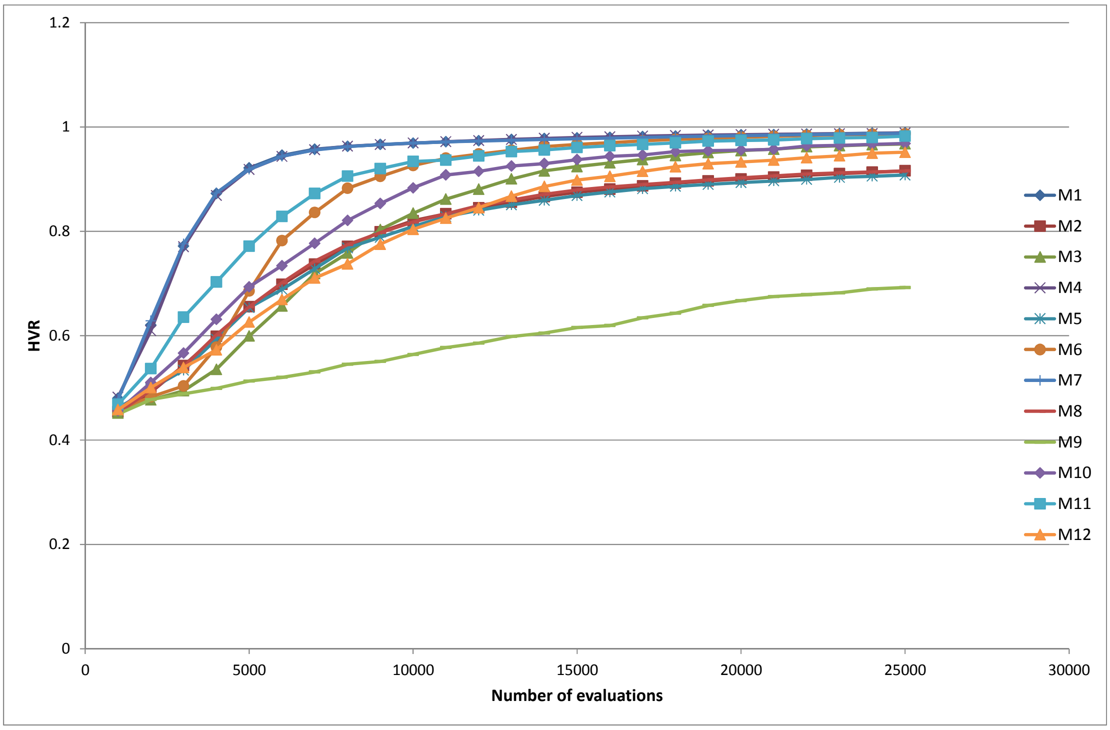

Figure A-2 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem ZDT3

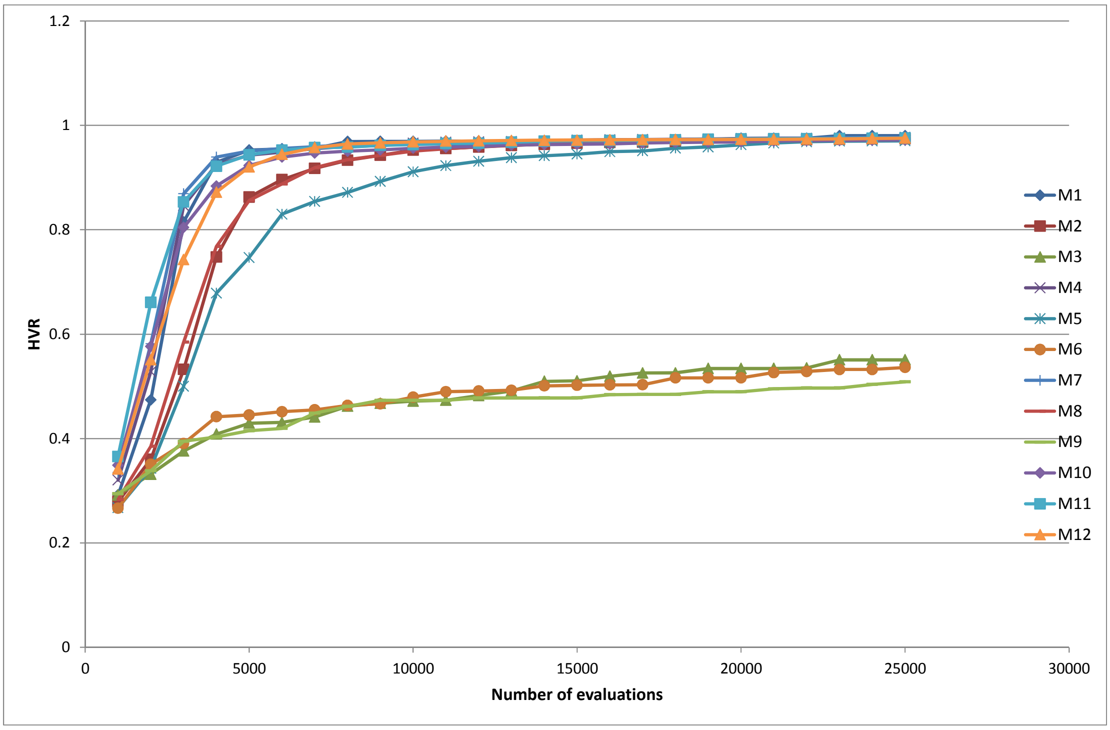

Figure A-3 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem ZDT4

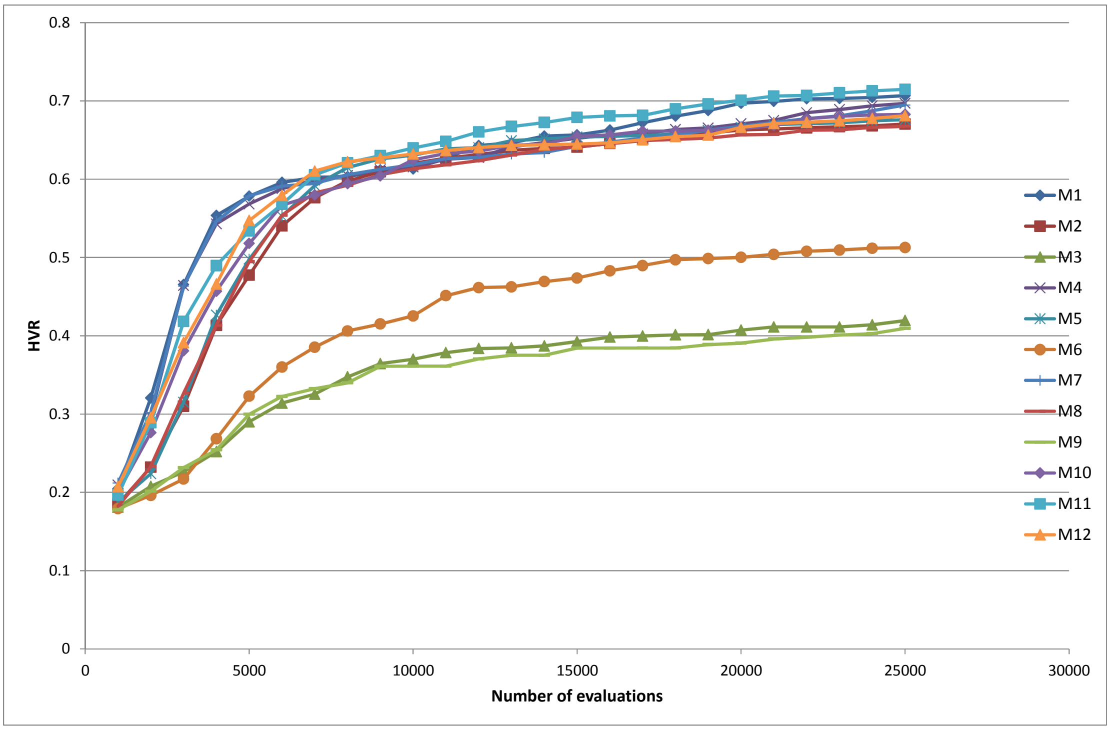

Figure A-4 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem ZDT6

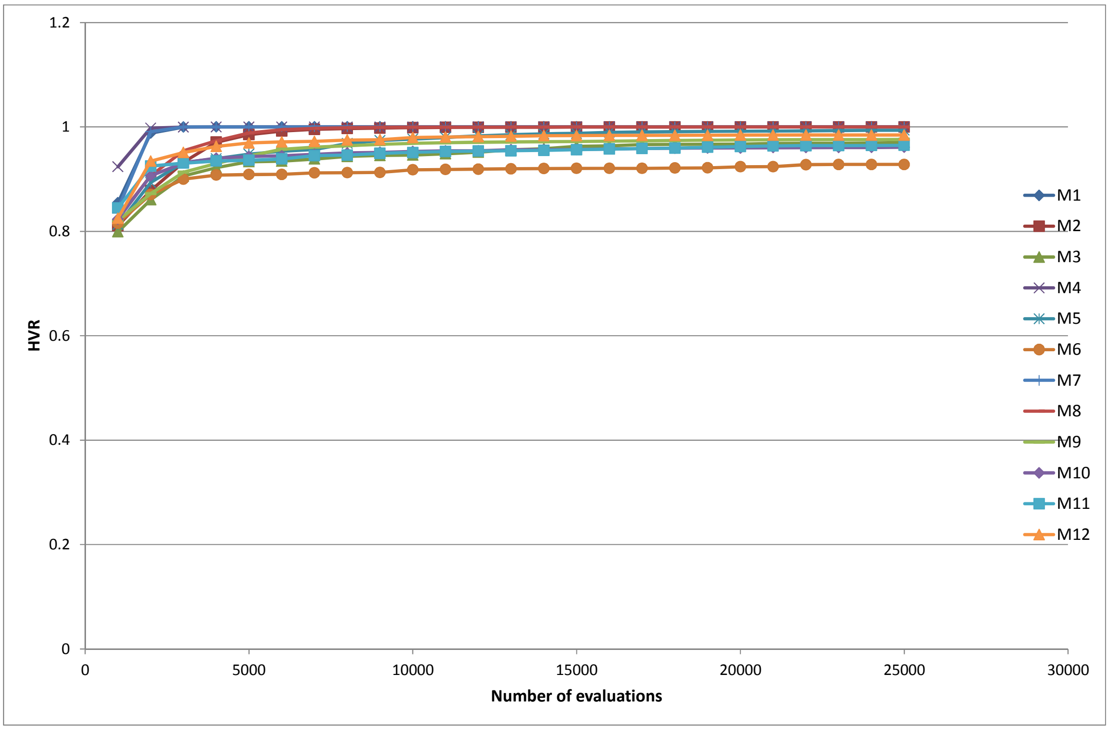

Figure A-5 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem DTLZ1

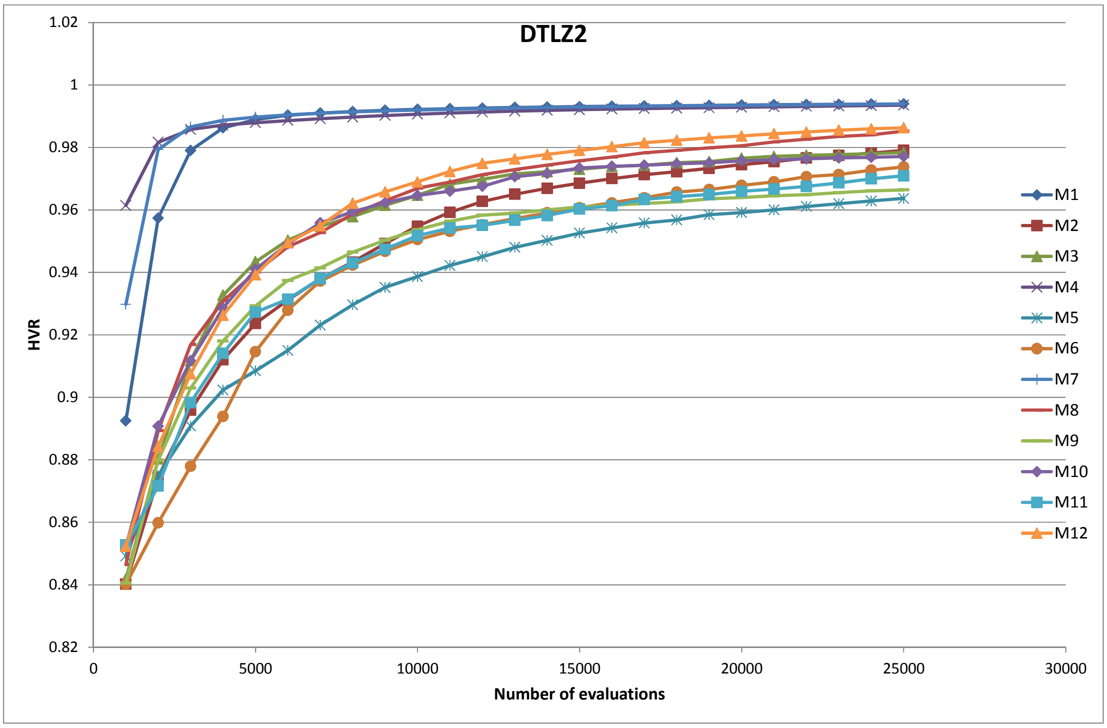

Figure A-6 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem DTLZ2

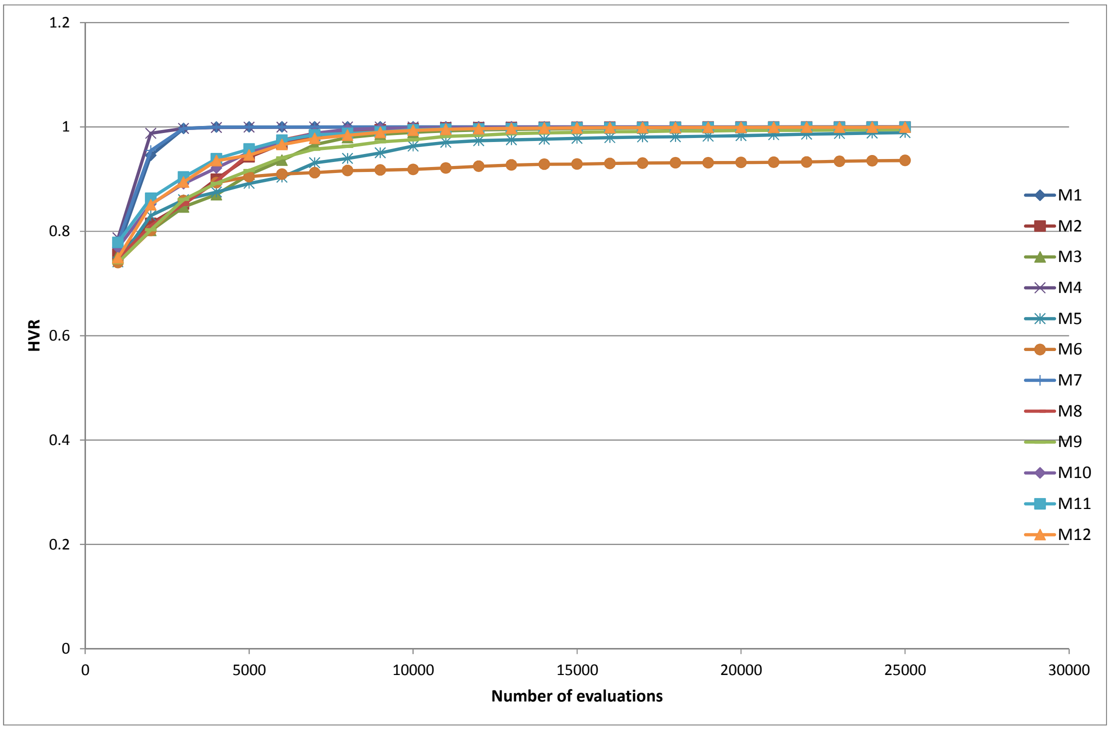

Figure A-7 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem DTLZ3

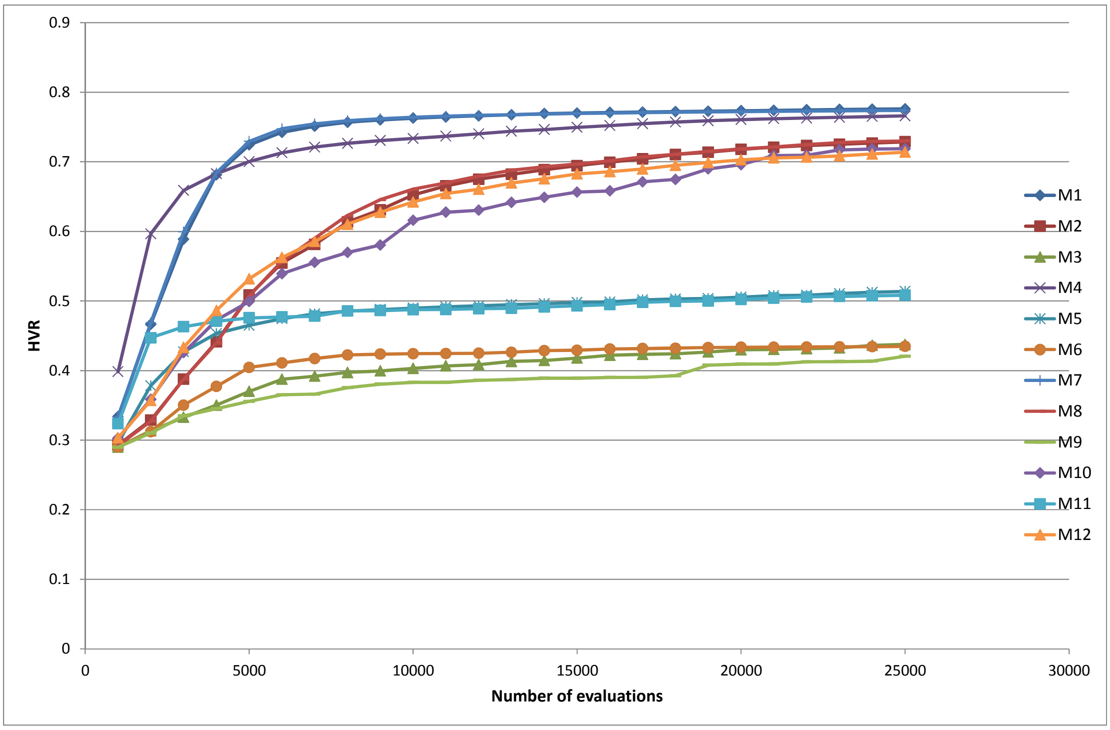

Figure A-8 HVR measure for twelve variants as a function of the number of evaluations for the benchmark problem DTLZ6

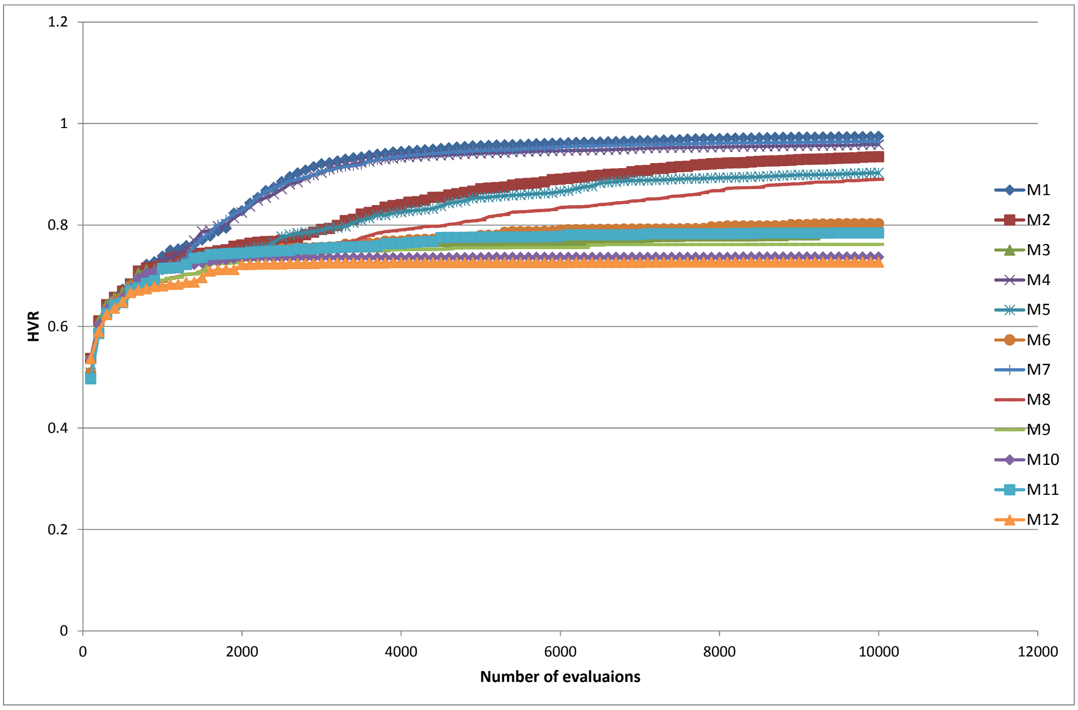

Figure A-9 HVR measure for twelve variants as a function of the number of evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

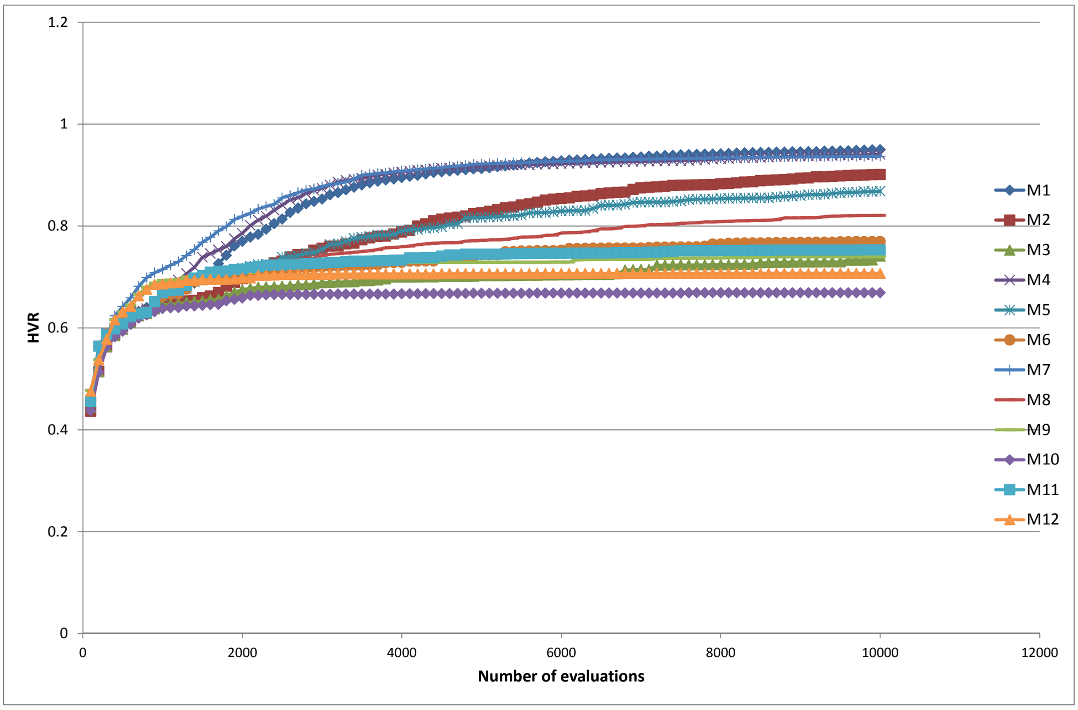

Figure A-10 HVR measure for twelve variants as a function of the number of evaluations for the Canberra case study minimizing three objectives: present worth cost, restriction frequency and the fraction of time that storage is less than 20%

In Figures A-11 to A-20 the hypervolume measure versus the number of evaluations for M1, M4, M7, M13, M14 and M15 variants is presented for each of the ten problems. These variants differ in terms of the amount of pheromone to be put on the non-dominated so-far solutions. The M1, M4 and M7 variants update pheromone based on the value calculated using Eq. (5.21) while M13, M14 and M15 update pheromone based on the inverse of objective function values. It was observed that the difference between the maximum and minimum amount of pheromone given by Eq. (5-21) was small – the pheromone deposit was effectively constant.

The results for the benchmark problems show little difference among the six variants. This is because the scale of the objective function values is not very large, so the scale of the inverse of the objective function values is small. However, in the Canberra problems, there are large differences in the scale of the objective function values. As a result, greater differences in the performance of the variants are observed. Figure A-19 shows that all three variants using the inverse objective method, M13, M14 and M15, are worse than M1, M4 and M7. This difference is starkly greater for the three-objective Canberra problem shown in Figure A–20.

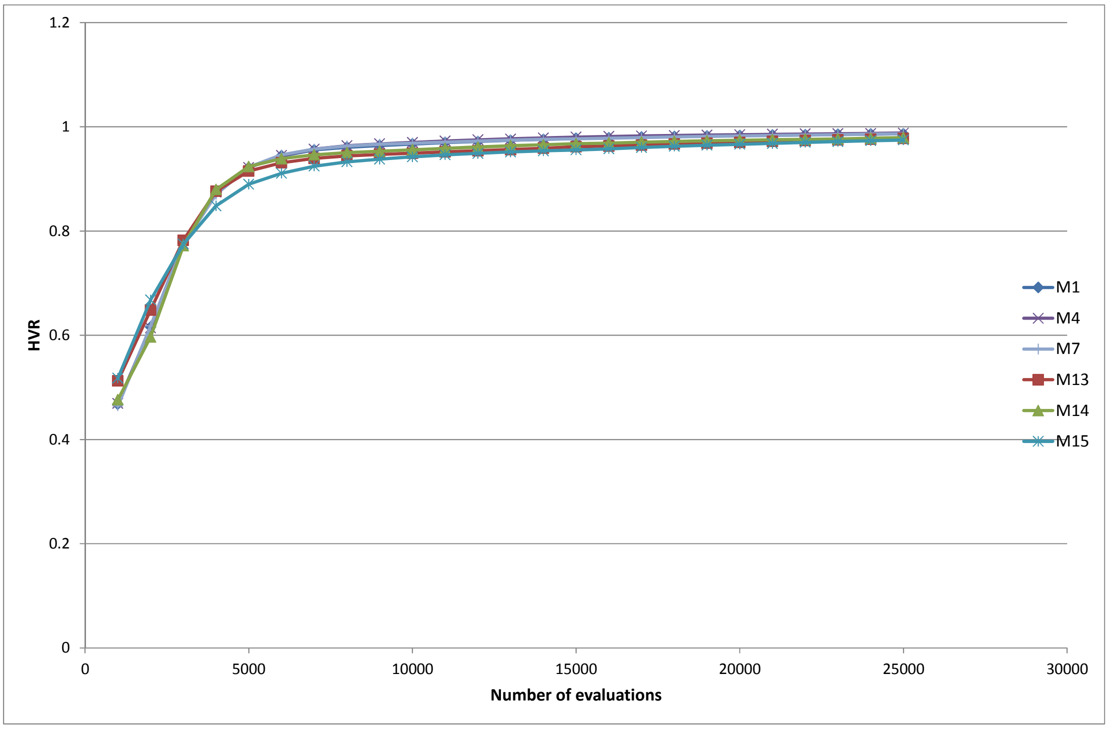

Figure A-11 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem ZDT1

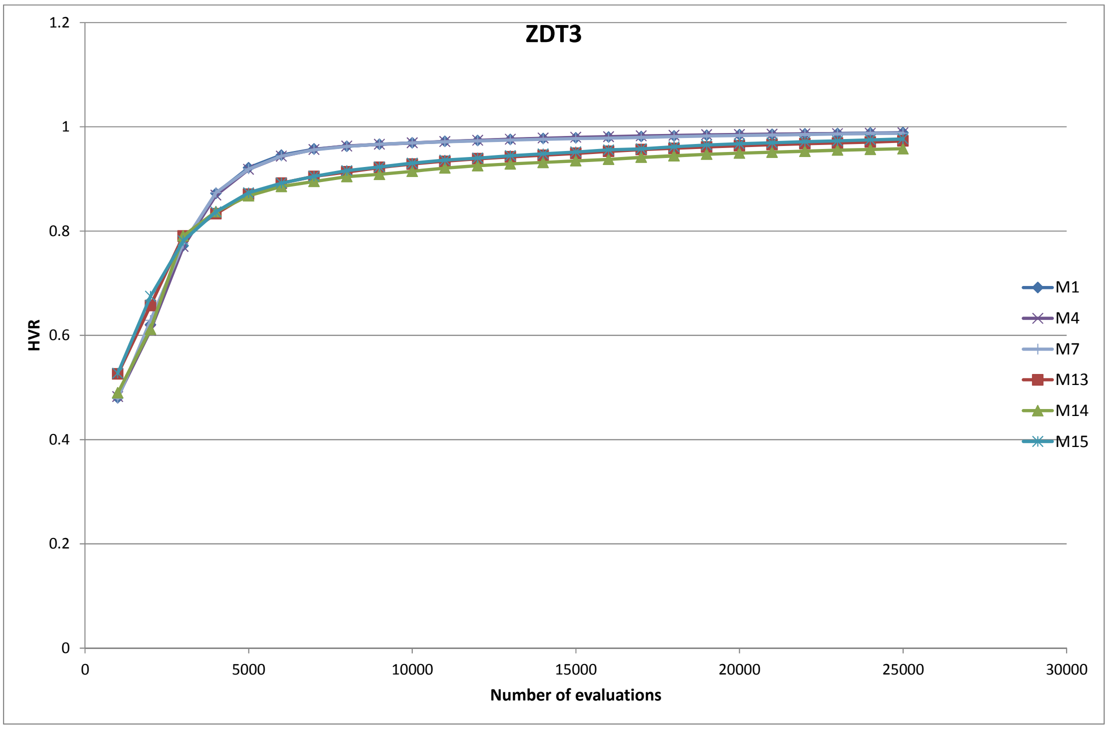

Figure A-12 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem ZDT3

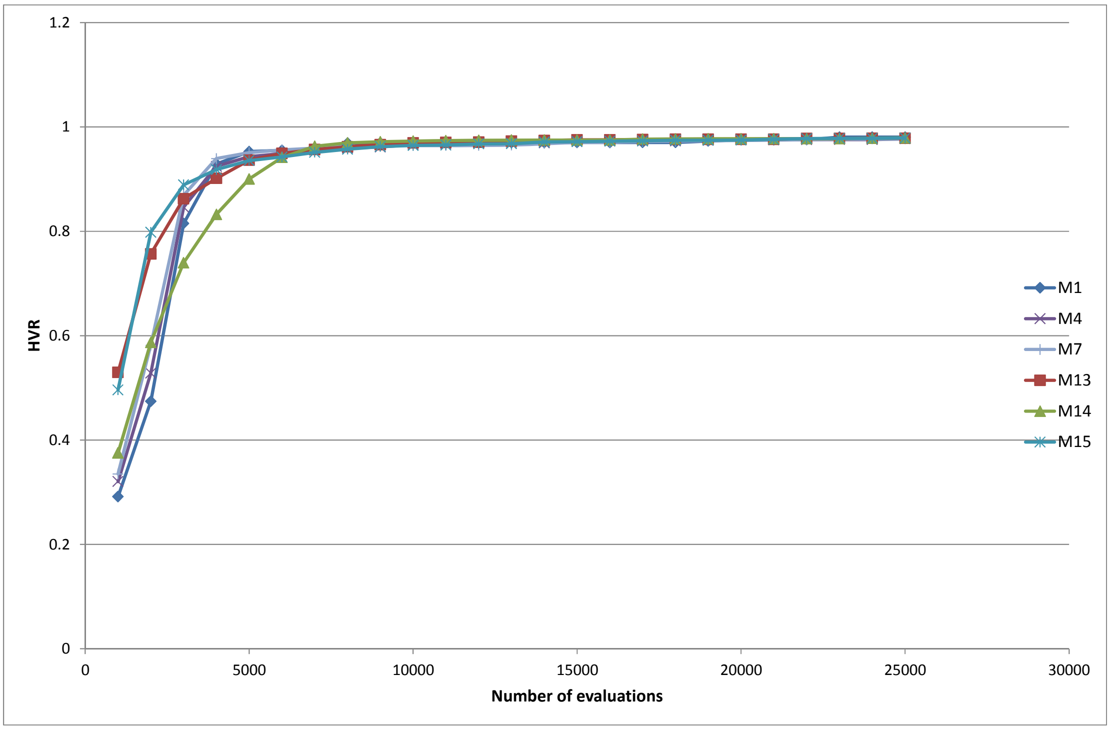

Figure A-13 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem ZDT4

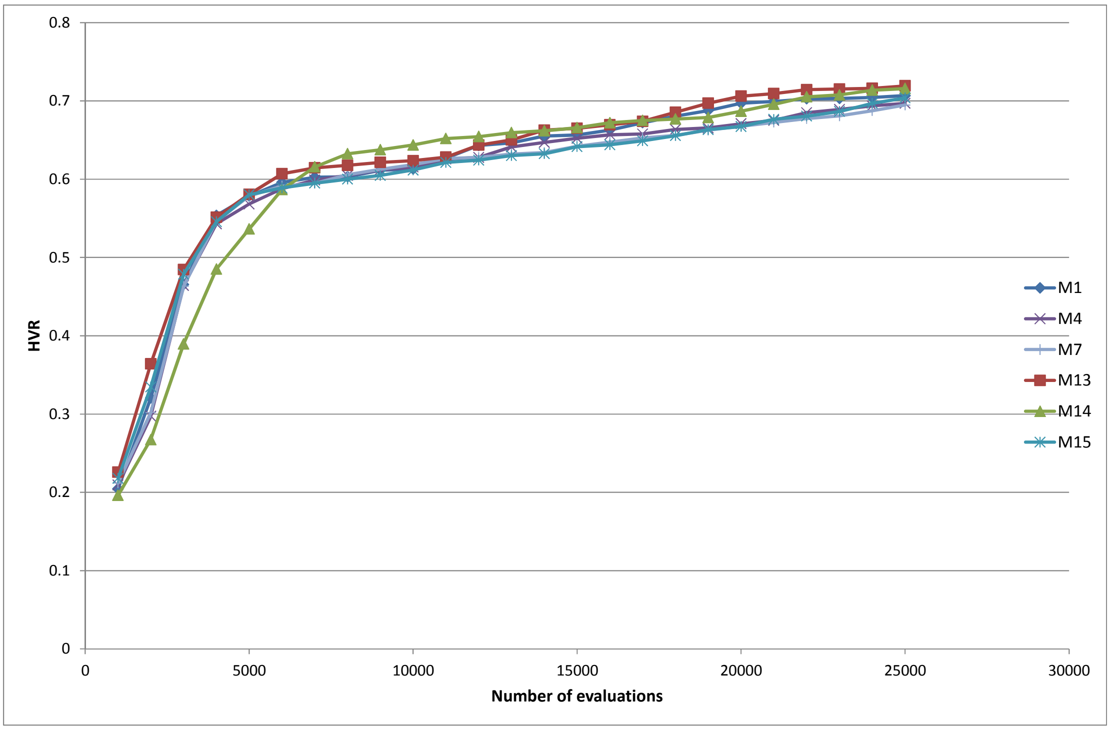

Figure A-14 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem ZDT6

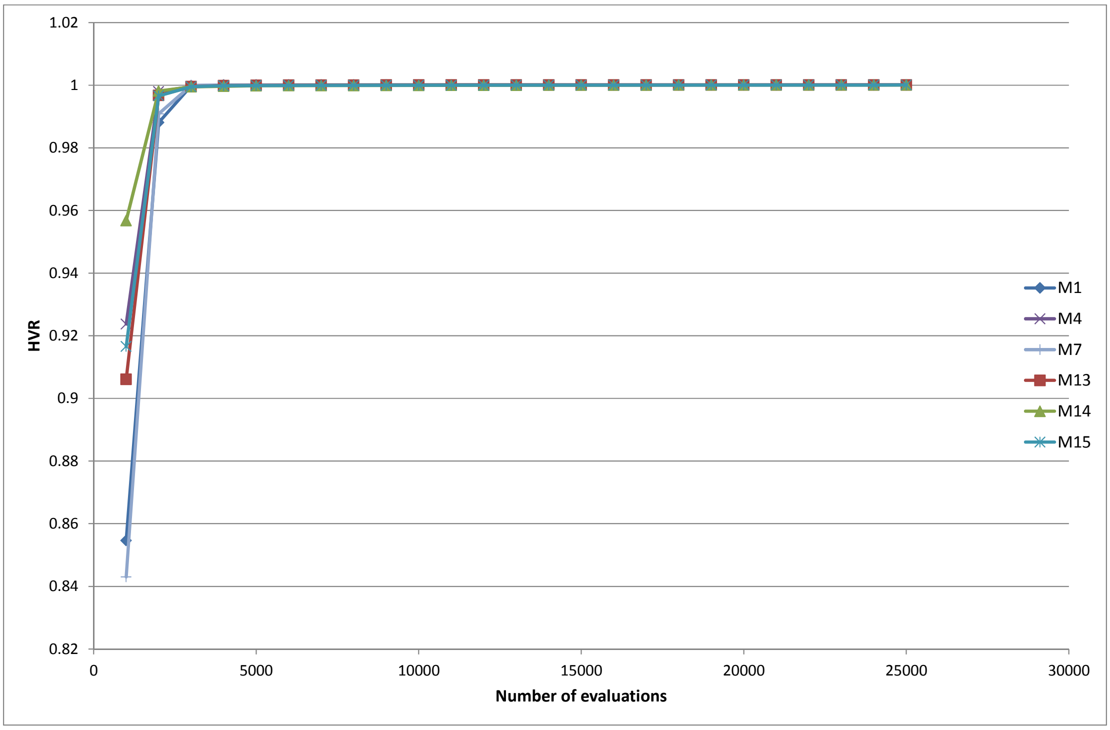

Figure A-15 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem DTLZ1

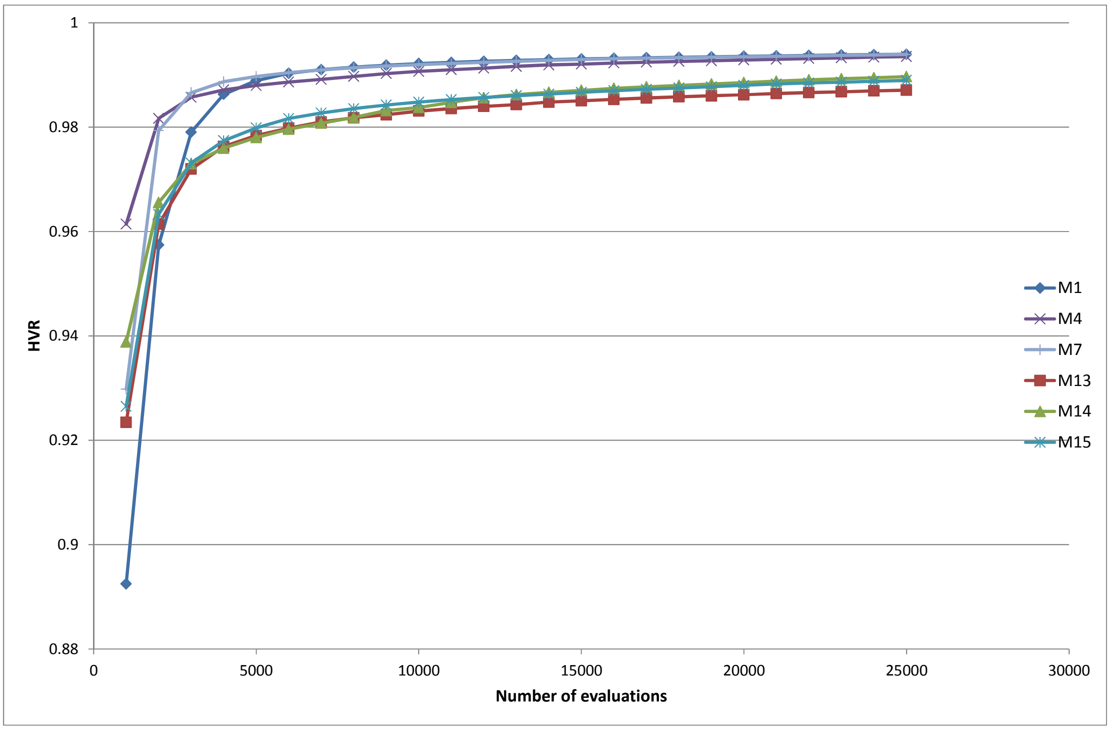

Figure A-16 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem DTLZ2

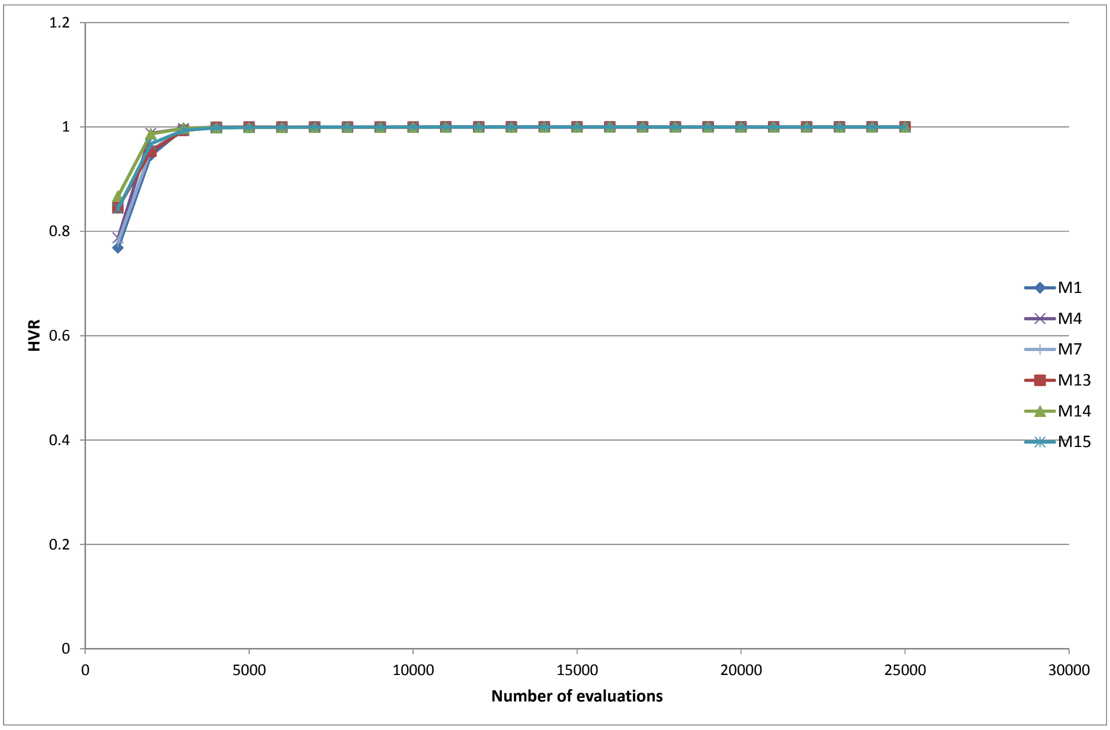

Figure A-17 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem DTLZ3

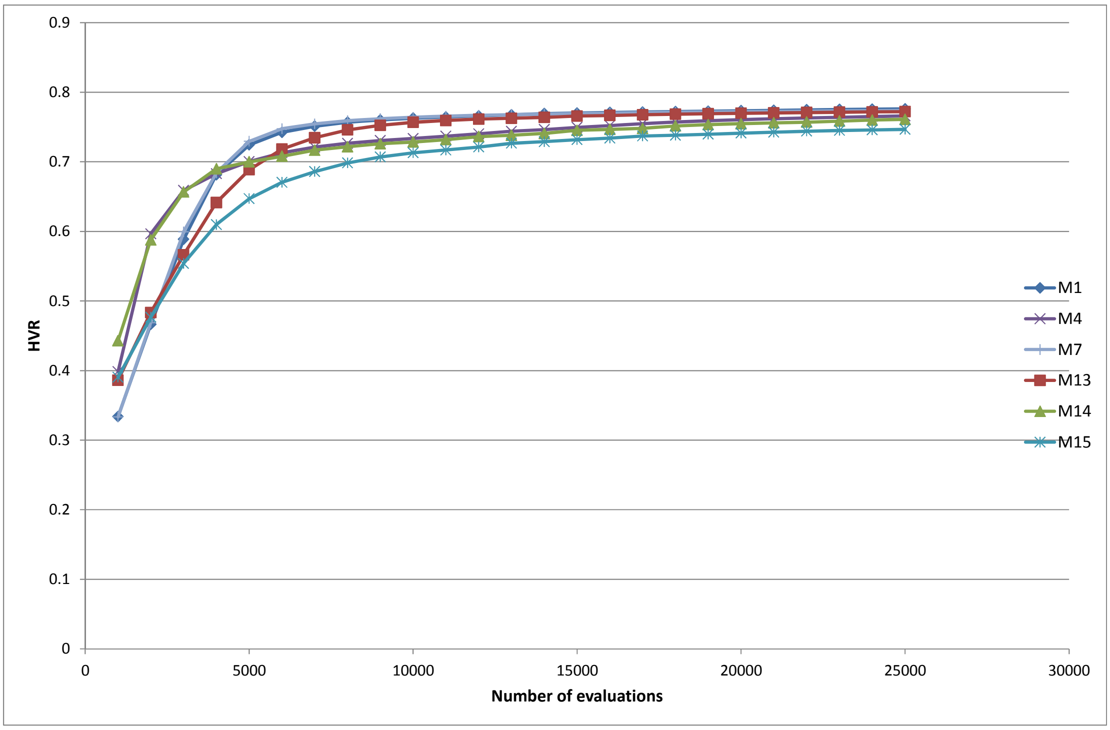

Figure A-18 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the benchmark problem DTLZ6

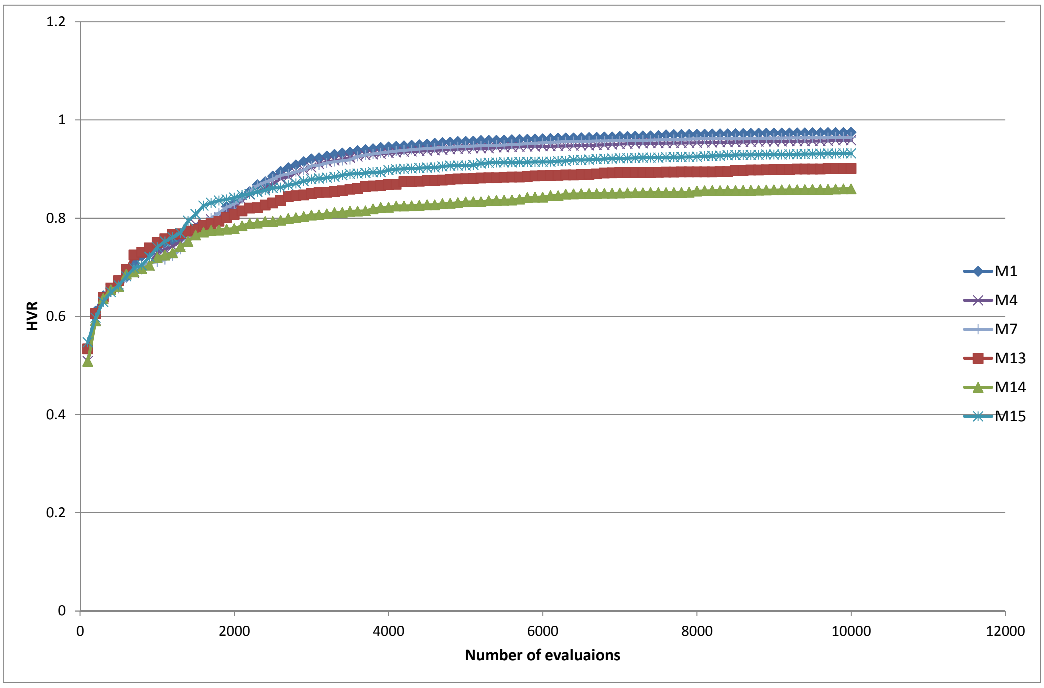

Figure A-19 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

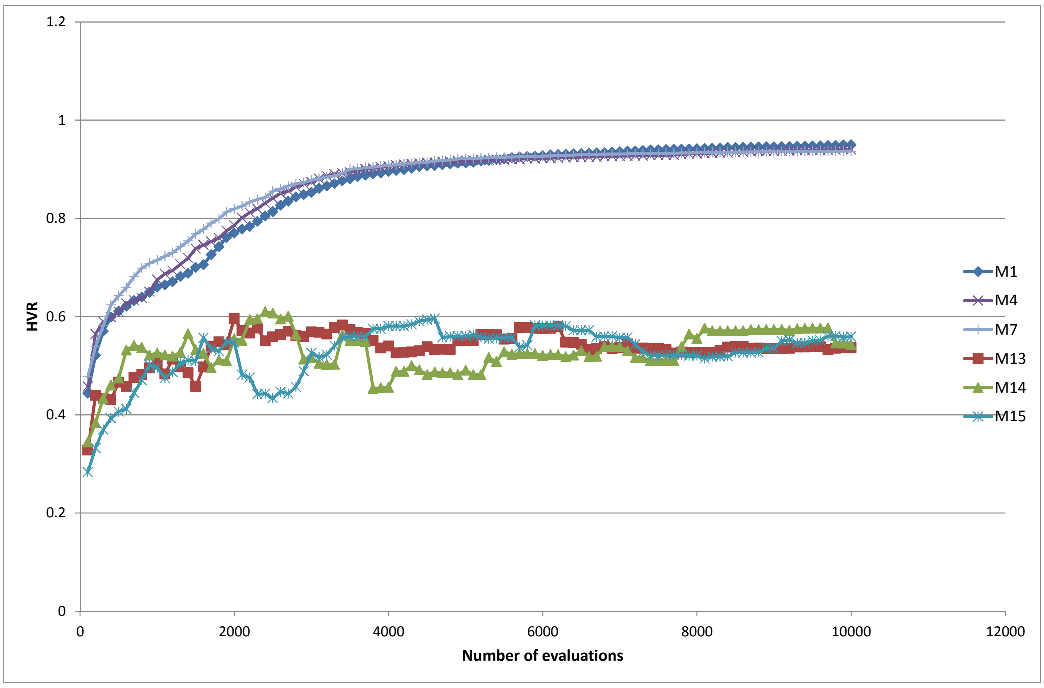

Figure A-20 HVR measure for six variants (M1, M4, M7, M13, M14 and M15) as a function of the number of evaluations for the Canberra case study minimizing three objectives: present worth cost, restriction frequency and and the fraction of time that storage is less than 20%

The results presented so far consistently show that the M1, M4 and M7 variants outperform the other variants. However, these variants depend on the objective function values in either the pheromone updating system or in the transition rule. To assess the significance of being dependent on objective function values, the sixteen variant, MOACO-State, is proposed with a single pheromone matrix and constant value for updating pheromone on non-dominated solutions found so-far. In Figures A-21 to A-30 the hypervolume measure of these three variants is compared against MOACO-State for the ten problems.

For the benchmark problems, the performance of M1, M4 and M7 is similar to that of MOACO-State, although MOACO-State would be judged to be superior. This finding is in agreement with (Martínez et al. 2007) who concluded that the updating process is more important than the number of pheromone or heuristic matrices. For the Canberra problems, MOACO-State has once again similar performance similar to the other variants, though in the three-objective problem, M4 and M7 perform noticeably worse for 3000 or fewer evaluations. MOACO-State is judged superior to M1, M4 and M7 for two reasons:

1) Because MOACO-State applies constant pheromone C regardless of the number and scale of objectives, it is considered to be more robust.

2) Because the primary interest in this chapter is identifying MOO methods that perform well for a limited number of evaluations, MOACO-State is judged to have a slight edge over the other variants, particularly for the three-objective Canberra problem.

For these reasons, MOACO-State is taken as representing the best fusion of existing MOACO algorithm features.

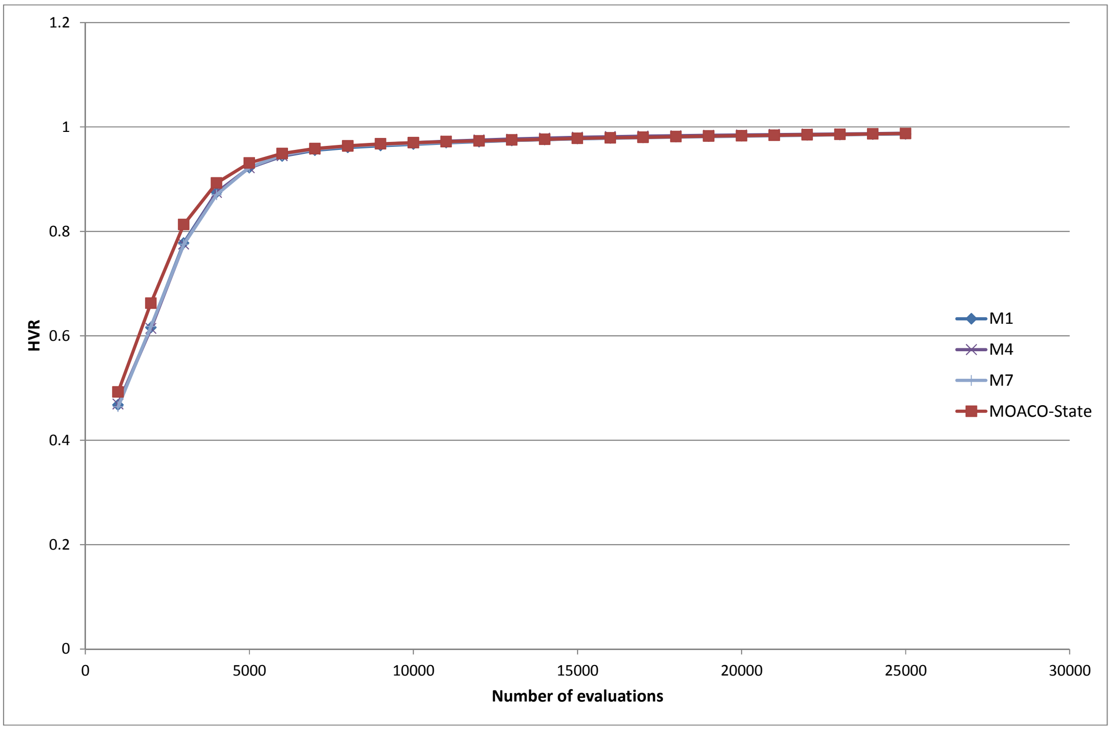

Figure A-21 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem ZDT1

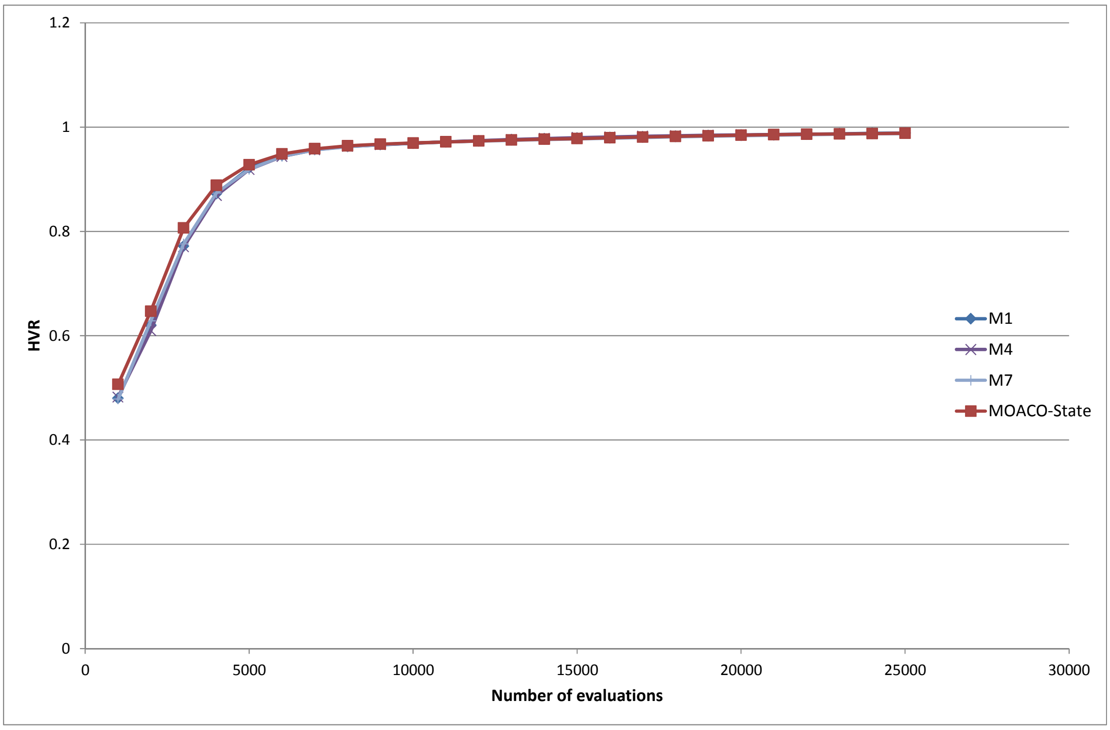

Figure A-22 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem ZDT3

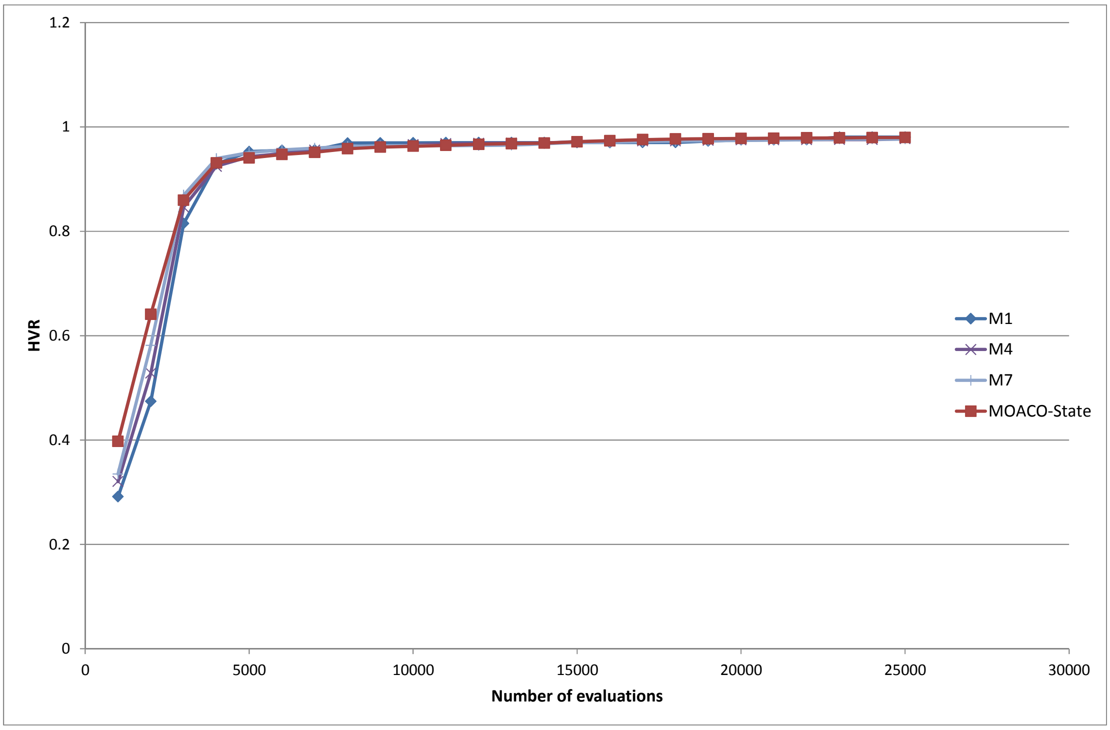

Figure A-23 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem ZDT4

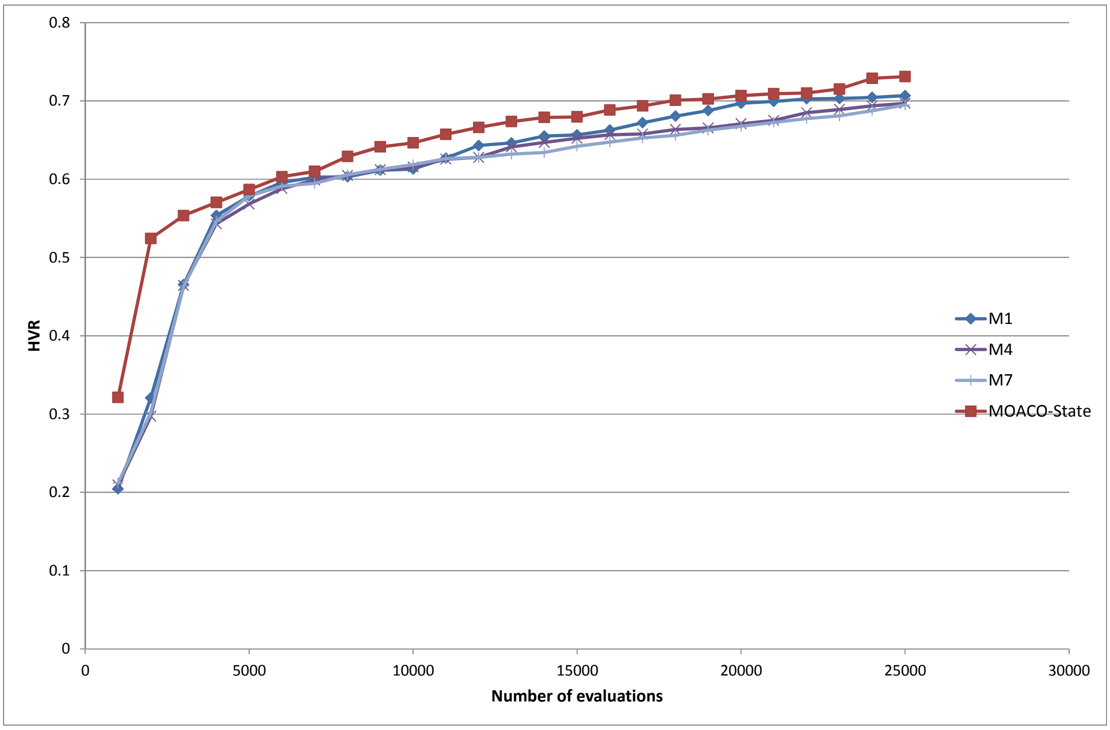

Figure A-24 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem ZDT6

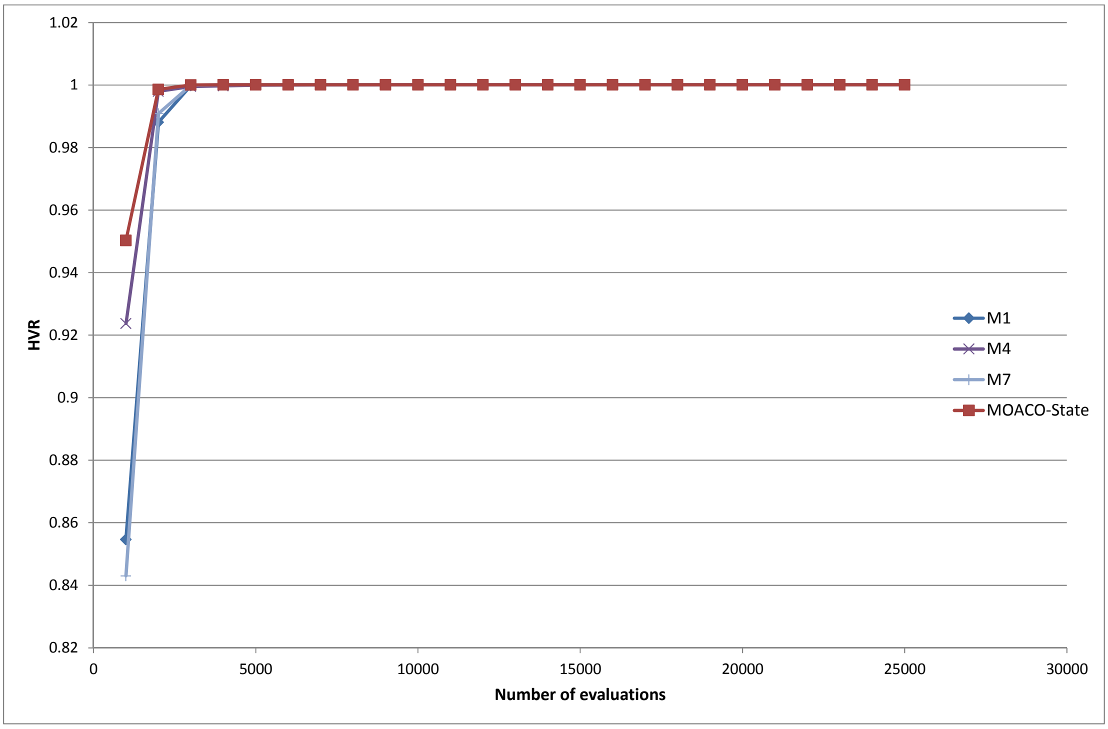

Figure A-25 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem DTLZ1

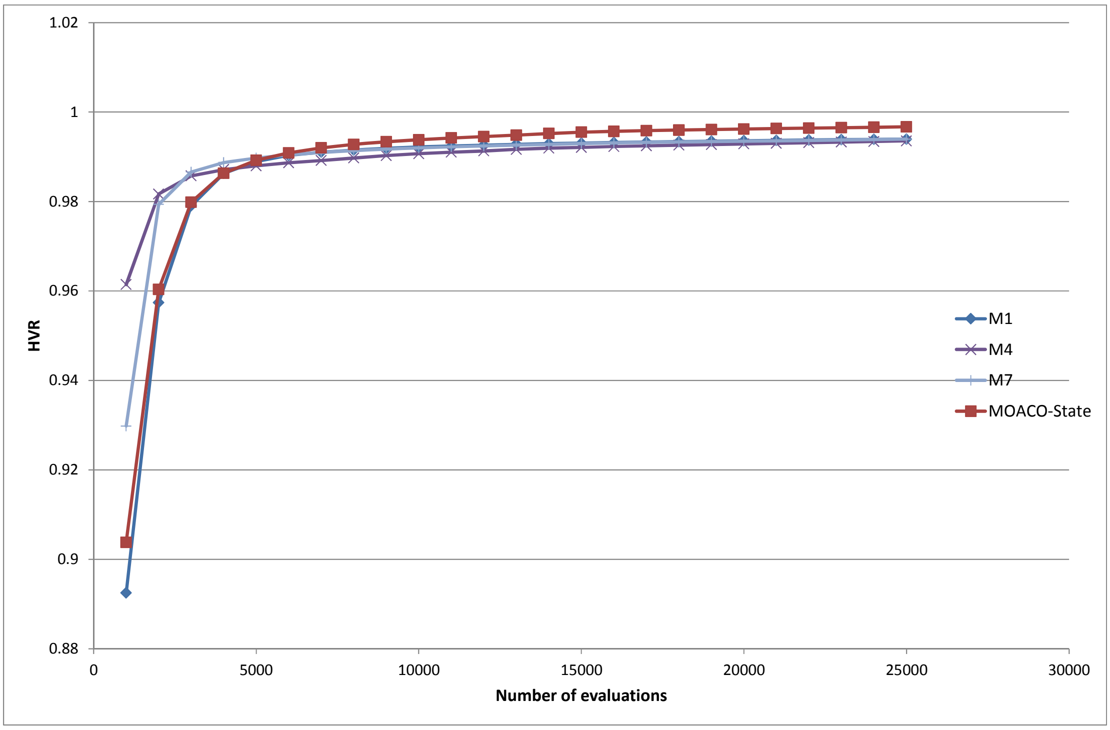

Figure A-26 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem DTLZ2

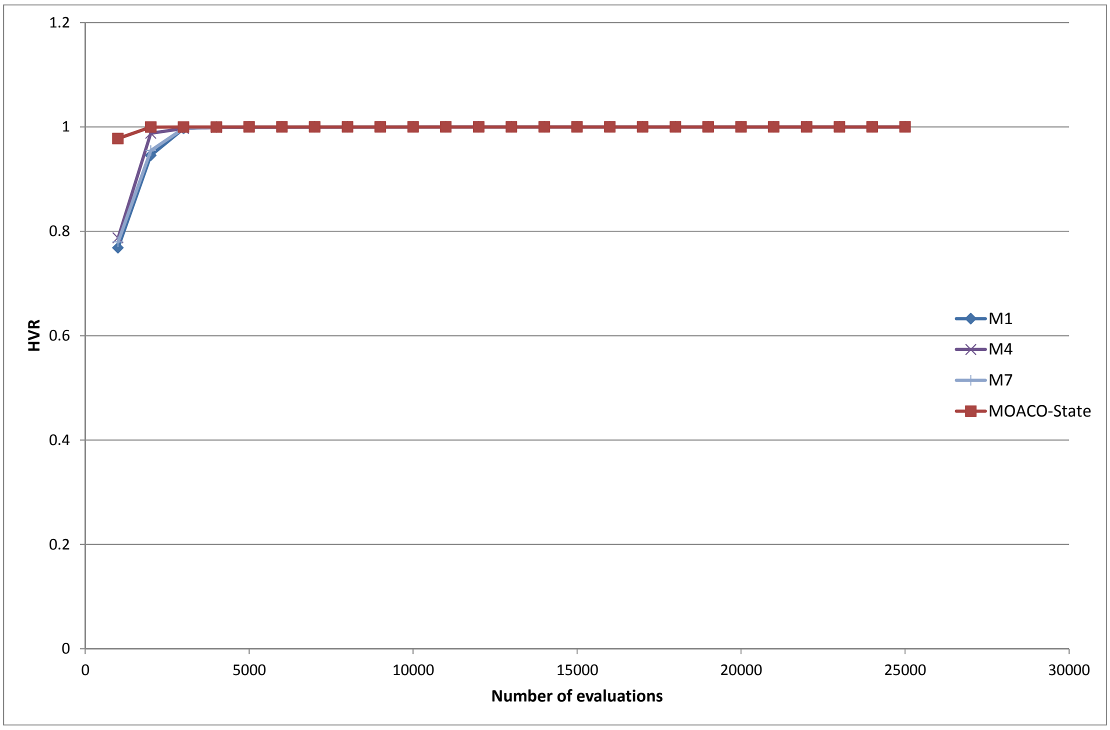

Figure A-27 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem DTLZ3

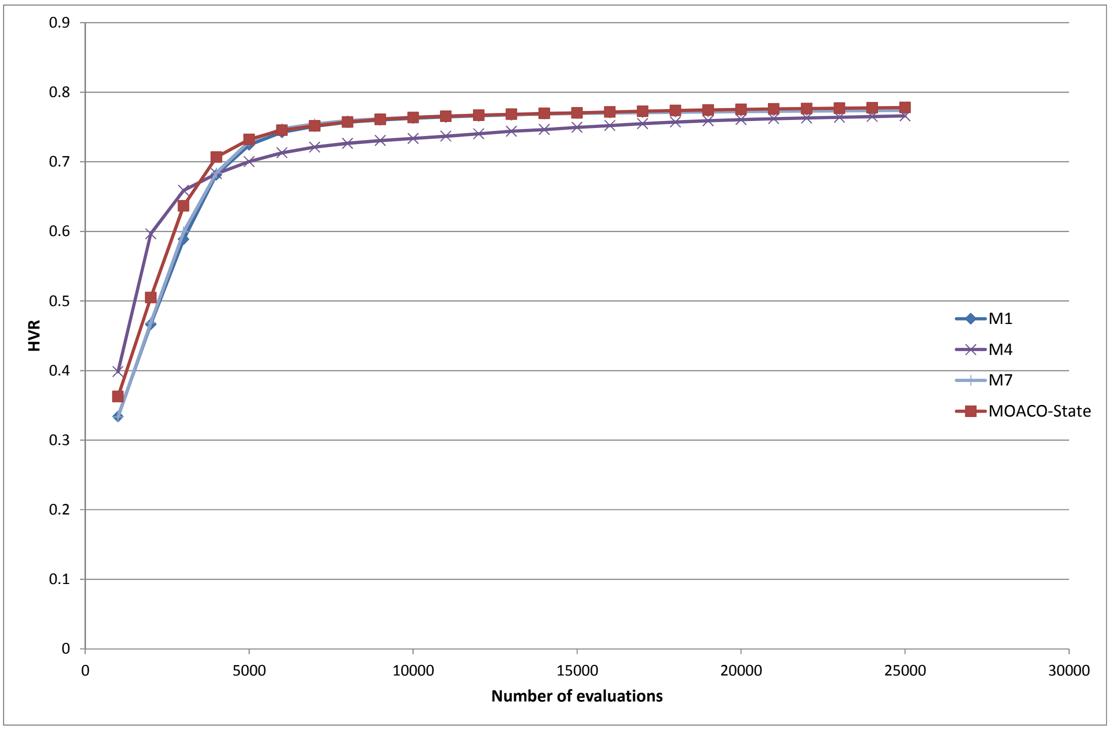

Figure A-28 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the benchmark problem DTLZ6

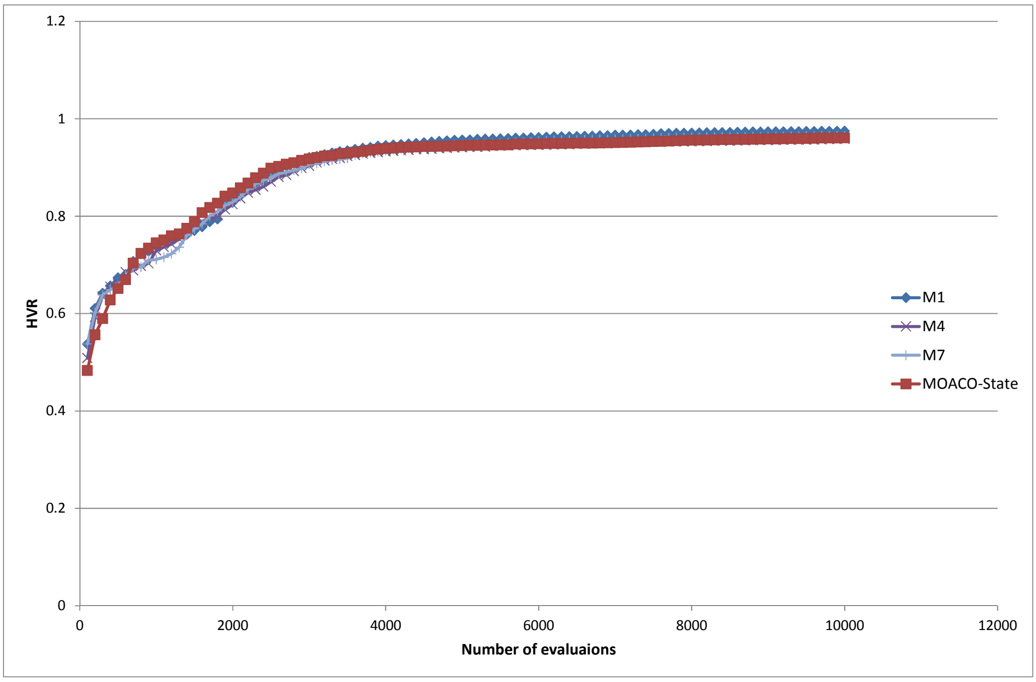

Figure A-29 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the Canberra case study minimizing two objectives: present worth cost and restriction frequency

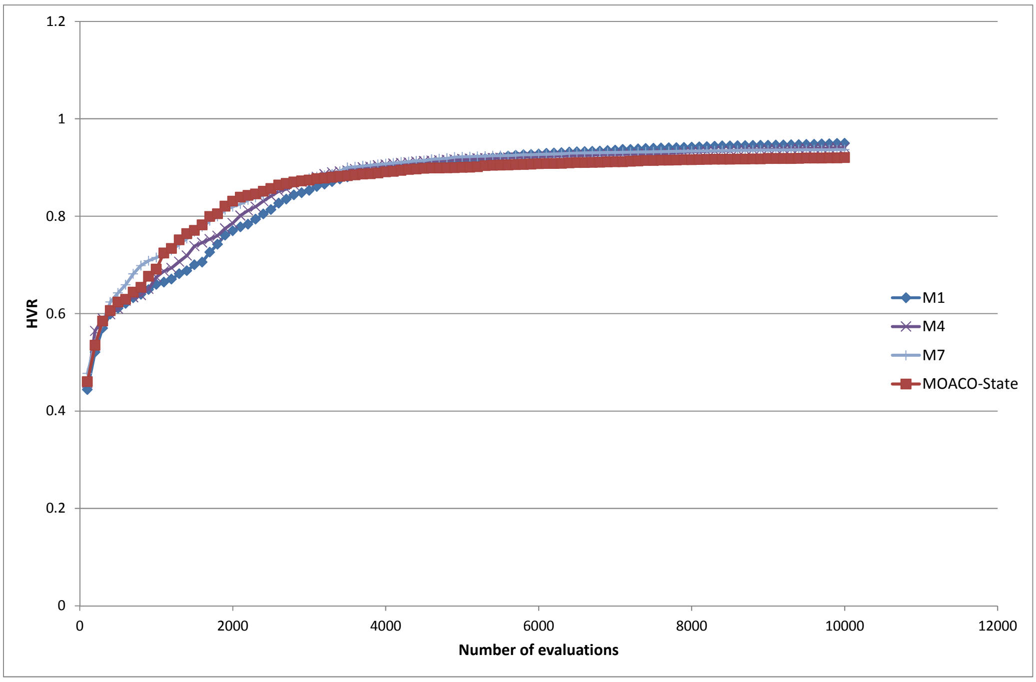

Figure A-30 HVR measure for three variants (M1, M4 and M7) and MOACO-State as a function of the number of evaluations for the Canberra case study minimizing three objectives: present worth cost, restriction frequency and the fraction of time that storage is less than 20%

Chapter 6

Conclusions and Future Work

6-1 Introduction

The overall aim of this thesis was to produce multi-objective optimization methods of greater practical relevance to urban water resource management. In pursuit of this aim, three specific objectives were set:

1. Identify and address the shortcomings of existing multi-objective optimization methods in urban water resources planning and operation;

2. Investigate the application of multi-objective optimization to scheduling and scaling of options to efficiently and equitably manage the challenge of population growth; and

3. Investigate the efficiency of multi-objective optimization search methods in the urban water resources applications.

This chapter reflects on the contributions made in this thesis. It summarizes the main findings with regard to each objective and then explores future research directions.

6-2 Summary and Conclusions

In this section, the rationale for each of the three objectives is revisited followed by a discussion of the major findings and their significance.

6-2-1 Moving Towards More Practical Multi-objective optimization Methods For Urban Water Resource Systems

Urban water management requires making decisions in the presence of conflicting objectives. The management of drought security in urban water supply is typically tackled using a mix of short-term options that manage the immediate response to drought and long-term options that control the risk of triggering the drought contingency plan. However, the maximization of drought security conflicts with other important objectives such minimizing economic cost and environmental impacts. This, along with the potentially very large number of solutions available to a water agency, makes multi-objective optimization a potentially very useful decision-support tool.

Review of past studies on the use of multi-objective optimization in water resource applications identified several shortcomings in a number of practically important areas. These include the following: 1) focusing exclusively on either long-term (or infrastructure) options or on short-term options (such as operation rules or drought contingency plans) may lead to sub-optimal solutions; 2) the use of short hydro-climate forcing data time series in simulation models to evaluate drought security can produce solutions that make the system highly vulnerable to severe drought; and 3) the *a priori* setting of environmental constraints may hide trade-offs between environmental, economic and security factors that are of considerable interest to decision makers.

These shortcomings have been addressed by a new multi-objective methodology. It exploits the ability of evolutionary algorithms to handle complex nonlinear objective functions and to interface with complex simulation models. An important contribution is the explicit treatment of drought security. The constraint is imposed that no unplanned shortfalls in demand may occur during the simulation – an unplanned shortfall would occur when the storages "run dry" or when there are limitations on transfer capacity, resulting in a failure to meet the minimum water needs specified in the drought contingency plan. Therefore, for an N-year simulation, the optimization produces a solution capable of dealing with a drought with an expected return period of N years. By using stochastically generated hydro-climate inputs, it is possible to consider very high levels of drought security.

A case study based on the headworks system for Australia's largest city, Sydney, demonstrated the practical significance of these shortcomings and, importantly, the ability of the new approach to deal with these shortcomings in a practicable manner. The following conclusions can be drawn from this case study:

1) Optimizing either operating rules or infrastructure options runs the risk of producing significantly inferior solutions. In the Sydney case study, there was strong interaction between some of the operating rule and infrastructure decisions. This highlights the importance of embedding an adequate simulation model in the optimization framework to ensure that joint optimization of operating rules and infrastructure options is possible.

2) The Sydney case study demonstrated the very considerable sensitivity of Pareto-optimal solutions to the expected return period of the severest drought. Indeed where high levels of drought security are expected, the use of historic or short stochastic hydro-climate records is fundamentally flawed leading to so-called “optimal” solutions that render the “optimized” system highly vulnerable to severe drought.

3) Environmental flow constraints are typically imposed on allocations within urban the water resource systems. The *a priori* imposition of such constraints runs the risk of missing potentially good solutions. It was shown in the Sydney case study that system performance was sensitive to the level of environmental flow constraints. Translating such constraints into objectives, difficult as it may be, provides a rich set of trade-offs between economic, social and environmental factors.

6-2-2 Efficient and Equitable Scheduling of Options to Cater For Future Changes

In the face of urban population growth and the accompanying growth in water demand, the performance of the urban water resource system is expected to deteriorate over time. This will result in the need to intervene and adapt the system to the changing conditions.

The scheduling capacity expansion problem seeks to identify the optimal schedule for changes to operating rules and infrastructure. In past studies, this problem has been largely tackled by minimizing the total present worth of capital, operational and rationing costs. A significant drawback of minimizing the total present worth cost is that it is likely to produce solutions that lead to more severe and frequent rationing in the future. Such a solution is likely to be socially unacceptable.

A new multi-objective formulation for the scheduling capacity expansion problem is developed to overcome this shortcoming while addressing the need to explicitly deal with drought security and jointly optimize operating and infrastructure decisions. The formulation enables the trade-off between cost and equity (the equal sharing of the burden of restrictions over the planning horizon) to be explored and deals with drought security by performing simulation using multiple replicates of future climate. A case study based on the headworks system for Australia's capital city, Canberra, was conducted to evaluate the merits of the new approach. The following conclusions can be drawn from this case study:

1) It was shown that minimizing total present worth cost can lead to more severe and frequent restrictions in the latter stages of the planning horizon. This unequal sharing of the burden of water shortages would be seen as a politically and socially-sensitive equity issue. A sensitivity analysis of the discount rate revealed that the higher the discount rate the greater the inequity across planning stages.

2) The importance of jointly scheduling both operating rule and infrastructure decisions was clearly demonstrated. By allowing decisions to adapt to the initial state of the system and to the growth in demand, the total present worth cost was reduced from \$462 m (when all decisions were made at the start of the planning horizon) to \$444 m (when decisions could be made at any of the three change points within the planning horizon). Indeed, in the Canberra case study, it was found that virtually all the benefit of scheduling decisions over the planning horizon could be attributed to the scheduling of operating rule decisions.

3) It was shown that formulating the scheduling capacity expansion problem as a multi-objective problem enabled the trade-off between cost and equity to be explored. The core idea was to introduce an objective which seeks to minimize the difference in the cost of restrictions between the planning stages. This produced a much richer, more relevant set of solutions for a decision maker to consider.

6-2-3 Computationally Efficient Multi-objective Optimization Methods

Computationally expensive simulation models are typically used in urban water resource applications to evaluate system performance. Simulation run times can range from less than a minute to over thirty minutes. These long run times are considered an impediment to the practical usage of MOO methods in urban water management.

The final objective of this thesis was to identify MOO methods that can produce approximate Pareto-optimal solutions with a limited number of objective function evaluations. Three benchmark MOO methods, NSGA-II, εMOEA and SMPSO, were selected among existing MOO methods for comparison.

The good performance of ant colony optimization (ACO) in single objective problems motivated investigation of its suitability for multi-objective optimization in the context of urban water management. A review of past work identified a number of shortcomings in existing MOACO methods with the principal one being the problem-specific nature of the algorithms. In this thesis, three MOACO methods were developed to address these shortcomings and were compared against the three benchmark methods using four urban water resource test problems.

The comparison of the six MOO methods using three metric that assessed the convergence, diversity (hypervolume ratio) and consistency ($I_{\varepsilon+}$) of solutions revealed that none of the methods was superior but that two of the methods, NSGA-II and MOACO-State, were inferior to the other methods. The EMOACO-I algorithm was found to be the best method in terms of the convergence and hypervolume ratio metrics but other methods produced better $I_{\varepsilon+}$ metrics. Out of the three benchmark methods, none emerged superior – εMOEA was ranked first for three of the four urban water resource test problems and SMPSO was ranked first for the fourth problem.

6-3 Future Research Directions

While this thesis has made a number of significant contributions that make multi-objective optimization more relevant and practicable in urban water resources management, there remain many opportunities for further advancement.

The case studies for Sydney and Canberra mainly focused on decisions associated with the headworks systems. However, there are considerable and practically important opportunities to extend the scope of these studies to include a much richer decision space. For example, the characterization of the drought contingency plan could be extended considerably. The case studies in this thesis limited rationing to domestic water use, when, in fact, there could be many more stages in a drought contingency plan that would impose severer rationing on all water sectors. It was shown that the total present worth cost for Sydney was sensitive to the choice of drought return period. Likewise it is expected there is similar sensitivity to the way the drought contingency plan is managed in a severe drought. Another opportunity lies in linking headworks or centralized solutions with decentralized solutions (Chiou et al., 2007; Daigger and Crawford, 2007; Daigger, 2009) that harvest water at the local urban scale. In the Canberra case study rainwater tanks were included as one of the infrastructure options. However, the scope of decentralized options is far greater including local stormwater harvesting and non-potable water substitution using grey and blackwater sources. Apart from the policy relevance of including decentralized options in Australian urban water supply, there are considerable technical challenges dealing with sensible treatments of scale over time and space.

The Fourth Assessment Report of the Intergovernmental Panel for Climate Change indicates that climate change is likely to impact on water resources around the world (Rosenzweig et al., 2007). In urban areas that are already vulnerable to drought, a drying climate is likely to significantly compound the stress arising from population growth (O'Hara and Georgakakos, 2008). Although the multi-objective scheduling approach developed in Chapter 4 was motivated by the need to cater for population growth, it is intrinsically suited to exploring the added stress of potentially drying climates. Because the approach allows scheduling of decisions over time, it provides a capable tool to explore adaptive management strategies provided, of course, a sufficiently rich decision space is used. However, the challenges, both technical and computational, would be considerable.

Ensuring practicable computational turnaround times for multi-objective optimization applications in urban water resources remains a formidable challenge. Generally there are three strategies to deal with computationally-intensive optimization problems: 1) use parallel computing; 2) adopt more efficient optimization methods; 3) and use meta-models to approximate the mapping between the decision and objective function spaces. The third option was not explored in this thesis but has received considerable attention in recent years. There is a strong case for exploring this option in the context of urban water resources.

The opportunity to develop new multi-objective ACO algorithms was limited by time constraints. There is clear scope to parallelize the ACO algorithms along the lines reported by Stützle (1998), Manfrin et al. (2006), Koshimizu and Saito (2009) and Chen et al. (2012). There is considerable scope to exploit the use of tabu lists to prevent ants exploring infeasible solution spaces. There is scope to better tune ACO parameters. A number of studies have investigated the issue of tuning and the adoption of adaptive tuning methods (Arabshahi et al., 1996; Zecchin et al., 2005; Stützle et al., 2010; Randall, 2004; Pellegrini et al., 2012). Finally, there is scope to improve the performance of the EMOACO and EMOACO-I algorithms with respect to the $I_{\varepsilon+}$ measure by dealing with the clustering behaviour of ants as they sample non-dominated routes.

In conclusion, while there is considerable scope for further work, this thesis has made several original and significant contributions to produce multi-objective optimization methods of greater practical relevance to urban water resource management.

References

Abbaspour, K. C., R. Schulin and M. T. van Genuchten (2001). "Estimating unsaturated soil hydraulic parameters using ant colony optimization." Advances in Water Resources, **24**(8): 827-841.

Afshar, A., F. Sharifi and M. R. Jalali (2009). "Non-dominated archiving multi-colony ant algorithm for multi-objective optimization: Application to multi-purpose reservoir operation." Engineering Optimization, **41**(4): 313-325.

Ajami, N. K., G. M. Hornberger and D. L. Sunding (2008). "Sustainable water resource management under hydrological uncertainty." Water Resources Research, **44**(11): W11406.

Alaya, I., C. Solnon and K. Ghedira (2007). Ant Colony Optimization for Multi-Objective Optimization Problems. 19th IEEE International Conference on Tools with Artificial Intelligence, 2007, **1**: 450-457.

Alonso, J. M., F. Alvarruiz, D. Guerrero, V. Hernández, P. A. Ruiz, A. M. Vidal, F. Martínez, J. Vercher and B. Ulanicki (2000). "Parallel Computing in Water Network Analysis and Leakage Minimization." Journal of Water Resources Planning and Management, **126**(4): 251-260.

Andrews, E. S., F. I. Chung and J. R. Lund (1992). "Multilayered, Priority-Based Simulation Of Conjunctive Facilities." Journal of Water Resources Planning and Management, **118**(1): 32-53.

Angus, D. (2007a). Population-Based Ant Colony Optimisation for Multi-objective Function Optimisation. Progress in Artificial Life, **4828/2007:** 232-244.

Angus, D. (2007b). Crowding Population-based Ant Colony Optimisation for the Multi-objective Travelling Salesman Problem. IEEE Symposium on Computational Intelligence in Multicriteria Decision Making, 333-340.

Angus, D. and C. Woodward (2009). "Multiple objective ant colony optimisation." Swarm Intelligence, **3**(1): 69-85.

Arabshahi, P., J. J. Choi, R. J. Marks, II and T. P. Caudell (1996). "Fuzzy parameter adaptation in optimization: some neural net training examples." Computational Science & Engineering, IEEE, **3**(1): 57-65.

Arthington, A., R. E. Tharme, S. O. Brizga, B. J. Pusey and M. J. Kennard (2004). Environmental Flow Assessment With Emphasis on Holistic Methodologies. The second international symposium on the management of large rivers for fisheries. R. L. Welcomme and T. Petr. Kingdom of Cambodia, RAP publication, **2**: 37-66.

Baran, B. and M. Schaerer (2003). A Multiobjective Ant Colony System For Vehicle Routing Problem With Time Windows. Proceedings of the 21st IASTED International Conference, Innsbruck, Austria.

Barreto, W. J., Z. Vojinovic, R. K. Price and D. P. Solomatine (2007). Approaches to Multi-Objective Multi-Tier Optimization in Urban Drainage Planning. 7th International Conference on Hydroinformatics, Nice, FRANCE.

Basagaoglu, H. and H. Yazicigil (1994). "Optimal capacity-expansion planning in multiaquifer systems." Journal of Water Resources Planning and Management, **120**(6): 836-856.

Bellman, R. E. (1957). Dynamic programming, Princeton University Press, Princeton, NJ.

Bertsekas, D. P. and P. Tseng (1988). "Relaxation Methods for Minimum Cost Ordinary and Generalized Network Flow Problems." Operations Research, **36**(1): 93-114.

Braga, B. P. F., J. G. L. Conejo, L. Becker and W. W.-G. Yeh (1985). "Capacity Expansion of Sao Paulo Water Supply." Journal of Water Resources Planning and Management, **111**(2): 238-252.

Brendecke, C. M., W. B. DeOreo, E. A. Payton and L. T. Rozaklis (1989). "Network Models of Water Rights and System Operations." Journal of Water Resources Planning and Management, **115**(5): 684-696.

Broad, D. R., H. R. Maier and G. C. Dandy (2010). "Optimal operation of complex water distribution systems using metamodels." Journal of Water Resources Planning and Management, **136**(4): 433-443.

Bui, L. T., J. M. Whitacre and H. A. Abbass (2008). Performance analysis of elitism in multi-objective ant colony optimization algorithms. IEEE World Congress on Computational Intelligence, 1633-1640.

Butcher, W. S., Y. Y. Haimes and W. A. Hall (1969). "Dynamic Programing for the Optimal Sequencing of Water Supply Projects." Water Resources Research, **5**(6): 1196-1204.

Cai, X., D. C. McKinney and L. S. Lasdon (2001). "Solving nonlinear water management models using a combined genetic algorithm and linear programming approach." Advances in Water Resources, **24**(6): 667-676.

Cardoso, P., M. Jesus and A. Márquez (2003). MONACO - Multi-Objective Network Optimisation Based on ACO. X Encuentros de Geometría Computacional, Seville, Spain.

Chang, L.-C. and F.-J. Chang (2009). "Multi-objective evolutionary algorithm for operating parallel reservoir system." Journal of Hydrology, **377**(1-2): 12-20.

Chang, L.-C., Y.-W. Chen and M.-S. Yeh (2009). "Optimizing system capacity expansion schedules for groundwater supply." Water Resources Research, **45**(7): W07407.

Chen, L., J. McPhee and W. W. G. Yeh (2007). "A diversified multiobjective GA for optimizing reservoir rule curves." Advances in Water Resources, **30**(5): 1082-1093.

Chen, L., H.-Y. Sun and S. Wang (2012). "A parallel ant colony algorithm on massively parallel processors and its convergence analysis for the travelling salesman problem." Information Sciences, **199**: 31-42.

Chiou, R. J., Chang, T. C., & Ouyang, C. F. (2007). Aspects of municipal wastewater reclamation and reuse for future water resource shortages in Taiwan. Water Science and Technology, **55**(1-2), 397-405.

Chung, F. I., M. C. Archer and J. J. DeVries (1989). "Network Flow Algorithm Applied to California Aqueduct Simulation." Journal of Water Resources Planning and Management, **115**(2): 131-147.

Cloke, P. S. and S. Samra (2009). Impacts of the 2001–2007 Drought and Climate Change on Security of Water Supplies in Country NSW. The 32nd Hydrology and Water Resources Symposium, Newcastle, Australia. H2009.

Coello Coello, C. A. (2006). "Evolutionary multi-objective optimization: A historical view of the field." Computational Intelligence Magazine, IEEE, **1**(1): 28-36.

Coello Coello, C. A., G. B. Lamont and D. A. V. Veldhuizen (2007). Evolutionary Algorithms for Solving Multi-Objective Problem, Springer.

Connarty, M. C. and G. C. Dandy (1996). Optimum Pricing and Capacity Expansion of Water Supply Headworks Using Genetic Algorithms. Hydrology and Water Resources Symposium (23rd : 1996) Hobart, Tas., 357-361.

Consoli, S., B. Matarazzo and N. Pappalardo (2008). "Operating Rules of an Irrigation Purposes Reservoir Using Multi-Objective Optimization." Water Resources Management, **22**(5): 551-564.

Cui, L.-J. (2003). Optimization of Urban Water Supply Headworks System Using Probabilistic Search Methods and Parallel computing. Civil engineering, University of Newcastle, Newcastle, **PhD**.

Cui, L. and G. Kuczera (2005). "Optimizing water supply headworks operating rules under stochastic inputs: Assessment of genetic algorithm performance." Water Resources Research, **41**(5): W05016.

Czyzżak, P. and A. Jaszkiewicz (1998). "Pareto simulated annealing - a metaheuristic technique for multiple-objective combinatorial optimization." Journal of Multi-Criteria Decision Analysis, **7**(1): 34-47.

Daigger, G. & Crawford, G. (2007). Enhancing water system security and sustainability by incorporating centralized and decentralized water reclamation and reuse into urban water management systems. Journal of Environmental Engineering Managegement, **17**, 1-10.

Daigger, G. T. (2009). Evolving Urban Water and Residuals Management Paradigms: Water Reclamation and Reuse, Decentralization, and Resource Recovery. Water Environment Research, **81**(8): 809-823.

Dandy, G. C., A. M. Edward and G. H. Bruce (1985). "Pricing and Expansion of a Water Supply System." Journal of Water Resources Planning and Management, **111**(1): 24-42.

Deb, K. (2001). Multi-objective optimization using evalutionary algorithms, John Wiley & Sons, Inc.

Deb, K. and H. Gupta (2006). "Introducing robustness in multi-objective optimization." Evolutionary Computation, **14**(4): 463-494.

Deb, K. and S. Jain (2002), Running Performance Metrics for Evolutionary Multi-Objective Optimization, *Rep. KanGAL Report No. 2002004*, Indian Institute of Technology.

Deb, K., M. Mohan and S. Mishra (2003a), A Fast Multi-objective Evolutionary Algorithm for Finding Well-Spread Pareto-Optimal Solutions, *Rep. KanGAL Report No. 2003002*, Indian Institute of Technology.

Deb, K., P. Zope and A. Jain (2003b). Distributed Computing of Pareto-Optimal Solutions Using Multi-Objective Evolutionary Algorithms. The Second Evolutionary Multi-Criterion Optimization (EMO-03) Conference. 8-11 April, Faro, Portugal.**,** 535-549.

Deb, K., A. Pratap, S. Agarwal and T. Meyarivan (2002a). "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II." IEEE Transactions on Evolutionary Computation, **6**(2): 182-197.

Deb, K., L. Thiele, M. Laumanns and E. Zitzler (2002b). Scalable multi-objective optimization test problems. The 2002 Congress on Evolutionary Computation, **1:** 825-830.

Dewson, Z. S., A. B. W. James and R. G. Death (2007). "A review of the consequences of decreased flow for instream habitat and macroinvertebrates." Journal of the North American Benthological Society, **26**(3): 401-415.

Doerner, K. and W. J. Gutjahr (2004). "Pareto Ant Colony Optimization: A Metaheuristic Approach to Multiobjective Portfolio Selection." Annals of Operations Research, **131**(1-4): 79-99.

Doerner, K. F., R. F. Hartl and M. Reimann (2003). "Are COMPETants more competent for problem solving? The case of a multiple objective transportation problem." European Journal of Operational Research, **11**(2): 115-141.

Dorigo, M. and G. Di Caro (1999). Ant colony optimization: a new meta-heuristic. The 1999 Congress on Evolutionary Computation, 1999. CEC 99, **2:** 1470-1477.

Dorigo, M. and L. M. Gambardella (1997). "Ant colony system: a cooperative learning approach to the traveling salesman problem." IEEE Transactions on Evolutionary Computation, **1**(1): 53-66.

Dorigo, M., V. Maniezzo and A. Colorni (1991), Positive Feedback as a Search Strategy, *Rep. 91-016*, Dipartimento di Elettronica, Politecnico di Milano, Milan.

Dorigo, M., V. Maniezzo and A. Colorni (1996). "Ant system: optimization by a colony of cooperating agents." IEEE Transactions on Systems, Man, and Cybernetics–Part B, **26**(1): 29-41.

Dorigo, M. and T. Stützle (2004). Ant Colony Optimization, The MIT Press, Cambridge, Mass.

Durillo, J. J. and A. J. Nebro (2011). "jMetal: A Java framework for multi-objective optimization." Advances in Engineering Software, **42**(10): 760-771.

Durillo, J. J., A. J. Nebro, F. Luna, C. A. Coello Coello and E. Alba (2010). "Convergence speed in multi-objective metaheuristics: Efficiency criteria and empirical study." International Journal for Numerical Methods in Engineering, **84**(11): 1344-1375.

Dziegielewski, B., W. R. Mee, Jr. and K. R. Larson (1992). "Developing a long-term drought plan for Phoenix." Journal of American Water Works Association, **84**(10): 46-51.

Engelbrecht, A. P. (2006). Fundamentals of computational swarm intelligence, Wiley.

Erlanger, P. and B. Neal (2005). Framework for Urban Water Resource Planning, WSAA.

Erlenkotter, D. (1973). "Sequencing Expansion Projects." Operations Research, **21**(2): 542-553.

Erlenkotter, D. and R. R. Trippi (1976). "Optimal Investment Scheduling with Price-Sensitive Dynamic Demand." Management Science, **23**(1): 1-11.

Eskandari, H., C. Geiger and G. Lamont (2007). FastPGA: A Dynamic Population Sizing Approach for Solving Expensive Multiobjective Optimization Problems. In Evolutionary Multi-Criterion Optimization. S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu and T. Murata, eds. Springer Berlin / Heidelberg, **4403:** 141-155.

Evenson, D. E. and J. C. Moseley (1970). "Simulation/Optimization Techniques for Multi-Basin Water Resource Planning." JAWRA Journal of the American Water Resources Association, **6**(5): 725-736.

Farmani, R., W. Godfrey and D. A. Savic (2006). "Evolutionary multi-objective optimization of the design and operation of water distribution network: total cost vs. reliability vs. water quality." Journal of Hydroinformatics, **8**(3): 165-179.

Fonseca, C. M. and P. J. Fleming (1993). Genetic Algorithms for Multiobjective Optimization: Formulation, Discussion and Generalization. Genetic Algorithms: Proceedings of the Fifth International Conference. San Mateo, CA**,** 416–423.

Fonseca, C. M., L. Paquete and M. López-Ibáñez (2006). An Improved Dimension-Sweep Algorithm for the Hypervolume Indicator. IEEE Congress on Computation, 2006. CEC 2006, 1157-1163.

Fulkerson, D. R. (1961). "An Out-of-Kilter Method for Minimal-Cost Flow Problems." Journal of the Society for Industrial and Applied Mathematics, **9**(1): 18-27.

García-Martínez, C. (2004). An Empirical Analysis of Multiple Objective Ant Colony Optimization Algorithms for the Bi-criteria TSP. ANTS 2004, LNCS 3172. M. D. e. al., Springer-Verlag Berlin Heidelberg**,** 61–72.

Geist, A. (1994). PVM: Parallel Virtual Machine: a Users' Guide and Tutorial for Networked Parallel Computing, MIT Press, Cambridge, Mass.

Gillig, D., B. A. McCarl and F. O. Boadu (2001). "An Economic, Hydrologic, and Environmental Assessment of Water Management Alternative Plans for The South Central Texas Region." Journal of Agricultural and Applied Economics, **33**(01): 20.

Grant, T. R. and P. D. Temple-Smith (2003). "Conservation of the Platypus, *Ornithorhynchus anatinus* : Threats and Challenges." Aquatic Ecosystem Health & Management, **6**(1): 5-18.

Gravel, M., W. L. Price and C. Gagné (2002). "Scheduling continuous casting of aluminum using a multiple objective ant colony optimization metaheuristic." European Journal of Operational Research, **143**(1): 218-229.

Grossman, D. and D. H. Marks (1977). "Capacity expansion for water resource facilities: a control theoretic algorithm for nonstationary uncertain demands." Advances in Water Resources, **1**(1): 31-40.

Hadka, D. and P. Reed (2011). "Diagnostic Assessment of Search Controls and Failure Modes in Many-Objective Evolutionary Optimization." Evolutionary Computation, 1-31.

Horn, J., N. Nafpliotis and D. E. Goldberg (1994). A niched Pareto genetic algorithm for multiobjective optimization. IEEE World Congress on Computational Intelligence, **1:** 82-87.

Hsu, N. S., W. C. Cheng, W. M. Cheng, C. C. Wei and W. W. G. Yeh (2008). "Optimization and capacity expansion of a water distribution system." Advances in Water Resources, **31**(5): 776-786.

Huang, V. L., P. N. Suganthan and J.J. Liang (2006). "Comprehensive learning particle swarm optimizer for solving multiobjective optimization problems." International Journal of Intelligent Systems, **21**(2): 209-226.

Hydrologic Engineering Center (1998). HEC-5 Simulation of Flood Control and Conservation Systems, User's Manual, Version 8, U.S. Army Corps of Engineers, Davis, CA, USA. Hydrologic Engineering

Ilich, N. (2009). "Limitations of Network Flow Algorithms in River Basin Modeling." Journal of Water Resources Planning and Management, **135**(1): 48-55.

Iredi, S. (2001). Bi-Criterion Optimization with Multi Colony Ant Algorithms. The First International Conference on Evolutionary Multi-Criterion Optimization. E. Z. e. al., Springer-Verlag Berlin Heidelberg 2001**,** 359-372.

Jaimes, A. L. P. and C. A. C. Coello (2007). "MRMOGA: a new parallel multi-objective evolutionary algorithm based on the use of multiple resolutions." Concurrency and Computation: Practice and Experience, **19**(4): 397-441.

Jenkins, M. W. and J. R. Lund (2000). "Integrating yield and shortage management under multiple uncertainties." Journal of Water Resources Planning and Management, **126**(5): 288-297.

Kandil, A. and K. El-Rayes (2005). "Parallel Computing Framework for Optimizing Construction Planning in Large-Scale Projects." Journal of Computing in Civil Engineering, **19**(3): 304-312.

Kasprzyk, J. R., P. M. Reed, B. R. Kirsch and G. W. Characklis (2009). "Managing population and drought risks using many-objective water portfolio planning under uncertainty." Water Resourorces Research, **45**: W12401.

Kennedy, J. and R. Eberhart (1995). Particle swarm optimization. IEEE International Conference on Neural Networks, 1995.

Kennedy, J. F., J. Kennedy, R. C. Eberhart and Y. Shi (2001). Swarm intelligence, Morgan Kaufmann Publishers.

Kennington, J. L. and R. V. Helgason (1980). Algorithms for Network Programming, John Wiley & Sons, Inc. New York, NY, USA.

Khaliquzzaman and C. Subhash (1997). "Network Flow Programming Model for Multireservoir Sizing." Journal of Water Resources Planning and Management, **123**(1): 15-25.

Kim, S. K. and W. W. G. Yeh (1986). "A Heuristic Solution Procedure for Expansion Sequencing Problems." Water Resources Research, **22**(8): 1197-1206.

Kim, T., J.-H. Heo, D.-H. Bae and J.-H. Kim (2008). "Single-reservoir operating rules for a year using multiobjective genetic algorithm." Journal of Hydroinformatics, **10**(2): 163-179.

Kim, T., J.-H. Heo and C.-S. Jeong (2006). "Multireservoir system optimization in the Han River basin using multi-objective genetic algorithms." Hydrological Processes, **20**(9): 2057-2075.

Knowles, J. (2006). "ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems." IEEE Transactions on Evolutionary Computation, **10**(1): 50-66.

Knudsen, J. and D. Rosbjerg (1977). "Optimal Scheduling of Water Supply Projects." Nordic Hydrology, **8**(3): 171-192.

Ko, S.-K., D. G. Fontane and J. W. Labadie (1992). "Multiobjective Optimization of Reservoir Systems Operation." Journal of the American Water Resources Association, **28**(1): 111-127.

Kollat, J. B. and P. M. Reed (2006). "Comparing state-of-the-art evolutionary multi-objective algorithms for long-term groundwater monitoring design." Advances in Water Resources, **29**(6): 792-807.

Koshimizu, H. and T. Saito (2009). Parallel ant colony optimizers with local and global ants. International Joint Conference on Neural Networks. IJCNN 2009.

Kuczera, G. (1987). "On maximum likelihood estimators for the multisite lag-one streamflow model: Complete and incomplete data cases." Water Resources Research, **23**(4): 641-645.

Kuczera, G. (1992). "Water supply headworks simulation using network linear programming." Advances in Engineering Software, **14**(1): 55-60.

Kuczera, G. (1993). "Network Linear Programming Codes for Water-Supply Headworks Modeling." Journal of Water Resources Planning and Management, **119**(3): 412-417.

Kuczera, G., L. Cui, R. Grilmore, A. Graddon, S. M. Mortazavi and C. E. Jefferson (2009). Addressing the Shortcomings of Water Resource Simulation Models Based on Network Linear Programming. H2009, The 32nd Hydrology and Water Resources Symposium, Newcastle, Australia, Engineers Australia.

Kuczera, G. and G. Diment (1988). "General Water Supply System Simulation Model: WASP." Journal of Water Resources Planning and Management, **114**(4): 365-382.

Kumar, D. and M. Reddy (2006). "Ant Colony Optimization for Multi-Purpose Reservoir Operation." Water Resources Management, **20**(6): 879-898.

Labadie, J. W. (2004). "Optimal Operation of Multireservoir Systems: State-of-the-Art Review." Journal of Water Resources Planning and Management, **130**(2): 93-111.

Labadie, J. W., D. A. Bode and A. M. Pineda (1986). "Network Model For Decision Support In Municipal Raw Water Supply." JAWRA Journal of the American Water Resources Association, **22**(6): 927-940.

Laumanns, M., L. Thiele, K. Deb and E. Zitzler (2002). "Combining Convergence and Diversity in Evolutionary Multiobjective Optimization." Evolutionary Computation, **10**(3): 263-282.

Liang, Q., L. E. Johnson and Y.-S. Yu (1996). "A Comparison of Two Methods For Multiobjective Optimization For Reservoir Operation." Journal of the American Water Resources Association, **32**(2): 333-340.

López-Ibáñez, M., L. Paquete and T. Stützle (2004). On the Design of ACO for the Biobjective Quadratic Assignment Problem. In Ant Colony, Optimization and Swarm Intelligence, **3172/2004:** 214-225.

López-Ibáñez, M. and T. Stützle (2010). The impact of design choices of multiobjective ant colony optimization algorithms on performance: an experimental study on the biobjective TSP. Proceedings of the 12th annual conference on genetic and evolutionary computation. Portland, Oregon, USA, ACM**,** 71-78.

Lund, J. R. (1987). "Evaluation and Scheduling of Water Conservation." Journal of Water Resources Planning and Management, **113**(5): 696-708.

Luss, H. (1982). "Operations Research and Capacity Expansion Problems: A Survey." Operations Research, **30**(5): 907-947.

Mahmoud, M. R. (2006). "High dimension dynamic programming model for water resources expansion projects." Engineering Optimization, **38**(3): 371-389.

Maier, H. R., A. R. Simpson, A. C. Zecchin, W. K. Foong, K. Y. Phang, H. Y. Seah and C. L. Tan (2003). "Ant Colony Optimization for Design of Water Distribution Systems." Journal of Water Resources Planning and Management, **129**(3): 200-209.

Manfrin, M., M. Birattari, T. Stützle and M. Dorigo (2006). Parallel Ant Colony Optimization for the Traveling Salesman Problem. Lecture Notes in Computer Science M. Dorigo, L. Gambardella, M. Birattari, A. Martinoli, R. Poli and T. Stützle, Springer Berlin / Heidelberg, **4150:** 224-234.

Mariano-Romero, C. E., V. c. H. Alcocer-Yamanaka and E. F. M. b (2007). "Multi-objective optimization of water-using systems." European Journal of Operational Research, **181**: 1691-1707.

Martínez, C. G. a., O. Cordón and F. Herrera (2007). "A taxonomy and an empirical analysis of multiple objective ant colony optimization algorithms for the bi-criteria TSP." European Journal of Operational Research, **180**: 116–148.

Matalas, N. C. (1967). "Mathematical assessment of synthetic hydrology." Water Resources Research, **3**(4): 937-945.

Matalas, N. C. and M. B. Fiering (1977). Water-resource System planning. Climate, Climatic Change, and Water Supply, Natl. Acad. of Sci., Washington, D.C., 99–110.

McMullen, P. R. (2001). "An ant colony optimization approach to addressing a JIT sequencing problem with multiple objectives." Artificial Intelligence in Engineering, **15**(3): 309-317.

Morin, T. L. (1973). "Optimal Sequencing of Capacity Expansion Projects." Journal of the Hydraulics Division, **99**(9): 1605-1622.

Morin, T. L. and A. M. O. Esogbue (1971). "Some Efficient Dynamic Programing Algorithms for the Optimal Sequencing and Scheduling of Water Supply Projects." Water Resources Research, **7**(3): 479-484.

Mortazavi N., S. M., L.-J. Cui and G. Kuczera (2009). Comparison of Genetic Algorithm and Ant Colony optimization methods for optimization of short-term drought mitigation strategies. Hydroinformatics in Hydrology, Hydrogeology and Water Resources. Hyderabad, India, IAHS Publ. **331,** 80-90.

Mousavi, H. and A. S. Ramamurthy (2000). "Optimal design of multi-reservoir systems for water supply." Advances in Water Resources, **23**(6): 613-624.

Nainis, W. S. and Y. Y. Haimes (1975). "A Multilevel Approach to Planning for Capacity Expansion in Water Resource Systems." IEEE Transactions on Systems, Man and Cybernetics, **SMC-5**(1): 53-63.

Nakashima, M., J. Harry G. Wenzel and J. E. Downey Brill (1986). "Water Supply System Models with Capacity Expansion." Journal of Water Resources Planning and Management, **112**(1): 87-103.

Nebro, A., J. Durillo, C. Coello Coello, F. Luna and E. Alba (2008). A Study of Convergence Speed in Multi-objective Metaheuristics. Parallel Problem Solving from Nature – PPSN X. G. Rudolph, T. Jansen, S. Lucas, C. Poloni and N. Beume, Springer Berlin / Heidelberg, **5199:** 763-772.

Nebro, A. J., J. J. Durillo, J. Garcia-Nieto, C. A. Coello Coello, F. Luna and E. Alba (2009). SMPSO: A new PSO-based metaheuristic for multi-objective optimization. IEEE symposium on computational intelligence in multi-criteria decision-making, 66-73.

New South Wales Dept. of Environment, C. C. and Water (2010). 2010 Metropolitan water plan summary / Dept. of Environment, Climate Change and Water ,Sydney : Dept. of Environment, Climate Change and Water.

Nicklow, J., P. Reed, D. Savic, T. Dessalegne, L. Harrell, A. Chan-Hilton, M. Karamouz, B. Minsker, A. Ostfeld, A. Singh, E. Z. A. T. C. o. E. C. i. Environmental and W. R. Engineering (2010). "State of the Art for Genetic Algorithms and Beyond in Water Resources Planning and Management." Journal of Water Resources Planning and Management, **136**(4): 412-432.

O'Laoghaire, D. T. and D. M. Himmelblau (1974). Optimal Expansion of a Water Resources System, Academic Press Inc.

O'Hara, J. and K. Georgakakos (2008). "Quantifying the Urban Water Supply Impacts of Climate Change." Water Resources Management, **22**(10): 1477-1497.

Pacheco, P. S. (1997). Parallel Programming With MPI, Morgan Kaufmann Publishers.

Palmer, M. A., C. A. Reidy Liermann, C. Nilsson, M. Flörke, J. Alcamo, P. S. Lake and N. Bond (2008). "Climate change and the world's river basins: anticipating management options." Frontiers in Ecology and the Environment, **6**(2): 81-89.

Pellegrini, P., T. Stützle and M. Birattari (2012). "A critical analysis of parameter adaptation in ant colony optimization." Swarm Intelligence, **6**(1): 23-48.

Perera, B. J. C., B. James and M. D. U. Kularathna (2005). "Computer software tool REALM for sustainable water allocation and management." Journal of Environmental Management, **77**(4): 291-300.

Petts, G. E. (2009). "Instream Flow Science For Sustainable River Management." JAWRA Journal of the American Water Resources Association, **45**(5): 1071-1086.

Pierro, F. d., S.-T. Khu, D. Savic and L. Berardi (2009). "Efficient multi-objective optimal design of water distribution networks on a budget of simulations using hybrid algorithms." Environmental Modelling & Software, **24**(2): 202-213.

Pierro, F. d., K. Soon-Thiam, S. Dragan and B. Luigi (2009). "Efficient multi-objective optimal design of water distribution networks on a budget of simulations using hybrid algorithms." Environmental Modelling & Software, **24**(2): 202-213.

Poli, R., J. Kennedy and T. Blackwell (2007). "Particle swarm optimization." Swarm Intelligence, **1**(1): 33-57.

Randall, D., M. H. Houck and J. R. Wright (1990). "Drought Management of Existing Water Supply System." Journal of Water Resources Planning and Management, **116**(1): 1-20.

Randall, M. (2004). Near Parameter Free Ant Colony Optimisation Ant Colony Optimization and Swarm Intelligence. Lecture Notes in Computer Science. M. Dorigo, M. Birattari, C. Blum, L. Gambardella, F. Mondada and T. Stützle, Springer Berlin / Heidelberg, 3172: 262-285.

Rani, D. and M. M. Moreira (2010). "Simulation-optimization modeling: A survey and potential application in reservoir systems operation." Water Resources Management, **24**(6): 1107-1138.

Razavi, S., B. A. Tolson and D. H. Burn (2012). "Numerical assessment of metamodelling strategies in computationally intensive optimization." Environmental Modelling & Software, **34**: 67–86.

Reddy, M. and D. Kumar (2006). "Optimal Reservoir Operation Using Multi-Objective Evolutionary Algorithm." Water Resources Management, **20**(6): 861-878.

Reddy, M. J. and D. N. Kumar (2007a). "Multi-objective particle swarm optimization for generating optimal trade-offs in reservoir operation." Hydrological Processes, **21**(21): 2897–2909.

Reddy, M. J. and D. N. Kumar (2007b). "Multiobjective Differential Evolution with Application to Reservoir System Optimization." Journal of Computing in Civil Engineering, **21**(2): 136-146.

Reyes-Sierra, M. and C. A. C. Coello (2006). "Multi-Objective Particle Swarm Optimizers: A Survey of the State-of-the-Art." International Journal of Computational Intelligence Research, **2**(3): 287–308.

Richter, B. D., A. T. Warner, J. L. Meyer and K. Lutz (2006). "A collaborative and adaptive process for developing environmental flow recommendations." River Research and Applications, **22**(3): 297-318.

Rosenberg, D. E., R. E. Howitt and J. R. Lund (2008). "Water management with water conservation, infrastructure expansions, and source variability in Jordan." Water Resources Research, **44**(11): W11402.

Rosenzweig, C., G. Casassa, D. J. Karoly, A. Imeson, C. Liu, A. Menzel, S. Rawlins, T. L. Root, B. Seguin and P. Tryjanowski (2007). Assessment of observed changes and responses in natural and managed systems. Climate Change 2007: Impacts, Adaptation and Vulnerability Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. M. L. Parry, O. F. Canziani, J. P. Palutikof, P. J. v. d. Linden and C. E. Hanson, Cambridge University Press**,** 79-131.

Rubinstein, J. and L. Ortolano (1984). "Water Conservation and Capacity Expansion." Journal of Water Resources Planning and Management, **110**(2): 220-237.

Salas, J. D., C. Fu, A. Cancelliere, D. Dustin, D. Bode, A. Pineda and E. Vincent (2005). "Characterizing the Severity and Risk of Drought in the Poudre River, Colorado." Journal of Water Resources Planning and Management, **131**(5): 383-393.

Salewicz, K. A., D. P. Loucks and C. Gandolfi (1991). "IRIS: Interactive River Simulation program." Environmental Software, **6**(1): 5-10.

Santana-Quintero, L., N. Ramírez-Santiago, C. Coello, J. Luque and A. Hernández-Díaz (2006). A New Proposal for Multiobjective Optimization Using Particle Swarm Optimization and Rough Sets Theory. Parallel Problem Solving from Nature – PPSN IX. T. Runarsson, H.-G. Beyer, E. Burke, J. Merelo-Guervós, L. Whitley and X. Yao, Springer Berlin / Heidelberg, **4193:** 483-492.

Scott, A. and T. Grant (1997), Impacts of water management in the Murray-Darling Basin on the platypus (Ornithorhynchus anatinus) and the water rat (Hydromys chrysogaster) *Rep.*, CSIRO Land and Water, Canberra.

Shelokar, P. S., V. K. Jayaraman and B. D. Kulkarni (2002). "Ant algorithm for single and multiobjective reliability optimization problems." Quality and Reliability Engineering International, **18**(6): 497-514.

Shiau and Wu (2007). "Pareto-optimal solutions for environmental flow schemes incorporating the intra-annual and interannual variability of the natural flow regime." Water Resources Research, **43**: W06433.

Shiau, J.-T. (2009). "Optimization of Reservoir Hedging Rules Using Multiobjective Genetic Algorithm." Journal of Water Resources Planning and Management, **135**(5): 355-363.

Shih, J.-S. and C. ReVelle (1994). "Water-Supply Operations During Drought: Continous Hedging Rule." Journal of Water Resources Planning and Management, **120**(5): 613-629.

Shih, J.-S. and C. ReVelle (1995). "Theory and Methodology Water supply operations during drought: A discrete hedging rule." European Journal of Operational Research, **82**: 163-175.

Shourian, M., S. Mousavi and A. Tahershamsi (2008). "Basin-wide Water Resources Planning by Integrating PSO Algorithm and MODSIM." Water Resources Management, **22**(10): 1347-1366.

Sigvaldson, O. T. (1976). "A Simulation Model for Operating a Multipurpose Multireservoir System." Water Resources Research, **12**(2): 263-278.

Srinivas, N. and K. Deb (1994). "Multiobjective Optimization Using Nondominated Sorting in Genetic Algorithms." Evolutionary Computation, **2**(3): 221-248.

Stützle, T. (1998). Parallelization Strategies for Ant Colony Optimization. Fifth International Conference on Parallel Problem Solving from Nature, Springer-Verlag**,** 722-731.

Stützle, T. and H. Hoos (1997). MAX-MIN Ant System and local search for the traveling salesman problem. IEEE International Conference on Evolutionary Computation**,**309-314.

Stützle, T. and H. Hoos (1996), Improving the Ant System: A Detailed Report on the MAX-MIN Ant System, *Rep. AIDA-66-12*, Technical university of Darmstadt.

Stützle, T. and H. H. Hoos (2000). "MAX–MIN Ant System." Future Generation Computer Systems, **16**(8): 889-914.

Stützle, T., M. López-Ibáñez, M. Dorigo, J. J. Cochran, L. A. Cox, P. Keskinocak, J. P. Kharoufeh and J. C. Smith (2010). A Concise Overview of Applications of Ant Colony Optimization. Wiley Encyclopedia of Operations Research and Management Science, John Wiley & Sons, Inc.

Stützle, T., M. López-Ibáñez, P. Pellegrini, M. Maur, M. M. d. Oca, M. Birattari and M. dorigo (2010), Parameter Adaptation in Ant Colony Optimization, *Rep. TR/IRIDIA/2010-002*.

Suen, J. P. and J. W. Eheart (2006). "Reservoir management to balance ecosystem and human needs: Incorporating the paradigm of the ecological flow regime." Water Resources Research, **42**(3).

T'Kindt, V., N. Monmarché, F. Tercinet and D. Laügt (2002). "An Ant Colony Optimization algorithm to solve a 2-machine bicriteria flowshop scheduling problem." European Journal of Operational Research, **142**(2): 250-257.

Tharme, R. E. (2003). "A global perspective on environmental flow assessment: emerging trends in the development and application of environmental flow methodologies for rivers." River Research and Applications, **19**(5-6): 397-441.

Thyer, M., A. J. Frost and G. Kuczera (2006). "Parameter estimation and model identification for stochastic models of annual hydrological data: Is the observed record long enough?" Journal of Hydrology, **330**(1-2): 313-328.

Toscano-Pulido, G., C. Coello Coello and L. Santana-Quintero (2007). EMOPSO: A Multi-Objective Particle Swarm Optimizer with Emphasis on Efficiency. Evolutionary Multi-Criterion Optimization. S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu and T. Murata, Springer Berlin / Heidelberg, **4403:** 272-285.

Tu, M.-Y., N.-S. Hsu, F. T.-C. Tsai and W. W.-G. Yeh (2008). "Optimization of Hedging Rules for Reservoir Operations." Journal of Water Resources Planning and Management, **134**(1): 3-13.

Tu, M.-Y., N.-S. Hsu and W. W.-G. Yeh (2003). "Optimization of Reservoir Management and Operation with Hedging Rules " Journal of Water Resources Planning and Management, **129**(2): 86-97.

UN-Water and FAO (2007), Coping with Water Scarcity Challenge of Twenty-first Century, *http://www.fao.org/nr/water/docs/escarcity.pdf, last visit 07/05/2012*

UN (2011), World Population to reach 10 billion by 2100 if Fertility in all Countries Converges to Replacement Level, *Rep.*, United Nations, Department of Economic and Social Affairs, Population Division.

Van Veldhuizen, D. A. and G. B. Lamont (2000). On measuring multiobjective evolutionary algorithm performance. Evolutionary Computation, **1:** 204-211

Voivontas, D., G. Arampatzis, E. Manoli, C. Karavitis and D. Assimacopoulos (2003). "Water supply modeling towards sustainable environmental management in small islands: the case of Paros, Greece." Desalination, **156**(1-3): 127-135.

Watkins Jr, D. W. and D. C. McKinney (1998). "Decomposition methods for water resources optimization models with fixed costs." Advances in Water Resources, **21**(4): 283-295.

Wilchfort, O. and J. R. Lund (1997). "Shortage management modeling for urban water supply systems." Journal of Water Resources Planning and Management, **123**(4): 250-258.

Wurbs, R. A. (1993). "Reservoir-System Simulation and Optimization Models." Journal of Water Resources Planning and Management, **119**(4): 455-472.

Wurbs, R. A. (1995). Water Management Models: A Guide to Software, Prentice Hall.

Yang, C. C., L. C. Chang, C. H. Yeh and C. S. Chen (2007). "Multiobjective planning of surface water resources by multiobjective genetic algorithm with constrained differential dynamic programming." Journal of Water Resources Planning and Management, **133**(6): 499-508.

Yeh, W. W. G. (1985). "Reservoir Management and Operations Models: A State-of-the-Art Review." Water Resources Research, **21**(12): 1797-1818.

Zagona, E. A., T. J. Fulp, R. Shane, T. Magee and H. M. Goranflo (2001). "RIVERWARE: A Generalized Tool For Complex Reservoir System Modeling." JAWRA Journal of the American Water Resources Association, **37**(4): 913-929.

Zecchin, A. C., A. R. Simpson, H. R. Maier and J. B. Nixon (2005). "Parametric study for an ant algorithm applied to water distribution system optimization." IEEE Transactions on Evolutionary Computation, **9**(2): 175-191.

Zitzler, E., K. Deb and L. Thiele (2000). "Comparison of multiobjective evolutionary algorithms: Empirical results." Evolutionary Computation, **8**(2): 173–195.

Zitzler, E., M. laummans, L. Thiele, C. M. Fonseca and V. Grunert da Fonseca (2002). Why Quality Assessment Of Multiobjective Optimizers Is Difficult. GECCO 2002**,** 666-674.

Zitzler, E. and L. Thiele (1998). Multiobjective Optimization Using Evolutionary Algorithms - A Comparative Case Study. Parallel Problem Solving from Nature — PPSN V**,** 292-301.