



# **Mathematical Applications for Conservation Ecology: The Dynamics of Tree Hollows and the Design of Nature Reserves**

Ian R Ball

Submitted in partial fulfilment of the requirements of a PhD.  
University of Adelaide,  
Department of Applied Mathematics

# Table of Contents

TOC	i
Abstract	iii
Acknowledgments	iv
Declaration	v
Papers arising from work	vi

## **Part 1: HOLSIM the tree hollow simulator**

### **Seeing the Hollows for the Trees**

Section 1	Introduction	1
Section 2	The Model	5
Section 3	The Scenarios	23
Section 4	Results	
4.1	Null run (Introducing the Scenarios)	25
4.2	Proportional Harvesting	27
4.3	Retaining a Fixed Number of Trees with Hollows	31
4.4	Retaining Trees with or without Hollows	34
4.5	Retained Tree Mortality	37
4.6	More Detailed Regimes	40
Section 5	Sensitivity Analysis	45
Section 6	Conclusions	57

**Part 2: The Nature of Reserve Design and  
the Design of Nature Reserves**

Section 1: Introduction	63
Section 2: Literature Review	66
Section 3: Formal Problem Statement	81
Section 4: Methods used in this Thesis	93
Section 5: Data sets	105
Section 6: Comparison of Different Solution Methods.	111
Section 7: Cost Threshold	123
Section 8: Incremental Reserve Design	134
Section 9: Fragmentation	141
Section 10: Spatial Rules	150
Section 11: Conclusions	158

**Bibliography**

Tree Hollow Dynamics Bibliography	171
Nature Reserve Design Bibliography	176

## Abstract

The first part of this thesis describes a deterministic computer model for simulating forest dynamics. The model predicts the long term dynamics of hollow bearing trees which occur in a single species (monotypic) forest stand under an array of different timber harvesting regimes over a time scale of centuries. It is applied to a number of different timber harvesting scenarios in the mountain ash (*Eucalyptus regnans* F.Muell.) forests of Victoria, south-eastern Australia. The results have far-reaching implications for forest management. These include: 1) When the harvest rotation time is 100 years or less, a supply of trees with hollows cannot be ensured by only retaining trees which already have hollows; and 2) When some retained trees are lost through logging related mortality, the effect on the number of trees with hollows is exaggerated. For instance if half of the retained trees are lost via logging related mortality, it is not sufficient to double the number of trees retained in order to maintain the same number of hollow bearing trees.

The second part of the thesis looks at a number of new mathematical problems in the design of nature reserve systems. The basic problem is to select a set of sites from a list to try to meet the representation requirements of a set of species which occur on these sites for as small a cost or number of sites as possible. After comparing a number of methods for solving basic problems a number of new problems are introduced. These include: Fixing the cost or size of the reserve system and then trying to maximise species coverage; Building a reserve system up in stages with species requirements incremented - this is another way of controlling the size or cost of the reserve system and also allows the cost - biodiversity trade-off to be examined; Controlling fragmentation of the reserve system by trying to minimise the boundary length as well as the size of the reserve; and introducing spatial requirements for each species. These requirements are that individual patches with the species be large enough to be viable on their own, as well as having some patches separated by a great enough distance that they are not all likely to be destroyed in a local catastrophe.



## Acknowledgments

This work was in part funded by Environment Australia. I would like to thank Andrew Taplin from Environment Australia, David Lindenmayer, Drew Tyre, Hugh Possingham (of course) and Donna, Callum and Ethan.

*Now small fowls flew screaming over the yet yawning gulf; a sullen white surf beat against its steep sides; then all collapsed, and the great shroud of the sea rolled on as it rolled five thousand years ago.*

Herman Melville, *Moby Dick*

# Declaration

This work 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 available for loan and photocopying.

Ian Ball

Date 9-3-2000

## Publications arising from this work

### Already Published

Ball, I. R., Lindenmayer, D. B., Possingham, H. P., (1997) "A tree hollow simulation model for forest managers: The dynamics of the absence of wood in trees." Proceedings International Congress on Modelling and Simulation (MODSIM 97) Eds: McDonald, A. D., and McAleer, M.

Ball, I. R., Possingham, H. P., and Lindenmayer, D. B., (1996) "Modelling of retained trees in logged forests" (ANCA: Canberra)

Ball, I. R., Lindenmayer, D. B., Possingham H. P. (1999) "A tree hollow dynamics simulation model for forest managers: Seeing the hollows for the trees." *Forest Ecology and Management*, **123**, 179-194.

### In Press


Ball, I. R., Smith, A., Day, J. R., Pressey, R. L., Possingham, H. "Comparison of mathematical algorithms for the design of a reserve system for nature conservation: An application of genetic algorithms and simulated annealing." In press: *Journal of Environmental Management*.

Possingham, H., Ball, I. R., Andelman, S. (2000) "Mathematical methods for identifying representative reserve networks" in *Quantitative methods for conservation biology*. Ferson, S., and Burgman, M. (eds). Springer-Verlag, New York.

### Submitted

Leslie, H., Ruckelshaus, M., Ball, I. R., Possingham, H. P. "Using siting algorithms in the design of marine reserve networks", *submitted to Ecological Applications*

# Part 1



**HOLSIM the tree hollow simulator:  
Seeing the Hollows for the Trees.**



# Section 1

## Introduction

The potential impacts of timber harvesting operations on the conservation of biodiversity has become a major planning and policy issue in the management of Australian Eucalypt forests (Resource Assessment Commission, 1992; Commonwealth of Australia, 1992; Department of the Environment, Sports and Territories, 1995). In a typical unharvested forest there will be trees representing all stages in the life cycle including trees with hollows. Tree hollows include fissures, and cracks in trees, holes and hollow branches, and they are occasionally termed cavities in the literature (Lindenmayer *et al.* 1993). These hollows are extremely important in the maintenance of animal diversity as a wide of animal species use them. In Australia alone it is estimated that approximately 400 vertebrate species need hollows (Ambrose, 1982). Trees with hollows are a characteristic component of forest ecosystems, not only in Australia, (Scotts, 1991; Lindenmayer *et al.*, 1993) but in many places around the world (Newton, 1994). They provide a key habitat component for a wide array of vertebrate and invertebrate taxa (Gibbons and Lindenmayer, 1996) and provide nesting and denning sites as well as places for animals to roost and perch (Lindenmayer and Franklin, 1998) for a wide range of purposes including shelter, protection and breeding

Species which use hollows in trees are considered to be among those most vulnerable to the impacts of timber harvesting (McIlroy, 1978; Recher *et al.*, 1980; Lindenmayer *et al.*, 1990a, 1991a, 1991b; Scotts, 1991; Gibbons and Lindenmayer, 1996). This is because hollows suitable for occupation by wildlife may take several hundred years to develop (Ambrose, 1982; Saunders *et al.*, 1982; Lindenmayer *et al.*, 1991c; 1993; Gibbons 1994), this and the type of logging operations used, as well as the interval between harvesting events, may prevent or severely impair the development of trees with hollows (Lindenmayer *et al.*, 1990b; Recher, 1996; Gibbons and Lindenmayer, 1996). Logging typically results in substantial changes to forest structure (Shaw, 1983; Recher, 1985a, 1985b, 1995; Lindenmayer *et al.*, 1990b; Lindenmayer 1992a, 1992b, 1994; Ough and Ross, 1992), such as a smaller size-class distribution (Recher, 1985a; Lindenmayer, 1995). Such changes typically lead to a reduction in the numbers of hollow-bearing trees (McIlroy, 1978; Smith, 1985; Mackowski, 1988; Lindenmayer *et al.*, 1991c). In Australian Eucalypt forests the short period of time between

logging operations (typically 40-120 years) (Government of Victoria, 1986; van Saane and Gordon, 1991; State Forests of New South Wales, 1994) severely impedes the recruitment of hollow bearing trees. This has a negative effect on the large number of vertebrate and invertebrate taxa that are obligate or facilitate hollow-using species (Gibbons and Lindenmayer, 1996). As hollow ontogeny in Australian eucalypt trees takes a long time, rectifying hollow shortages in wood production forests will take a long time (Lindenmayer and Possingham, 1995; Recher, 1995).

Tree retention strategies can help avoid shortages in hollow bearing trees and mitigate the impact of logging on hollow dependent fauna. (Gibbons and Lindenmayer 1996). Such approaches may enable those stand attributes which are essential for wildlife to be maintained within logged forests. These key attributes can then be perpetuated through many harvest rotation cycles to maintain essential habitat components for wildlife. Consequently tree retention strategies could help maintain the populations of those species that might otherwise be permanently eliminated from logged areas, or even lost from the entire wood production forest. Studies of a range of different types of eucalypt forest have highlighted the value of retained trees in facilitating the recolonisation of logged sites by birds (for example Loyn *et al.*, 1980; Recher *et al.*, 1980; Smith, 1985; Kavanagh and Turner, 1994 ).

A detailed review of a range of aspects of tree retention strategies has indicated that many key issues remain to be resolved (Gibbons and Lindenmayer, 1996). For example, in most Australian wood production forests, there is insufficient field data to determine the number, type and spatial arrangement of trees that should be retained (Gibbons and Lindenmayer, 1996). Furthermore, effective ways to protect trees during and after logging operations have yet to be determined (Gibbons, 1994). The long-term effectiveness of tree retention strategies remains unknown. The long term nature of hollow dynamics makes it difficult to assess the efficacy of tree retention strategies in the field. It also means that it is difficult for forest and wildlife managers to forecast the stand conditions which will occur well into the future that result from management actions taken now. The problem of visualising and understanding the long-term changes to forest architecture, particularly through successive harvest rotations, is complex and the value of good forest models becomes obvious.

The aim of this section of work is to develop a computer simulation transition matrix model which explores and forecasts the long-term changes in stand structure. It is hoped that such a model will overcome the current difficulty of predicting the impact of different harvesting

scenarios on hollow dynamics. This model is different from many other stand models that have been developed because it is focussed on the development of hollows and simulates forest conditions well after growth rates have peaked and trees have entered prolonged periods of senescence. The model is of interest from a modelling perspective because its structure is determined by the data that is available or readily obtainable, and the questions of direct interest to managers. For example, stand development is modelled in a phenomenological fashion, ignoring key driving processes like water, light, and nutrients. This is because data on the influences of these processes on the long term growth of trees and formation of hollows is rare or non-existent for most tree species. This is the first model to explicitly include hollow ontogeny and measures of tree hollows for managed stands.

The model is applied using data from a detailed case study of Mountain Ash. *Eucalyptus regnan* forest of the central highlands of Victoria, south-eastern Australia. Mountain Ash typically occurs as a single stand (Ashton, 1976), making it inappropriate to use any of the many forest succession models (see for example Acevedo *et al.*, 1996; Lin *et al.*, 1996; Osho 1996). Mountain Ash stands are fast growing and regenerate only after a disturbance such as fire or harvesting (Attiwill, 1994), gap formation and regrowth do not generally occur on Mountain Ash stands and hence gap models such as JABOWA (Botkin *et al.*, 1972) and FORET (Shugart, 1984) were not considered.

There are a large number of generic vegetation models based on ecological mechanisms, such as VAFS/STANDSIM (Roberts 1996), TREEDYN3 (Bossel 1996) and ForM-S (Oja *et al.*, 1995) and numerous yield growth models (see for example Rayner and Turner, 1990; Dale *et al.*, 1985). The yield growth models have been criticised because they do not extrapolate beyond the management regimes for which they have been parameterised (Voit and Sands, 1996). A generic vegetation model could have been modified to include the key mechanisms and data of interest into its structure but it was believed to be superior to let the data determine the optimal structure of the model. A forest succession model which was designed to look at arboreal marsupial habitat quality is EDEN (Pausas *et al.*, 1997). This model includes a predictor for the abundance of hollow bearing trees dependent upon stand species composition and a number of site related parameters.

This investigation is limited to tracking populations of trees with hollows in logged sites. No attempt has been made to include the abundance of hollow-dependent animals in logged sites as a function of the availability of these trees, even though strong relationships between the

abundance of trees with hollows and the presence and abundance of hollow-dependent animals have been established from a range of studies in Mountain Ash forests (for example Smith and Lindenmayer 1988; 1992; Lindenmayer *et al.*, 1990b, 1990c, 1991c, 1994, 1995). This is because the dynamics of animal distribution is considerably more complex than a simple measure of available habitat (see for example Tyre 1999).

This model will be able to answer the following types of questions: What are the long term effects of different logging strategies on the number of trees with hollows in a managed stand? For example, if 10% of a stand were left after a harvest would we get 10% of the number of hollows in an untouched forest? How would the harvest rotation length effect this? How would the effects be different if we retained a set number of trees rather than a percentage, or if we selected the trees to be retained in some fashion? How great an effect does tree loss after the harvest, due to increased exposure of retained trees to the elements, have on the number of tree hollows in the stand?

The model was tested with a number of different timber harvesting scenarios in Mountain Ash forests. It was parameterised using information generated from past studies of Mountain Ash forests (see Table 1.1) including: rates of stand growth (Dahl, 1940; Cunningham, 1960; Ashton 1975a-c; 1976; 1981), typical stocking rates of trees (Burgman *et al.*, 1995; Ambrose, 1982; Banks, 1993), patterns of cavity ontogeny (Ambrose, 1982; Lindenmayer *et al.*, 1993), and patterns of decay and collapse among hollow-bearing stems (Lindenmayer *et al.*, 1990a; 1997).

Although the work has focussed on a single forest type in one region, the model was constructed with a generic modelling framework that could readily be used to examine other forests dominated by different species of trees, and characterised by different rates of growth, mortality and silvicultural practices (for example see studies of Blackbutt, *Eucalyptus pilularis*, forest by Mackowski, 1984; Horne *et al.*, 1991).

This is the first model to explicitly include tree hollows. It has been designed around parameters which were readily available for Mountain Ash and which could be realistically obtained for other commercially important tree species. The scenarios which are examined in this thesis include: logging cycles in which a proportion of the trees on a stand are retained between each harvest; cycles where a set number of trees are retained; cycles where retained trees are of particular classes and ones where trees are retained from any class; retained tree schemes when there is increased mortality for the trees left standing after a harvest.



## Section 2

### The Model

#### General

The program models a single-species (monotypic) forest stand and simulates the composition of the forest with respect to the size and life stage of the trees. It was parameterised to model a stand of Mountain Ash forest. The model is general and with different input data it could be used for a different species of tree. It was designed to explore the impact of different harvesting regimes and tree retention systems on the number of hollow-bearing trees in a stand. It is assumed that the number of hollows is only a function of tree size and tree form and that by modelling the dynamics of the trees in a stand we can predict the dynamics of the number of hollow bearing trees using data relating number of hollows for trees of different classes. In this thesis a default stand area of 25 hectares has been used throughout, although this is controlled by the input parameters and is not a limitation on the model. The location of individual trees and the spatial distribution of classes of trees within the stand is not included in the model. For this reason the stand area is simply a scaling parameter which can be set to any area.

The length of the time step used in the simulation can be varied; we used a time step of five years. In each time step, the forest changes through a series of processes, each of which is described below. These processes are growing the trees, computing the decay of healthy trees, recruitment of new trees, tree death, the decay of damaged trees, and changes to a stand associated with timber harvesting operations. The program does not allow for any catastrophes (eg wildfires or storms) or long term environmental changes, and is completely deterministic.

For the purposes of this study, we assumed recruitment in mountain ash forests only occurs after a disturbance (logging), at which point new trees fill all the available space from the seed bank. Stands of mountain ash regenerate well after disturbances like wildfire and logging (Attiwill, 1994). Because the only disturbance allowed in the simulation is logging, a simulation without logging would be unrealistic as we have little information about the dynamics of this species when there is no disturbance for hundreds of years (Therefore, scenarios without the effects of timber harvesting were not pursued).

## Forest Structure and Tree Growth

Trees are categorised into states determined by size and condition. Because the model is deterministic, the number of trees in each state is the expected or average number of trees for that state. Thus the number of trees is not an integer value.

Each tree in the forest has two defining attributes, 1) its size class, and, 2) its form describing the condition, or life stage, of the trees (seen below in Figure 2.2). The size class category represents the size of the trees and the tree form represents advancing stages of senescence. The two categories are correlated to a degree in that the larger trees will tend to be in a higher form of senescence. Although it is still possible for small trees to be in an advanced stage of senescence due to the effects of self thinning or random effects of weather, and it is possible for quite old and large trees to have largely escaped the ravages of time.

There are 50 distinct size-classes where the largest size-class contains mature trees of approximately 250 years of age as well as trees older than this. It is scaled so that a healthy tree (one unaffected by decay and with no hollows) will increase on average one size-class in each time step (5 years). Before decay begins, the age of the tree will roughly be its size-class times the length of a time-step. The relationship between the average diameter at breast height (DBH) of a size-class and the average age of the size-class is calculated by a formula given by Ashton (1976):

$$\text{Log } Y = 1.02 \text{ Log } X, \quad (2.1)$$

or equivalently

$$Y = X^{1.02}, \quad (2.2)$$

where Y is the DBH in centimetres and X is the age in years this relationship is shown in Figure 2.1




Figure 2.1 Relationship between size class and DBH. The diameter at breast height is equal to size class of the tree divided by 5 and raised to the power of 1.02. Note that size-class 50 is the highest size class so that trees do not continue increasing in DBH indefinitely.

The form of trees ranged from zero to eight reflecting advancing stages of senescence (Figure 2.2). The forms are defined as follows (after Smith and Lindenmayer, 1988): Form 0, living tree without hollows; Form 1, mature living tree with hollows; Form 2 mature living tree with a dead or broken top; Form 3, dead tree with most branches still intact; Form 4, dead tree with 0 - 25% of the top broken off, branches remaining as stubs only; Form 5, dead tree with the top 25-50% broken away; Form 6, dead tree with the top 50-75% broken away; Form 7, solid dead tree with greater than or equal to 75% of the top broken away; Form 8 hollow stump.




Figure 2.2 Tree forms 0 through 8. Tree form 0 is a living tree which has not formed any hollows; Form 1 is a mature living tree with hollows; Form 2, mature living tree with a dead or broken top; Form 3, dead tree with most branches still intact; Form 4, dead tree with 0 - 25% of top broken off, branches remaining as stubs only; Form 5, dead tree with top 25 - 50% broken away; Form 6, dead tree with top 50 - 75% broken away; Form 7, solid dead tree with greater than or equal to 75% of the top broken away; Form 8 hollow stump (Modified from Smith and Lindenmayer, 1988).

## Transition Processes

During each time step, the distribution of trees in the forest changes through a series of transition processes. The transition processes cause a fraction of trees in each state to flow into a variety of other states. These processes include: senescence (using the form transition matrix shown in Table 2 from Lindenmayer *et al.*, 1998), mortality through overcrowding, density independent mortality, growth, recruitment, and harvesting. The processes that cause trees to change state are summarised in Figure 2.3 and described below.




Figure 2.3. Schematic diagram of transition processes. Trees are divided into size classes from 1 to 50 and tree forms from 0 to 8. The main processes are; overcrowding mortality which transfers trees into tree form 3 (dead trees with most branches intact, see Figure 2) or removes them from the system; growth increases the size class of trees and growth at a reduced rate increases trees of forms 1 and 2 (living trees with hollows and living trees with missing crown); density independent mortality moves trees from form 0 to form 1 (living trees without hollows to living trees with hollows); senescence is the general process which increases the forms of trees until they eventually collapse and are removed from the system. Note that mortality and senescence moves trees to the right and growth moves trees downwards - the only two valid directions.

## Senescence and Changing Tree Form

Senescence is the first process applied in each time step. It includes the senescence of already damaged trees (tree forms 1...8) and density independent mortality (the mortality of healthy trees which is not caused by overcrowding) (see Figure 2.4). These transitions are independent of the composition of the forest in the model.

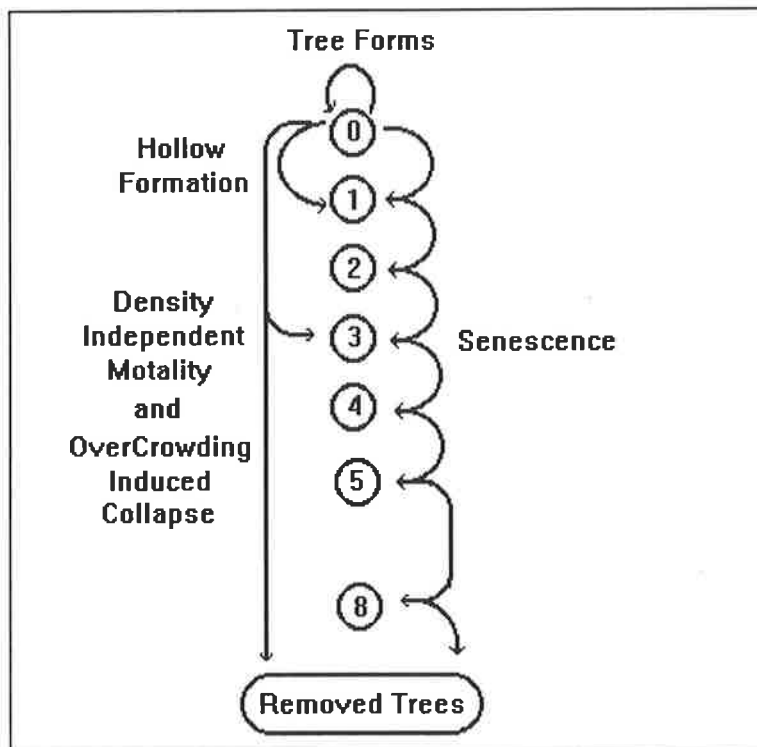


Figure 2.4. Senescence processes. All trees start in tree form 0 and will increase their tree form until they are eventually removed. Hollow formation moves trees from form 0 (living with no hollows) to form 1 (living with hollows). Overcrowding either removes trees completely (through immediate collapse) or transfers them to form 3 (dead trees with most branches intact).

Mature trees (at or above size class 18) can form hollows and hence be transferred from tree form 0 to tree form 1 (mature living trees with hollows). The rate at which they are transferred is a linear function of stem basal area and hence a quadratic function of size class. No trees form hollows at size class 18 and 5% at size class 50, in any given time step (Figure 2.5). This value was estimated by calibrating the model with actual forests of known age (Lindenmayer *et al.*, 1993).

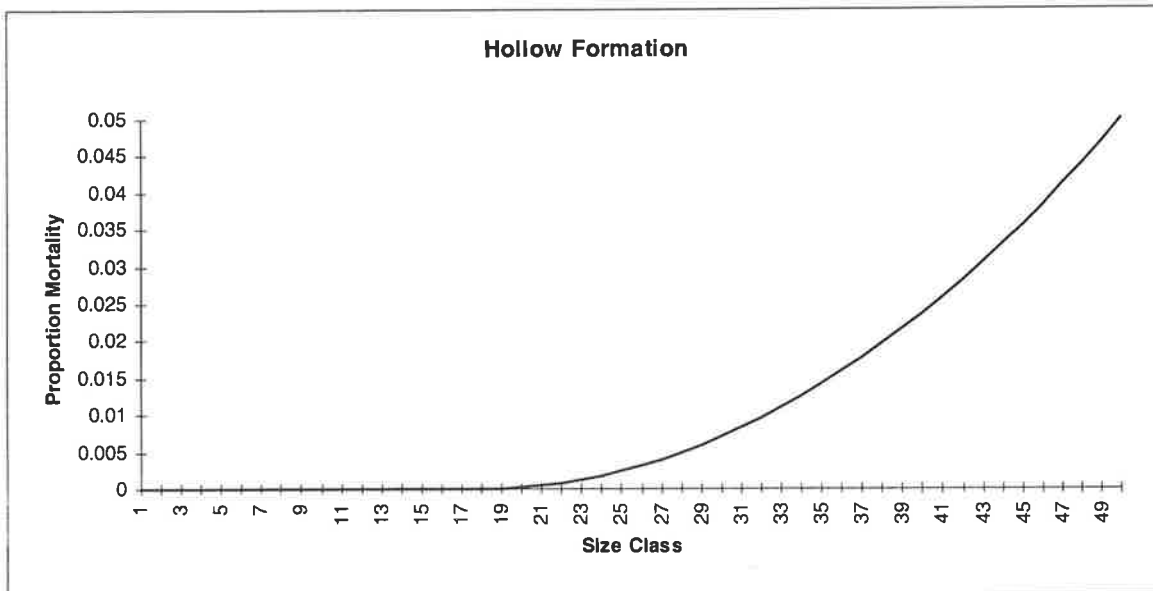


Figure 2.5. Hollow formation quadratic curve. The relationship between the size class of a tree and the proportion which forms hollows in a given time step. This is a quadratic function of the size class. Proportions are given by the square of (size class - 18)/32 times 0.05. Trees of size class 18 or less rarely form hollows (Lindenmayer *et al.*, 1993) so only the 32 size class between 18 and 50 may do so based on the basal area or the square of the size class.

Trees which are already higher than tree-form 0 have started the decay process and the form transition matrix is used (Table 2.1). This transition matrix was taken from Lindenmayer *et al.*, (1996). It was derived by monitoring the condition of more than 2,000 Mountain Ash trees for varying periods of time between 1983 and 1993. In 1983-84 hollow bearing trees in 32 sites of 3 hectares each were recorded, marked and plotted on a map. In a follow up survey in 1988, the form and whether these trees were still standing was recorded and another 497 sites were then measured. A follow up survey was conducted in 1993 and the combined data of tree form transitions has been scaled to become the transition frequency matrix of Table 2.1.

The transition frequency matrix is applied to the trees in forms 1 through to 8 in each size class. The transition frequencies are not dependent upon the size classes of the trees. The form of a tree can only ever increase or stay the same. This transition matrix is the key to the model. It comes from the data which is the hardest to collect for the model, as the transitions are slow processes. The estimates used in this thesis derive from a single, albeit detailed and large scale, survey for Mountain Ash.

Density independent mortality affects living trees (tree forms 0, 1 and 2). Trees below a cut-off size class collapse when they die through self-thinning (Ashton, 1981) or density independent mortality and are removed from the simulation. Trees at, or above, the cut-off size class will die, but not immediately collapse. These trees are placed into tree form 3 (dead trees with most branches intact, see Figure 2.2). The cut-off size class, which is the size-class below which trees immediately collapse, has been set at 18 which corresponds to trees approximately 90 years old. This means that trees that die before this size class will not become hollow-bearing stems (based on work by Ambrose, 1982). The rate of density independent mortality is 1% of trees in form 0 (living stems with and without hollows) as illustrated in the transition matrix given in Table 2.1.

New	Form 0	Form 1	Form 2	Form 3	Form 4	Form 5	Form 6	Form 7	Form 8	Removed
Old										
Form 0	0.99	0	0	0.01	0	0	0	0	0	0
Form 1	0	0.99	0.01	0	0	0	0	0	0	0
Form 2	0	0	0.87	0	0.02	0.04	0.02	0	0	0.06
Form 3	0	0	0	0.21	0.43	0.21	0	0	0.07	0.07
Form 4	0	0	0	0	0.38	0.29	0.07	0.09	0.02	0.16
Form 5	0	0	0	0	0	0.53	0.26	0.11	0.03	0.08
Form 6	0	0	0	0	0	0	0.29	0.22	0.2	0.3
Form 7	0	0	0	0	0	0	0	0.43	0.36	0.21
Form 8	0	0	0	0	0	0	0	0	0.58	0.42

Table 2.1. Transition Matrix. (Adapted from Lindenmayer *et al.* (1996)). Each row is the old tree form and the columns are the new tree forms. The numbers are the proportion of the row's tree form which are transferred into the column's tree form at each time step. For example, 99% of tree form 1 trees remain in tree form 1 and 1% are transferred to tree form 2. Each row adds to 1, with round-off error, because all trees will end up in one of the ten states.

## Growth

In the model, growth is parameterised using information on the rate of growth in mountain ash gathered by Ashton (1976). This data was used as it is relatively complete and is readily available in the published literature and is compatible with the phenomenological modelling approach pursued in this thesis. Living trees (tree form 0, 1 and 2) grow at each time step. It is possible for a tree to stay in the same size class in a time step or to increase either one, two, or three size classes in a single time step. It is not possible for a tree to grow more than three size classes, which would be a faster rate of growth than is known to occur in mountain ash. In addition, it is not possible for the size class of a tree to decrease. The proportion of trees of forms 0, 1 and

2 which move to the different size classes is shown in Figure 2.6. This is based upon a normal curve with a mean of one size class and a standard deviation of 0.25 size classes. The rates of growth of trees in tree forms 1 and 2 are less than the rate of growth of trees of form 0.

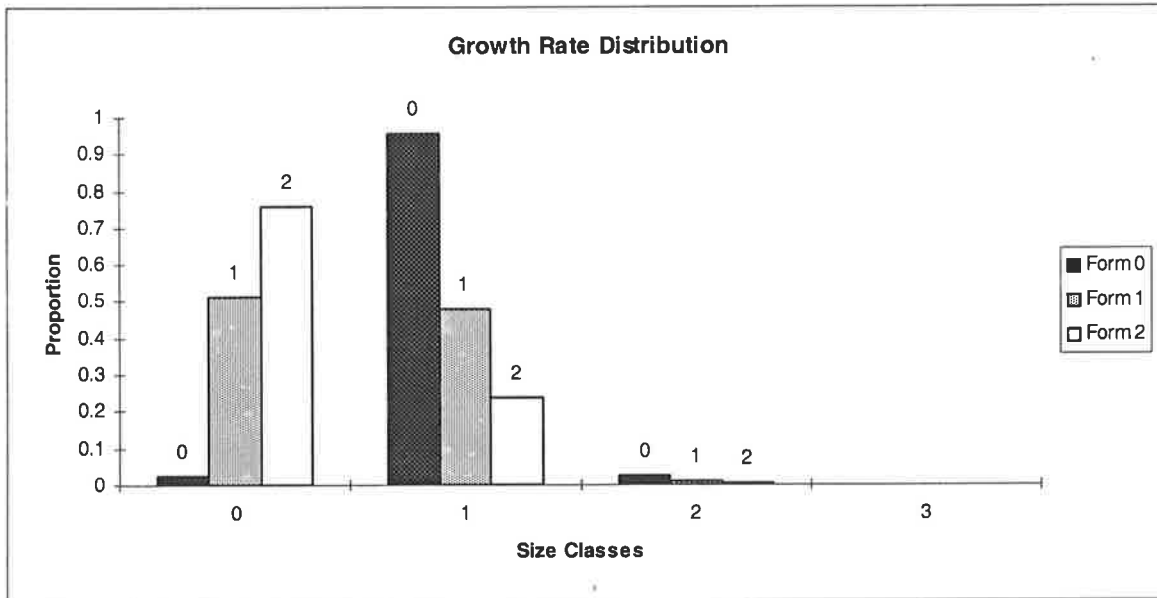


Figure 2.6. Growth rate distribution. The proportion of trees (of different tree forms) that increase 0,1 or 2 size classes in each time step.

The rates of growth of trees in tree forms 1 and 2 are derived from the rate of growth of tree form 0. Neither tree form 1 nor form 2 can increase more than three size classes in one time step in the model. The proportion of trees in form 1 (living trees with hollows) which increase a single size class in a time step is half that of trees in tree form 0. The proportion of stems which increase two size classes is one quarter that of tree form 0 trees increasing the same amount. In a similar manner, the proportion of form 2 trees (trees with a missing apical crown) which increase one size class is half that of form 1 trees and the proportion of form 2 trees which increase two size classes is half the proportion of form 1 trees which increase two size classes.

Deriving the growth parameters for this model from observations of a forest is not straight forward. Trees which grow faster than average will become the dominant trees in an even-aged stand and hence have a higher survival probability. Those which grow more slowly tend to be “crowded out” and collapse and hence will not be observed. For this reason it is necessary to calibrate the growth parameters to agree with observed forests. For the purposes of this study the standard deviation for growth of tree form 0 trees was estimated using field information



from Lindenmayer (unpublished data) to create realistic dynamics. The standard deviation was adjusted so that the forest stand would be similar in age structure to actual forest stands. The sensitivity analysis carried out on this parameter in Section 5 show it to have relatively little effect on the main results.

## Crowding, Recruitment and Self Thinning

“Crowding” is a term used to describe the number and size of trees in a stand. The level of crowding of trees in the forest is computed for the purposes of calculating self-thinning and, when appropriate, recruitment. It is assumed that each tree in each state uses a particular amount of some limiting resource such as space, light, or water. By assuming that limiting resources are a linear function of space, measures of the availability of space can be used to calculate the rate of self thinning, without loss of generality or the need to specify exactly which resource is limiting. When the amount of a resource that trees are using is greater than the available resource, then the stand is overcrowded.

Dahl (1940) included a general formula for the maximum stocking of a plot. It is given as a function of tree species and size and follows a general form from Reineke (1933). Once an initial measure is taken from an even-aged stand of a given tree species, it is then possible to determine the relative space required by trees of that species for all sizes. This is because this is a three parameter model where the parameters are tree size, number of trees of that size which can fit on a stand and a tree species parameter  $k$ . Knowing the number and size of trees on an even-aged stand allows us to determine the species specific parameter  $k$ . The formula for the number of possible trees of Mountain Ash of a given size is

$$\log_{10} N = K - 1.605 \log_{10} D. \quad (2.3)$$

Where  $D$  is the diameter measured at breast height (DBH) of a tree in inches,  $N$  is the maximum number of trees of the given DBH per acre and  $K$  is a species dependent parameter. Ashton (1976) estimated  $K$  to be 4.26.

The inverse of  $N$  is the amount of space that each tree requires. This equation then becomes:

$$S(D) = \frac{C_a}{10^{-4.28}} \left( \frac{D}{C_{iCm}} \right)^{1.605}, \quad (2.4)$$

where  $S(D)$  is the space in hectares each tree requires.  $D$  is measured in centimetres and the conversion factors  $C_a$  and  $C_{iCm}$  converts acres to hectares and inches to centimetres respectively.

The constants can be combined to give:

$$S(D) = 10^{-K} D^{1.605}. \quad (2.5)$$

For *E. regnans*  $K$  is 4.023,  $D$  is now DBH in centimetres and  $S(D)$  is the space in hectares each tree requires. For example,  $S = 0.051$  hectares for a  $D$  of 50 centimetres and  $S = 0.154$  hectares for a tree diameter of 1 metre. The relationship between the number of trees per hectare for different diameters is shown in Figure 2.7. Without specifying the limiting resource, it is impossible to calculate how much of the resource is used by trees in various states of decay.

Having determined the relative space requirements for different sized trees, it is next important to determine the relative space requirements for trees of different tree-forms. It should be noted that trees of higher tree-forms, which include dead trees, do not necessarily have spatial requirements as such, but they prevent other trees from using resources associated with that space. For the purposes of this study, tree-forms 1 to 8 use, or prevent other trees from using, progressively less of the limiting resource. For example, we assumed a form 5 tree uses 20% of the resource used by a form 0 tree of the same diameter. The values used in this study are displayed in figure 2.8. When the ratio of decayed to healthy trees is high and the forest is near or at its maximum-crowding, then errors in the estimation of resource use by these damaged and dead trees may be significant. However, this situation rarely occurs in the model.




Figure 2.7. Maximum number of trees per hectare (calculated from Dahl 1940). Note the Y-axis has a logarithmic scale.




Figure 2.8. The amount of space used or required by trees of different forms. These are given as a proportion of the amount of space a tree of form 0 uses.

Overcrowding can occur after the growth phase of stand development and initiates the self-thinning phase. In self-thinning, trees are removed (ie collapse or die) from the forest until the stand is no longer overcrowded. Trees of the smallest size class are removed first, followed by the next smallest and so on. Trees in size-class 18 or higher might not immediately collapse, but might be transferred to tree-form 3 (dead trees with most branches still intact, see Figure

2.2). The frequency with which these trees will be transferred to tree form 3 rather than collapsing is based upon a quadratic function of the tree's size-class (hence, it is approximately linearly dependent upon the basal area of the trees). Because the function is set so that trees in size class 18 or lower have a 100% chance of collapsing and 100% of trees in size class 50 will remain standing in form 3 after they die through overcrowding (Figure 2.9). When overcrowding effects trees in larger size classes, trees in form 2 will be the first ones transferred to tree form 3 followed by form 1 trees and then those in form 0, until the forest is no longer overcrowded.




Figure 2.9. Proportion of trees of different size classes which move from form 3 (dead trees with most branches intact) upon death. The proportion is a quadratic function of size class and is the square of (size class - 18) divided by 32. This is based on Lindenmayer *et al.*, (1996) and shows that trees with smaller diameters are more likely to collapse.

Recruitment occurs in Mountain Ash only after a disturbance (Attiwill, 1994). The only source of disturbance in this model is timber harvesting, because the impacts of wildfires have been excluded. During the recruitment phase, trees fill out the first size-class (trees less than 5 cm) so that the forest is again at maximum stocking rate.

## Harvesting

Timber harvesting is managed through the harvesting sub-model, which allows us to simulate most of the typical harvesting regimes used in Australian forests. There are two types of

harvesting. In the first type the user specifies the number of trees of a particular quality to be retained. An example might be to retain 5 'good' trees and 3 'mediocre' ones on every hectare, where the characteristics of 'good' and 'mediocre' trees are specified in the submodel. For example, 'good' could be all trees between size classes 26 and 50 (the maximum size class) which are in tree forms 0 or 1 and 'mediocre' could be defined as trees in tree form 0,1 and 2 (i.e, living stems) and between size classes 16 and 18. The second type of instruction is the removal of a proportion of trees in any given state or range of states. A simple example of this prescription is to remove 100% of trees below a certain size which have no hollows. A number of harvesting instructions can be combined to make a harvesting operation and any number of different harvesting operations can be included in a rotation period. The operations in a rotation period are repeated for the length of the simulation. To take a specific example from scenario 4.4: A harvest states that the rotation length is 100 years with a fifty year delay before the rotation starts. At the end of the rotation length a number of trees are marked for retention. Up to ten trees per hectare are retained with tree forms of 1 or higher (ie hollow bearing trees). All size classes can be retained starting with the largest trees and continuing until either 10 trees per hectare have been set aside or down to size class 18 (the minimum size which can form hollows). 100% of the trees not marked for retention are removed. Harvested trees are completely removed from the system and no count is kept of them.

When trees are set aside for retention and the rest of a stand is harvested, there is an increased rate of mortality for those retained stems. This higher mortality rate is caused both by the increased exposure of the trees to changed wind, temperature, and other environmental features after logging and the effects of high intensity post-logging burns used to promote regeneration. Mortality rates can be set in the program as the proportion of the retained trees which will collapse within a single time step. The default setting for this is zero, which means that there is no additional mortality or accelerated rates of collapse for retained trees. The results of field observations (Lindenmayer *et al.*, unpublished data) indicates that we would expect 50% or more of retained trees to collapse soon after logging. Because this mortality could vary tremendously depending on how the stand is treated, this parameter is varied in the simulations.

## Hollows

The dynamics of three types of hollows in trees can be followed in this model. This is especially important for fauna species that prefer specific hollow types. These types of hollows, for mountain ash trees, are holes, fissures, and hollow branches. For each tree form and size class, there is an expected number of each of these hollows per tree. This can be used to calculate the expected number of each type of hollow for the entire stand. The values for mountain ash are displayed in Figure 2.10 and they are based on extensive statistical analysis by Lindenmayer *et al.* (1993). The default output of the model is the number of trees which have hollows, rather than the total number of different types of hollows summed across all hollow-bearing stems of given types within the stand. The program also has the option of allowing the examination of each of the different types of hollows individually.




Figure 2.10. Number of different types of hollows in form 1 trees (living with hollows) of different size classes. Adapted from Lindenmayer *et al.* (1993).

## Multiple Coupe System

The model was designed for a single stand of trees, but it is possible to use it to explore a multiple coupe system. To do this, the user treats each coupe in the system separately. The results from simulations for individual stands can then be combined to produce the results of the multiple coupe system.

As an example, consider a two coupe system with a 100 year rotation time on each coupe, staggered so that timber harvesting occurs every 50 years. The first coupe is simulated with the first harvest starting fifty years after a previous clear cut. The second coupe is simulated with the first harvest starting at one hundred years after a previous clear cut. The two outputs are then combined either by graphing statistics for the two coupes on the same graph, or by adding the appropriate values ( such as the number of trees with hollows in the two coupes) and graphing or extracting the relevant statistics.

## Output

HOLSIM allows the user to specify two types of output. It can produce a time series of the number of hollow bearing trees in the stand, or the number of a particular type of cavity (ie. a hole, hollow branch or fissure). Additionally, it can extract the minimum, average, and maximum number of hollow bearing trees over a pre-defined period. For the scenarios described below, we extracted the “steady state” statistics on the hollow dynamics by setting the statistical extraction period to 650 to 749 years. This was ample time for each of the scenarios to achieve a stable cycle of number of trees with hollows and predicts hollow dynamics if a harvesting system is implemented over a long time.

## Assumptions and Limitations

This model is an abstraction and implementation was limited by the information available on Mountain Ash stands in particular, and forest processes in general. The main limitations to the model are that the stand consists of a single species of trees, that the effects of climate and foliage damaging fauna are constant with time, that there is no spatial complexity to the stand, and that there are no random catastrophes.

The primary limitation of the model is that the modelled stand consists of a single species of trees, which is appropriate in the study because stands of Mountain Ash are generally monotypic. The effects of other species on the system are assumed to be constant or insignificant and presumably limited to non-tree species (such as shrubs and insects). The main interactions between species in a multiple species stand would be between different species of trees. This can not be simulated by the model

Interactions with other species, such as insects, other plants and animals are not included. These factors are of marginal importance to the main objective of this study. If the effects of such interactions are small, or they vary over time by a small magnitude, then they do not need to be included explicitly.

Long term climatic changes also are not included in the model. Each five year time step is assumed to be the same as the last. Natural variation of the climate is allowed for, to some extent, by the averaging nature of the model - the assumption that the scenario results don't relate to an individual stand but to the average stand. Current theories on climatic change predict an average raising of the temperature on the Earth within the time scale of a simulation. However, the effects of such a change on the processes included in the model (growth, hollow formation, collapse) are not well understood in general, nor for Mountain Ash trees in particular.

Spatial complexity is not included in the model. Even though for mountain ash there is information on the effects of micro-climatic influences on the presence of tree hollows and other measures (Lindenmayer et al., 1993), it was not appropriate to include them in the model. If complete information of the position of every tree and the geography of every hectare were included, then on a practical level the model would become unworkable without promising the addition of any further insights. Too much information would be needed for any given stand before any modelling could be undertaken. Detailed spatial information might be useful if the stand contains gross geographic features, such as half of the stand being on a steep slope and the other half being on flat ground. In this case, the model could be used by modelling two stands independently, using different input parameters, and combining the results.

The central concept underpinning the model is that the simulation gives a single deterministic result for given inputs. Stochastic effects (such as climatic stochasticity mentioned above) are averaged out of the model. This approach does not give an indication of the level of variance associated with a result. Including such stochastic effects would involve a qualitative change, and would require an entirely new model. Given reasonable consistency in the ways in which forest stands grow, this complication is considered unwarranted.



Catastrophes were not included in the model. Fires are an important part of Australian forest ecology (Williams and Gill 1995) and excluding them places a large constraint on the applicability of the model. However, as the model represents an average forest, the inclusion of fires would be very difficult. Wildfires have a substantial effect on an individual stand and occur at unpredictable time intervals. It is, however, possible to study the effects of a particular fire event at a specific time. If one were interested in observing the after-effects of a large fire on a stand, then the simulation could be commenced after the point of the fire. Alternatively, if the aim was to model a stand after a fire and a given tree harvesting regime, the fire could be simply included as a special type of harvest.

There was insufficient information available in the literature to derive some parameters for mountain ash. The solution to this was to estimate the parameters and set them to values to make the stand conform to data on real stands. Tests of the sensitivity of the model on these parameters were completed. Facilities to test the sensitivity of the model to most of the parameters are included in the software.

Because of lack of data it was necessary to estimate the rate at which trees formed hollows, and, the relative amount of the limiting resources which dead and dying trees used (or prevent other trees from using). The variance in tree growth were estimated as described above due to theoretical difficulties with existing data. In addition, values for the rate of collapse of trees which have been retained following logging operations was left as an input parameter to be varied and its effects explored.

The process by which mountain ash forms hollows is quite complicated and the average rate at which hollows are formed is not precisely known. In the model, this was the rate at which trees move from tree form 0 (healthy trees with no hollows) to tree form 1 (living trees with hollows). This rate was estimated using the known number of trees which have hollows of given size classes.

There are data to determine the amount of space (which is used as the limiting resource) which trees require in different size classes. However, the amount of space which different tree forms use in various states of decay is not well known. In this study, dead trees in the advanced states of decay were assumed to occupy or use a small amount of the limiting resource (Figure 2.8).

The proportion of trees in these higher tree forms tended to be very small and hence the model is relatively insensitive to these values.

The rate of tree growth in the program was set so that the average tree will grow by one size class in a single time step. Rate of growth information is readily available for many tree species but the variation in this rate of growth is difficult to derive in a form appropriate for use in this program. Natural variation in the rate of tree growth is compounded by the effects of crowding by neighbouring trees. The variance of growth in this simulation for Mountain Ash was estimated so that the model would be consistent with the known composition of stands at various ages. Variance of growth is one of the input parameters open to easy exploration and sensitivity analysis.

Trees which are retained during a timber harvesting operation are likely to collapse within a few years. This is caused both by sudden exposure of the tree to wind and the likelihood of the trees being burnt in a high intensity regeneration fire (Lindenmayer *et al.*, 1990). The proportion of trees which collapse after harvesting is a user-defined input variable.

## Section 3

### Scenario Description

#### Introduction

Scenarios were simulated to explore a number of different harvesting strategies, particularly tree-retention strategies. Tree-retention scenarios were explored to answer such questions as what kind of trees need to be retained (Scenario 4.3) and what impact do different levels of post logging mortality have on hollow availability (Scenario 4.5). Other scenarios included proportional harvest strategies (Scenario 4.2) and more complex cutting strategies (Scenario 4.6). The theme throughout this section is to look at how effective different strategies are at retaining hollow bearing trees under different conditions. In brief the scenarios are:

Scenario 4.1 examines a null run of the model with no harvesting included. The aim of this is to see how many trees with hollows the program will achieve in the absence of logging.

Scenario 4.2 Looks at proportional harvesting schemes. These are where a certain proportion of all types of tree are removed at each harvest.

Scenario 4.3 looks at a harvesting regime where all trees which have hollows are retained and all the rest are removed. This section dispels the idea that retaining all the trees which have hollows would in itself be sufficient to preserve the habitat qualities of a forest.

Scenario 4.4 Here a set number of trees are retained, with preference for trees with hollows where present but including some which have not yet formed hollows if they aren't enough hollow bearing trees available.

Scenario 4.5 The previous regime is repeated with additional mortality included for retained trees to reflect mortality due to increased exposure to the weather and to post logging regenerative burns.

Scenario 4.6 is a more complex regime where the trees are first thinned and later cleared. The purpose of this scenario is to test the simulation of complex regimes.

In each case the scenario was run for a total of two hundred time steps (one thousand years) for a single species stand of Mountain Ash. This was always adequate time to allow the long term dynamics of the simulation to be seen, which was the reason for projecting so far into the future.

Two main forms of output were examined for each scenario. The first was a graph depicting the number of trees with hollows per hectare as time progresses for each scenario and the second was the minimum, maximum and average number of trees with hollows per hectare after the number of trees with hollows has reach some form of equilibrium. Here equilibrium is meant to mean that the number of trees with hollows is either constant or changes through a fixed and repeating cycle (a dynamic equilibrium). It does not mean that the forest itself has reached a state where it no longer changes. In discussing the various scenarios shown below, the main measure used was the minimum number of trees with hollows in the 25 hectare stand. The minimum was used because it is the factor likely to be limiting the number of cavity dependent organisms using hollows in logged coupes (reviewed by Gibbons and Lindenmayer 1996), and also because it is convenient to have a single measure for comparisons between scenarios.

In each scenario the stand started as a newly cleared and regenerated stand. There was a 50 year delay before the beginning of the first harvest rotation in all scenarios which had harvests. The harvests generally occurred at the end of the rotation cycle (except for Scenario 4.6 with 2 separate harvests in each cycle). Thus the first harvest occurs when the stand is either 100 or 150 years old in all scenarios. It is believed that in all scenarios which come to a steady number of steady cycle of trees with hollows the initial condition of the forest does not affect the expected long term outcome.

## Scenario 4.1

### No logging

#### Description

The aim of this scenario was to examine a forest which has not been subject to logging. This is done so that we may better see the impacts of logging on the forest from an absolute perspective. It is done to give an indication of how the forest would evolve, with the understanding that as recruitment only occurs after a tree harvest event, the scenario is not realistic over a longer time scale (eg > 400 years).

#### Results




Figure 4.1.1 The dynamics of trees with hollows when there is no logging (Scenario 4.1). Forest starts as a newly cleared regeneration stand.

The results from this run are displayed in Figure 4.1.1. The forest commences as an even aged stand after a clear cut and it is not harvested after that. Because the forest only regenerates after a disturbance some form of disturbance is normally included in each scenario (which is normally a logging event) in order to get sensible results over a long time scale. Hence over a long time scale the dynamics seen above will not occur.

## Discussion

Figure 4.1.1 demonstrates the effect of the main process on the system. Because the simulation starts from a newly cut forest there are no trees with hollows in the system until after the year 100. At that point there is a sudden increase in hollow-bearing trees. This occurs from the density independent mortality coupled with the effects of overcrowding. These effects kill trees but some of the dead trees remain standing and these trees are hollow bearing.

Dead trees remain standing for only a relatively short length of time (Tree form 3 trees will remain standing for 5.3 time steps or 28.2 years on average as computed from Table 2.1) and the number of trees available to enter this state must be decreasing as the overall number of trees decreases due to over-crowding. Thus the number of trees which form hollows due to these processes peak much earlier than those which form hollows through the second effect..

The second effect is the transfer of trees into tree form 1, the formation of hollows in living trees. These trees will remain standing for a much longer period than dead trees. Hollow formation is much slower for living trees but as they are longer lived, they contribute to the number of trees with hollows greatly, as seen in Figure 4.1.1.

The maximum number of trees with hollows per hectare is quite high and is much higher than occurs in any of the other scenarios examined in this thesis. It is not the highest possible number because of the lack of new tree recruitment, experimentation has shown this theoretical maximum to be just above 20 trees per hectare.

## Scenario 4.2

### Proportional Harvesting

#### Description

One method of ensuring that the harvested stand retains some habitat quality is never to completely clear the stand make every harvest a thinning operation. This scenario explores that harvesting approach. In this scenario the proportion of trees removed at each tree harvesting event is constant for all size classes and tree forms and it is constant for all tree harvesting events. This proportion varies between 10 and 90%. Two rotation lengths are examined, 50 year and 100 year rotation.

#### Results




Figure 4.2.1. Number of trees with hollows. 50% trees removed from all size classes and tree forms. 100 year rotation with first harvest at 150 years.

Figure 4.2.1 displays a time series showing the number of trees with hollows, in this case the proportion of trees removed is 50% over a 100 year rotation period. Typical of this scenario there is a period of uneven oscillations leading to a more steady cycle after about 400 years. Even so by 250 years it has reached the maximum.

The down spikes correspond to the harvesting event. After each harvest the number of trees with hollows is reduced by half due to the thinning operation. The number of trees with hollows increases immediately after each harvesting operation, this is because there are plenty of mature trees left in the system, which are of sufficient size to develop hollows.

% removed	Rotation Length	
	50 Years	100 Years
10	14.4	16
20	10.1	12.6
30	6.7	9.6
40	4.2	7.1
50	2.4	4.9
60	1.2	3.1
70	0.5	1.8
80	0.15	0.8
90	0.02	0.2

Table 4.2.1 The minimum number of trees with hollows which occurred once the system had reached a steady state. Taken after 650 years.




Figure 4.2.2. Graph of data presented in table 4.2.1. Number of trees with hollows per hectare under different thinning percentages for two different rotation lengths.



## Discussion

Table 4.2.1 and Figure 4.2.2 summarise the results from this scenario. It is immediately clear that the effects of increasing the percentage of trees removed on the number of trees with hollows is non-linear. If 50% of the trees are removed at each harvest then the number of trees with hollows is a third or less the number of trees when only 10% are removed, depending upon the harvest rotation length.

It is also clear that the effect of decreasing the rotation length depends upon the percentage of trees being retained. When the percentage removed is low then the effect is small, it is greatest when 40% of trees are removed in a harvest. At higher percentages it is proportionally greater but smaller in absolute terms (as there are fewer trees with hollows in either case). For example when 10% of trees are removed the number of trees with hollows is 0.9 times as high in the 50 year rotation as in the 100 year rotation, at 50% it is 0.49 times as high and at 90% it is 0.1 times as high. This means that the greater the number of trees being removed the more critical is the rotation length on the number of trees with hollows.




Figure 4.2.3. The proportion of trees which would remain after 4 consecutive harvests if no trees were removed for any reason between harvests.

The reason that this occurs is that trees are subjected to more than one thinning before they are large enough to become hollow bearing trees. With a 50 year rotation length a tree will experience two thinning events before it is large enough to become a hollow bearing tree, most

hollow-bearing trees will have experienced more than two thinning events before they became hollow bearing - this gives the non-linear effect. Figure 4.2.3 demonstrates this principle. This figure shows how many trees would remain after 4 consecutive harvests of different proportions. For example, if 10% of the trees were removed for four consecutive times there would be 67% of the original amount left. This distribution is non-linear with a greater relative effect occurring with higher proportions removed. If we increased the number of consecutive harvests, which would correspond to decreasing the rotation time in the scenario, the curve would be sharper and the non-linear effect would be stronger. In the scenario, between each harvest other processes are also active, including self-thinning mortality and other forms of mortality. However, the principle demonstrated in Figure 4.2.3 is still clear.

## Scenario 4.3

### Retaining a Fixed Number of Trees with Hollows.

#### Methods

This is the first of the scenarios to examine the retention of classes of trees (size classes and tree forms) for the maintenance of trees with hollows. In this scenario *all* trees which are hollow-bearing are retained at each harvest event and the remaining trees are removed. The question that this scenario addresses is how successful is this strategy of retaining all suitable habitat.

Three rotation times are examined, 50 year, 100 year and 150 year rotation. The last is considerably longer than typical rotation times (Gibbons and Lindenmayer 1996) but is included for completeness. The rotations are timed so that the first harvest occurs 200 years after a clear-cut. This allows the forest to become quite mature before the effects of harvesting are applied.

#### Results

Time series of each of the three simulations are displayed in Figures 4.3.1, 4.3.2 and 4.3.3.




Figure 4.3.1. All trees without hollows are logged each 50 years with the first logging event at 200 years.




Figure 4.3.2: All trees without hollows are logged each 100 years.




Figure 4.3.3: All non-hollow bearing trees are removed each 150 years.

## Discussion

It is immediately obvious from these results that the strategy of retaining all hollow-bearing trees does not work when the rotation length is 100 years or less. Only when the rotation cycle is greater than 100 years do we see a stable cycle forming.

The reason is that it takes close to 100 years before trees start to form hollows. When the rotation length is 50 years there is no opportunity for new trees to grow large enough to become suitable habitat trees. The trees which are already hollow-bearing remain but collapse eventually and no new hollow-bearing trees replace them. Those which are already dead collapse relatively quickly and account for the steep downward slope on the first two graphs immediately following the first harvesting event. Trees which are living but have hollows die and collapse relatively slowly and account for the long shallow slope on the graph.

In the model, trees start to form hollows after 90 years but not enough trees form hollows to stem the loss when the rotation length is 100 years. At 150 years there is enough time for new trees to form hollows in between harvests to start replacing those which have collapsed in the interim.

In these simulations there is no extra post-logging mortality of retained trees, which makes this scenario idealistic. Including such additional mortality would make the situation obviously worse for the retention of hollow bearing trees.

Retaining only those trees with hollows is clearly a very poor strategy, and difficult as well, as it relies on forest managers being able to identify all of the hollow bearing trees. When it works at all is when the time between harvests is long enough for the forest to mature. In order to have a rotation time of less than 150 years a different strategy must be employed. In the previous scenario (Scenario 4.2) a proportional tree removal strategy was examined, which did lead to a stable cycle of trees with hollows. The next scenario shall examine the retention strategy where the number of trees being retained at each harvest is fixed, but the retained trees aren't constrained to be already hollow-bearing.

## Scenario 4.4

### Retaining trees with or without hollows

#### Methods

In this scenario a number of trees per hectare were retained at each tree harvesting event. When there were not enough trees with hollows to fill the retention quota, trees without hollows were included in the quota, starting with the largest stems. This number was varied from 1 tree up to 35 trees per hectare, with a rotation length of either 50 or 100 years long. As in the previous scenario, retained trees were not subjected to additional mortality, which might be the case where they are sufficiently protected against increased exposure to weather and any post-logging regenerative burns. This might correspond to a situation where effective wind-breaks are added to protect retained trees and where logging debris was dispersed without burning.

#### Results




Figure 4.4.1. Time series showing the number of trees with hollows per hectare. 10 trees per hectare are retained (with preference for hollow-bearing trees). There is a 50 year rotation which starts 50 years into the simulation.

A time series is displayed in Figure 4.4.1. This is the number of hollow-bearing trees per hectare in the system as a function of time. The first harvest occurs in year 100 at the end of the 50 year rotation cycle. The number of hollow-bearing trees rises smoothly until reaching a

stable level. The number of hollow-bearing trees is approximately two thirds of the number of trees being retained. All of the time series displayed a similar time series with the exception of the extreme case of retaining 35 trees per hectare. In this extreme case the system settled into a slow cycle.

The number of trees with hollows per hectare which results from each simulation (taking the minimum number of trees occurring after the system has reached a stable cycle or equilibrium) is displayed in Figure 4.4.2. The final two points occur when the number of trees per hectare retained are often greater than the number of trees per hectare standing (all trees are retained).




Figure 4.4.2. Summary of results from Scenario 4.2. Each point is the minimum number of trees with hollows per hectare which occurs after the simulation has stabilised. The X-axis is number of trees per hectare being retained.

### Discussion

There is a very strong and linear relationship between the number of trees being retained and the number of these trees which are hollow-bearing. For the 50 year rotation cycle approximately 64% of the trees which are retained are hollow-bearing and for the 100 year rotation this is closer to 67%. These proportions are stable across all the values except for the extreme points. What this means is that one third of the trees being retained are not hollow-bearing yet but are being kept so that they can grow large enough to replace trees which are hollow-bearing as those trees collapse.

The number of hollow bearing trees is very similar for both the 50 and 100 year rotation cycle. This indicates that as long as the retained trees are protected against increased chances of collapse or damage, then the rotation cycle is not critical. The next scenario will consider the case where there is additional mortality to the retained trees due to the effects of tree harvesting and increased exposure.

The number of trees per hectare that were being retained was increased through each simulation to see the point at which the linear relationship in Figure 4.4.2 no longer held. The relationship held up to the point where nearly every tree on the stand was being retained.

The important conclusions from this scenario and the previous one are that it is not sufficient to retain only trees which are already hollow-bearing. Trees which are not yet hollow-bearing must also be put aside as potential replacements for the others as they decay and collapse. This scenario indicates that one in three trees which are being retained should be non-hollow-bearing trees of the largest sized cohort. This value depends critically on the decay rate of hollow-bearing trees and the rate of formation of hollows in trees. If these rates were different then the same qualitative conclusions would apply but the recommended number of non-hollow-bearing trees to retain would be different (this is explored in Section 5). When there is no post-logging mortality of the retained trees the rotation length is not important to the number of hollow bearing trees.



## Scenario 4.5

### Retained tree mortality

#### Methods

In this scenario as in Scenario 4.4, a quota of trees are retained which include some which are hollow-bearing and some which are not. The retained trees are subjected to retained-tree mortality. This is the mortality due to increased exposure to the elements and also due to logging practices such as the post-logging regenerative burns which are often used. A standard 10 trees per hectare is retained. Rotation time is either 50 years or 100 years and the retained-tree mortality rate varies between 0 and 90%.

The retained tree mortality is applied in the same five year time-step as harvesting, for the sake of convenience. In the case of mortality due to post-logging regenerative burns this is accurate, in the case of mortality due to increased exposure, this is an approximation. There is evidence that retained trees tend to suffer twice the normal mortality due to increased exposure also indications are that mortality due to post-logging regenerative burns is very high, much higher than 50% (Lindemayer, *pers. comm.*).

#### Results

The results from this series of simulations are displayed in Figure 4.5.1 and they show that the minimum number of hollow bearing trees tapers off quickly as mortality increases. This decline is more dramatic when the rotation time is 50 than 100 years. The greatest absolute difference occurs when the retained tree mortality is around 20% (1.588 trees per hectare). The relative difference is always increasing, for example when the proportion is 40% there are 3 times as many hollows in the forest with a 100 year rotation then there are in the 50 year rotation forest, when the proportion is 90% then there is 33 times as many trees in the longer rotation forest.




Figure 4.5.1. Graph of the results from Scenario 4.5. The minimum number of trees with hollows per hectare which occurs in a rotation cycle after a steady cycle has been reached. Mortality proportion is the proportion of trees which are to be retained but collapse immediately following a tree harvesting event.

## Discussion

It was not surprising that the number of trees with hollows decreases as the mortality rate increases. The interesting result is that this change is not linear. For example, when there is 50% post logging mortality, the number of hollow bearing trees remaining is less than half the number which occurred when there is no additional post logging mortality. This is particularly evident in the case where a 100 year rotation was implemented and it is much stronger when the rotation length is only 50 years.

The reason that this occurred was that with each cut, a proportion of all retained trees were removed (particularly the largest stems). When there was limited mortality, these trees persisted through several harvesting events. However, with added mortality, few trees may grow large enough for hollows to form. Replacement is limited because of the length of time it takes for a tree to start forming new hollows.

## Conclusions

The important conclusions in this scenario are that it is important to reduce the amount of retained-tree mortality and that the benefits from reducing retained-tree mortality is non-linear, or out of proportion to the mortality reduction. The level of retained-tree mortality is particularly important to the length of the rotation cycle. At shorter rotation cycles the effects are much more dramatic, particularly at mortality levels of 50% or more. If the retained-tree mortality is low enough then the effects of shortening the harvest rotation length is less damaging.

## **Scenario 4.6**

### **More Detailed Regimes.**

#### **Methods**

This scenario was designed to examine a more complex harvesting regime. The software allows for any number of different types of harvests of arbitrary complexity to repeat within a defined rotation cycle. In this case two types of harvest occur within the rotation cycle. The first is a thinning operation which removes a given proportion of trees from every size-class and tree-form. This operation occurs 50 years into the rotation cycle in all the simulations in this scenario. The second operation occurs 100 years into the rotation cycle (at the end of the cycle). At this event the trees are all removed except for a given number of retained trees.

In the first set of simulations 50% of all trees are removed in the thinning harvest and the number of trees retained at the clearing harvest is varied between 1 and 25 trees per hectare. In the second set of simulations this number of trees retained is set at 10 trees per hectare and the percentage of trees removed in the thinning harvest is varied between 10% and 90%.

## Results




Figure 4.6.1. A time series showing the number of trees with hollows throughout the simulation. Harvest event 1 clears 50% of all trees at 50 years, harvest event 2 clears all but 10 trees at end of the 100 year rotation cycle.

Figure 4.6.1 shows a time series of a forest where 50% of all trees are cleared after 50 years and all but 10 retained trees are cut at the end of the 100 year rotation cycle. There is a 50 year delay before the start of the first cycle so the thinning harvest occurs on the century and the clear cut with retained trees occurs 50 years into the century. The thinning cut is where the number of hollow bearing trees decreases and the cut with retention marks a change in slope of the number of trees forming hollows.




Figure 4.6.2. A summary of the results for the complex cutting strategy. At 50 years 50% of all trees are removed, at 100 years a number are retained and the rest removed.

Figure 4.6.2 is the summary of the first set of simulations where the proportion of trees removed at the thinning harvest is kept constant at 50% and the number of trees retained at the final clearing harvest is varied between 1 and 25 with no extra post-logging mortality. The points fall upon a nearly straight line with the number of hollow bearing trees being approximately 15% of the number of trees being retained on the second harvest.




Figure 4.6.3. The percentage of trees removed at year 50 harvest varies and number retained at 100 year harvest is fixed at 10 trees per hectare.

Figure 4.6.3 is the summary of the second set of simulations. Here the proportion of trees removed at the thinning harvest is varied between 10% and 90%. The number of trees retained at the final harvest is always 10 trees per hectare and there is no extra post-logging mortality. The number of trees with hollows decreases non-linearly as the harvest proportion increases.

## Discussion

It is natural to consider this scenario in conjunction with Scenario 4.5. where there is one harvest and post logging mortality. This scenario could be thought of as the same except with a fifty year delay between the clearing harvest and the post-logging mortality. If we look at the results this is what can be seen. Comparing Figure 4.6.3 with Figure 4.5.1 we see that we get a very similar curve to the 100 year rotation curve in Figure 4.5.1. The curve is the same shape but the values are all slightly higher, due to the fact that in this scenario there is a delay between the final harvest and the next thinning operation whereas in the previous scenario the harvest is immediately followed by tree removal due to post-logging mortality.

The time series in Figure 4.6.1 shows some interesting dynamics. It is not surprising that the thinning operation causes a decrease in the number of trees with hollows (half of them are removed). What is interesting is that at the point where the second harvest occurs the slope increases, which means that the rate at which trees are forming hollows increases. This is due to the sudden drop in resource competition and the subsequent decrease in self-thinning related mortality. Only a few of the retained trees will already be hollow bearing (1.5 trees per hectare in this simulation), many of the rest will be of the latest cohort. These retained trees will not be subjected to overcrowding mortality after the harvest and will tend to remain living when forming hollows. The collapse rate for trees increases dramatically when the tree dies, even if it is left standing and the greatest cause of tree mortality is self-thinning.

In this time-series the minimum number of trees with hollows once a cycle establishes is 1.52 trees per hectare, this compares favourably with the number from Scenario 4.5 where there is 50% post-logging mortality. In the latter case the minimum number of trees with hollows once a cycle established is 1.084 trees per hectare. The fifty year delay before the clearing with retained trees and the 50% mortality (in one case caused by logging practices and the elements

in the other by a second tree harvest) allows some of the retained trees which did not have hollows, to form them.

Figure 4.6.2 shows a line which is very close to being straight. Here approximately 15% the number of retained trees survive as hollow bearing. Without the 50% removal harvest this value would be 67% as shown in Figure 4.4.1. It is interesting that in both cases the ratio is the same for different numbers of trees per hectare being retained.



## Section 5

# Sensitivity Analysis

### General

In this section I look at the sensitivity of the model and some of the results to the values of different parameters. To what extent would the results be different if the parameter values were different?

It is expected that the quantitative results of the model will be different when the parameters are different, although the relative effects of changing one or another parameter could well be insignificant. To what extent will the quantitative results of the model change as the parameters are varied through a sensible range? Will this also effect the qualitative results ? If the parameters were different would the conclusions about the relative effects and worth of different tree retention schemes remain the same or would they be completely different.

I have taken a single tree retention scenario as the baseline model upon which to test the sensitivity of the results to the parameters. In this scenario 10 trees per hectare are retained at each harvest. The retained trees are not restricted to those trees which are currently hollow bearing. The harvest rotation length is 50 years. There is no extra post-logging mortality associated with each harvesting event. This scenario features in section 4.4. With the parameter values used in this model the long term minimum number of trees with hollows is approximately 6.4 trees per hectare.

The approach used was to test the parameter set against the baseline scenario and change the parameters over a range which extends beyond the range of sensible parameter choices. This allowed the effect of extreme parameter changes to be examined. The parameters examined in this section include the growth parameters, the decay parameters including the density independent mortality rate and hollow formation rate, the spatial requirements for trees of different classes and the length of the delay before the first harvest.

Tree growth is controlled by three parameters: the average number of size classes advanced in a single time step, the maximum number and the standard deviation in the number of size classes advanced in a single time step. The standard deviation associated with growth is the least well established of these parameters, although the value used is a sensible one based on standard Mountain Ash forest dynamics. For this reason the standard deviation is the growth parameter which is tested for sensitivity here. The significance of the average rate of growth of trees is also considered.

Two decay constants were estimated from known typical forest structures. These are the hollow formation rate, which transfers trees from tree form 0 to tree form 1 in the model (Living non hollow bearing to living hollow bearing), and the density independent mortality rate, which is the rate of tree loss (to tree form 3) due to factors other than overcrowding. The sensitivity of both of these constants are considered here independently.

The “spatial requirements” of trees of different sizes is well defined from censuses of even-aged Mountain Ash stands of different ages. However, it is not known the amount of space (or the resource for which space is taken as a surrogate) which is used by or occupied by damaged and dead trees. To determine this, either the limiting resource must be delineated and studied, or the number of damaged and dead trees must be censused in maximally stocked stands. The first method is more suitable for a process model rather than a phenomenological model. The second is a better method for refining these parameters, although it would be a time consuming and relatively expensive task. In this thesis, reasonable looking parameter values were used with the understanding that the number of dead and damaged trees on a stand is usually very low and the belief that their impact on overcrowding is slight. 5 different sets of parameters were used to test this assumption with the levels of space occupation and usage ranging from very low to very high. This test will ascertain whether the assumption was correct and also help determine the value of attempting to gain further information on the true value of these parameters..

In all of the scenarios in this thesis there is a fifty year delay from the start of the simulation to the start of the first harvest rotation cycle. This is to give the stand some time to grow. The final sensitivity test in the section is on this harvest rotation delay. Would the long term dynamics be different if we started with a more mature or with a younger stand?

## Standard Deviation of Growth

The parameters which control the growth of trees in the system include the mean and standard deviation of the rate of growth of trees. The mean growth is the number of size classes which a tree will advance on average at a single time-step. This is well measured and the model is designed so that on average a tree will advance one size class per time step. If the mean growth were changed then the results would be different because younger trees would become hollow bearing and overcrowding effects would be more intense. In fact it would have the effect of speeding up the simulation.

The effects of changing the standard deviation are not *a priori* obvious. The standard deviation effects the spread of trees in growth (although the maximum is still capped at three size classes per time step). Increasing the standard deviation would likely increase the speed at which the forest matures and decrease the time needed before trees start to form hollows. This is because a higher standard deviation would mean that some trees grow much faster than average. It is these trees which will become dominant and the smaller trees would be the ones dying due to overcrowding. Thus a high variation in growth would tend to accelerate the growth of the stand.

Whereas the mean growth rate is well understood and parameterised from excellent data (Ashton, 1976), the standard deviation is not well measured in the field, largely due to the fact that the slower growing trees quickly die off - making it hard to see the effect of growth variation in terms of tree numbers and sizes. For this reason the standard deviation is chosen as a prime candidate for a sensitivity test.




Figure 5.1 The sensitivity to the standard deviation of growth. The results are the minimum number of trees with hollows when retaining 10 trees per hectare with no post logging mortality and a fifty year rotation length.

Figure 5.1 demonstrates the sensitivity of the test result on the standard deviation of tree growth. The effect associated with altering this parameter on the scenario is small, although the number of trees with hollows in the long run does decrease very slightly for very high standard deviations. A higher standard deviation will cause trees to grow quicker, the slower trees being crowded out of the system. The main effect of this is that the time needed for a tree to grow large enough to form hollows will decrease as will the time needed for them to grow fast enough so that they do not immediately collapse upon mortality. These effects aren't large enough to significantly change the number of hollow bearing trees in the retention scenario. Dead trees make good nesting sites but they collapse quickly enough that they aren't as important a feature in the system as living trees which are hollow bearing. This importance of dead trees to the stand is examined in depth later in this section.

The effect of decreasing the time until trees are large enough to start forming hollows is not terribly great because hollow ontogeny is still a very slow process. This is examined in the next section where the maximum hollow formation rate is varied.

## Maximum Hollow Formation Rate

Living trees which have a size class greater than 18 will start to form hollows. The rate at which they form hollows follows a linear function of basal area, thus it is a quadratic function of size class set so that the rate is zero at size class 18 and a set maximum rate at size class 50. In all of the scenarios in this thesis the maximum hollow formation rate was 0.05 per time step as a proportion of living trees.

This value was set so that the number of hollows occurring on an untouched stand corresponded with what was known to occur in nature (Lindenmayer, *pers. comm.*). As such it was considered a prime parameter upon which to do a sensitivity analysis. This is particularly true as this hollow formation rate acts on living trees to give living hollow bearing trees, which are considered to be the type of hollow bearing trees which are most significant to the number of hollow bearing trees on the stand (as seen in Scenario 4.1). They are thought to be significant because they are the least decayed hollow bearing trees and hence take the longest time to collapse.

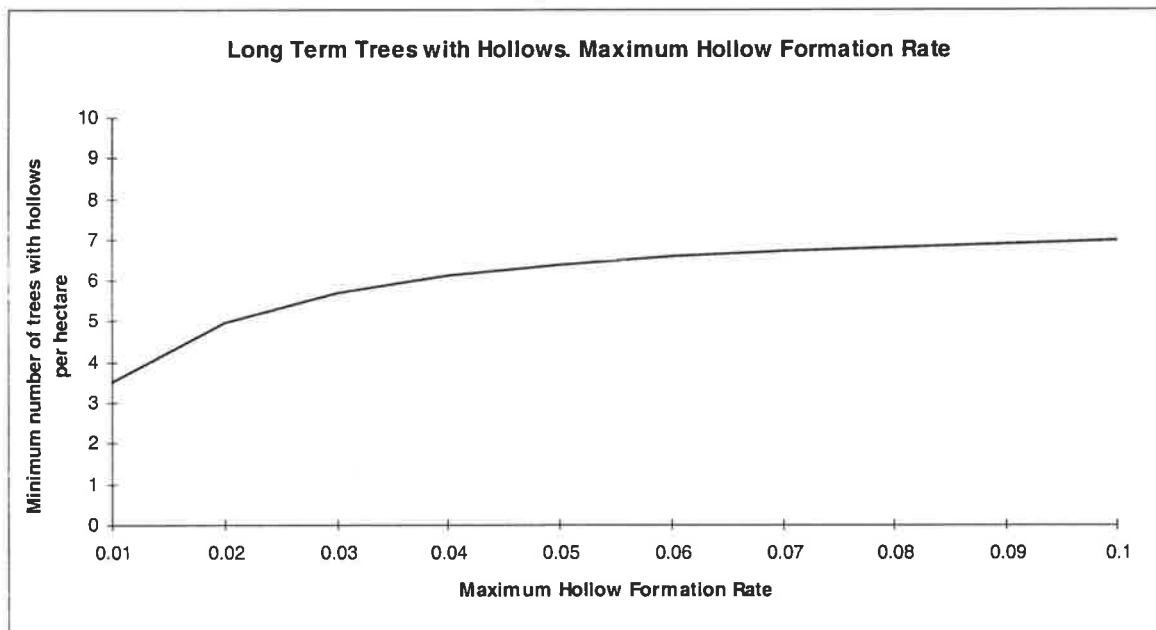


Figure 5.2. The long term number of trees with hollows per hectare on the stand. The maximum hollow formation rate is varied between 0.01 and 0.1. The value used throughout this thesis is 0.05.

Figure 5.2 shows this to be a relatively sensitive parameter, which lends confirmation to the idea that living trees which have hollows are more important to the number of hollow bearing trees, than dead ones. Interestingly the sensitivity of the model to the parameter decreases at higher values. The value used in the model was 0.05 and this is seen to be in a less sensitive region of parameter space. The response to this value becomes increasingly important when the constant is below a value of 0.03.

For hollow formation there is a period in which trees are too small to form hollows, after which there is a quadratic increase in the number of trees forming hollows (due to normal hollow ontogeny and not mortality). This helps explain why the number of trees with hollows asymptotes to a value below 10 even for very high hollow formation rates. The maximum is limited by the length of time before trees can form hollows and the mortality and senescence rates of hollow bearing trees.

### Density Independent Mortality

Density independent mortality is the transfer of trees from tree form 0 to tree form 3. The trees must be large enough not to collapse upon mortality. Trees which enter tree form 3 will senesce and collapse relatively quickly, although they are hollow bearing (which they aren't in tree form 0). For this reason it is not immediately clear what the effect will be of increasing the density independent mortality on the number of hollow bearing trees, but at higher mortalities the effect must unquestionably be detrimental to the number of hollow bearing trees on the stand.

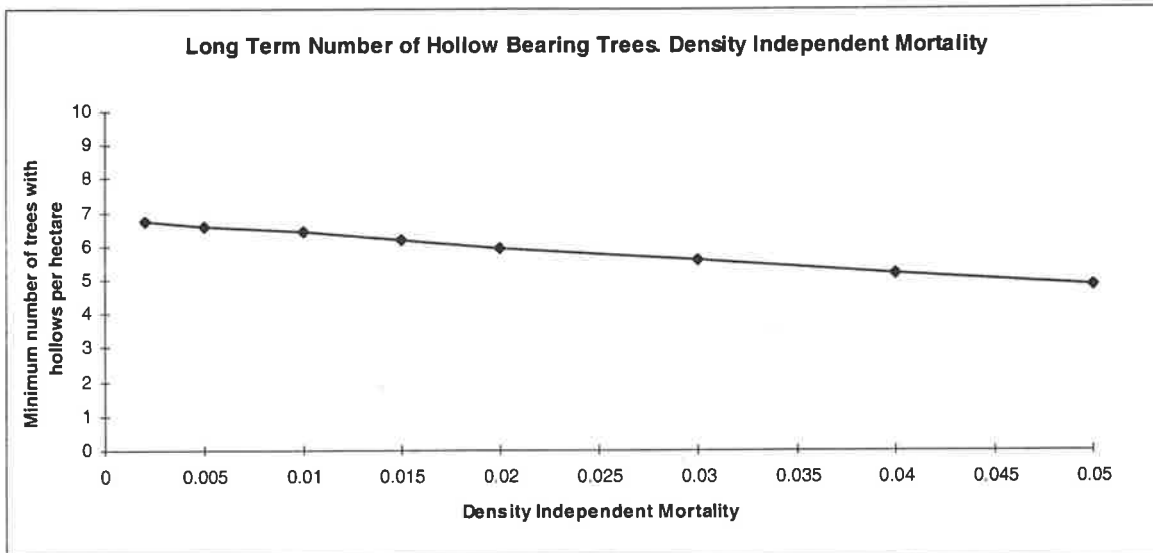


Figure 5.3. The long term number of hollow bearing trees per hectare. Density independent mortality is varied from 0.005 up to 0.05 trees in being transferred from tree form 0 (healthy with no hollows) to tree form 3 (dead trees).

These dynamics are close to linear. The greater the number of trees being transferred from tree form 0 (living with no hollows) to tree form 3 (dead trees) the lower the number of trees with hollows on the stand. At first this looks paradoxical as when a tree dies it will automatically have hollows in this model. However, this parameters shows us the relative rapidity with which dead trees decay and collapse. Hollow bearing trees which are not living are a relatively transient feature of the landscape in comparison with living hollow bearing trees.

The impact of this type of mortality on the long term number of trees with hollows is relatively slight over the parameter range of 0.005 to 0.05. The number of hollow bearing trees changes by 1.9 tree per hectare as the magnitude of the parameter changes by a factor of 10. This could only be significant if the parameter value used in this thesis was particularly inaccurate.

### Area Requirements for different tree forms

Resource use is assumed to be a linear function of space in this model. Without specifying what the limiting resource is we instead set an amount of space which is either used by each type of tree or which each type of tree prevents others from using. Determining the amount of space required by trees of different size classes is achievable by looking at normal stocking levels of differently aged stands of even-aged trees. The number of trees in such stands

immediately gives the space requirements of trees of different ages and the space requirements of trees of different size classes is derived from this.

Determining the amount of space required by or used by trees of higher tree forms is more problematic. Of course dead trees do not have a space requirement as such but they still have a space which they prevent other trees from using. It does not seem improbable that this be the case for the limiting space-related resource of the trees no matter what that resource may be. The amount of space which these trees prevent others from using is given as a proportion relative to tree form 0 space requirements (living trees without hollows). The absolute requirement will rise as the size class of the tree rises but will still follow this proportion of the space requirement of tree form 0 trees.

In all the scenarios of this thesis the proportional space requirement was set to 1 for tree form 1 trees (Which are living trees with hollows) slightly lower for tree form 2 trees, which are still living but are now missing their crowns, and significantly lower for tree form 3 and higher (which are all dead trees). The resource (or area) spoilage by collapsed trees are not included in this model.

In this sensitivity analysis 5 different cases were considered, extra low, low, normal, high and extra high requirements. The extra low area requirement has each tree form (except tree form 1) require a twentieth the area of a living tree. The low area requirement was similar to the normal but with slightly lower values, high had higher values following the same pattern and extra high had equal area requirements for all tree forms. The two extreme cases were included after the results of the low and high cases were seen.

In all five cases tree form 1 trees have the same requirements of similarly sized tree form 0 trees. These two classes only differ in the presence of tree hollows and in no other way. Tree hollows represent a form of damage to the tree but the trees are otherwise healthy and, significantly, they still bear an undamaged crown. Tree form 2 trees, although still living, exhibit more significant damage. Whereas it is expected that these trees will have very similar requirements to tree form 0 trees they have also been lowered in the most severe case of this sensitivity analysis to test the significance of this belief. The different cases are displayed in Table 5.1 and the results from the simulations in 5.4.



## Relative Area Requirements for different tree forms

Tree Form	Xlow	Low	Normal	High	Xhigh
1	1.00	1.00	1.00	1.00	1.00
2	0.05	0.80	0.90	1.00	1.00
3	0.05	0.20	0.30	1.00	1.00
4	0.05	0.15	0.25	0.50	1.00
5	0.05	0.10	0.20	0.45	1.00
6	0.05	0.05	0.15	0.40	1.00
7	0.05	0.05	0.10	0.35	1.00
8	0.05	0.05	0.05	0.30	1.00

Table 5.1 Tree form relative area usage for 5 different causes. This is the amount of space that each tree form uses relative to tree form 0 trees.




Figure 5.4. The area requirements for trees with different tree forms. These are the results from simulations using each of the 5 cases from table 5.1.

It was not expected that changing these requirements would have a great effect on the resulting number of trees with hollows in the retention scenario because it was hypothesised that the

impact of higher tree forms (particularly the dead trees) was low due to low numbers and fast decay rates of these trees. It is still surprising that the results are the same to 3 significant figures.

This shows quite clearly that the results are insensitive to the decisions made about these higher tree form trees. It suggests that the impact of higher trees forms (certainly in terms of overcrowding) is not significant it also suggests that the decision on how many different classes of tree forms to divide these trees into was quite arbitrary. A much smaller number of tree form types would do equally well for this species.

### Delay before start of harvest rotation cycle

In all of the scenarios in this thesis the harvest rotation cycle, if any, begins after the forest has aged for 50 years. Harvests occur at the end of the rotation cycle with the exception of Scenario 4.6 in which there is a harvest in the middle and one at the end of the rotation cycle, hence the first harvest occurs on a stand which is 100 years old or older in all cases. The initial delay before the harvest starts will obviously effect the speed with which the stand reaches an ultimate steady state or cycle in the number of hollow-bearing trees. But hopefully it will not effect what the final state actually is




Figure 5.5 The long term number of trees with hollows with different delays before the harvest rotation begins. Note that the harvest occurs at the end of the fifty year rotation cycle so even with zero delay the forest is harvest occurs at least fifty years after the start of the scenario.

The delay does not significantly effect the long term number of hollow bearing trees on the stand. There is a very slight effect as the delay increases through 100 years and beyond. In these cases the stands are very mature when first harvested and have an abundance of hollow bearing trees. All the retained trees in the first harvests of these simulations would be hollow-bearing and as they collapse they would be replaced with a mix of hollow-bearing and recruitment trees which have yet to form hollows. This means that they are slower to reach the equilibrium value - slower to the extent that they still differ from the other simulations after 1,500 years. They are, however, within one decimal place of the equilibrium at this stage.

## Conclusions

The results from the test scenario have proven to be insensitive to many of the parameters tested here. The scenario tested was chosen because it is one of the central scenarios in this thesis. This is an important point because I believe that it is quite possible to get a greater effect by changing the selected parameters by using different scenarios. For example if instead of retaining the ten best trees one where instead to retain trees only over a certain size class then increasing the standard deviation in growth could be expected to have a larger impact. This is because increasing the standard deviation would allow trees to reach this size class more quickly.

The main result of the baseline model is relatively insensitive to most of the parameter changes. This is the result that in the long run approximately 64% of retained trees will have hollows and the others will not yet be hollow bearing. Although this value is relatively stable it is still believed that the qualitative results from this thesis, such as that not all retained trees will or should have hollows, is more important than the specific quantitative results. The value is certainly indicative as a starting point for tree retention but it is not expected to be as accurate in the real as in the modelled world.

This sensitivity analysis points to hollow ontogeny rates as an important parameter for research. The analysis suggested that it might be sufficient to determine the lower bound these values. The results are indicative that the study of hollow ontogeny from a process oriented viewpoint is worthwhile.

It would be possible to extend this sensitivity analysis. Although it is not feasible to systematically test over the entire parameter space, the space could be examined by taking the results from a great many tests where all the parameters have been varied at random. This would still be a large undertaking and would involve some complexities in deciding how to vary the parameters between each test and in the analysis of a parameter space with a large number of dimensions.

# Section 6

## Conclusions

### Implications of the Scenarios

There are a number of results from the scenarios which have implication for forest management of Mountain Ash forests and similar forests.

- When the harvests are purely thinning operations, with a given percentage of all trees removed, the harvest rotation length becomes very important to the number of trees with hollows (Scenario 4.2). Figure 4.2.2 shows that the effect of removing a percentage of trees is magnified by a shorter rotation time with proportionately greater impacts when the percentage is high.
- Scenario 4.3 demonstrates that it is not sufficient to retain only those trees which already have hollows unless a very long (greater than 100 years) rotation cycle is in place. It is necessary to retain trees which are not yet hollow bearing in order to replace those which collapse.
- When the retained trees are *completely* protected from the normal factors which increase mortality during and following a harvesting event and retained trees are not restricted to those already possessing hollows, as in Scenario 4.4, the minimum number of trees with hollows is independent of the rotation length. In this case the number of trees with hollows which persist on the stand is simply proportional to the number which are retained. This is shown in Figure 4.4.2.
- Additional mortality associated with logging on retained trees can not be countered simply by reserving more trees so that the same number survive. This is because of the non-linear effect of the additional mortality seen in Scenario 4.5. For example, when there is 50% mortality of retained trees with hollows persisting in the stand, a greater proportion of the trees retained must be 'recruitment' trees, which do not yet have hollows, but are being retained to replace those which do as they age, collapse, and suffer tree harvesting related mortality (Figure 4.5.1).
- A short rotation time can greatly exaggerate the effect of additional post logging mortality (Figure 4.5.1). This implies that introducing measures to reduce post logging mortality

could well be the most cost-effective way of ensuring the persistence of adequate numbers of trees with hollows. Such measures could include the retention of trees in clumps to decrease effects of increased exposure to weather and clearing brush from the base of retained trees before applying any post logging regeneration burns.

- When proportional thinning is combined with tree retention in a complex harvesting schedule (Scenario 4.6) the results are very similar to the retention of trees with post-logging mortality added, although the number of hollow bearing trees does not fall as low when there is a period of time between the retention cut and the thinning operation.

In general it is harder to preserve trees with hollows than one might naively expect. Retaining all trees with hollows is not good enough, unless other trees are also retained. A post logging mortality of 50%, for example, will give less than a third the number of trees with hollows than when there is no post logging mortality, if the rotation time is 100 years (Figure 4.5.1). If the rotation time is 50 years then post logging mortality has an even greater effect. When harvesting involves removing a proportion of all trees then the number of trees with hollows can taper off dramatically with the proportion removed. For example, the number of trees with hollows for a fifty year rotation and 50% of trees removed is less than one sixth the number of trees with hollows when only 10% of the trees are removed (Table 4.2.1).

Because the model is structured around a phenomenological description of the world, the results can only be expected to be true in a qualitative sense and this within the bounds of the parameters explored. For this reason the detailed sensitivity analysis of Section 5 is very important.

## Sensitivity Analysis

The sensitivity analyses in Section 6 gives confidence in the robustness of these qualitative results. In particular, many of the hard to determine parameters, such as the space or resource requirements which dead trees prevent others from using, are not at all important to the results which derive from this study.

The most sensitive parameter was the maximum hollow formation rate parameter. Whilst there is no reason to believe that the qualitative results could be affected if this parameter were

incorrect, the quantitative value could well be significantly different. This suggests that the study of hollow formation in Mountain Ash trees is a worthwhile pursuit.

The natural rate of mortality (density independent mortality) is also understandably important, although the general findings from the sensitivity analysis indicate that the parameters are well enough understood to justify the findings arising from the scenarios described in this thesis.

## Further Work

Areas which could be explored using HOLSIM include the study of multiple coup systems and also the effect of fires on logging regimes. It is not difficult to manage a multiple coup system so that the overall minimum number of trees with hollows is some given amount. This could be well effected by staggering tree harvests so that if the number of trees with hollows follows a cycle on each coup then the cycles are staggered so that some kind of average number of trees with hollows always exists. This would only be important in a retention scheme where the realistic assumption of additional post-logging mortality is included.

Having a sufficient number of trees with hollows over a large area is not necessarily sufficient to preserve species which use tree hollows. If local hollow numbers occasional get very low then local populations might go extinct, the problem then becomes a metapopulation problem where local extinctions and recolonisations are important. HOLSIM in this framework would give you a function describing how many trees with hollows there are in each individual coup at each given point in time and then some other model could be built around this which takes the biology of the study organism into consideration with particular attention paid to migration rates. This could be done numerically using a model such as ALEX (Possingham and Davis, 1995) where the habitat value varies according to the number of hollow bearing trees found in HOLSIM, or alternatively the function describing the number of trees with hollows discovered in HOLSIM trials could be mimicked by mathematical equations which could be put into a mathematical metapopulation model (such as Day and Possingham, 1995).

The effect of fire on a coup system can be easily modelled by running a simulation up to the point of a fire, applying the fire (which would remove a great number of trees and move many others to a more degraded or higher tree form) and then following this with a continuation of

the simulation possibly with a different harvest system. Many questions could be posed. The immediate effect of a fire could be to either destroy the majority of the trees and hence reduce the number of trees with hollows, or to damage so many trees that in the immediate aftermath there are a great number of hollow bearing trees. Where the forest has attained a steady cycle (or value) of trees with hollows, the question could be asked as to how quickly this cycle is once again attained following a fire with either a continuation of the same harvest rotation or with a modified one. One could also ask whether there were some straight forward tactic to help a decimated forest stand, quickly to become more stable in terms of hollow bearing trees. Such a tactic might be to start with a larger than normal number of trees being retained and then drop down to the normal number after a few rotations.

Having looked at this model for Mountain Ash it seems natural to do a similar thing to other similar tree types such as Blackbutt. The model might also be extendable for more complex forest interactions particularly for mixed forest stands.

HOLSIM is a purpose built model to study a particular facet of the problem of retaining hollow bearing trees for the purposes of species conservation. It was devised largely with a knowledge of the data available, and with the data that could be gathered. It included those processes which were necessary at a level of detail that was commensurate with the detail known for all processes. For example, there would be no point in mixing highly detailed models of tree growth based on nutrient, aspect, slope and climactic knowledge for Mountain Ash when there was only a phenomenological understanding of hollow formation and tree decay. The detail was appropriate to the problem at hand.


There are many models for tree growth, many of them are highly detailed focussing on the processes of growth and micro-climate conditions. HOLSIM is unique in looking at hollow bearing trees and modelling of tree hollows. Because tree hollows are a limiting factor on many species and are used or are important for many more species, this makes this modelling project a significant and novel piece of work.

The conclusions are qualitatively correct although the quantitative results should be held with a degree of caution. It is clear, for example, that when retaining trees it is necessary to include



trees which do not yet have hollows - it would be rash, however, to decide that this study requires that of every ten trees retained 3.6 trees on average must be hollow-less.

## **Part 2**



### **The Nature of Reserve Design and the Design of Nature Reserves**

# Section 1

## Introduction

This part of the thesis looks at the design of nature reserve systems. This involves the allocation of sites to a reserve status which is assumed to protect all the species occurring on these sites. In this thesis sites are considered indivisible and the problem is to select which sites should be included in the reserve system from a list which contains information on the distribution of species across the sites. In the advanced problem, information on the spatial layout of the sites is also used. This thesis extends the theory of nature reserve design to problems of maximising biodiversity for a fixed area and problems involving spatial complexity both for the reserve system as a whole and for the individual requirements of the species being conserved. New algorithms have been developed to solve these problems. All of these issues are examined through a series of worked examples.

There have been a number of approaches to the design of nature reserve systems, both in the way the problem has been defined, and in the methods used to find solutions for the problem. These are reviewed in the literature review of Section 2. The third section describes formally the new problems which are examined in this thesis. The fourth section details the methods which are used to solve these problems and the fifth section details the data sets on which the problems are applied. The main methods used in this thesis are simulated annealing and a number of iterative heuristics. The former almost always gives better resulting reserve systems and the latter has the undying appeal of simplicity and speed.

There are a number of quite separate questions which are addressed in this thesis. These are tackled in Sections 6 through 10. Section 6 compares the methods described in Section 4 across a number of the traditional representation problems. These are non-spatial problems where a certain number of occurrences or area is to be preserved for each conservation value. A conservation value is an item to be preserved. It could be a species, or a vegetation type or landform type or anything else of a similar nature.

The aim of the problem is to ensure adequate representation of conservation values for the least cost or using the least area. This can be reversed to a problem of maximising representation for

a fixed area or cost. Although the algorithms used in this thesis are not designed specifically for this 'problem reversal' they can be used to approach it. In Section 7 a new fast way of integrating the two problem types is explored.

Section 8 looks at a new and different creation process for reserve systems. In this approach a reserve system is generated meeting some nominal target level. This reserve system is then used as the basis for building a reserve system with a slightly higher target level. This incremental reserve building is repeated with higher and higher targets. This is a new approach to reserve design and could be used as another way of reversing the problem, by increasing targets until a fixed cost is met. It is also a method which could be of use in scheduling the addition of new sites in a situation where the entire reserve system cannot be wholly created in a single period of time. It would be expected that a reserve built incrementally to some high target level would not be as good as one built from scratch using the high target level. In this section we will see whether, and to what extent, this is the case.

Fragmentation is an important issue for reserve design and there has been a little work done on spatial reserve design issues previously, most of which has focussed on single species reserve design. A simple way of controlling fragmentation which has not been used before is to control the boundary length of the reserve. The shorter the boundary length the more compact and less fragmented is the reserve system, for a fixed area. This is a new way of incorporating spatial considerations and it is examined in Section 9.

Section 10 looks at new ideas on spatial issues. In this section, spatial restrictions for species are included. These restrictions include what are termed the aggregation and separation rules. The aggregation rule states that the target conservation value can only exist on an area of the given size or smaller. With this rule, unconnected fragments of a conservation value are not counted toward their goals. The separation rule can be included to attempt to provide some genetic diversity within a conservation value. This states that there must be at least three locations which are separated by a given distance on which the conservation value is preserved. The aim of the separation rule is both to increase the diversity within a species and also protect the conservation value from localised disasters (risk spreading).

There are two data sets on which all these problems are tested and they are described in detail in Section 5. The final chapter reviews the findings on this part of the thesis and concludes with speculation on further directions in which this research could be pursued.

## Section 2

# Literature Review

### Introduction

A key concept to the preservation of biodiversity is the establishment of a reserve system. A reserve system is a series of protected areas, or reserves, which retain as great an amount of biodiversity as possible. In this thesis we are interested in the process of selecting which sites should make up a reserve system. The aim being to adequately preserve a set of conservation values. Conservation values are species or vegetation types or any other biodiversity feature that can be captured in a reserve system. The problem of what constitutes adequate representation is not examined here, it is assumed that experts can give some estimation of conservation value requirements.

The design of reserve systems has involved several approaches of increasing sophistication. The most common method can be termed the *ad hoc* approach, in which sites are added to the reserve system for a variety of reasons, such as that they are available and inexpensive to make reserves, that they preserve some particular conservation value such as a species or vegetation type which is seen as desirable, or they contain some conservation value not associated directly with biodiversity such as a remote wilderness value or desirable geological feature or perhaps just because the sites are all that's left after the rest of area has been allocated to other uses. Another method is the use of scoring systems which is an attempt to produce an objective method for the design of reserve systems. Scoring systems allow individual sites to be objectively scored with respect to a number of criteria but in considering sites individually rather than the worth of the reserve system as a whole, problems occur. The worth of an individual site depends critically on what is already in the reserve system. When the reserve system as a whole is considered it has usually occurred with scant regard for the spatial arrangement of sites within the reserve system.

This review will begin by looking at the scoring approach and then lead into the concept of comprehensive representation. It will then look at more advanced methods of designing reserve

systems and the addition of important components such as spatial characteristics and advanced measures of species requirements.

### Scoring Systems for Biological Conservation

A key requirement for the objective design of reserve systems is the need to identify areas which require preserving, either because they contain threatened taxa or because they have high species richness. The identification of such areas is one of the crucial steps in the construction of a reserve system and it is also useful in pin-pointing regions which are under threat. This is an important use of scoring systems. It is a separate issue from the reserve design problem covered in this thesis which does not include considerations of the level of threat to individual sites.

Scoring systems were adopted in order to produce objective results, to objectify the methods used in identifying threatened areas or to encode subjectivity so as to make it appear objective, depending upon who is to be believed. These systems have the advantage of making the basis for reserve system decisions very clear and explicit by identifying the criteria which were seen to be important and defining measures for these criteria. Two good reviews of the scoring system approach are Magules and Usher (1981) and Smith and Theberge (1986)

There were a great number of scoring systems based on a wide range of quality measures ranging from social/political measures (such as Gehlbach 1975 who also included many non-social/political measures) through concepts such as 'naturalness' or what might be termed as wilderness value (Brown and Hickey, 1990) and also including strict biological measures. These last are targeted at ensuring the preservation of biodiversity by focussing on rare or endemic species and on areas of high biodiversity. Some strictly biological measures include the Purdie *et. al.* (1986) PDI index (Priority, Density index), Rabe and Savage (1979) for marine biological measures, and Goldman (1975) with a mathematical combination of rarity, richness and global extent of species and locations.

The basic idea behind these methods is to score each site independently of all other sites depending upon a number of different criteria. The individual criteria scores are combined to give the site a score which is the site's value for a reserve system. The simplest and most common way of combining the different criteria is to add them together, possibly with some

multiplicative factor for individual criteria. A reserve could then be constructed by choosing a given number of the highest scoring sites.

Generating the score for a site was often difficult because of some of the subjectivity that was inherent in many of the measures. For example, Margules and Usher (1984) tested different scoring methods and found different results were obtained from each of 9 different experts who attempted to apply the same scoring method on a single data set. The goal of scoring systems was to produce a method for reserve selection which was objective. Perhaps for this reason the focus in conservation biology journals was placed primarily on biological objectives with the recognition that social and political objectives lay in a different sphere. The use of expert opinion has the pitfall of possibly being subjective but it has the advantage that information can be pooled from a variety of sources in a manner which is difficult for a mathematical index to use. The use of questionnaires for experts was found to be useful by Sinden and Windsor (1981), for example, to determine both the scarcity and abundance measures for species.

Scoring systems are useful for identifying locations which are under threat. They have been used to locate global hotspots and regional areas which deserve attention (Prendergast *et. al.*, 1993). Also on a smaller scale they are a useful tool to identify areas of concern such as specific locations impacted by proposed projects (See Rossi and Kuitunen, 1996, for example).

## Comprehensiveness

Comprehensiveness is the measure of how well the reserve system as a whole preserves biodiversity. This issue was first described by Kirkpatrick (1983). Thus the reserve system should consist of a collection of sites which work best together rather than a collection of sites which are individually good without regard to what is already in the system.

The scoring procedures described above do not address the criteria of comprehensiveness. The ten highest scoring sites in a region, for example, could conceivably be identical with regard to which conservation values are contained on them. A reserve system consisting of these ten sites would be no better than a reserve system consisting of only one of these sites, assuming that the conservation values are adequately represented on each of the ten sites. Prendergast *et.*



*al.* (1993) showed this effect by looking at the distribution of several groups of species in Great Britain. If the best 5% of locations were taken, measured solely on the number of species they contained, then only 50% of the species under consideration were captured.

This highlights the fact that the value of a site for a reserve system is not based solely on intrinsic factors such as species richness. The value of a site must be based on what that site adds to the reserve system, which in turn depends upon what is contained on all the other sites in the reserve system. It is possible to measure the value of an individual site for a reserve system but this value will change every time a site is added to or removed from the reserve system.

### Iterative Scoring

A simple method for handling the comprehensiveness problem is to score each site according to criteria based on what that site is adding to the reserve system. The highest scoring site is added to the reserve system and then the remaining sites are re-scored based on what they add to the updated system. This continues until some stopping criteria is met, such as the adequate representation of a suite of conservation values. This general principle was first described in Kirkpatrick (1983) and an example given by Kirkpatrick and Harwood (1983) on a Tasmanian wetland. In this example, sites were ranked according to species richness, the number of species on a site which are not yet adequately represented within the reserve system. The goal of this test problem was to preserve a single occurrence of each species.

The general principal was also outlined by Thomas and Mallorie (1985) when considering the conservation of butterflies in Morocco. These iterative methods are obviously more successful at producing efficient solutions than the approach of taking all the highest scoring sites. This has been experimentally tested by Pressey and Nichols (1989a) on two data sets taken from New South Wales wetlands as well as by Prendergast *et. al.* (1993) in Great Britain and Kershaw *et. al.* (1994) in South Africa. In South Africa, Rebelo and Siegfried (1992) used the iterative heuristics to examine the spatial layout of various reserve designs and reserve design strategies.

An algorithm which selects reserves on the basis of their richness is what is called a greedy algorithm in the field of operations research. It is a commonly used algorithm of first choice for

a variety of optimisation problems but it is known to be poor. In the field of reserve design it is not surprising that many of the rare species do not appear on the richest sites (Prendergast *et. al.*, 1993). The sites tend to be species rich because they may contain relatively common species (relatively common within the data set). Conservation values which do not appear on the richer sites will still be included in a reserve system but they are not included until much later in the iterative algorithm, after most of the constraining choices of site selection have been made.

Although the rarity of species was already a component of scoring procedures, it was in 1988 that Margules *et. al.* introduced the use of rarity of conservation values in an iterative selection procedure. The use of rarity was also tested by two of the authors in a data set based in semi-arid New South Wales (Pressey and Nicholls, 1989b). Mixtures of greedy and rarity measures became the basis for many of the heuristic reserve selection algorithms as used, (for example, see Sætersdal *et. al.*, 1993). The iterative heuristics were reviewed by Pressey *et. al.* (1993). This was followed by a test of 30 different iterative heuristics (Pressey *et. al.*, 1997) which were all different combinations of richness and rarity measures. Most recently Nantel *et. al.* (1998) used greedy based heuristics to study problems where the aim was to minimise land use conflicts.

The iterative methods produced solutions which minimised the number of sites but the problems contained no spatial information and the resulting reserve systems were usually very fragmented (Bedward *et. al.* 1992). A simple way of avoiding a large degree of fragmentation was to increase the likelihood that sites which were adjacent to existing reserves be selected. This was done by Nicholls and Margules (1993) where they incorporated whether or not a site was adjacent to existing reserves into the selection process. Testing this method was done on a coastal hardwood area in south-eastern New South Wales. A similar approach was also used to study Fynbos communities on the Agulhas Plain in South Africa by Lombard *et. al.* (1997).

An early advance on these solutions methods worked around the problem of determining the location of hard to locate species (such as fauna species) by looking at the probability that the target species occurred on each site (Margules and Nicholls 1987). This was done using the linear regression system of GLIM (McCullough and Nelder 1983). The reservation targets were then based on the expected number of occurrences for that species.

The reserve selection problem is a variant of the set covering problem of operations research (Possingham *et. al.* 1993). Some of the techniques specific to the set covering problem have found their way into the nature reserve literature such as row and column reduction (Possingham *et. al.* 1993) which can reduce the size of the problem. They showed that with simple problems the size reduction can be considerable. These were presented by Camm *et. al.* (1996) to a wider audience. Some iterative heuristics already incorporated a form of column reduction, sites which contained the sole example of some species were included in the reserve in the initiation stage of the heuristic.

### Optimality of Heuristics

In recognising that the problem is of the set covering problem type and recognising its roots in operations research it was also recognised that the iterative heuristics described above would tend to produce sub-optimal results. This was a point made clearly by Underhill (1994) who presented a simple case where the greedy algorithm fails to find the optimum value on a data set small enough that the optimal answer can be seen at a glance by the reader. Such a problem forms Example 3.1 in Section 3 below. Underhill (1994) presented the linear programming formulation of the problem suggesting that this was the appropriate method for tackling this class of reserve design problems.

### Linear Programming and the Mathematical Formulation of the Problem

There have been a number of papers outlining the representation problem in terms of mathematical programming. Initially by Cocks and Baird (1989, 1991), and most recently by Williams and Reville (1997) who make a strong case in favour of this approach. The obvious benefit that mathematical programming gives is that the result is a true optimum in the sense that the targets are met for the least possible cost of the reserve system. For this reason, reserve selection has been tackled a number of times using linear programming (for example, Saetersal *et. al.* 1993, Csuti *et. al.*, 1997, Church *et. al.*, 1996).

The problem can be simply formulated as a linear program. Let the total number of sites be  $m$  and the number of different conservation values to be represented be  $n$ . Then the set of all sites can be described as the  $m \times n$  matrix  $A$  whose elements  $a_{ij}$ , are

$$a_{ij} = \begin{cases} 1 & \text{if conservation value } j \text{ occurs in site } i \\ 0 & \text{otherwise} \end{cases} \quad (2.1)$$

for  $i = 1, \dots, m$  and  $j = 1, \dots, n$ .

Let  $X$  be an array of dimension  $m$  with elements  $x_i$ , given by

$$x_i = \begin{cases} 1 & \text{if site } i \text{ is included in the reserve} \\ 0 & \text{otherwise} \end{cases} \quad (2.2)$$

for  $i = 1, \dots, m$ .

Then the problem is to

$$\text{minimise } \sum_{i=1}^m x_i \quad (2.3)$$

$$\text{subject to } \sum_{i=1}^m a_{ij} x_i \geq 1 \quad \text{for } j = 1, \dots, n \quad (2.4)$$

where  $a_{ij}, x_i \in \{0,1\}$ .

This formulation assumes that only presence absence information is available and that the requirements is to capture at least one representation of each species. If abundance information were available then this could be included in the matrix  $A$  by allowing  $a_{ij}$  to take on values other than 0 and 1. The constraint equation (2.4) could be changed in a straight forward manner to allow for different target levels of representation for each conservation value. Once the problem has been formulated, linear programming packages can be applied to achieve an optimal solution if the problem is small enough.

The requirement of optimality was criticised by Pressey *et. al.* (1996). They made the point that optimality in terms of reserve size or cost is only one of a number of desirable features for a reserve selection process. The draw-back of mathematical programming techniques is that they are limited by the detail or complexity of the problem (Garey and Johnson, 1979). They can produce results in reasonable time for the simple forms of the problem (such as the minimum representation problem) but have difficulty solving problems with more realistic targets (Day *pers. comm.*) and have not been tested on problems which include any spatial complexity.

Optimality is obviously a desirable feature for a reserve selection process. Of equal, if not greater, importance is that the selection process has the ability to produce a range of good solutions, that the process be reasonably quick, and that it be flexible (Pressey *et. al.*, 1996). A reasonably quick selection method can be repeated with changes to the data set, the conservation values' target level of representation and with proscriptions as to which sites may or may not be included in the reserve system. This allows a land-use manager to use the method in an interactive and useful manner. This is particularly important as the final reserve will be decided by its value for conservation as well as other social factors, and also because the information on the distribution of conservation values in the data set is often open to revision. These reasons also dictate the property of flexibility. The greater and more different are the types of question that the selection process can answer, the stronger is that process.

Traditional mathematical programming methods are limited in the complexity which they can effectively handle. This makes it a relatively weak approach with respect to speed, flexibility and the ability to generate alternative solutions. The iterative algorithms are much stronger in this regard. They are easy enough to implement and its fairly easy to apply numerous different iterative algorithms to each problem. Two other methods have also been put forward to tackle these problems: Genetic Algorithms and Simulated Annealing.

### Alternative Algorithms

Both genetic algorithms and simulated annealing are general problem solving techniques which have proven to be strong and adaptable in a range of other problems (see Vidal 1993 for a large

review of applications). They usually produce good results and can be applied to problems of complexity where exact optimisation methods will fail.

Both these types of algorithms have been applied to the problem of nature reserve design (Csuti *et. al.* 1997). The genetic algorithms have had an uncertain success but simulated annealing has proved to be a strong and promising method (as will be shown in Section 6 of this part of the thesis). Of the modern methods simulated annealing appears to be the most promising (along with adapted iterative heuristics) to solve more complex and interesting problems.

## Comparison of Approaches

A number of authors have compared the many different approaches to solving reserve design problems to see which techniques are the best. The clear superiority of iterative heuristics over simple scoring selection in terms of producing comprehensive solutions was shown both by Pressey and Nicholls (1989a) and by Kershaw *et. al.* (1994). In terms of measuring the success of different types of iterative heuristics on different types of problems, there have been a few studies (eg Pressey and Nicholls, 1989b, Rebelo and Siegfried 1992), but the most comprehensive is by Pressey *et. al.* (1997) who worked on 30 different heuristics to find that algorithms based on rarity criteria but with an element of richness produced the strongest and most robust results.

A few authors have tested different types of algorithms. Sætersdal *et. al.* (1993) and Willis *et. al.* (1996) looked at solving problems both with linear programming and with heuristics. Willis *et. al.* (1996) stated that linear programming was better. In their test case on the Cape Floristic Region of South Africa both methods found the optimal solution, but here linear programming has the advantage that the resulting solution is known to be optimal. In another test of iterative heuristics, simulated annealing and linear programming conducted using terrestrial vertebrate data from Oregon U.S.A, Csuti *et. al.* (1997) felt that linear programming was the best method on problems simple enough for it to solve and that the iterative heuristics were useful for producing quick solutions and solutions to difficult problems.

In a study comparing alternatives to linear programming, in particular simulated annealing, iterative heuristics and genetic algorithms, Ball *et. al.* (in press) found simulated annealing to be the best of the algorithms for producing robust and efficient solutions with iterative heuristics still being useful for producing rapid solutions. Genetic algorithms were less successful, producing useful results but taking longer than simulated annealing and producing less efficient results. This work forms the basis of Section 6, which includes and extends this paper.

### Difficulties with Reserve Design

The difficulty with defining the problem in terms of maximising biodiversity is that it is hard to measure biodiversity. In the above problem description we have blithely looked at maximising the number of species without worrying about the relative value and requirements of species. In truth, not all species are equal, and their requirements are obviously different. Possingham and Andleman (in press) describe a method for ensuring an equity between what is reserved for different fauna species. An equivalent rule for ensuring equity in the preservation of flora species has also been prepared (Burgman *et. al. in press*).

It is not necessarily true that every species is of equal importance. In order to try to capture the idea of true biodiversity coverage there have been approaches to measure the relative input to biodiversity of different species using such methods as cladistics. The problem could then be seen as trying to span the cladistic tree Vane-Wright *et. al.* (1991). There are many different possible measures of phylogenetic diversity, Krajewski (1994) discussed 7 different measures and concluded that they were useful for measuring biodiversity but needed some standardisation to become truly practical. The issue examined in this research is whether biodiversity should be measured by genetic diversity or whether phylogenetic diversity should be used. Phylogenetic diversity can theoretically be used to measure both diversity between and within species. Arguments for and against phylogenetic diversity measures can be conveniently found in the back-to-back papers of Faith (1992) and Crozier (1992).

These approaches can set the relative targets for each species as well as weighting their relative importance. Generally speaking, they do not change the fundamental problem of reserve design in terms of selecting which sites to allocate to the reserve. The concept of spanning the

cladistic tree means that some species could be replaced with others and this would add a second dimension in terms of comprehensiveness, but the other methods simply put relative weights on each of the conservation values.

### Where have all the species gone?

The other difficulty in using species as a basic unit is that they aren't easy to survey. Usually the full range of species on a site is not known and even if a study is restricted to a small set of species such as those recognised as rare and endangered, it is not always possible to reliably say whether they are present or absent on any particular site. There are a number of alternative methods which people have used to circumvent this problem.

One way of determining where species are, is to measure easily measurable quantities such as environmental variables and then to use these to predict whether or not a species is likely to be at that location using a linear regression method such as GLIM (Nicholls, 1989). This gives probabilities for species being at locations and makes the selection problem slightly more complex. This has been used, by both Margules and Nicholls (1987) and Cocks and Baird (1989). Another similar method is Mackey *et. al.*'s (1989) use of BIOLCIM and PATN to guess at the underlying species richness for an area that has had only a preliminary survey, and McKenzie *et. al.*'s (1989) case study of the Nullarbor Region of Australia.

An alternative method is to use the location of indicator groups. These are groups which are either easy to survey or are fairly well surveyed already. The idea is that by determining the distribution of species in the indicator group you can preserve them and hopefully capture other groups which share a similar distribution. The few tests of the main assumption that some groups can act as indicators of others have failed to produce good results. The correlation between indicator groups and other groups is lacking as has been shown convincingly by Prendergast *et. al.* (1993).

Another alternative to indicator groups is using vegetation or landform types. These are easier to survey and are more likely to capture the underlying biodiversity. They have been used extensively and, in fact, form the basis of the reservation policy in at least one country (JANIS



1997). Vegetation and landform types are the conservation values which are used in the data sets in this thesis.

A new approach to this problem is the use of environmental representation (Belbin 1993). With this approach, we try to capture as much of the environmental variability as possible in the belief that this will capture as much biodiversity as possible at the same time. This is similar to the use of landform types although using data which is more readily available. Also the data is of a different type, being continuous measures rather than discrete categories or types.

The use of either landform, vegetation types or indicator groups is completely compatible with the reserve selection algorithms described above. Rather than capturing a number of species, one is capturing a range of attributes or conservation values (which might also include non biological attributes) and the underlying problem type is the same. With environmental representation the data is now continuous and could be turned into a similar problem by assuming that capturing a site will capture a range of environmental conditions (a temperature range or range around the average temperature) for instance. This way of using environmental representation has not been tested yet. Instead it has been used in a problem reversal where the aim becomes to catch as wide a range of environmental conditions as is possible for some level of cost. Thomas and Mallorie (1985) found this to be a more effective method than species conservation when testing a number of iterative heuristics on the selection of sites to preserves butterflies in Morocco.

### Current and Future Work

Irreplaceability (and substitutability) are concepts which are relatively new, they were first outlined in Pressey *et. al.*'s (1993) review of reserve selection methods. A thorough description of the concept is provided in Pressey *et. al.* (1994). Irreplaceability has basically two forms. The first is a method for scoring sites based on the idea of how limited the selection of sites would be if that site were lost. This is similar to a rarity measure in that sites holding the rarest species are obviously the most irreplaceable but it also depends upon how many rare species there are on the given sites. Using this measure of irreplaceability we can easily devise an iterative heuristic for problem solving. Lombard *et. al.* (1997) used this concept to find 'core' sites on the agulhas plain of South Africa.

The second measure of irreplaceability is done by looking at a sample of solutions to a simplified problem and determining how often each site appears. Each reserve system is a single solution and the question becomes - how many solutions become unachievable if a given site was made unavailable for selection. By looking at the reduction in solution space we can get a measure of irreplaceability for each site. The drawback with this approach is that the solution space for problems of any reasonable size or complexity becomes uncontrollably large. The assumption that every solution should be equally weighted is also problematic.

There has only been one sensitivity analysis of these methods done so far (Freitag and Van Jaarsveld, 1998). We know that the final solution depends not only on the quality of the method but also the quality of the input data but this has only been looked at by Andelman and Meir (in press) so far. It is bound to be the focus of further work. In determining the sensitivity of procedure for reserve design we can determine the robustness of various methods and also direct survey work to locations where the greatest improvements can be had.

The reserve selection problem is often described as a set covering problem: Identify the minimum set of sites required to adequately represent all species. Often the actual problem is one where the best representation for a given amount of money (or given area of land) is desired. This is called the maximal coverage problem (MCP) in operations research. Such 'reverse' problems are relatively new in the literature. Camm *et. al.* (1996) described this problem mathematically and this was followed by a case study looking at greedy and linear programming solutions by Church *et. al.* (1996). In earlier work Sætersdal *et. al.* (1993) looked at a simple form of the MCP, Belbin's (1993) use of environmental partitioning was also a form of the MCP and this was upgraded by Faith and Walker (1996) using a system they described as a 'pattern based approach' to environmental representativeness. It has been tackled using mathematical programming since (Church *et. al.*, 1996). This is an interesting form of the problem, It is a problem which can easily be tackled using such methods as mathematical programming and simulated annealing. Iterative heuristics can also be quickly devised to try to tackle them, the greedy form being an obvious starting point.

There has been little work done at including spatial concerns to the reserve selection problem. The complaint that traditional methods in ignoring space produce very fragmented solutions has been levelled against the iterative heuristics method (Initially by Bedward *et. al.*, 1991) and

some work has been done to address the problem although spatial separation, which is a feature of fragmentation, has also been lauded for the benefit of reduction of risk due to localised disaster (Rebello and Siegfried, 1990). One of the strengths of the heuristics approach is that it is very simple to produce variations on them to try new things (see Pressey *et. al.*, 1997 for 30 tested variations). The spatial configuration problem has been addressed by the addition of an adjacency constraint, described and tested by Nicholls and Margules (1993) as a solution to fragmentation and also used by Lombard *et. al.* (1996) to find what were termed core sites.

An alternative approach to determining a reserve with a good spatial configuration has been to take the boundary length of the reserve system as a whole as something that should be minimised. In particular a reserve system with as low a boundary length/ reserve area ratio as possible is one which is compact and hence desirable. Section 9 of this thesis presents a new and superior measure of compactness.

The use of SLOSS theory has been examined in this aspect of reserve design with little obvious success. Margules *et. al.* (1982) looked at the ideal shape of reserve systems, Lahti and Ranta (1985) tested the SLOSS principle by looking at reserving big or little habitats on the Finnish peatland reserves. Lomolino (1994) looked at reserving sites going from largest to smallest and vice versa. Margules and Stein (1989) also looked at the basic question from a practical view point on data from the *Eucalyptus* forests in south-eastern New South Wales.

## Dynamic Allocation

The above mentioned problems are extensions of the older work but there is a closely related problem which is also worth considering, the dynamic reserve allocation problem. In this problem rather than designing an entire reserve system from scratch you are faced with the prospect of securing land for the reserve as both the land and the funds for purchasing the lands become available. This method has been originally identified in Possingham *et. al.* (1993) and described in detail in Possingham (*in press*).

A lot of the focus here has been on alternative problem types and refinements of the normal problem definition. Defining the problem is as important as considerations of solution methods, although the two are obviously intertwined. The basic concept of what constitutes an

ideal reserve system is refined through advances both in the methods for designing them and in the way in which the problem is defined. The next section describes in detail all of the new problems which are introduced in this thesis. These include the boundary length measure concept which has been described here and also alternatives which are related to the maximal coverage problem as well as some new spatial concepts.

## Section 3

# Formulation of the General Reserve Design Problem

### Problem definition

The problem is to design a nature reserve system which will achieve adequate representation of all the conservation values as efficiently as possible. Conservation values can include features of the landscape or climate types, vegetation types, individual species or any other measurable feature which is desired for the reserve system. In the data sets which are used in this thesis they consist of landform types and vegetation types as described in Section 5. The reserve system consists of a number of locations, which are termed sites, selected from a pool of available sites. Sites are considered the primary unit of selection and they cannot be further subdivided. The sites within a reserve system need not be contiguous.

This can be formulated as a problem by including the objective of producing an adequate reserve system for a minimum cost or area. For the simplest case of preserving one instance of each conservation value for the minimum number of sites this problem can be expressed with equations 2.1 through 2.4 in the previous section, although this doesn't include a decision of what constitutes the conservation values which are to be preserved. There are many additional criteria which could also be included in the definition of a good reserve as we saw in Section 2, such as the reserve being as compact as possible.

The process of defining the problem clearly and explicitly is as important as finding solutions (ie reserve systems) and refining solution methods. This process includes determining what criteria we can and should include into the objective and also determining what data can be used and how it should be used. Some obvious steps in this definition process include the definition of what conservation values should consist of and how we can define the criteria of adequate representation in such a way that is both clear and explicit and in a way which is amenable to solution.

The simplest definition of adequate representation is that one occurrence, or representation, of each conservation value is required, or that some number of occurrences is required. A more detailed requirement is that a given area of each vegetation or landform type or given population size or occupied habitat area for a species be included in the reserve. The ideal, in terms of generating a solvable problem, is for the requirements for each conservation value to be independent of the requirements of the other conservation values and that they be no more complicated than necessary.

Setting the operational definition of adequate representation is a balance between using available information and knowledge, and the complexity and tractability of the definition. In this thesis I have looked at increasing the complexity of conservation values requirements beyond what has been done before to include spatial considerations. These spatial requirements include a separation and an aggregation rule. The aggregation rule helps to decrease the fragmentation of a conserved conservation value. It helps to ensure that the conservation value occurs in large clumps and stops it reaching its target with small fragments. For example, if we want to preserve 100 hectares of some species habitat we can use the aggregation rule to ensure that this isn't done by preserving 200 half hectare lots. The aggregation rule states that only habitat which appears above some conservation value specific aggregation target of abundance can count towards meeting that conservation value's target.

The separation rule can be used when we want to ensure that our conservation value is not concentrated around a single location. This can be useful to protect against local catastrophes as well as increasing the genetic variation captured when preserving species and vegetation types. It works as a risk spreading strategy (see, for example, Lindenmayer and Possingham, 1995). The separation rule, when it is active, requires that the given conservation value appears in at least three sites which are separated by the required distance. The separation rule can be used in conjunction with the aggregation rule.

When the spatial rules are in use the quality of a reserve for a conservation value is no longer simply a linear combination of the quality for the conservation value of the sites in the reserve. This increases the complexity greatly, although in a manner which methods such as simulated annealing can still handle. In the formulation used in this thesis the constraints are combined into the objective function as well.

## Objective Function

The objective function is a function which gives a numerical value on a partial or whole reserve. A partial reserve is a reserve which does not satisfy all of the target criteria. The objective function is the primary method used to measure the quality of a reserve system, although it is not actively used by all of the algorithms. It is used by the greedy heuristic, the iterative improvement algorithm and simulated annealing.

The objective function (Equation 3.1) consists of two sections; the first is a measure of the 'cost' of the reserve system and the second a penalty for breaching various criteria. These criteria can include a cap on the 'cost' of the reserve system and always includes the target representation level for each conservation value. In this objective function the lower the value the better the reserve system.

$$\sum_{Sites} Cost + BLM \sum_{Sites} Boundary + \sum_{ConValue} CVPF \times Penalty + Cost Threshold Penalty(t)$$

(3.1)

Here the cost is some measure of the cost, area, or opportunity cost of the reserve system. It is the sum of the cost measure of each of the sites within the reserve system. 'Boundary' is the length (or possibly cost) of the boundary surrounding the reserve system. The constant BLM is the boundary length multiplier which determines the importance given to the boundary length relative to the cost of the reserve system. The next term is a penalty given for not adequately representing a conservation value, summed over all conservation values. CVPF stands for 'conservation value penalty factor' and is a weighting factor for the conservation value which determines the relative importance for adequately reserving that particular conservation value. The penalty term is a penalty associated with each under-represented conservation value. It is expressed in terms of cost and boundary length and is roughly the cost and additional modified boundary needed to adequately reserve a conservation value which is not adequately represented in the current reserve system. The cost threshold penalty is only used when there is a fixed maximum cost, an advanced problem (seen in Section 7). The cost threshold penalty is a penalty applied to the objective function if the target cost is exceeded. It is a function of the cost and possibly the boundary of the system and in some algorithms will change as the algorithm progresses (which is the  $t$  in the above formula).

This formulation is different from the traditional one (Equations 2.3 and 2.4) in that the constraints (of adequate representation of each species) are included in the objective function as penalties for not meeting the constraints. This is done to facilitate the solution methods used in this thesis. It allows them to search through parts of the solution space which would otherwise be considered invalid. The penalties are such as to encourage the methods to eventually design reserve systems which suffer no penalty, and hence would meet all the constraints in the old formulation.

The conservation value penalty and the cost threshold penalty are examined in some detail here. Many of the heuristic algorithms used here do not use either of them, instead they will use their own method to design a reserve system which meets all of the representation requirements, and hence would have no conservation penalty under the objective function formulation.

### Cost of the Reserve System

Typically in the literature, the cost of a reserve system has been taken as either the number of sites which make up the reserve (Eg Lombard *et al.*, 1997) or the area of the reserve system (Eg Bedward *et al.* 1991, Pressey *et al.*, 1997). This surrogate for the cost of the reserve system is the element which is typically to be minimised. The important feature of these measures is that the cost of the reserve is simply the sum of the costs of each conservation value. This is a reasonable assumption and if actual costs obey this assumption then the methods used in this thesis will work identically whether area or cost is used.

Another way of measuring cost is the opportunity cost of the system for alternative industries, such as wood production. With whichever measure is used the assumption is made that the cost of a site for the reserve does not depend upon what other sites are reserved. The cost of the reserve system is simply the sum of the cost measures of each site within the reserve system.

### Boundary Length and Fragmentation

It is desirable for the reserve system to be not too fragmented (see Bedward *et al.*, 1991, for example). One simple measure for fragmentation is the boundary length of the reserve system.



For a reserve system of any fixed size, the lower the boundary length the lower the level of fragmentation. Included in the objective function is the boundary length of the system. The boundary length of the reserve system can be expressed mathematically as:

$$\text{Boundary Length} = \sum_{i \in \text{sites}} x_i \left( bl_i + \sum_{j \in \text{sites}} bl_{ij} (1 - x_j) \right) \quad (3.3)$$

here  $i$  is summed over all sites.  $x_i$  is the status of site  $i$ , it is 1 if site  $i$  is in the system and 0 otherwise.  $J$  is also summed over all sites.  $bl_{ij}$  is the boundary length between sites  $i$  and  $j$  (this is symmetric so that  $bl_{ij} = bl_{ji}$ ).  $bl_i$  is the fixed (non-removable) boundary for site  $i$ . This is the boundary between this site and land which is not available for reservation.

The costs associated with a boundary can vary depending upon the type of boundary. For this reason it might be desirable to replace the boundary length with boundary cost. This has no effect on the methods of solution. Once a boundary measure is included, whether it be length or cost, the objective function becomes a non-linear function of the sites in the system. This means that methods based on linearity, such as linear programming, will no longer work without some doctoring.

The boundary length modifier is useful to equate the boundary length to the cost of the reserve system. It is necessary because the measures are in different units but it is also useful because it controls the relative importance of size and shape, or area and fragmentation, of the reserve system.

### Conservation Value Penalty

In Equation 3.1 the constraints of the problem that each conservation value be adequately represented were included in the form of the conservation value penalty. This will allow algorithms such as simulated annealing to work more effectively at finding good reserve systems.

The conservation value penalty is the penalty given to a reserve system for not adequately representing conservation values. It is based on a principle that if a conservation value is below

its target representation level, then the penalty should be the cost for raising that conservation value on its own up to its target representation level. To take an example: if the requirement was to represent each conservation value by at least one instance then the penalty for not having a given conservation value would be the cost of the least expensive site which holds an instance of that conservation value. If you were missing a number of conservation values then you could produce a reserve system that was fully representative by adding the least expensive sites containing each of the missing conservation values. This would not increase the objective function value for the reserve system, in fact, if any of the additional sites had more than one of the missing conservation values, then the objective function value would decrease.

It would appear to be ideal to recalculate the penalties after each change had been made to the reserve system. However, this would be time consuming and it turns out to be more efficient to work with penalties which change only in the simplest manner from one point in the algorithm to the next.

A greedy algorithm is used to calculate the cheapest way in which each conservation value could be represented on its own and this forms the base penalty for that conservation value. We add together the cheapest sites which would achieve our representation target. This approach is described in the following pseudo-code:

- I. For each site calculate a value-per-unit-cost value.
  - A. Determine how much of the target for the given conservation value is contributed by this site.
  - B. Determine the economic cost of the site.
  - C. Determine the boundary length of the site.
  - D. The overall cost is the economic cost + boundary length x BLM (Boundary Length Multiplier)
  - E. value-per-unit-cost is the value for the conservation value divided by the overall cost.
- II. Select the site with the lowest value-per-unit-cost. Add its cost to the running cost total and the level of representation for the conservation value to the representation level total.
  - A. If the level of representation is close to the target then it might be cheaper to pick a 'cheap' site which has the required amount of the conservation value regardless of it's value-per-unit-cost.

III. Continue adding up these totals until a collection of sites which adequately represent the given conservation value has been found.

IV. The penalty for the conservation value is the total of the costs (including boundary length times boundary modifier) of the sites in this collection.

Thus, if one conservation value were completely unrepresented then the penalty would be the same as the cost of adding the simple set of sites, chosen using the above code, to the system, assuming that they are isolated from each other for boundary length purposes. This value is quick to calculate but will tend to be higher than optimum. There will often be more efficient ways of representing a conservation value than that determined by a greedy algorithm, consider the following example.

Example 3.1: Conservation Value A appears on a number of sites, the best ones are:

Site Number	Cost	Amount of A represented
1	2.0	3
2	4.0	5
3	5.0	5
4	8.0	6

The target for A is 10 units. If we use the greedy algorithm we would represent this with sites 1, 2, and 3 (selected in that order) for a total cost of 11.0 units. Obviously if we chose only sites 2 and 3 we would still adequately represent A but our cost would be 9 units.

This example is a simple case where the greedy algorithm does not produce the best results. The greedy algorithm is rapid and produces reasonable results. The program will tend to overestimate and never underestimate the penalties when using a greedy algorithm. It is undesirable to have a penalty value which is too low because then the objective function might not be improved by fully representing all conservation values. If there are some conservation values which need not be fully represented this should be handled entirely by use of the conservation value penalty factor, which is described below. It is not problematic to have penalties which are higher than they absolutely need to be, sometimes it is desirable. The boundary cost for a site in the above pseudo-code is the sum of all of its boundaries. This

assumes that the site has no common boundaries with the rest of the reserve and hence will again tend to overestimate the cost of the site and the penalty.

The penalty is calculated and fixed in the initialisation stage of the algorithm. It is applied in a straight forward manner - if a conservation value has reached half of its target then it scores half of its penalty. The problem with this is that you might find yourself in a situation where you only need a small amount to meet a conservation value's target but that there is no way of doing this which would decrease the objective value. If we take Example 1 once again, then the penalty for conservation value A is 11 units (see above). If you already have sites 1 and 4 in the nature reserve then you have 9 units of the conservation value and the penalty is  $11.0 \times (10 - 9)/10 = 1.1$  units. So the species attracts a penalty of 1.1 units and needs only 1 more unit of abundance to meet its target. There is no site with a cost that low - the addition of any of the remaining sites would cost the reserve system much more than the gain in penalty reduction.

This problem can be fixed by setting a higher CVPF (Conservation Value Penalty Factor) for all conservation values. The CVPF is a multiplicative factor for each conservation value, described below. Two other methods would be to alter the penalty on the fly so that the penalty is always the greedy-cost of adding that conservation value to the system, or of ignoring the problem and hoping that it is not serious. Altering the penalty structure during the operation of the algorithm would slow it down considerably. It would, however, produce more accurate penalties which might improve the performance of an algorithm, but it might well not. The option of doing nothing is catered for in the software produced for this thesis by having a user defined cut-off score for when a conservation value is deemed to be missing. This score does not effect the running of the program, merely its reporting. An algorithm might be run with a cut-off score of 98%. Then, if a conservation value is 98% or more of its target, it will not be counted as missing. If none of the conservation values are below this user defined cut-off score then the reserve system might be adjudged as meeting its requirements. If the targets have been set somewhat arbitrarily then it might well not be necessary to exactly meet all of them as long as all the conservation values get sufficiently close to their targets.

It is quite possible that the target for a conservation value is set higher than can possibly be met. In Australia where the JANIS (1997) requirements state that 15% of the pre-European area of each forest ecosystem type should be reserved, we can easily have targets which are

larger than the current area of some forest ecosystems. When this is the case, the algorithm will scale up the penalty so that if, for example, it costs 100 units to reserve all the remaining area of a given ecosystem but that represents only half of the target the initial penalty will be 200 units. This means that if you get half-way to your target then the penalty for that conservation value will be half the maximum penalty, no matter how high the target or whether it is a feasible target.

### Spatial effects on the conservation value penalty

When calculating the initial penalty for a conservation value which has spatial requirements a different method is used. In this case a very simple estimation is made for conservation values which are subject to an aggregation and to a separation rule. Sites which contain the conservation value are collected at random until the collection meets the target and the spatial requirements for the conservation value. At this point there are superfluous sites in the collection and these are removed. The remaining sites are then scored and this is the penalty for that conservation value. Here the greedy method has been replaced with an iterative improvement method.

A conservation value which has a spatial aggregation rule has a second target value which is the smallest amount of contiguous patches which will count to the main target. A patch is a group of contiguous sites, on each of which the given conservation value occurs. The second target should be something like the minimum viable population size for a species or the area required for such a minimum viable population. If a group of contiguous sites contain a conservation value but not as much as the minimum clump size then the reserve system is penalised for that conservation value as if none of those sites contained the conservation value. For each conservation value the effective representation is:

$$\text{Effective Representation}_i = \sum_{j \in \text{Sites}} \begin{cases} a_{ij} & \text{if patch} \geq \text{Aggregation Target}_i \\ 0 & \text{Otherwise} \end{cases} \quad (3.4)$$

Here 'patch' is the amount of conservation value  $i$  in the target and site and the contiguous patch in the reserve system. The 'Aggregation Target' for conservation value  $i$  is set separately

for each conservation value. Because this summation is over all sites, every site in a patch of sufficient size for the conservation value will be counted once and only once.

The separation rule is handled in an even simpler way. Algorithms which use this option look through all the sites for the given conservation value and determine if there are enough of them at the required separation distance from each other. This separation distance must be specified for each conservation value which follows the separation rule. The program determines whether there are one, two, or three or more sites which are mutually separated by at least this distance. The maximum number of sites in valid patches which are mutually separated by the given distance is the separation count of that conservation value. Three is the required separation count and hence no penalty is applied if the separation count is three or more. If the separation count is less than three then an additional penalty is added which is 0.2 times the maximum penalty of that conservation value. It is 0.5 times this maximum conservation value penalty for a separation count of 1 (the lowest). The values of 0.2 and 0.5 were selected because they proved to work well in trials. The method for calculating the separation count for a particular conservation value can be described as follows:

- I. Make a list of all the sites containing the target conservation value. If the aggregation rule is in effect then this list should only contain those sites which contain the conservation value and meet the rule. Set the nominal separation count to 0.
- II. Pick the top site from the list and then search down all the sites in the list below this one until one is found which is the same distance or further than the required separation distance.
- III. If there are no sites of this distance than the separation count remains 0.
- IV. If a site is found of the required distance then increment the separation count and search for a site lower down on the list which is distant from both selected sites by the required distance.
- V. If one is found then increment the separation count.
- VI. This process can continue for as high a separation count as desired.

The above process gives the method by which the separation count is determined. It could be extended beyond a count of three if required without loss of generality but for the purposes of this thesis a separation count of three is as high as is needed.

## Conservation Value Penalty Factor

The conservation value penalty factor (CVPF) is a multiplicative factor which can be unique to each conservation value. It is primarily based on the relative worth of that conservation value but it includes a measure of how important it is to get it fully represented. The actual effect that it will have varies between the algorithms which use the objective function. If it is below 1 then the algorithm might well refuse to add a site to preserve that conservation value if there are no other conservation values on the site.

An algorithm might well fall slightly short in the representation of species (see Example 3.1), getting close to but not at or above the target value. To ensure that each conservation value (which can meet its target) meets the target it can sometimes be desirable to set the CVPF at a much greater value than 1.

## Cost Threshold Penalty

The cost threshold penalty has been included to make it possible to look at a reverse version of the problem. The reversal of the problem would be to find the reserve system which has the best representation for all conservation values constrained by a maximum cost for the reserve system. The cost threshold penalty is included to try to tackle this problem using the same framework built to solve the traditional problem. It works by applying a penalty to the objective function if the cost of the system has risen above the desired threshold. The threshold is based around the cost of the system only, but could be modified to include the modified boundary length as well. It runs differently for the two main classes of algorithms used in this thesis as will be seen in Section 4. With iterative heuristics it supplies the stopping point. Sites are added to the reserve system until this threshold is reached and then the algorithm stops. With algorithms which use the objective function, such as simulated annealing, the cost threshold penalty is applied whenever the system goes beyond the threshold. The penalty is such that the system should be forced down to the threshold by the time the algorithm has finished.

## Alternative Problem Ideas

This thesis does not look at the case where the requirements for different conservation values are dependent on each other, but this is a case which might occur in a practical application. For example suppose there is some species of small marsupial which occurs right across the large

area over which reserve planning is to occur. The species can be easily divided up into several distinct regions and the species within these regions could be considered sub-species. We might desire that some particularly rare sub-species be included in the reserve in a large enough population and that we capture at least three distinct sub-species from the rest of the distribution with some total population requirements for the species as a whole. So, in this example, the requirements are to put aside 1,000 hectares of occupied species habitat which must include at least 100 hectares of the rare sub-species and representation from at least three other sub-species. There are a number of ways of approaching this, one might be to divide up the sub-species into separate conservation values, one for the rare sub-species and maybe one or maybe more for the other sub-species. Then the requirements for this set of conservation values are clearly inter-dependent. The main solution methods in this thesis could be applied to this type of problem although this is not tested in this thesis, nor has this type of problem appeared in the published literature.

Another logical extension of this problem is to have more management classes. Currently there are two management states which a site can be in. It could either be included in the reserve system or not included in the system. However, the methods used in this thesis could be applied to the case where there are many management classes, each of which has a different effect for different conservation values. More extension are considered in Section 11.



## Section 4

### Solution Methods

#### General

Given the objective function and possibly some constraints, the next task is to find the best (or some very good) solutions. In this section I describe the methods which will be used in the rest of the thesis to find solutions to the problems defined in Section 3. The main methods for problem solving in this thesis are the simulated annealing and iterative heuristic algorithms. Also used is iterative improvement, an algorithm which is conceptually halfway between an iterative heuristic and simulated annealing and which is usefully used in conjunction with other methods. Iterative heuristics algorithms, or more simply heuristics, is the term used for 'rule of thumb' algorithms throughout this thesis, these are of the type of simple iterative heuristics for reserve selection described in Section 2. Simulated Annealing is the primary method used in the thesis in conjunction with iterative improvement. In Section 6 the relative effectiveness of these algorithms has been tested and also compared with genetic algorithms and linear programming.

All of the methods in this section have the aim of designing a nature reserve system which will achieve the adequate representation of a number of conservation values as efficiently as possible. Conservation values can include features of the landscape or climate, types of vegetation, individual species or any other measurable feature which is deemed to be desirable for inclusion in a nature reserve system. In the data sets which are used in this thesis they include landform types, vegetation types and fauna species. The reserve system consists of the selection of a number of locations, which are termed sites, from a pool of available sites. Sites are considered the primary item for selection and they cannot be further subdivided. The sites need not be adjacent and the reserve system could be comprised of a number of smaller contiguous nature reserves across the region. The aim of producing an efficient reserve system is basically that of achieving a reserve system which is as inexpensive as possible, whilst retaining adequate representation of all conservation values. Determining the cost of a reserve

system is not entirely straightforward and often a surrogate measure is used such as the number of sites in the system or the area used by the reserve system.

### ***Iterative Improvement***

Iterative improvement is an old method which can produce good solutions quickly, it has largely been supplanted by simulated annealing of which it is a subset. It is included here as an aid to the other algorithms and as a baseline against which to test their effectiveness, the expectation being that they should do better than the average iterative improvement run.

In pseudo-code the algorithm is:

- I. Initialise the algorithm with a random reserve system.
- II. Choose a site at random which hasn't been tested (see III) since the last change.
- III. Test the objective function change if the site status were changed, by either adding it to or removing it from the reserve system.
- IV. if the change decreases the value (or cost) or the system then accept the change.
- V. return to step II until all sites have been tested and none yield a change.
- VI. Final reserve is a local optima.

The initial reserve system can be any collection of sites, or even no sites at all. If iterative improvement is being run on its own then the initial solution will be an empty reserve system, otherwise it can take the final output of another algorithm as its initial reserve and work from there to ensure that a local optima is achieved.

At each iteration the algorithm selects a site at random and tests the objective function to see if adding or removing the site will improve the reserve system. It will consider adding the site if it is not already part of the reserve system and removing it if it already is part of the reserve system. It continues to add and remove sites which improve the objective function until the system reaches a stage where no improvement will come from adding or removing any of the sites to the reserve system, a local minimum.

The main strength of iterative improvement (over the heuristic algorithms described below) is that the random element allows it to produce multiple solutions. On average the solutions might be poor, but if it can produce solutions quickly enough then it may produce some very

good ones over a great many runs. It is theoretically possible, albeit unlikely, to reach the global minima by running iterative improvement starting from either an empty reserve or a situation where every site starts in the reserve system.

### ***Simulated Annealing***

Simulated annealing is based on iterative improvement with stochastic acceptance of bad changes to help avoid getting stuck prematurely in a local minima. The implementation used in this thesis will run for a set number of iterations. At each iteration a site is chosen at random which might or might not already be in the reserve system. The change to the value of the reserve system which would occur if this site were added or removed from the system is evaluated just as it was with iterative improvement. This change is combined with a parameter called the *temperature* and then compared to a uniform random number. The site might then be added or removed from the system depending on this comparison.

The temperature starts at a high value and decreases during the algorithm. When the temperature is high, at the start of the procedure, then both good and bad changes are accepted. As the temperature decreases the chance of accepting a bad change decreases until, finally, only good changes are accepted. When the algorithm finishes it is followed by iterative improvement to ensure that the final solution is a local optimum.

In pseudo-code the algorithm is:

- I. Initialise the system.
  - A. Select a reserve system at random.
  - B. Set the initial temperature and number of iterations.
- II. Choose a site at random.
- III. Evaluate the objective function change if the site status were changed, by either adding it to or removing it from the reserve system.
- IV. If  $e^{\left(\frac{-change}{temperature}\right)} < \text{Random Number}$  then accept the change.
- V. Decrease the temperature.
- VI. Go to step II for a given number of iterations in the system.
- VII. Invoke the iterative improvement algorithm.
- VIII. Final reserve is a local optima.

There are two types of temperature decreasing schedules used in this thesis. One is fixed schedule annealing in which the annealing schedule (including the initial temperature and rate

of temperature decrease) is fixed before the algorithm commences, usually following some trial runs. The other is adaptive schedule annealing in which the algorithm samples the problem and sets the initial temperature and rate of temperature decrease based upon its sampling. There are a number of options which control how the simulated annealing algorithm works. These are discussed below.

### Selecting the Initial Reserve System

The initial reserve system is selected at random with each site having an equal probability of being selected. The probability can be set to any value between 0 (an empty reserve) and 1 (all sites starting in the reserve). There is no theoretical reason to set the probability to any particular level.

### Number of Iterations

The number of iterations determines how long the annealing algorithm will run. It will always give a final answer but the longer the run the more likely it will be that a better answer (lower objective function value) will arise. Setting the number of iterations involves resolving a trade-off. The number of iterations needs to be quite large but if it is too large then the amount of improvement for the extra time might be too poor and it might be more profitable to run the algorithm repeatedly for a shorter time and select the best from the resulting multiple solutions. Both the fixed and adaptive annealing schedule require the number of iterations to be set before commencing an annealing run.

### Temperature

The initial temperature and temperature decreases are set prior to the commencement of the algorithm for a fixed annealing schedule. The number of temperature decreases is set prior to the commencement of the algorithm for both the fixed and adaptive annealing schedules.

The number of temperature decreases is important for technical reasons. If it is set too high then there is a chance that the final temperature might not be what was expected due to round-off error. A value of 10,000 was found to be adequate (when the total number of iterations was

above 10,000). A smaller value might also be required where the difference between the initial and final temperature is small, again because of potential round-off error.

The initial temperature should be set so that almost any change is accepted. The decreases should be set with an eye on the final temperature, the final temperature should be low but not zero. When the temperature is too low then the algorithm will begin to behave like an inefficient iterative improvement algorithm at which point it should terminate, followed by the iterative improvement algorithm using the final result from simulated annealing as its seed reserve. The difference between the simulated annealing algorithm with a zero temperature (where it only accepts good changes) and the iterative improvement algorithm is that the iterative improvement will not attempt to change the same site twice unless the reserve system has changed between the two tests. The simulated annealing algorithm does not keep track of rejected changes and may test the same site multiple times.

### Adaptive Annealing Schedule

The adaptive annealing schedule commences by sampling the system a number of times (number of iterations/100). It then sets the target final temperature as the minimum positive (i.e. least bad) change which occurred during the sampling period divided by ten. The initial temperature is the same as the maximum positive change. This is based upon the adaptive schedule of Conolly (1990) it is a simplified form based on detailed initial tests.

### Fixed Schedule Annealing

With fixed schedule annealing the parameters which control the annealing schedule are fixed for each implementation of the algorithm. This is done typically by some trials of the algorithm with different parameters for a number of iterations which is shorter by an order of magnitude to the number to be used in the final run. Once good parameters are set for the short runs longer runs can be implemented. The trial runs consist of sample runs where the results are measured to see which parameters do best as well as individual runs where the progress of the reserve is followed throughout the run. Following the detailed progress of a run can give some indication of the level of progress made at each temperature and how appropriate the final temperature is.

Fixed annealing schedules which have been carefully set are generally superior to adaptive annealing schedules. Adaptive annealing is advantageous as it does not require a skilled user to apply the algorithm and because it is quicker. It is faster in terms of the processing time required as there is much less in the way of initial runs, it is considerably faster in terms of the time which the user must apply in running the algorithm. For this reason it is considered to be a serious alternative. The land-use designers and managers who will be using this approach will tend to use the standard options and automatic methods to a large extent so that the ability of adaptive annealing to design reserve systems is very important, although obviously not definitively important. Adaptive annealing is also important for broad investigations, tests and trials on the system which would precede the more careful and detailed use of a fixed schedule annealing algorithm.

### General Process for setting a fixed schedule

The two key control parameters are the initial and final temperature. Although technically the simulated annealing program uses the initial temperature and cooling factor, the initial and final temperature also define the schedule. If the final temperature is too low, then the algorithm will spend a lot of the time in a local minimum having effectively reached its conclusion long before the end of the iterations. If the final temperature is too high, then many of the important near-minimum states will not be explored and the resultant reserve system will largely be delivered by the iterative improvement algorithm which follows the annealing schedule. An initial temperature which is too high will cause the system to spend too much time in the high temperature equilibrium and hence a smaller proportion of time at useful temperatures.

A good way to get a feeling for the annealing schedule is to watch the annealing algorithm in action keeping track of the current temperature and the current value of the reserve system. This was done by displaying these values every time the temperature decreased, which was 10,000 times in one application of the algorithm. Observing this decrease indicates what proportion of time is spent with the system sitting unproductively at the equilibrium value for the current temperature and hence indicate whether either the initial or final temperature parameters should be different. A complementary method is to do many short simulations with a variety of parameter choices and compare the quality of resulting solutions to see what reasonable parameter values look like.

Generally these trial runs are done with a much smaller number of iterations than the final runs. The parameter choices have to be altered to take this into consideration. In general this means raising the initial temperature parameter when the number of iterations is increased. The parameter is raised so that the system will still be spending the same number of iterations or more at the lower temperatures but will allow a larger and more gradual cooling curve.

## **Heuristics**

Heuristics is the general term for a class of algorithms which has historically been applied to the nature reserve problem (Pressey *et al* 1993). They spring from an attempt to automate the process of reserve selection by applying a logical rule based on a common sense approach. There are a few main types of heuristics upon which the others are variations. These are the greedy heuristic, the rarity heuristic, and irreplaceability heuristics.

All of the heuristics add sites to the reserve system sequentially. They start with an empty reserve system and then add sites to the system until their stopping criteria is reached. The stopping criteria is always that no site will improve the reserve system but the definition has two slightly different meanings as will be seen. The heuristics can be followed by an iterative improvement algorithm in order to make sure that, and see if, none of the sites added have been rendered redundant by later additions.

### **Greedy Heuristics**

Greedy heuristics are those which attempt to improve a reserve system as quickly as possible. The heuristic adds whichever site has the most unrepresented conservation values on it. This heuristic is usually called the richness heuristic and the site with the most unrepresented conservation value on it is the richest site. It has the advantage of making immediate inroads to the successful representation of all conservation values (or of improving the objective score) and it works reasonably well.

Greedy heuristics not only create a set of sites for a reserve system but also an ordering as well, so that if there aren't the resources to obtain or set aside the entire reserve system then you can

'wind back' the solution - the resultant reserve system will still be good for its cost. From this perspective they can be quite powerful.

These heuristics can be further divided according to the objective function which they use. The two used in this thesis have been termed the Richness and the Pure Greedy heuristic. These are described below.

### **Richness**

Here each site is given two scores; The value of the site and the cost of the site. The cost is a combination of the site's inherent cost and the change in modified boundary length. The other score is based on the amount by which all the conservation values are increased toward their targets. If a conservation value occurs on the site and it has not met its target yet, then the amount by which it approaches its target is added to this score. This is either the amount of that conservation value on that site or the current shortfall for the target, whichever is the smaller amount. The actual value is taken as a proportion of the target for that species, which means that it does not matter what units the conservation values are measured in or whether different conservation values are measured in compatible units. The richness of the site is then just the contribution it makes to representing all conservation values divided by its cost.

### **Pure Greedy**

The pure greedy heuristic values sites according to the change which they give to the objective function (Equation 3.1). This is similar to, but not the same as, the richness heuristic. When using the conservation value penalty system, which is used with simulated annealing, the pure greedy heuristic has a few differences from the richness algorithm. It might not continue until every conservation value is represented. It might turn out that the benefit for getting a conservation value to its target is outweighed by the cost of the site which it would have to add. Also the pure greedy algorithm would be able to contain all the complexity added to the objective function, notably with regard to a conservation value aggregation and segregation rule.





## Rarity Algorithms

The greedy method could easily be driven by the distribution of relatively common conservation values. The first sites added would be those which have the largest number of conservation values and probably many of these conservation values are fairly common in the data set. The rarity algorithms work on the concept that a reserve system is best designed by selecting the initial reserves based on the distribution of the rarest conservation values. Here, rarest means that the conservation value appears on the fewest sites, it does not mean that the conservation value is necessarily the rarest or most endangered globally (ie beyond the data set).

Many rarity algorithms were explored, although they all worked in a similar manner. These have been titled; Maximum Rarity, Best Rarity, Average Rarity and Summed Rarity. The rarity of a conservation value is the total amount of it across all sites. For example, with forest ecosystems this would be the total amount available in hectares. There is a potential problem here where the rarities for different conservation categories could be of different orders of magnitude or measured in different units. For example, vegetation types measured in hectares and species measured in individuals. This is circumvented in most of the following algorithms by using the ratio of the abundance of a conservation value on a particular site divided by the rarity of the conservation value. Because the abundance and the rarity of the conservation value are of the same units, this produces a non-dimensionalised value.

### Maximum Rarity

This method scores each site according to the formula:

$$\frac{\text{Effective Abundance}}{(\text{Rarity} \times \text{Site Cost})} \quad (4.1)$$

This is based upon the conservation value of the site which has the lowest rarity. The abundance is how much of that conservation value is on the site capped by the target of the conservation value. For example: suppose that the site's rarest species occurs over 1,000 hectares across the entire system. It has a target of 150 hectares of which 100 has been met. On this site there are 60 hectares of the conservation value. The cost of the site is 1 unit (including

both opportunity cost and boundary length times the boundary length modifier). Then the effective abundance is 50 hectares (the extra 10 does not count against the target). And the measure is  $50 / (1000 \times 1) = 0.05$  for this site.

Note that the maximum rarity is based upon the rarest species on the site and that rarity on its own is a dimensioned value. For this reason this algorithm is not expected to work well and is expected to work particularly poorly where more than one type of conservation value is in the data set (for example, forest ecosystems and fauna species).

The maximum rarity value is calculated for each site and the one with the highest value is added to the reserve with ties being broken randomly.

### **Best Rarity**

This is very similar to the maximum rarity heuristic described above. The same formula is used:

$$\frac{\text{Effective Abundance}}{(\text{Rarity} \times \text{Site Cost})} \quad (4.2)$$

Although the conservation value upon which this is based is the one which has the best (Effective Abundance / Rarity) ratio and not the one with the best rarity score. This avoids the dimensioning problem but is otherwise similar.

### **Summed Rarity**

This takes the sum of the (Effective Abundance/ Rarity) for each conservation value on the site. It further divides this sum by the cost of the site. Thus there is an element of both richness and rarity in this measure. Here it is possible for a site with many conservation values on it to score higher than one with a single, but rare, conservation value. The formula here is:

$$\frac{\sum_{con.val} \frac{\text{Effective Abundance}}{\text{Rarity}}}{\text{Cost of Site}} \quad (4.3)$$

## Average Rarity

This is the same as the summed rarity but the sum is divided by the number of conservation values being represented on the site (Ie those with an effective abundance > 0, not those which occur on the site). By dividing by this number the heuristic will tend to weigh more heavily for rarer than common conservation values. The formula here is:

$$\frac{\sum_{con.val} \frac{\text{Effective Abundance}}{\text{Rarity}}}{\text{Cost} \times \text{Number of Con. Vals.}} \quad (4.4)$$

## Irreplaceability

This is the latest idea for this type of heuristic (Pressey *et. al.*, 1994). Irreplaceability captures some of the elements of rarity and greediness in a new way. Irreplaceability works by looking at how necessary each site is to achieve the target for a given conservation value. This is based on the concept of the buffer for the conservation value. The buffer is the total amount of a conservation value minus the target of the conservation value. If the target is as large as the total amount of the conservation value available then it has a buffer of zero and every site which holds that conservation value is necessary for the reserve system. The irreplaceability of a site for a particular conservation value is:

$$Irrep.= \begin{cases} \frac{(\text{Buffer} - \text{Effective Abundance})}{\text{Buffer}}, & \text{Buffer} > 0 \\ 0, & \text{Buffer} = 0 \end{cases} \quad (4.5)$$

Note that if the buffer is zero then the irreplaceability is 0. If a site is essential then this gives a value of zero. A value close to one indicates that there are many other sites suitable for representing this conservation value. There are two ways in which the irreplaceability is used to generate scores for sites:

### **Product Irreplaceability**

The irreplaceability for each conservation unit is multiplied together to give a value for a site between 0 and 1 with 0 meaning that the site is essential for one or more conservation values. This number is subtracted from 1 so that a high value is better than a low value. The value of this product cannot be higher than the value for an individual conservation value and as such it is similar to the rarity heuristics in that it will tend to select sites based on their holding of hard-to-represent conservation values.

### **Summed Irreplaceability**

The irreplaceability for each conservation value on the site is subtracted from 1 to produce a value between 0 and 1 where a high valued site is necessary for conservation purposes and a low valued one isn't. These values are summed across all conservation values on the site. This summation means that the quantity of conservation values is important and it is related to the product irreplaceability heuristic in the same way that the summed rarity heuristic relates to the best rarity heuristic.

In this section we've looked at the methods which can be used to generate reserve systems using the problem formulations of the previous section. The next section will introduce two data sets on which these problems and methods are tested. The two data sets are of different sizes and have different qualities. They are suited to different types of problems and for testing the robustness of the methods across different data sets.

## Section 5

### Data Sets

#### General

Two data sets are used throughout the rest of this thesis. Only one contains spatial information and many of the problems described in Section 3 can be applied to this one alone, the other problems can be applied to both data sets, which they generally are. By applying the problems to two data sets with different sizes and different characteristics we can see how robust the methods from Section 4 are and whether and to what extent they are sensitive to idiosyncrasies in a particular data set.

In this section we look at both the data sets and their defining characteristics. This will allow us to judge the extent to which the results from the following sections might be determined by the characteristics of the data sets. We will also see what the general characteristics of conservation value distributions are for the two types of conservation values used in this thesis.

The main characteristic measures for the data sets are the number of conservation values and sites in the set. Other measures are taken from the conservation value versus site matrix where the rows are sites and the columns are conservation values and the values in the matrix are the distributions of conservation values across the sites. From this matrix we derive the sparsity, the distribution of site richnesses and the distribution of conservation value frequencies. The sparsity is the proportion of non-zero entries in the matrix, the distribution of site richnesses is based on the number of non-zero entries in each row and the distribution of conservation value frequencies or rarities is based on the number of non-zero entries in each column.

Both the data sets used in this thesis are from Australia. One is from western New South Wales (NSW), the other is from the Northern Territory (NT). The NSW data set consists of presence absence data with a large number of conservation values and a sparse distribution matrix. The NT data set consists of abundance information of fewer conservation values on a more densely distributed matrix as well as the spatial position of each site.

## New South Wales

The data set comes from the Western Division of New South Wales, a semi-arid region dominated by leasehold grazing land and occupying about 325,000 km<sup>2</sup>. There are 248 conservation values which are landform types. The total number of the sites is 1,885. On each site the presence or absence of each landform type is recorded but not the amount. Each site is assumed to have the same cost and reserve systems are scored by number of sites and amount of missing conservation values.

The presence (and absence) of each landform type is recorded for each site. There were 8,798 recorded presences which gives the matrix of landform type versus site a sparsity of 1.88%. That is to say that 1.88% of the possible positions on this matrix contained a 1.

The richness of a site is the number of different conservation values on a site. The lowest number of conservation values on a single site was 1 and the greatest was 18. Figure 5.1 shows the distribution of species richness across the data set. Most sites have between 2 and 5 conservation values.

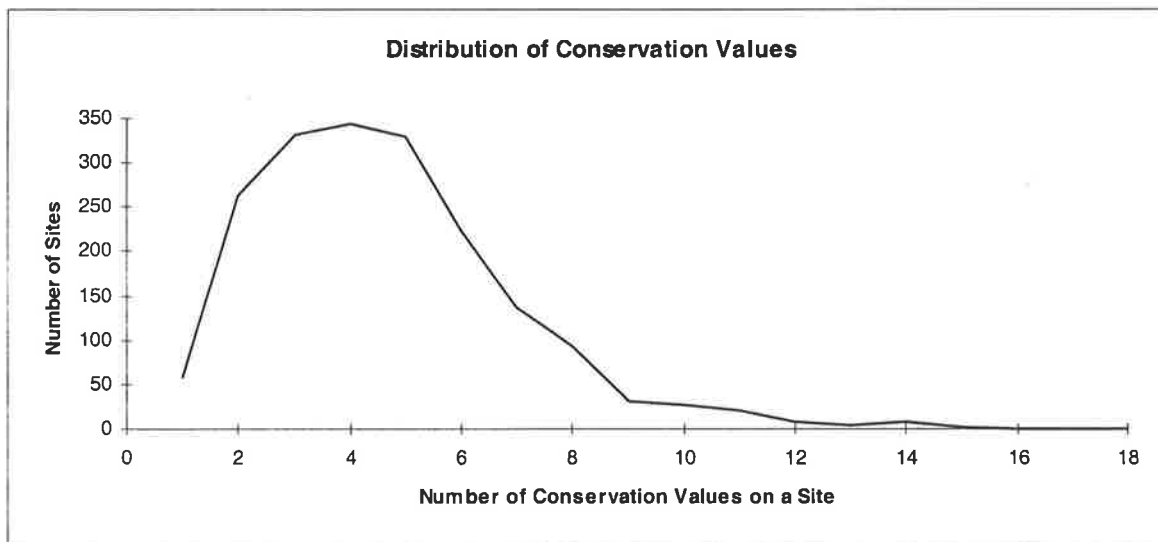


Figure 5.1 A graph of the distribution of site richness. The least rich site had 1 conservation value and the most rich had 18 conservation values.

Figure 5.1 shows the distribution of site richness. Most sites have between 2 and 6 conservation values with few having 9 or more. Most sites contain only a few different landform types.

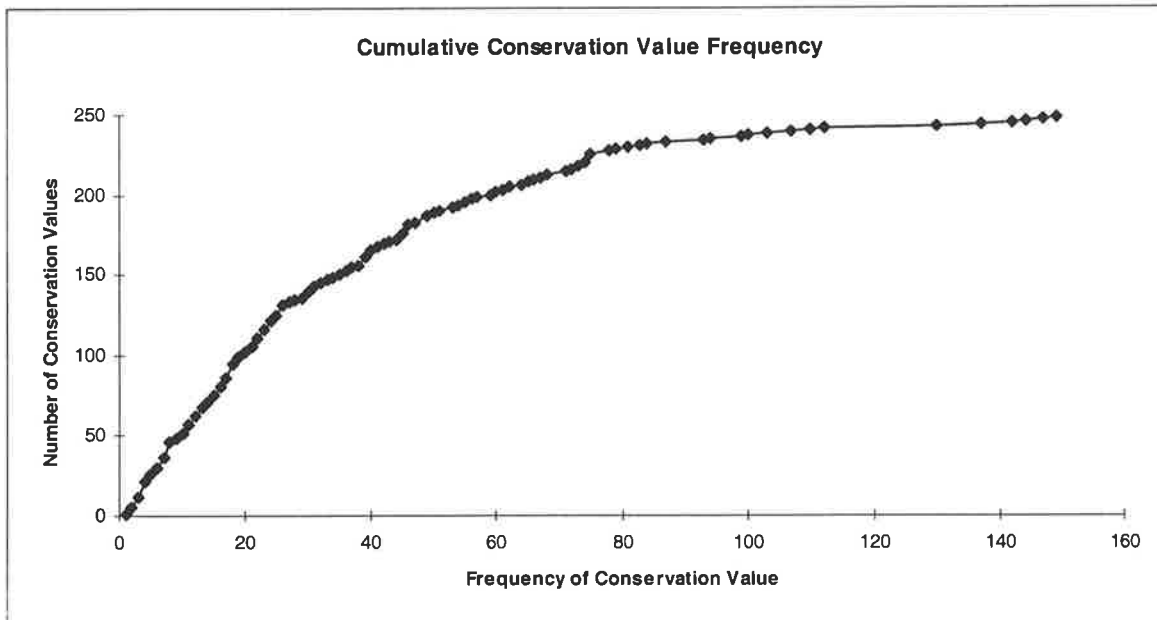


Figure 5.2 A cumulative frequency graph of the frequency of conservation values. The x-axis is the number of sites which a conservation value appears on. The y-axis is the cumulative number of conservation values of that frequency.

Figure 5.2 shows the frequency distribution of the conservation values. The rarest conservation value in the data set appears on a single site, the most common appears on 149 sites. The distribution of the species from rare to relatively common is fairly evenly spread.

The areas for the sites in this data set are not known, nor are the costs. Arbitrarily the costs have been identically set to 1 so that the aim is always to minimise the number of sites in the reserve system. Also the relative value of each landform is assumed to be identical.

## Northern Territory

The data set consists of the northern territory divided up into 1,958 sites. The entire territory has been divided up into a square grid of  $\frac{1}{4}$  degrees in size. The grid is 36 sites wide and 60 sites high. Sites are numbered sequentially from the top left, beginning with the top left corner of cell 1 at 11.00 S 129.00 E. So cell 37 is at 11.25 S 129.00 E (Decimal). Only land based cells are included.

The sites in this system form a regular square grid and the boundaries between each site are with its neighbours on each side with the exception of Melville Island which was considered

isolated from surrounding sites. Each site was considered to be a square and of equal area, those on the coast would be irregularly shaped but this was ignored. The coastline is a boundary which has no cost, or alternatively is not considered a boundary. Land boundaries on the edges of the system are non-removable boundaries.

When examining the boundary length of a reserve we are interested in the length of the perimeter of the reserve. Other elements such as roads through the reserve system are closely related to the perimeter in terms of their effect on the reserve system. It would seem valuable to include measures of such elements along with a measure of the perimeter. This does not occur in the Northern Territory data set but could be included in the fixed boundary cost for the sites.

There are 112 different vegetation types recorded in the data set and these are the only type of conservation values in the data set. The amount in hectares of each of these vegetation types is recorded for each of the sites. The sparsity of the distribution matrix in terms of presence/absence of vegetation types is 4.57%, which is more than twice as high as that of the New South Wales data set. All sites have identical costs and all vegetation types are given the same relative value for meeting their conservation targets.




Figure 5.3: Distribution of site richness in Northern Territory data set.



Figure 5.3 shows the distribution of conservation value richness in this data set. The richness of sites varies between sites which contain a single vegetation type and a site which contains 22 different vegetation types. Most sites contain between one and five different vegetation types.

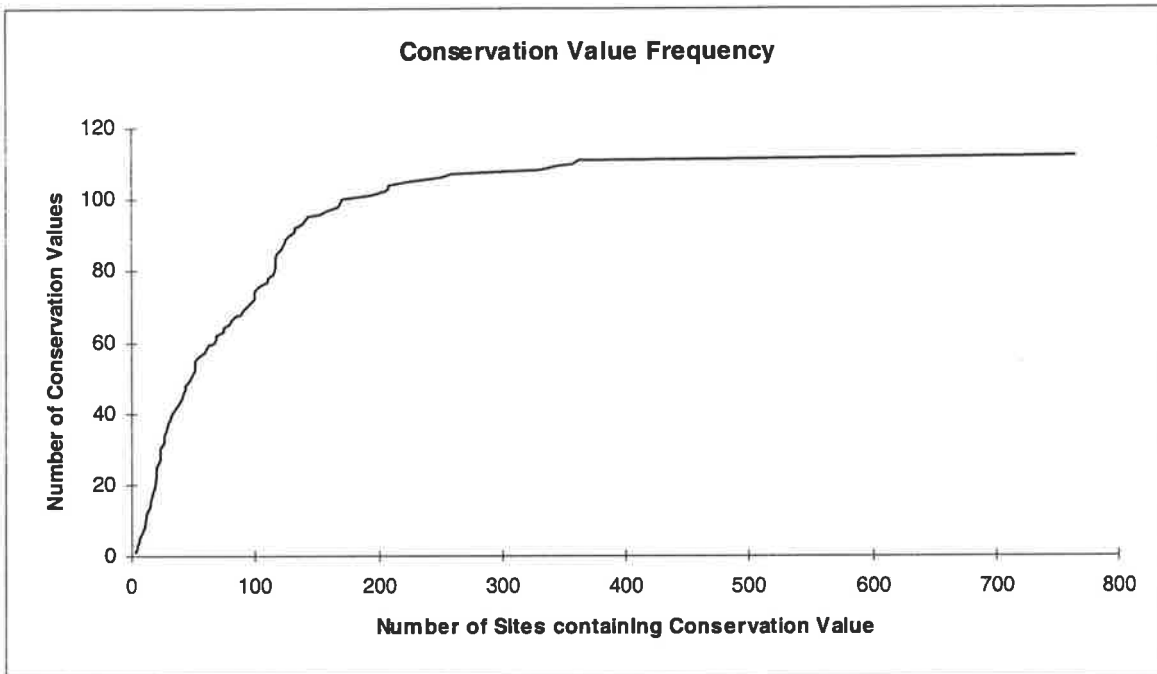


Figure 5.4. Cumulative frequency graph showing how common or rare are the different conservation values.




Figure 5.5. Cumulative abundance graph showing the total abundance for the different conservation values.

Figure 5.4 shows the frequency distribution as the number of occurrences for each conservation value. The rarest conservation value occurs on 3 sites and the most common occurs on 764 sites. The Northern Territory data set contains abundance data as well. Figure 5.5 is the abundance distribution which is similar to Figure 5.4 but using overall abundance rather than number of occurrences. The abundances of the conservation values varies between 191 Ha and 210,329 Ha. If we compare Figure 5.4 with Figure 5.5 we can note that the high end of the spectrum is more spread out with regard to abundance than it is with regard to occurrences.

The two data sets contain different types of information, occurrence data versus abundance data, and also are different in size and characteristics. The NSW data has fewer sites but more conservation values, the distribution matrix is more sparse and there are fewer conservation value rich sites. In the following sections we will see how robust are the different methods to the different data sets and how the results of the different problem are affected by differences in the data sets. The simplest problem is to try to preserve a set amount or set number of occurrences of each conservation value regardless of any spatial considerations. In the next section this problem is applied to both data sets.

## Section 6

# Algorithm Comparison

### Introduction

The aim of this section is to compare the quality of solutions which the different algorithms produce and attempt to state which algorithms work best under different circumstances and for different types of problems. Two data sets are used in this test, the NSW data set, which consists of presence/absence information and the NT data set consisting of the areas of vegetation types present on each site.

### **6.1 Algorithm comparison on New South Wales data set**

The New South Wales data-set consists of presence/absence records of a number of conservation values on a number of sites. Two types of problem are examined for this data set, occurrence problems and percentage problems. The occurrence problems require that either 1, 2 or 5 occurrences of each conservation value must be included in the reserve. The percentage problem is that 5% of the total number of occurrences for that species be included in the reserve. The percentage problem would not normally be applied to presence/absence data as it is here, however, it is a useful test giving a different pattern to how many occurrences must be reserved for each conservation value. This work includes and expands on Ball *et al* (in Press), in which genetic algorithms and linear programming were also tested for the single occurrence problem on the New South Wales data-set.

## Results

Minimum Value of Reserve Size					
Method	Problem				
	1 Occurrence no lt. imp.	1 Occurrence	2 Occurrences	5 Occurrences	5% rule
Richness	63	61	120	328	99
Greedy	63	62	126	334	102
Max Rarity	103	78	154	381	138
Best Rarity	172	72	150	375	139
Ave Rarity	119	82	155	378	150
Sum Rarity	57	57	114	316	98
Prod Irrep	57	56	114	317	98
Sum Irrep	57	57	113	316	99
It Imp	-	<b>77</b>	<b>147</b>	<b>380</b>	<b>130</b>
Adaptive Anneal	-	54	113	312	98
Genetic Algorithm	64	-	-	-	-

Table 6.1. Summary of results from Western Division NSW data. The value given is the minimum valued reserve from ten runs of the algorithm. In all but the first column an iterative improvement algorithm is run following the named algorithm to ensure that the final value is a local minima. The Genetic Algorithm result is taken from Ball *et al* (In Press) which only looked at the 1 occurrence problem. The iterative improvement algorithm is included within the adaptive annealing method and hence there is no value in the first column for this method.

The results from the experiment are presented in Table 6.1. The single occurrence problem was repeated for the heuristics in two ways. In the first (Column 1 of Table 6.1) the heuristic was run on its own, in the second (Column 2) it was followed by an iterative improvement algorithm working with the resulting reserve system to find the nearest local minima. The iterative improvement algorithm was also run on its own and is presented in the table in bold face because it is the baseline method, which other algorithms should be able to beat. Each of the main types of heuristic are separated by a line. These types are the greedy based algorithms, richness based algorithms, irreplaceability based algorithms and iterative improvement and adaptive annealing. It is known that for the 1 occurrence problem on this data set the optimum value is 54 sites (Day, *per. Comm*). This was proved by using a linear

programming package on this problem. The optimum value is not known for the other problems.

The data presented here is the *score* of the best scoring reserve of 10 repeats. This is the cost plus the penalty for any missing species. For the single occurrence and the 5% solution the score is identical to the cost. However, because one of the landforms appears on a single location it is not possible to reserve every landform more than once, and so in the two occurrence problem the score is one unit more than the cost for all algorithms. For the 5 occurrence problem the score is quite a bit above the optimum cost because there are 21 landforms which cannot be fully represented and they contribute to the score in the form of a non-removable penalty of 31 units. As we are interested only in the relative success of each algorithm it is not important whether we take the score or the cost.

The genetic algorithm result was taken from Ball *et al* (*in press*). The results were poor and the long implementation and running time of this algorithm were not supportive of this approach. Genetic algorithms have a slower running time because a number rather than a single solution is moved through the solution space, because the objective function value of new solutions need to be calculated from scratch rather than just calculating the change in value of a small perturbation of a single solution.

A strength of genetic algorithms is that it can search various aspects of the solution space independently (Beasley, *et al.* 1993), but there is no *a priori* reason to believe that the control parameters (the reserve status of each planning unit) in the reserve design problem can be considered sufficiently independent. The huge range of different variations on genetic algorithms (see for example Belew and Vose, 1997), indicate that it is certainly possible that one or another variation could perform well. However, because the results on the occurrence problem were poor and initial trials on the percentage problem were also poor (A. Smith, *pers. comm.*) this approach was not pursued for the purposes of this thesis.

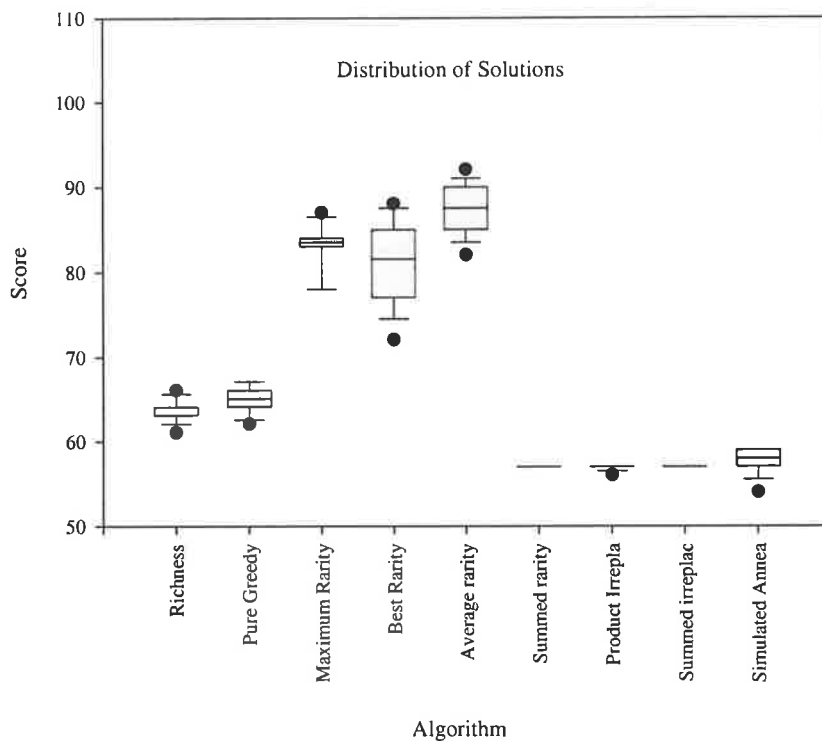


Figure 6.1: The distribution of solutions of different algorithms on the single occurrence problem.

All algorithms were repeated ten times and the best result (the smallest reserve) was taken as that algorithm's result. The algorithms produced reserve designs with a scattering of scores although the better algorithms had a much tighter scattering. Figure 6.1 shows the level of scattering for each of the main algorithms.

## Discussion

The heuristics used here add reserves to an initially empty system but they never remove them. Hence it is important to follow them with an iterative improvement schedule to make sure that the final result is a local minimum. The iterative improvement algorithm used in this way will simply remove any redundant sites. Looking at the first two columns of Table 6.1 we can see the effect of this. The rarity algorithms are aided to the greatest extent and possibly the greedy algorithms as well. The results of the product irreplaceability algorithms do not appear to be

aided by the iterative improvement algorithm. It is interesting to note that the iterative improvement algorithm improves the poorest results the most, these poorer heuristics have a greater tendency to make earlier selections redundant by later selections.

Note that running the iterative improvement algorithm following the heuristics will remove all the redundant sites. Redundant sites are ones which can safely be removed as they are no longer necessary for meeting the targets of *any* of the conservation values. Because they are redundant they do not effect the way in which the heuristics run and hence there is no benefit in checking for redundant sites whilst the heuristic is running rather than at its termination.

Of the heuristics in this problem the irreplaceability algorithms did the best followed by summation rarity and then the greedy based heuristics. The greedy based heuristics, particularly the richness heuristic did particularly well in the 5% representation problem.

In all cases adaptive annealing produced the best result. In all cases except for the five occurrence problem the score is very close to its nearest rival of the two irreplaceability heuristics. In the 5 occurrence problem it produces a result which is noticeably better.

With the exception of the summation rarity algorithm the rarity heuristics produced very poor results - poorer in general than simple iterative improvement. These are the heuristics which depend only on the rarity measure and do not take into account the richness quality of the sites. The rarity based algorithm which did do well was the summation rarity algorithm in which richness can play as great a role as rarity, particularly in the later stages of the algorithm. The problem with rarity algorithms is that once all the rare conservation values are collected the selection procedure will be based on the relative rarities of common conservation values. Once this point is reached the rarity based heuristics are adding sites almost at random. These results show how useful it is for the richness quality to be taken into account. The fact that the summation rarity heuristic outperforms the greedy based heuristics shows the importance that rarity has as well.

It is interesting that both the richness based algorithms do better under the 5% problem than the others. Here the target for each conservation value is based upon its frequency of occurrence with rare conservation values having low targets and common ones having higher

frequencies. The greedy based algorithms do better at meeting the targets of the more common conservation values. In this problem these more common conservation values play a greater role because they have higher targets.

Figure 6.1 displays the variation in solution scores for the algorithms applied to the single occurrence problem. The summed rarity and the summed irreplaceability have no variation, they only produced solutions of a single score. Product irreplaceability similarly has a very low variability. This figure shows that these results are very close to the simulated annealing results. If the average solution were taken they would produce superior solutions to simulated annealing but because the best of ten were taken the simulated annealing algorithms found an optimal solution. Because of the low variation of the best heuristic algorithms, this tends to suggest that repeating them is not worthwhile. A more worthwhile approach with the heuristic algorithms would appear to be using as many different types of heuristics as possible on each problem. Simulated annealing obviously benefits from being repeated, although another alternative that would be worth pursuing is to increase the iteration time of simulated annealing and run it fewer times. Determining the best balance between a long annealing schedule and creating multiple solution is not simple, particularly as having multiple solutions might be beneficial in adding flexibility to the reserve design process.

## ***6.2 Algorithm comparison on Northern Territory data-set***

The Northern Territory data set contains abundance measures for the conservation values, and it lends itself well to looking at percentage problems. The problems which are examined in this test are the 1 and 5 occurrence problems, and the 5% and 10% problem. The 5% and 10% problem are to preserve 5% and 10% of the total abundance of those conservation values.



## Results

Method	Problem			
	1 Occurrence	5 Occurrence	5%	10%
Richness	33	173	126	218
Greedy	33	171	114	205*
Max Rarity	47	209	145	252
Best Rarity	45	207	169	305
Ave Rarity	47	218	140	251
Sum Rarity	32	168	110	199
Prod Irrep	32	168	137	224
Sum Irrep	31	168	135	221
It Imp	<b>43</b>	<b>203</b>	<b>123</b>	<b>219</b>
Adaptive Anneal	30	165	115	207
Fixed Schedule			109	202

Table 6.2: Summary of comparison of algorithms on the Northern Territory data. The value given is the minimum valued reserve from ten runs of the algorithm. The value is the score of the system which is identical to the number of sites in the system with the single exception of the greedy heuristic in the 10% problem. Here one conservation value has not met it's target.

This table shows the results of each of the different algorithms on four problems in the northern territory data set. The 205 result for the greedy algorithm included one vegetation type not achieving its target. In this case a reserve could be designed with the same value which does represent that species, the number of sites in the system would be greater but the value would be the same as there would be no missing conservation value penalty.

## Discussion

The adaptive annealing schedule produced the best results for both the single and the 5 occurrence problem. They produced poorer results for the 5% and the 10% problems. The greedy based algorithms did much better for this than the NSW data set. The irreplaceability algorithms did well with the occurrence problems and poorly with the percentage problems. Again the rarity algorithms with the exception of summation rarity, did poorly on all of the tests. They did more poorly than the iterative improvement algorithm did on its own.

Because the adaptive annealing schedules did so poorly a fixed schedule annealing algorithm was also tested to see the extent to which the problem was the annealing algorithm or the particular annealing schedule used. The fixed schedule annealing had a number of initial runs at a low number of iterations to find a suitable annealing schedule and then the algorithm was run with a higher number of iterations to produce higher quality results. The indications are that the adaptive annealing schedule produces relatively poor results in this case. This is due to a poor annealing schedule which appears to be due to the wider range of values which occur when dealing with the full abundance problem. When dealing only with occurrence problems the values which occur are always integer values covering a relatively small range and the schedule is adequate.

It is surprising that the irreplaceability algorithms do more poorly than the iterative improvement algorithm in the percentage problems. Quite possibly this indicates that rarity measures are less effective for data sets containing abundance information as opposed to simple presence/absence information. However, it is always possible that this is just due to idiosyncrasies in the Northern Territory data set.

The summed rarity does well in both types of problems. This appears to be the best mix of rarity and greedy measures of all of the heuristics tested here. That the greedy component is important is evidenced by the good quality of solutions from the greedy based algorithms for the 5% and 10% problems. Possibly the finer detail of information in the abundance problems suits greedy based algorithms to a greater extent.

### ***6.3 Time comparisons of different algorithms***

This section includes time comparisons of the different algorithms for two different scenarios: The single occurrence problem applied to the NSW data set and the 5% problem applied to the Northern Territory data set. In each case the iterative improvement algorithm followed the application of the algorithm to ensure a local minima. In each case the algorithm was repeated ten times, the same number as used in Sections 6.1 and 6.2.

All the test runs were carried out on a Toshiba 440CDX notebook running a Windows NT operating system. Only the relative running times have any significance, the absolute running times are primarily effected by the type and speed of the computer on which they are run.

## Results

Running Time of Algorithms		
	Problem	
Method	1 Occurrence NSW Data set	5% NT data set
Richness	8' 30"	12' 37"
Greedy	6' 8"	5' 32"
Max Rarity	7' 22"	5' 1"
Best Rarity	13' 58"	8' 0"
Ave Rarity	8' 16"	5' 19"
Sum Rarity	4' 25"	4' 22"
Prod Irrep	4' 18"	4' 50"
Sum Irrep	4' 20"	4' 40"
It Imp	58"	38"
Adaptive Anneal	103'33"	89' 40"
Genetic Algorithm	*	

Table 6.3. This is the running times of the different algorithms on two of the problems described in Sections 6.1 and 6.2. All the test runs were carried out on a Toshiba 440CDX notebook running a Windows NT operating system. The genetic algorithm took half again as long as simulated annealing in Ball *et al* (*in press*).

## Discussion

These results are primarily important for their relative value and also for a rough measure of the magnitude of their value. Cross testing this on desk top computers have given running times which are approximately a quarter of these running times. It is believable that newer computers will be able to run them faster as well.

The iterative improvement algorithm was the fastest. The others were all followed by an iterative improvement algorithm but would obviously still be an order of magnitude longer had they not been. The different heuristics took between about 4 minutes and 13 minutes to run 10 repetitions. The adaptive annealing algorithm took over an hour. The different heuristics varied considerably in the length of time they took to run but they were still of relatively short duration.

The length of time it takes for an adaptive annealing run can be controlled by changing the number of iterations. In this implementation the number of iterations was 1,000,000 but when time permits a value of 10,000,000 is also worthwhile. The longer the run the better the solutions are on average.

A sensible categorisation of run time which corresponds to how an approach could be used is: Very Brief (~1 minute), Short (~5 minutes), Medium (~1 hr) and Long (hrs/days). This would correspond with how the algorithm could be used by a reserve system designer. Any method which was of medium or long duration would have to be carefully set up and run, possibly overnight. This would make sense where the number of scenarios to be explored were small and the desire for good results was great. With algorithms which take a short time, scenarios can be posed quickly and modified quickly in a much more interactive fashion. Run times which are very brief would give very quick answers. For simple questions, a simple examination of the results and the posing of another scenario would each take longer than the actual run length for such algorithms.

The algorithms basically fell into three main time classes. The iterative improvement was Very Brief, the heuristics were Short, particularly if run only once. Adaptive annealing was either Medium or Long depending upon the number of iterations chosen. Taken with the quality of the results they produce (Tables 6.1 and 6.2), and the amount of time available, this can indicate when one or another approach is desirable.

## Conclusions

These results show some interesting differences between occurrence problems and percentage problems. The refined rarity heuristics such as the irreplaceability algorithms and the

summation rarity algorithms do very well with occurrence problems and the greedy based algorithms do better with the percentage algorithms. Most of the rarity based algorithms do very poorly. They clearly demonstrate the importance of mixing rarity with greedy or richness measures for doing heuristics.

Of the heuristics the summation rarity is the most robust and produces very good results under a wide variety of conditions. However, one of the chief characteristics of the heuristics is their speed and it would appear wise to keep many of the heuristics available and test many of them on each problem or data set. If the number of heuristics were to be pared down the obvious ones to include are the greedy heuristic, the summation rarity heuristic and the two irreplaceability heuristics.

The annealing algorithm does well under all conditions and this test proves its strength at solving a range of problems. Weaknesses in the adaptive annealing schedule are apparent when looking at abundance data. It was easy to improve the schedule through trial and error and it would have been easy to produce an adaptive annealing method which performed well under the abundance situation but it would not necessarily work well for the occurrence problem.

The difference in the abilities of the heuristics on different types of problems leads to the conclusion that it is often worthwhile testing a number of different heuristics for a given problem to see which one is useful for that particular problem. The strength of simulated annealing is that it does well, often the best, under all of the problem types which we've considered. Having used more than one data set with different types of measures and different types of conservation values gives confidence in the conclusions but the more data sets that are tested the stronger would be these conclusions. Ideally these tests would be repeated over a larger number of data sets and on data sets which used different types of conservation values, such as actual species, and at different scales.

The use of the algorithms in an interactive reserve system design atmosphere should be modified by the length of time the algorithms use. To produce the best results a designer would use the best algorithms independent of how long they took to run (as long as run time was no longer than a week or so). However, quicker algorithms which produced poor results could still

be very useful for understanding the data and testing different design ideas and scenarios in preparation for a final application of a slower and better algorithm.

# Section 7

## Cost Threshold

### Introduction

Designing the most efficient reserve system which achieves all conservation value representation requirements is the ideal form of the nature reserve problem. An alternative approach is to try to produce the best reserve system for a fixed amount of area or site cost. What was a constraint becomes the objective and vice-versa. This is a reversal of the problem, and it is then not necessarily straight forward to define what is meant by the best reserve system. The best reserve system should maximise the representativeness of all conservation values, preferably without sacrificing one conservation value for another.

One way to approach this problem is using a cost threshold. A maximum cost is set on the reserve and the algorithms attempt to produce a reserve design which best achieves its objectives subject to that maximum cost. With the heuristics this means that they continue to add sites to the system until the cost threshold has been met. For both iterative improvement and simulated annealing the cost threshold works using penalties (described in Section 4) so the algorithms will never exceed the threshold.

In the problems in this section the cost is the number of sites in the reserve, but the cost could just as easily refer to the cost of the reserve or some combination of cost and boundary length.

It is to be expected that the greedy algorithms would perform well under a tight cost threshold. A greedy algorithm is obviously optimal in the degenerate case where only a single site is allowed in the reserve. We have seen from the algorithm comparison in Section 6 that they work poorly relative to algorithms which include a rarity component for the non-threshold problem. We would expect that for a threshold that was much lower than the minimum reserve size (which is 54 sites for the NSW data set) the greedy based heuristics would do the best and that they would do relatively worse when the threshold was closer to this optimal size.

Likewise simulated annealing can do no better than a greedy algorithm for the extreme case of adding a single cost threshold and it is known to do better in the case where there is no cost threshold. Does it do better or worse than the greedy algorithms? and if so when? Do we see the same pattern between the success of the summation rarity and the irreplaceability heuristics as we did in Section 6?

### **7.1 Cost threshold: Single Occurrence Problem**

The data set is the New South Wales data set and the problem is to reserve 1 occurrence of each conservation value. It has been found previously that the optimum reserve size is 54 sites. For this test the cost threshold has been set to 20, 30, 40 and 50 sites. The best solution will be one which represents the most number of conservation values for the given cost. For this reason the number of missing conservation values is the main score which is used for the test.

Three of the heuristics that were tested in the previous chapter have not been tested here. These are the maximum rarity, best rarity, and average rarity heuristic. They have not been included in this test because they produced very poor results in the algorithm comparison test. Iterative improvement was found to produce very poor results as well, but it is included to provide a lower baseline.

There are three cost threshold parameters - the threshold and two shape control parameters. The threshold is used by all the algorithms and it is the cost which should not be exceeded. The heuristics use this as a stopping point and simulated annealing along with iterative improvement uses this as a point beyond which a penalty is applied. The penalty depends upon how far beyond the threshold the reserve system is and on the two control parameters according to the following formula:

$$\text{Cost Threshold Penalty} = \text{Amount over Threshold} \times (Ae^{Bt} - A) \quad 7.1$$

Here  $t$  is the proportion of the run, it starts at 0 and ends at 1 at the termination of the run.  $A$  and  $B$  are the control parameters.  $B$  controls how steep the curve is (a high  $B$  will have the multiplier varying little until late in the run).  $A$  controls the final value. A high  $A$  will penalise



any excess of the threshold greatly, a lower  $A$  might allow the threshold to be slightly exceeded. The multiplier starts at 0 when  $t$  is zero. Both  $A$  and  $B$  require some experimentation to set. In this test these pre-trial experiments were carried out with the most stringent cost threshold (20) and once the values were found they were kept throughout the trial.




Figure 7.1. The cost threshold penalty as applied during the simulated annealing runs. The x-axis is a proportion of the run length. The parameters were  $A = 10,000$  and  $B = 0.1$ . The curve is almost indistinguishable from a straight line.

The cost threshold system in simulated annealing works more crudely when using the adaptive annealing variation. The cost threshold system works by applying an ever increasing threshold penalty on the system. This is to allow the system to explore parts of the solution space above the cost threshold at an early stage of the annealing schedule and yet keep it to the cost threshold toward the end of the run. The way that the adaptive annealing system works is to sample the solution space a number of times and then set a fixed annealing schedule from this. The sampling period is at the start of algorithm and works under an initial cost threshold penalty regime. For this reason there is a conflict with using adaptive annealing under a cost threshold system. It will still work and is included in this experiment as a test of its robustness.

A second type of cost threshold was tested, the fixed cost threshold system. In this system the same penalty is applied throughout the annealing process whenever the cost threshold is exceeded. This has the advantage for adaptive annealing in that the initial sampling can be

more accurate. The drawback is that with a large threshold penalty the system might find its search effectively restricted to some small part of the solution space.

A fixed annealing schedule was used under the cost threshold of 20. The initial and final temperature values were trialed with short runs to find a good combination and then this combination was used on the problem with a normal run length (1,000,000 iterations). This was done to measure how much better a refined method could do than the simple adaptive annealing.

## Results

<b>NSW Cost Threshold Test</b>				
	<b>Cost Threshold</b>			
<b>Algorithm</b>	<b>20</b>	<b>30</b>	<b>40</b>	<b>50</b>
Richness	83	47	26	13
Greedy	83	47	26	18
Sum Rare	118	74	37	11
Prod Irreplaceability	128	76	39	11
Sum Irreplaceability	129	76	37	11
Iterative Improvement	153	118	83	51
Adaptive Anneal	113	63	28	8
Fixed Thresh Anneal	112	65	27	4
Fixed Anneal	109			

Table 7.1. The missing conservation value penalty from the best of ten runs applying each algorithm. This reflects the quality of the solution with the lower the number meaning the closer to the target representation. Fixed annealing, which requires considerable pre-testing, was applied only to the cost threshold of 20 to see how close the annealing method could come to the greedy-based algorithms for this problem.

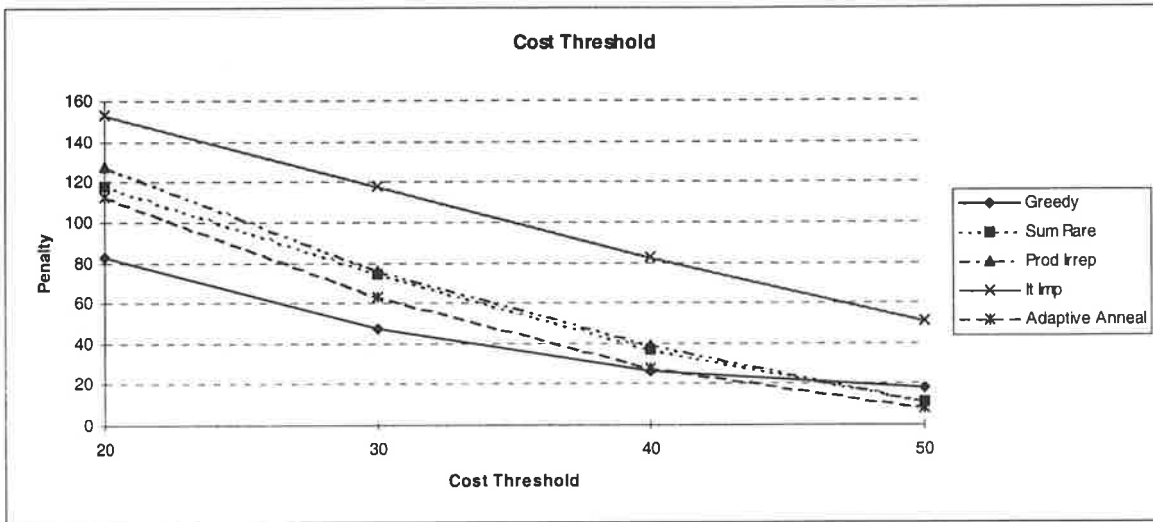


Figure 7.2. The results from table 7.1 displayed in a graph. For clarity the richness and the summed irreplaceability heuristics are excluded from the graph. Their results are not visibly different from the greedy and product irreplaceability heuristics respectively.

The results given are the best of ten runs. The values are the number of conservation values missing from the resultant solution. On the graph two algorithms have been excluded because their results are very close to those of similar algorithms and the graph would be less legible with them. The fixed annealing schedule was only used on the cost threshold 20 run and is not displayed on the graph either.

## Discussion

The results show that the greedy based heuristics do exceptionally well for the highly restrictive cost thresholds and their success tapers off only as the cost threshold approaches the cost of the optimum solution (54 sites capturing all conservation values). For the lower cost thresholds they do very well.

It is also interesting to note that they also do better than both the adaptive annealing schedule as well as the fixed annealing schedule. It is not surprising that rarity algorithms will do poorly at restricted solutions because they make the initial selections very carefully to create a good base for later choices but these initial selection are not expected to be a good reserve on their own. However, annealing is expected to work well on both unrestricted and restricted problems. It still does much better than the non-greedy based algorithms being very close to greedy when the cost threshold is 40 and better when the threshold reaches 50.

In the initial trial the Summed Rarity algorithm was almost indistinguishable from the two irreplaceability heuristics. Here it is significantly better under the lowest cost threshold and achieves similar results for higher cost thresholds. Perhaps this is an indication that it is more dependent upon the richness of the sites rather than the rarity of the conservation values. Rarity is a more sensible measure to use than irreplaceability when the threshold is below the minimum required amount. Irreplaceability measures how necessary each site is to achieve the targets for all conservation values but this is less relevant when it is impossible for each conservation value to achieve their targets. Rarity is similar in that it determines how difficult it is to collect each conservation value. As this is not based upon the target size (which is too high to be met under a low cost threshold) it will do better under a low cost threshold situation. It is expected that this difference would be greater the larger the targets for the conservation values. Thus if the targets for each conservation value are much higher than can be achieved for the given cost threshold, the irreplaceability heuristics' performance would slip significantly.

The reserves which are created by the adaptive annealing algorithm tend to do better when the threshold penalty is fixed, rather than using an increasing penalty. Although, the results are close enough that any benefit is not clearly significant. When the threshold was set at 50 the fixed threshold penalty version of adaptive annealing found an optimum solution. The smallest number of sites required to meet the targets is 54, hence any solution in which the number of missing conservation values + the cost threshold is 54 is also optimum. Fixed annealing was only applied to the most severe threshold of 20 and it is better than the other versions but still significantly worse than the greedy based heuristics.

The iterative improvement algorithm does the most poorly of all, which is to be expected. It is interesting to note that the penalties for the solution, as a function of the cost threshold, is a linear relationship, for all others it is a convex curve.

## Conclusion

From this study we can conclude that greedy based algorithms will do better when the threshold is below what is required, and that the lower the threshold the better will be the result relative to other methods. We note that the cost threshold problem is a type of maximal coverage problem. The typical greedy heuristic for a maximal coverage problem is the same for

the minimum set problem, the iterative improvement and annealing algorithms are also appropriate for the problem. The others are not designed for this type of problem but still perform tolerably well.

A natural extension of this work is to look at increased requirements and an increased cost threshold. Does the pattern continue when 2 occurrences of each conservation value are required and the cost threshold is higher?

## **7.2. Cost threshold: 2 Occurrence problem.**

### **Aims**

The aim of this experiment is to repeat the cost threshold experiment where the conservation requirements and the threshold are higher. A secondary aim is to examine the evenness of the solutions. How many conservation values are represented at least once before the algorithms begin to represent some a second time? This is an important issue for the cost threshold approach because the algorithms are not designed to aim for evenness - There is nothing explicit in the cost threshold approach which attempts to maximise evenness. Gaining an additional occurrence of any of the under-represented conservation values has the same benefit, whether it is the first occurrence for the conservation value or whether it is relatively well represented but still below its target. For this reason it is interesting to know to what extent evenness occurs naturally with this approach.

### **Methods**

A reduced number of algorithms were used in this study. Because a number of the algorithms work in very similar ways and produce very similar results this experiment limits itself to looking at the main different types of algorithms. These include the greedy, summation rarity, product irreplaceability and iterative improvement algorithms along with the adaptive annealing algorithm using the normal cost threshold method as described in Section 4.

Each conservation value has a target of two occurrences and the cost threshold is set at 60, 70, 80, 90 and 100. The smallest reserve from Section 6 which satisfies the two occurrence requirement is 113 sites. The missing conservation value penalty is taken as the base measure for the solutions. This penalty gives two points for every conservation value which is completely unrepresented and 1 point for every conservation value which is represented only

once. As well as this, the breakdown of conservation value representation is examined separately.

## Results

2 Occurrence Problem					
Algorithm	Cost Threshold				
	60	70	80	90	100
Greedy	96	71	48	43	24
Sum Rarity	156	115	78	50	24
Prod Irreplaceability	171	126	88	54	23
Iterative Improvement	221	184	183	100	79
Adaptive Annealing	132	98	58	28	12

Table 7.2. The missing conservation value penalty for the best of ten runs of each of the algorithms. Note that this is larger than the number of missing conservation values.




Figure 7.3. The results from Table 7.2 presented as a graph.

<b>2 Occurrences. Missing Species</b>										
<b>Cost Threshold</b>	<b>60</b>		<b>70</b>		<b>80</b>		<b>90</b>		<b>100</b>	
	<b>0</b>	<b>1</b>	<b>0</b>	<b>1</b>	<b>0</b>	<b>1</b>	<b>0</b>	<b>1</b>	<b>0</b>	<b>1</b>
<b>Greedy</b>	25	46	16	39	10	28	4	35	4	16
<b>Sum Rarity</b>	50	56	36	43	22	34	12	26	4	16
<b>Prod Irrep</b>	61	49	43	40	30	28	14	26	3	17
<b>It Improvement</b>	75	71	65	58	55	68	31	38	25	29
<b>Adaptive Anneal</b>	32	68	29	40	16	26	3	22	3	6

Table 7.3. showing the level of partial representation. This is the number of conservation values which have 0 or 1 representation in the best solutions.

<b>Representation Ratio</b>					
	<b>60</b>	<b>70</b>	<b>80</b>	<b>90</b>	<b>100</b>
<b>Greedy</b>	0.543	0.410	0.357	0.114	0.250
<b>Sum Rarity</b>	0.893	0.837	0.647	0.462	0.250
<b>Prod Irreplaceability</b>	1.245	1.075	1.071	0.538	0.176
<b>Iterative Improvement</b>	1.056	1.121	0.809	0.816	0.862
<b>Adaptive Annealing</b>	0.471	0.725	0.615	0.136	0.500

Table 7.4. Derived from Table 7.3, this is the ratio of those conservation values which have 0 representation divided by those which have 1 representation. A value of 0 would mean that all conservation values are represented at least once and would be ideal. The higher the value the poorer the result.

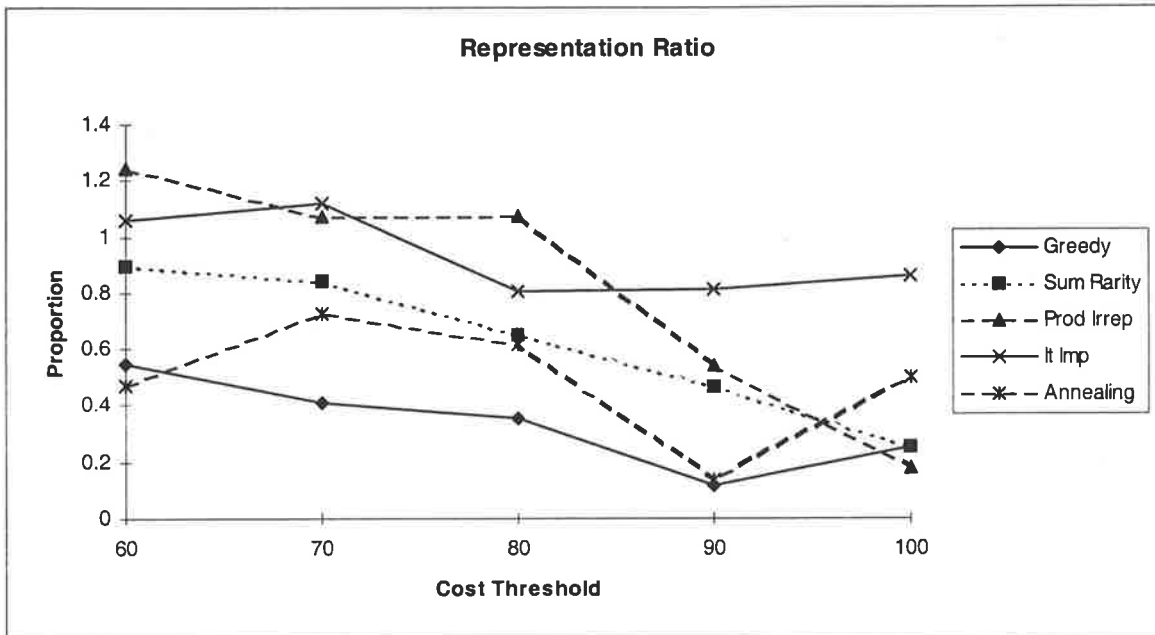


Figure 7.4. Graphical representation of Table 7.4.

The results are given in three different forms here. The first is the raw missing conservation penalty. This is a measure of how well the different algorithms do. The second is the number of conservation values represented not at all and only once. This indicates the evenness of the selection process. A single measure of the evenness is the ratio of the number of conservation values not represented divided by the number represented only once. This ratio looks at the split of conservation values as opposed to the overall number. It is desirable for as few conservation values as possible to be completely unrepresented and so a value of 0 is the best possible value and the higher the value the worse the result. When the threshold is 50 the number of missing species is low and the ratio varies widely and is probably misleading.

## Discussion

The trend of the previous experiment is repeated here. The greedy algorithm does very well for the lower cost thresholds and better than the other heuristic algorithm for all but the highest cost threshold tested. The adaptive annealing algorithm is better at a cost threshold of 90 and higher. It generally seems to do better for this problem than the last in comparison with the summation rarity and product irreplaceability algorithms, this could be due to the fact that the solutions are larger and there is more freedom in the solution space for the annealing algorithm to explore.



Looking at tables 7.2 and 7.4 we see that both the greedy and the adaptive annealing are similarly good at covering each conservation value once before conserving conservation values a second time. They are good in relation to the other algorithms but they still leave a number of conservation values unrepresented even with a high cost threshold. It is encouraging to see that this ratio is not terribly bad in general. Testing this approach on a 3 occurrence problem might make the trend clearer.

## Conclusions

When using a cost threshold that is significantly below the level required to meet our targets the greedy algorithm is the best of the heuristic algorithms and often is better than simulated annealing. When the cost threshold approaches the level required to meet all the targets of the conservation values the rarity based algorithms prove their worth and the benefits of careful selection of the initial sites become greater.

None of the algorithms are designed to encourage an evenness of representation of conservation values. Gaining the second representation of a conservation value is valued the same as getting the first representation of an as-yet unrepresented conservation value for all the algorithms. However, there is a tendency for both the greedy and the simulated annealing algorithm to perform well in terms of evenness of representation.

One way of getting very even representation for the occurrence style problem is to first design a reserve which represents each conservation value once, then to use this as the initial reserve for designing one which represents each conservation value twice and so on. This is the incremental reserve design method and it is the focus of Section 8.

## Section 8

# Incremental Reserve Design

### Introduction

The focus of this section is a particular manner of building up a reserve system using occurrence data where there is a target of many occurrences for each conservation value. I call this the incremental reserve design method. In this method we start with a target of a single occurrence of each conservation value, we then increase our target to 2 occurrences of each conservation value using the previous reserve as the basis for a new one, and we continue until we reach some final number of occurrences for each conservation value. This assumes that the final target number of occurrences is identical for each conservation value, but the approach could be modified for the more general case where each conservation value has a different target level of representation.

This method is of interest as an alternative or complementary method to that of using a cost threshold to obtain some level of representation for a set cost. With the incremental method we have a progression of reserves with increasing cost which satisfy increasing targets. Unlike the cost threshold method, when we pick a reserve near the desired cost the conservation values will tend to be more even in their representation of conservation values.

This is also an interesting experiment as it shows the rate at which the reserve has to increase in size to cope with an increase in the conservation value targets. Naively we might expect that the area would increase linearly with the targets. Alternatively there might be a drop-off in the extra area needed to meet higher targets because there are probably a number which are already represented above the current target.

Finally we can compare the incremental reserve design, where we base each new reserve on an older one with a lower target, with the case where the reserve for each new target is designed from scratch. Here we would expect that designing each new reserve from scratch would give a

better result. If this is the case, is there any trend in the difference between the two methods as the targets increase?

### **8.1 Incremental design on New South Wales data set.**

The first test is carried out on the New South Wales data set. This data set consists of presence/absence data and lends itself to targets of increasing numbers of occurrences. The incremental method is tested here against the method of designing each new reserve system from scratch. In the incremental method the target starts out as 1 occurrence for each conservation value. The reserve which results is used to start the design of a reserve with targets of 2 occurrences for each conservation value. This initial reserve is locked into the system and cannot be removed. This process of locking in each new reserve and increasing targets continues until the ultimate number of occurrences, in this case 5, is achieved.

Two algorithms are used in this test, the adaptive annealing algorithm and the greedy algorithm. In both cases they are tested using the incremental method and the design from scratch method for targets of 1 through 5 occurrences. In the New South Wales data there are a number of conservation values which appear fewer than five times and hence cannot meet the higher targets.

### **Results**

<b>Increasing Requirements</b>					
<b>Algorithm</b>	1 Occurrence	2 Occurrence	3 Occurrence	4 Occurrence	5 Occurrence
Anneal Inc.	54	112	168	221	273
Anneal Non Inc.	54	112	170	227	278
Greedy Inc.	62	118	177	232	282
Greedy Non Inc.	62	125	177	233	295
Miss	0	1	5	12	21
Amount	0	1	6	18	39

Table 8.1. Increasing number of occurrences of each conservation value required. The values presented are the number of sites in the system not including penalties for species which are not adequately represented in the system. All solutions meet targets where possible. The algorithms tested are adaptive annealing (Anneal) and the greedy heuristic (Greedy). 'Inc' indicates that the methods builds on the

previous reserve and 'total' means that it doesn't. 'Miss' is the number of species which cannot be met at that level (because there are fewer occurrences available than the target). 'Amount' is the inevitable shortfall for that level of occurrence due to conservation values which cannot have their requirements satisfied.




Figure 8.1. Graph of Table 8.1. This shows the results for simulated annealing. The triangle points are the total reserve design with the addition of the missing species penalty.

The results for this experiment are shown in Table 8.1. Note that at the higher target levels not every conservation value can be represented the required amount of times. For example, there is one conservation value which can only be represented once, 21 conservation values appear 4 or fewer times over all sites. The values presented in the main body of the table are of the number of sites in the reserve system.

The incremental algorithms build upon the reserve for the next largest target except for the 1 occurrence problem in which case they are identical to the 'total' algorithms. The 'total' algorithms are built from scratch. When an algorithm increments from a lower to a higher number of occurrences the initial reserve is fixed and non-removable.

There is an almost linear relationship in this problem between the required number of occurrences and the number of sites. The addition of the shortfall gives a curve with a slope which increases slightly with increasing target level (the triangle points in Figure 8.1).

## Discussion

Interestingly the size of the reserve increases almost linearly with the number of occurrences required. The increase in number of sites from one increment to the next varies between 51 and 58 new sites required. At each increment a number of conservation values are, presumably, already over-represented but the number of conservation values which are hard to represent increases and this larger number of conservation values which appear on few sites counterbalance any benefits gained from over-representing the more common species.

Surprisingly the incremental system gives better results than working out the total afresh. There are a few possible reasons why this might be the case. It might be an idiosyncrasy of this problem in particular. The best reserve increases linearly as the target number of occurrences increase, so it might be that working out the best value at each level is a simple process of adding a similar number of sites onto what is already there. There is, perhaps, no great savings to be had by redesigning the system from scratch.

This is very clear for the case of the greedy algorithm which sometimes does worse when designing from scratch than when incrementing the total. The difference between the two methods is that in the second method the greedy algorithm will often try to get duplicates of common conservation values before getting the first occurrences of less common conservation values. By incrementing the requirements some element of rarity selection comes into the selection process, which is beneficial in this case.

The annealing algorithm does more poorly as well, although the solutions are close. The incremental reserve design method has the benefit that much more annealing time is spent getting each solution. To get the five occurrence problem the number of annealing iterations has been five times as great as the design-from-scratch approach and the design-from-scratch approach has not discovered any greater efficiencies in this system.

This would be an interesting experiment to repeat on a different data set specifically to test whether the relationship between size and representation is linear and whether the sizes are similar for both the incremented and total design are so similar, particularly if there is a non-linear relationship between number of sites required and level of representation.

## Section 8.2: Incremental reserve design on the NT data set

This is a repetition of the experiment in Section 8.1 using the NT data set. Here only adaptive annealing is used with the targets increasing from 1 occurrence up to 5. The question that this hopes to address is whether the increase in required area is non-linear and what is the difference between designing a reserve system by scratch and designing one based on a reserve system which meets a lower target.

### Results

Northern Territory Incremental Reserve Design		
Target	Incremental Design	Non Incremental Design
1	32	32
2	59	49
3	85	72
4	111	91
5	141	111

Table 8.2. The number of sites in the smallest reserve meeting the designated target. In the incremental design the reserve system from the target 1 lower is 'locked in' before the algorithm commences.




Figure 8.2. A graph of the data in table 8.2

Table 8.2 and Figure 8.2 shows that the gap between designing a reserve from scratch for a given target and designing one based on a smaller reserve grows larger as the target increases.

There is no obviously non-linear increase in the number of sites required to meet ever-increasing targets.

## Discussion

This experiment shows a reverse result to the trial done on the New South Wales data set. Similar to the New South Wales test, the number of sites required to meet the targets increases in a nearly linear fashion with either the incremental or non-incremental design. The non-incremental design has a lower slope. It is better than the incremental design and the amount by which it is better increases with the target which is different from the New South Wales test (Figure 8.2). The inefficiencies of the incremental design method are compounded as the target is increased. This is different from the previous test in which the incremental method produces superior results. It could be that this does not happen with the Northern Territory data set because there are more conservation values and more occurrences in total (the occurrence matrix has a higher sparsity). This might increase the range of solutions of good quality for the algorithm to choose from and make the benefits from non-incremental design greater. If this is the case then the prediction would be, that if the problem were complex, perhaps using abundance rather than occurrence targets, then the non-incremental method is superior, if the problem is simple then incremental reserve design might suffice.

## Conclusions

The incremental reserve design system is another way in which a problem reversal can be tackled while using the same objective function and basic problem formulation. If a set amount of sites were available for reservation (for example, 250 sites in the NSW data set) then reserves could be designed using the increment system described here by increasing the target number of occurrences until we find the maximum target which can be met for all conservation values in the desired area (250 sites in our example). In this example, 4 occurrences can be reserved for 221 sites but 5 would require at least 273. So we use the 4 occurrence solution as our starting point and then run simulated annealing with a cost threshold of 250.

An alternative would be to use a fast algorithm such as the greedy algorithm to find roughly the level of representation that could be met. We then use a slower algorithm such as simulated annealing to produce a better result at the same representation and use this as our starting reserve system with the given cost threshold in place.

Here we have assumed that our targets are the same for every conservation value. If we were using abundance rather than occurrence data the differential targets could be handled elegantly. We would set relative abundance requirements and then increment the system by relative amounts for each conservation value, possibly with aggregation and separation restrictions.

It would be ideal to test this system on other data sets, given the differences in the results in this Section. Another critical element of reserve design is the spatial location of the sites within the reserve system. It would be interesting to see how spatial requirements would constrain the incremental method as the targets were increased. The next section introduces a spatial element with the inclusion of a measure of compactness (or inversely fragmentation) to the reserve design problem.



# Section 9

## Fragmentation

### Introduction

This section introduces spatial issues to the reserve design problem. In particular, I consider the boundary length of a reserve system. The smaller the boundary length the more compact the reserve system, for a given area. The boundary length and size of the reserve are the twin measures used here for reserve system quality. The Northern Territory data set is used for both tests in this section. All the problems were tackled with the adaptive annealing algorithm. In the first test I looked at the effect of changing the relative weight placed on area versus boundary length. The second looks at the effect of having a cost threshold limiting the area (but not the boundary length) of the system.

### **9.1 Varying the importance placed on boundary length**

The boundary length of the system is in different units of measurement than the area. Both could theoretically be converted to a cost measure perhaps using opportunity cost for the area and the actual cost associated with constructing and protecting boundaries for the boundary cost. In the method used in this thesis the boundary length is multiplied by a boundary length modifier, whose purpose is to correct for the differing units and also allow the control of the relative importance of boundary length and area. In this test the effect of increasing this boundary length modifier is examined to see the types of reserve systems which different choices make as well as to examine the trade-off between boundary length and area. The boundary length is combined with the cost of the system according to the objective function from Section 3 (Equation 3.1):

$$\sum_{Sites} Cost + BLM \sum_{Sites} Boundary + \sum_{ConValue} CVPF \times Penalty + Cost Threshold Penalty(t)$$

In this section there is no cost threshold. The conservation value penalty is calculated with the boundary length included so that the system will always be driven toward one which has no penalty and in which all the conservation values meet their targets.

BLM is the boundary length modifier. In this problem, it takes on the values: {0, 0.001, 0.01, 0.1, 1, 5, 10}. The adaptive annealing algorithm is repeated ten times at each value and from this the most compact and the smallest reserve system are recorded which could be different. The size in terms of the number of sites is taken from each reserve as is the boundary length. In the Northern Territory data set the sites are all equi-sized and hence given a size (or cost) of 1. The boundary between any two sites is also set at 1. Hence the boundary length of a reserve on its own will be a maximum of 4 or possibly less if it includes coastline.

Compactness is measured crudely by the total boundary length of the reserve system and more elegantly by the ratio of the boundary length divided by the area of the reserve. The ratio to ideal compactness is defined as the boundary length of the reserve system divided by the length of the perimeter of a circle which has the same area as the reserve system. This gives a ratio and hence a non-dimensionalised measure of the compactness of the system. A circle is chosen as the theoretically ideal compact shape, even though the Northern Territory data set consists of a square grid. Ignoring coastline effects, the theoretical most compact shape for this data set would be a square, as the perimeter of a square of fixed area is a fixed ratio (1.128) to the perimeter of a circle of the same area this means that the values but not the rankings are different by taking a circle rather than a square as the ideal shape. As already suggested, it is theoretically possible to design a reserve with a smaller boundary length than the circle of same area if the reserve includes a fair amount of coastline (which does not add to the boundary length).

The ratio to ideal compactness should always be higher than one, with the lower the ratio, the better the reserve. An alternative, which is not used in this thesis, is to take the inverse of this measure. This would give a value that was bounded in the interval [1,0) if we ignore the potential effect of the free boundary on the coastline.

## Results

BLM	Fewest sites		Best BL/Ideal	
	Sites	BL/Ideal	Sites	BL/Ideal
0	117	11.788	121	11.745
0.001	115	11.732	119	11.559
0.01	116	11.629	121	11.617
0.1	116	11.053	118	11.037
1	114	9.353	127	8.035
5	133	6.947	165	5.446
10	157	5.899	189	4.781
50	440	3.591	708	2.014
100	458	3.124	1259	1.208
500	385	3.35	1073	1.361
1000	473	3.217	1355	1.081

Table 9.1. Two reserves are reported for each set of 10 simulated annealing runs. The one with the fewest sites and the one with the best ratio of boundary length to ideal boundary length. BLM is the boundary length modifier. BL/Ideal is the ratio of the boundary length of the reserve system to the 'ideal' boundary length which is the perimeter of a circle with the same area as the reserve system.

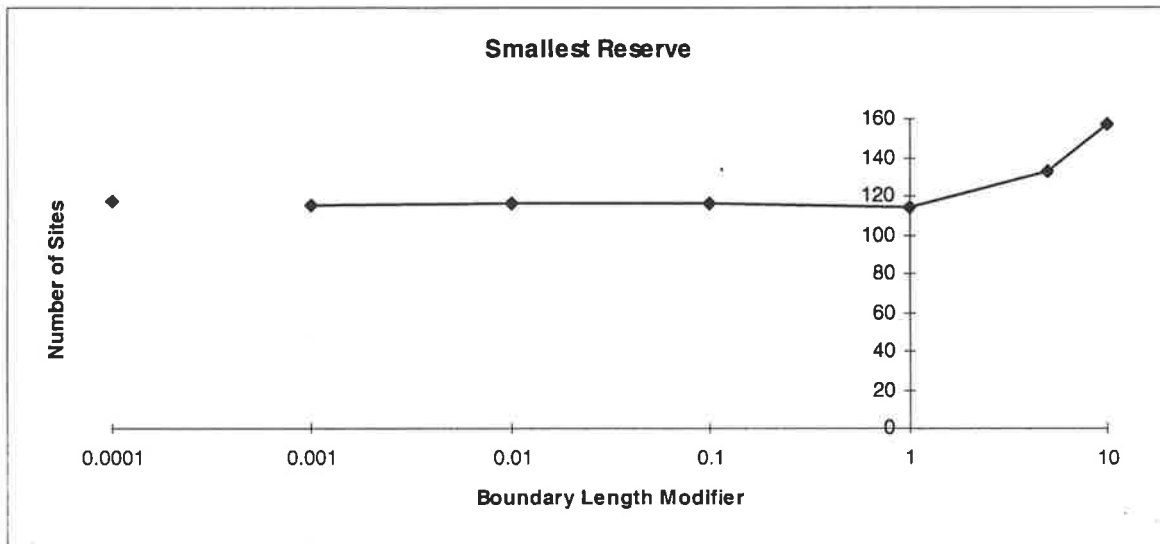


Figure 9.1a. Smallest Reserve looking at the number of sites for the reserve. Note that the x scale is logarithmic and that the isolated point on the left is the case where the boundary length modifier is




Figure 9.1b. Smallest Reserve looking at the ratio to the ideal compactness for the reserve. Note that the x scale is logarithmic and that the isolated point on the left is the case where boundary length modifier is 0. Also note that with this measure the ideal ratio is 1.

Figure 9.1 gives the values for the smallest reserve out of then ten runs for each value of the boundary length multiplier. The size of the reserve increases when the boundary length modifier gets larger than 1. The boundary length begins to decrease when the boundary length modifier increases beyond 0.1.




Figure 9.2a. Most Compact Reserve looking at the number of sites for the reserve. Note that the x scale is logarithmic and that the isolated point on the left is the case where the boundary length modifier is 0.




Figure 9.2b. Most Compact Reserve looking at the ratio to the ideal compactness for the reserve. Note that the x scale is logarithmic and that the isolated point on the left is the case where the boundary length modifier is 0. Also note that with this measure the ideal ratio is 1.

Figure 9.2 shows the same set of runs, although this time taking the most compact of the ten runs for each instance of the boundary length modifier. A similar albeit more exaggerated pattern to Figure 9.1 is apparent here.




Figure 9.3. Map showing sites included in the reserve when the boundary length modifier is 0. The 'X's are sites selected by the algorithm and the '.'s are the other sites in the data set. This reserve system consists of 117 sites with a boundary length of 452




Figure 9.4. Map showing sites included in the reserve when the boundary length modifier is 10. The 'X's are sites selected by the algorithm and the '.'s are the other sites in the data set. This is the most compact reserve for this boundary length modifier. This reserve system consists of 157 sites with a boundary length of 262




Figure 9.5. Map showing sites included in the reserve when the boundary length modifier is 1,000. The 'X's are sites selected by the algorithm and the '.'s are the other sites in the data set. This reserve system consists of 473 sites with a boundary length of 248

## Discussion

As is expected when the boundary length modifier increases the results will tend to be more compact but over a larger area. Without a boundary length check the reserve will tend to be quite fragmented (see Figure 9.3). This remains the case for even small boundary lengths but as it rises we see two interesting things. One is that the boundary length consistently decreases even as the sites are increasing a lot (Figure 9.5 has the smallest boundary length although the area is obviously huge). The other is the crossover point where the boundary length modifier is 1. At this point the reserve has not increased significantly in size but has become significantly more compact. Beyond that point both the size increases and the compactness decreases. This point is a desirable one, but it is only established by doing a full set of simulations.

The difference between the solutions can be clearly seen in Figures 9.3 through 9.5. Figure 9.3 looks completely fragmented, the boundary length is the greatest although the size is comparable with other, more compact, reserves. The biggest clump has only three contiguous sites in it and as they appear in a straight line it looks strung out. (Although in a square matrix any formation of three contiguous sites has the same boundary length). In Figure 9.4 the number of sites is considerably greater at 189. 60 extra sites are included but the boundary length of 233 is close to half of that of the smallest reserve (with boundary length of 452). The final figure shows the danger of valuing the shape base solely on the boundary length. Here a huge reserve is designed which retains all conservation values and fills up the interior with extra sites in order to avoid having unnecessary boundaries. This reserve has a boundary length of 141 (the smallest in this test) and 1,355 sites (the greatest in this test).

This last figure demonstrates the concept that even if the boundary length is the most important feature for a reserve system, it is often sensible to include even a token site cost in the objective function. The reserve system might not be quite so big if the grid were not square. For example, in a square grid of sites a large cross will have the same boundary length as a rectangle which has the same height and width as the cross, as seen in Figure 9.6.




Figure 9.6. In both reserves the boundary length is 30 units long. The reserve on the left has an area of 15 sites the one on the right has an area of 56 sites.

An alternative method is to pick a fixed size for the reserve and apply the cost threshold system to keep the reserve system to that size and search for the most compact solution. This is the subject of Section 9.2.



## 9.2. Cost threshold

### Aim

Following on from the above test, two runs are done with a cost threshold. Here the size of the reserve system is controlled by the cost threshold and a high degree of compactness is encouraged by applying a high boundary length modifier (BLM = 10).

Two cost thresholds were tested. The first is a cost threshold of 117 sites which is the number of sites that was selected when there was no boundary length multiplier. The second is a cost threshold of 120 sites which is very close but gives three more sites to play around with.

### Results

Cost Threshold			
PU	BL	BL/Sites	BL/Ideal
117*	452	3.863	11.788
117	440	3.761	11.475
120	450	3.75	11.588

Table 9.2. Results from fixed cost solutions. The first row is the result generated in the previous section with no boundary length modifier and without the cost threshold. The second row has a boundary length modifier of 10 and a cost threshold of 117. The third row has a boundary length modifier of 10 and a cost threshold of 120.

### Discussion

The results here are promising. With 117 sites a reserve is designed which is more compact than that designed with no boundary length modifier. The boundary length is 12 units shorter and the ratio of actual boundary length to idea boundary length is better.

When we increase the allowable size of the reserve to 120 sites the boundary length to site ratio decreases but interestingly the boundary length to ideal ratio increases. In this case the benefit of having three extra sites to try to decrease the boundary length are not clear. It is probably not enough, in this case, to gain any real advantage in clumping. The fact that one ratio improves as the other one becomes worse shows that the choice of an index of compactness is important. Mathematically this occurs because the first ratio is dependant upon the area of the reserve and the second upon the square root of the area, which is more natural when combining it with a length measure.

# Section 10

## Spatial Rules

### Aim

It is more desirable to have compact rather than fragmented reserves but it is also desirable to have a reserve in which the sites containing individual conservation values are separated by a distance which will protect them from localised disasters and encourage a greater amount of genetic diversity capture. The previous section looked at decreasing, or controlling, the level of fragmentation of the reserve system as a whole by trying to minimise boundary length as well as reserve area. In this section I look at applying spatial requirements at the level of conservation values. I use the spatial aggregation and separation rules of Possingham and Andelman (in press) where we want to ensure that individual conservation values are in local clumps which are large enough to ensure their survival as well as ensuring that these clumps will occur in at least three adequately separated locations.

The question which is asked here is how do the fragmentation and conservation value spatial requirements effect each other. When all the requirements are included in the targets, how does this effect the compactness and the size of the solutions? When only size or size and compactness are used how does this effect the conservation values? How many of them are good according to measures of aggregation and separation? Is this the same for the case where the boundary length modifier is zero, which gives us a fragmented reserve, as when it is 10 giving us a tightly clumped reserve system?

### Methods

In this section the requirements for conservation values now include spatial components and hence the conservation value penalty is calculated slightly differently. There are two distinct spatial components which are included, the aggregation rule and the separation rule. Changing

the conservation value penalty does not effect the other terms of Equation 3.1 and the aim is still to minimise either area or some combination of area and boundary length.

With the aggregation rule there is a minimum clump size (which is 10% of the target total). A clump for a conservation value is a collection of contiguous sites on which the conservation value occurs. A clump might be a single site but it could also be larger. If contiguous sites contain the conservation value then it is considered that the conservation value occurs in contiguity on these sites. This is an important simplification which decreases the problem size and also makes it possible to use the Northern Territory data set in which the spatial position of vegetation types within a site is not known. If a conservation value occurs on a clump in the reserve system which is smaller than the minimum clump size then that clump cannot count toward the conservation values total, if the clump is larger than this minimum clump size then it will count toward this total.

The minimum clump size has been set to 10% of the total representation target for each conservation value. This is a somewhat arbitrary value, but it will ensure that the most fragmented conservation value meets its target needing at most 10 separate clumps. In a real problem the minimum clump size would be set depending upon the individual requirements of the conservation values, although a simple percentage of the total target level would probably work well as an initial estimate.

The separation rule states that the conservation value under consideration must occur on at least three sites which are mutually separated by the stated distance (In our square grid this is a distance of 2.5 grid cells). When used in conjunction with the aggregation rule (as it is here) the three sites must each be part of a valid clump. It is possible that some or all of the separated sites are part of the same clump, if the clump were large enough. This is not problematic and would probably occur in the entirely desirable case where the entire region was reserved. Although the separate sites of the conservation value can occur in the one clump, this clump would have to be large enough to withstand localised disasters.

The separation distance was completely arbitrary, it was chosen because most conservation values were spread out enough that they could theoretically meet these requirements, but the distance was still high enough to ensure a reasonable level of spread. In this data set with the

given separation distance there are 11 conservation values which cannot meet their separation requirements.

The resulting reserve systems were scored for their cost (the number of sites in the reserve system), their boundary length, the number of conservation values which did not meet their targets when the aggregation rule was computed and the number which did not meet their targets when the separation rule was tested. The treatments differed by the boundary length modifier used and whether the aggregation and separation rule were also applied. The treatments tested were:

- Boundary length modifier 0, no conservation value aggregation or separation rule.
- Boundary length modifier 1, no conservation value aggregation or separation rule.
- Boundary length modifier 10, no conservation value aggregation or separation rule.
- Boundary length modifier 0, no conservation value separation rule. Conservation value aggregation rule applies.
- Boundary length modifier 0, conservation value separation and aggregation rules apply.

These treatments were applied to the Northern Territory data set using the adaptive annealing algorithm.

One might expect that as the boundary length modifier was increased the reserve system would become more compact and this would increase the chances of conservation values satisfying aggregation requirements and decrease the chance that the separation requirements were met. When applying the aggregation rule on its own it would be expected that the boundary length of the reserve be less than the case when no aggregation rule is applied (with a zero boundary length modifier). It would be expected that the application of the aggregation rule on its own would tend to work against the separation rule. The treatments used in this section are designed to test these predictions.

## Results

<b>Effects of Different Spatial Rules</b>				
<b>Treatment</b>	<b>Cost</b>	<b>BL</b>	<b>Agg Missing</b>	<b>Sep Missing</b>
BLM 0	112	424	15	73
BLM 1	117	293	13	80
BLM 10	204	227	5	70
Agg	112	446	0	77
Agg + Sep	122	446	0	11

Table 10.1: The effect of using different spatial rules on the cost, boundary length and the adequacy of representation of conservation values in a reserve system. 'Agg Missing' is the number of conservation values who do not meet the aggregation requirement, 'Sep Missing' is the number which do not meet the separation requirement. The 'Agg' treatment includes penalising the system for conservation values not meeting the aggregation requirement and the 'Agg + Sep' treatment also penalises for those not meeting the separation requirement. Both of those treatments have a boundary length modifier (BLM) of 0.




Figure 10.1. A map showing the reserve system generated without aggregation and separation rules and with a boundary length modifier of 0. The 'X's are reserved sites and the '.'s are the other sites in the data set. This reserve system consists of 112 sites with a boundary length of 424.




Figure 10.2 A map showing the reserve system generated without aggregation and separation rules and with a boundary length modifier of 1. The 'X's are reserved sites and the '.'s are the other sites in the data set. This reserve system consists of 117 sites with a boundary length of 293.




Figure 10.3 A map showing the reserve system generated without aggregation and separation rules and with a boundary length modifier of 10. The 'X's are reserved sites and the '.'s are the other sites in the data set. This reserve system consists of 204 sites with a boundary length of 227.




Figure 10.4 A map showing the reserve system generated with the aggregation rule but not the separation rule and with a boundary length modifier of 0. The 'X's are reserved sites and the '.'s are the other sites in the data set. This reserve system consists of 112 sites with a boundary length of 446.




Figure 10.5 A map showing the reserve system generated with both the aggregation rule and separation rules and with a boundary length modifier of 0. The 'X's are reserved sites and the '.'s are the other sites in the data set. This reserve system consists of 122 sites with a boundary length of 446.

Table 10.1 summarises the results for this test. The best reserve systems (lowest objective function value) are displayed in figures 10.1 - 10.5. In all cases, every conservation value met its primary reservation target. When the aggregation rule was in place then every conservation value also met its aggregation target. When the separation rule was in place every conservation value which could meet its separation requirement, did so.

## Discussion

As the boundary length modifier is increased we see the same pattern that occurred in Section 9. The reserve cost increases as the boundary length decreases. As this boundary length increases the number of conservation value which do not satisfy their separation requirements does not appear to change. The prediction was that this number would increase but the number which do not meet their separation requirements is already very high and is over 60% of the total number of conservation values. When the separation rule is not in place the algorithm produces very poor dispersal of conservation values.

As expected, the number of conservation values which do not meet their aggregation requirements decreases as the boundary length of the reserve decreases. Even though the

aggregation requirements were not enforced through a penalty, they tend to be met as the reserve system becomes more and more compact. This is probably because there is a tendency to collect conservation values from a single population when trying to minimise the boundary length of the reserve system. The fact that the area of the reserve systems also increases as the boundary length increases also improves the chances of conservation values meeting their aggregation requirements without the rule being explicitly used.

When the aggregation rule is in place all conservation values meet their aggregation requirement, as expected. Interestingly the reserve system is as fragmented as when there is no boundary length modifier (Figure 10.4). Aggregating the conservation values does not appear to aggregate the sites. When the aggregation rule is applied on its own, a similar number of conservation values fail the separation test as when the aggregation rule is not used at all.

The separation rule encourages aggregations of conservation values which are separated by enough distance so that a localised disaster will not take out the entire representation for that conservation value. When both the aggregation and the separation rules are in place all conservation values which can meet their requirements do meet their requirements. 11 of the conservation values cannot meet their separation requirements anyway. The size and boundary length of the reserve system is of the same order as when the aggregation rule alone is in place. This is against the prediction that the separation rule would tend to make the reserve system more fragmented. However, the reserve system for the aggregation rule on its own is already as fragmented as the one with no aggregation rule and no boundary length modifier (compare Figure 10.5 to Figure 10.1).

## Conclusion

From this test we can conclude that using the boundary length modifier to produce a compact reserve system will also tend to aggregate those sites containing individual conservation values. Oddly enough the reverse is not true. When we aggregate sites containing individual conservation values we do not get a more compact reserve system.

The separation requirement is met poorly independently of what the boundary length modifier is or whether the aggregation rule is in effect whenever the separation rule is not specifically applied. The conclusion is that one might be able to avoid directly applying the aggregation



rule without too poor a result with a high boundary length modifier, but that if separation is desired then there is no substitute to explicitly applying the separation rule.

Possible further work which could follow from this section is to repeat the treatments on a different data set and by testing the sensitivity of the results to the separation distances of the conservation values.

# Section 11

## Conclusions

### General

In this half of the thesis we have looked at a number of new problems involving the design of reserve systems and looked at the methods which work well for them. The comparison of the different algorithms was extensive involving a number of trials of two formulations of a non-spatial objective on two separate data sets. The formulations included both occurrence and proportional representation formulations. The work looked at a whole suite of new problems including ways of trying to get the best representation for a set cost and of adding spatial elements to the design problem.

The use of a cost threshold to control reserve design is new to this thesis and works quite well at providing reserves which are as representative as possible for a set cost. Incremental reserve design has not previously appeared in the published literature and it is proven to be a good means of examining features such as the relationship between reserve size and conservation value richness. Fragmentation and other spatial issues have always been important issues for the design of nature reserves, although previously they have not been examined in depth. Previous work on such issues has focussed on extensions to the iterative heuristics to encourage a level of clumping (see, for example, Willis *et. al.*, 1996, and Lombard *et. al.*, 1997). The use of the boundary length modifier is new and is a powerful addition to designing reserves with some spatial complexity. The application of the species separation and aggregation rule to the reserve design problem is new to this thesis. The rules go well beyond the simple species targets which are normally used. This has shown that quite complex information on the requirements of conservation values can be incorporated into these methods of reserve design.

In this section we will look at the implications of these additions to the reserve design problem. This will be followed by an examination of how these reserve selection methods could be used

within an actual reserve design process. There are a large number of directions in which this research could proceed. A few of these are examined at the end of this section.

## Discussion of Results

In Section 6 we saw the comparison of a number of algorithms for the design of reserve systems. The two main categories were iterative heuristics and simulated annealing. Of the heuristics there were three main categories of heuristic - richness based, rarity based and irreplaceability based heuristics. The irreplaceability heuristics were substantially different from the other types, although rarity and richness qualities were embedded in the irreplaceability measure.

Heuristics which include a mixture of both rarity and richness measures were generally the most successful. Possibly the reason for this lies in the distribution matrix for the data set. The distribution matrix is a large matrix where the columns are conservation values and the rows are sites. The values in the matrix are the amount of the conservation values on the sites, with an occurrence problem these values are zeros and ones. This matrix holds all the information for the non-spatial reserve design problems. The sum across each row is the richness for a particular site and the sum down each column is the rarity (or frequency) for that conservation value. It should not be surprising that both these properties are essential for designing a reserve system because the two measures span the matrix.

The irreplaceability heuristics and summation rarity heuristics did well under a variety of different problems. The greedy heuristic was also reasonably good and it was the best under some circumstances. It would seem that the best way of using these heuristics, which are very quick to apply, is to throw a battery of heuristics at a problem and look at the range of results rather than trying to choose the heuristic which should be used in all situations. It is also clear, however, that some of the heuristics, such as the maximum rarity heuristic, should not be considered for serious work.

One algorithm which performs particularly well under certain circumstances is the greedy algorithm. This algorithm produces reasonable results for many problems, although it is

generally not as good as some of the more sophisticated algorithms. However, it is clearly the best algorithm to be used under heavily constrained circumstances. When the number of sites is constrained to be far below the amount required to represent each conservation value adequately the greedy heuristic is clearly the best. We can see why this might be the case when we consider that the rationale behind the greedy heuristic is to find the best single improvement at each step possible. We have seen why this fast improvement method will not generally produce the best comprehensive design in Example 3.1. For a tightly constrained problem, however, it is clearly optimal.

The simulated annealing algorithm was the strongest of the algorithms tested in this thesis. It produced reasonably good results for every problem tested. It is the best method for most problems and, importantly, it is robust to the problem type. The only case where it did not produce the best results was in the highly constrained case of the problem in Section 7, where a greedy algorithm was the best. The main conclusion of the tests in Section 6 was that the simulated annealing algorithm was the best to use in general but that the iterative heuristics are a useful alternative, primarily for their speed.

## Cost Threshold

There were two new methods introduced in this thesis aimed at solving the inverse of the normal reserve design problem. This is the problem where the maximum area or cost of the reserve system is fixed and the aim is to maximise the level of representation of the conservation values for this area or cost. The incremental reserve design method could be used to tackle this type of problem but it was the cost threshold method which was specifically designed with this problem reversal in mind. The cost threshold keeps the reserve system to the given area or cost and appropriately set conservation value targets and penalties will drive these towards maximum representation.

The main problem with using the cost threshold method is that it does not include any weighting for the different conservation values. It is usually desirable that conservation values are relatively evenly balanced in their representation but the cost threshold method does not directly do that. For example, if the target for each conservation value were 20 occurrences,

then raising one conservation value from 0 to 20 occurrences would have the same benefit to the objective function as getting the first occurrence of 20 different conservation values.

Section 7 specifically addressed this concern and found that the method did leave an imbalance of representation of conservation values, but that this imbalance was not great. There seemed to be a natural balancing of representations in the test problem. It would be interesting to implement a non-linear measure of representation to the objective function for this problem. This could be done without substantially altering the current approach. A simple method would be to add an additional penalty if a conservation value were under half of its target and an extra penalty if it was completely unrepresented. This should drive the reserve system toward one in which the level of representation was more balanced between the different conservation values.

When the incremental design method was used, some non-intuitive results ensued. Particularly interesting was the fact that as the targets increased, the size of the reserve system increased roughly linearly (Figures 8.1 and 8.2). When a reserve system has met a given target level for all conservation values, there are usually some which are over-represented. This means that some conservation value are already at a higher level of representation and one would expect that for this reason one would require fewer new sites to raise all conservation values to that higher level of representation. However, it appears that this over-representation of some conservation values is balanced by the difficulty in finding the extra representation for the remaining conservation values. There are fewer instances of these remaining conservation values available for conservation as the target value is incremented.

An interesting fact which arose from this test was that the incremental design method can sometimes be superior to designing a reserve system from scratch for a given target size using simulated annealing. The reason that this is thought to be the case is that under highly constrained circumstances the benefits of design from scratch are not great and when design from scratch is done using simulated annealing a lesser number of iterations is preformed for higher targets than has been cumulatively used by the incremental system. This was the case for the New South Wales data set. The Northern Territory data set revealed the expected situation were the design from scratch method produced superior results.

The incremental design method can be used as a method for reversing the reserve design problem. If the cost threshold is only an approximate value then the design is done by incrementing the conservation value targets and applying the solution method until the cost of the reserve system surpasses the cost threshold. Then one of the last two incremented reserve systems is the solution. A more refined method would be to use the incremental method to calculate roughly the maximum even-target level for all conservation values which can be achieved for less than the cost threshold. Then the reserve could be redesigned from scratch using the target level just found, hopefully finding a smaller or cheaper reserve system (If the resulting reserve is larger than the incremented reserve system then the incremented reserve system is used). This reserve system is then used as the basis for the final reserve system in which the target level is incremented and the cost threshold is applied using the cost threshold penalty. The conservation value representation in this reserve system will be as even as it is possible to achieve.

## Fragmentation and Space

Spatial considerations are essential for the design of realistic reserve systems. Without them reserve systems could be very fragmented (as seen in Figure 9.3). Section 9 looked at a new and excellent way of reducing the fragmentation of a reserve system. This was to include the boundary length as a measure of the quality of the reserve. The problem then becomes a multiple objective problem but it is reduced to a single objective function by using a boundary length modifier to equate the boundary length and the size measure of the reserve system. This gives a powerful means for controlling the compactness of the reserve system and controlling the trade-off between compactness and size for the reserve system.

In the test case there was a 'golden point' in the compactness versus size trade-off. It occurred when the boundary length modifier was around 1. At this point the boundary length was significantly smaller without a significant area increase than when the boundary length modifier was zero. When the boundary length modifier increased, the boundary length decreased but the area of the reserve system started to get quite large. Whether there is generally such a 'golden point' for other data sets remains to be seen.

The aggregation and separation rules are also new to the reserve design problem. These are a spatially explicit representation requirements for conservation values. We saw in Section 10 that to some extent it is possible to do reasonably well at meeting aggregation requirements without explicitly implementing the aggregation rule. As the boundary length modifier increases, the compactness (as well as the area) tends to get greater. When this happens the representation for an individual conservation value will tend to be greater and tend to be gathered from fewer separate populations. This improves the conservation value's standing with regard to the aggregation rule, although the benefit became apparent in Section 10 only with a large boundary length modifier of 10. What this means is that we might expect that a compact reserve system is good in terms of conservation value aggregation even when conservation value aggregation is not explicitly included in the objective function.

Interestingly the reverse does not appear to be true. When the aggregation rule is in operation but there is no boundary length modifier, we do not get an improvement in compactness. Aggregating individual conservation values does not lead to an aggregation of sites in general within the reserve system.

The separation rule would seem to push solutions in the opposite direction to either a clumping mechanism such as the boundary length modifier, or the aggregation rule. No clear trend was evident, possibly due to the fact that so few conservation values met their separation rule when it was not explicitly implemented. If the separation rule is required then it must be set explicitly.

## Future Work

There are a large number of directions in which this work could proceed. Some ideas were described in Section 3. These included complex conservation value requirements and the use of multiple classification of reserves rather than the two classes of reserved and unreserved. There are a great many other avenues that could be pursued.

## Surrogates for Conservation Values

One of the main aims in reserve design is to capture as much biodiversity as possible. Often this means to capture as many different species and sub-species as possible. Because it is not generally easy to accurately establish the number of species on individual sites and their complete distribution, it is often convenient to try to use a surrogate measure for the species. In this thesis both landform type and vegetation type have been used. Other surrogates include environmental types (Belbin, 1993), indicator groups (Faith and Walker, 1996), umbrella species and keystone species (Simberloff, 1998).

In Australia a reserve requirement has been explicitly laid down based on vegetation types (JANIS, 1997). This requirement is that 15% of the distributions of each vegetation type which existed before the arrival of Europeans be retained in reserve systems. With an appropriate data set we can ask a few questions. To what extent does preserving vegetation types also preserve threatened species? In general how well do surrogates for species preservation actually preserve the species.

This question could be examined with a data set which included both the species' distributions and the surrogates' distributions. Linking the target representation of the surrogates with the requirements for the individual species is a useful problem to examine. Also the spatial rules can profitably be examined. Of particular interest, following Section 10, is the question of whether applying the separation rule to the surrogate for the species will also help the separation requirements of the species.

## Alternative Generator

The design of reserve systems goes beyond the automatic selection of sites by a computer algorithm. It encompasses many more subtle decisions and often there are many external constraints as to which reserve systems are acceptable or not. In this light the reserve system algorithms are aids to the designer of reserves - not replacing, but speeding up the designers ability to select which sites to include and exclude. In keeping with this way of looking at the problem, the generation of alternative solutions could be very useful.



Simulated annealing, and a number of the other algorithms examined in this thesis, already do this to some extent. Simulated annealing will generate a large number of different reserve systems. However, it is not known how similar these reserve systems are to each other. They could be nearly identical or completely different. It could be useful to try to change the algorithm so that it would produce a series of solutions, each one as different as possible from all previous solutions. Furthermore, it would be desirable to provide the algorithm with one or more solutions and then use it to generate solutions as different from these as possible.

It would be desirable to provide the algorithm with one or more solutions and then use it to generate solutions which were both as small as possible and also as different from all the provided solutions as possible. Alternatively the algorithm might produce a series of solutions each one as different as possible from all previous solutions.

As well as using alternative solutions to give an insight into the design of the reserve they can be used to get a measure of irreplaceability for sites within the reserve system. The simplest way that this could be done would be to generate a great many reserve systems and then determine in what proportion of the solutions each site was included in the reserve system. This proportion would give some indication of a site's irreplaceability. If a site appears in every solution then that is a strong indication that it is a necessary site. This would work well using simulated annealing in which the solutions are not as biased toward one type of solution or another unlike the iterative heuristics.

A simpler way of looking at irreplaceability is to generate a baseline reserve system and then see how much larger the reserve system would have to be if each of the sites in the system in turn were removed from the data set and others substituted into the solution. After the site were removed from the baseline solution the algorithm would be applied again to create a new reserve system, possibly just adding new sites in a greedy fashion or possibly designing one from scratch. The resulting reserve system should be the same size or larger than our existing reserve. The increase in area or cost is a good measure of the irreplaceability of a particular site. It would also be simple to quickly test to see if the site could be replaced at all, or whether it contains a necessary representation for one or more conservation values.

## Alternative cost measures

The primary cost measure used in this thesis was the number of sites in the reserve system. Area is another obvious cost measure which has been widely used. It should be possible to include explicit measures of the cost of the reserve system, either in terms of the actual cost of purchasing and maintaining sites or possibly in terms of the opportunity cost of the site for different land uses.

The opportunity cost for the reserve system as a whole will often be more than just the sum of the opportunity costs of the sites within the reserve system. The opportunity cost for any specific site could depend upon which other sites are in the reserve system. Could this non-linear measure be approximated by a simple opportunity cost associated with each site? Or could it be handled by fixing the opportunity cost for each site and then updating the opportunity costs routinely in the algorithm after the reserve system has changed by a sufficient amount? These ideas would be straight forward to test.

## Other ideas

One of the main difficulties of reserve design is the lack of knowledge about conservation value distributions. Not only is the knowledge scant for most species but a lot of it is of poor quality, with detailed knowledge generally available only for a handful of species. It would be useful to study the question of whether it is better for survey design to focus on one or two threatened species, and try to ascertain their distribution as accurately as possible, or whether it is better to cover as broad a range of species as possible. This is related to the issue of using surrogate species but a bit more subtle as all conservation value information would still be used. It is just the detail of knowledge which might shift from one group of conservation values to another.

A different formulation of the problem is to look at a similarity matrix of the sites within the data set with each other. From this large similarity matrix we'd try to select sites which are as dissimilar to each other as possible. Here we are trying to span the similarity matrix and this is a new and different way of describing the reserve problem, although similar in some regards to methods such as environmental representation (Belbin 1993) which tries to take as dissimilar environmental types as possible.

The design of marine reserves involves some different issues. Not only does it have more of the data collection problems than terrestrial reserve design has, but the way that sites are connected and defined is not as clear. Many of the species of interest can move very large distances and this makes concepts of site aggregation problematic. Instead of neighbouring sites whose boundary length is to be minimised we could use a similar mechanism to try to increase the connectivity of sites. Connectivity in this case is an allied measure to boundary length but is based, not on the simple geography of which site borders on which site, but includes an element of water flow. The connectivity matrix can include a large number of sites which are connected by different rates of fluid flow. Also the matrix need not be symmetrical. It is possible that the flow only ever occurs in one direction.

### The Software Spexan

Much of the software used to generate the results in this thesis were originally encoded into FORTRAN77 programs called SIMAN and ALGO. These were re-encoded into a C program called Spexan, which is an acronym of 'spatially explicit annealing'. This software included all of the algorithms described in Section 4 and was used on all of the reserve design problems in this thesis.

The software works by reading in the data set and a file containing all of the control information. The main input file is constructed using a separate program which has a graphic user interface, although it can be edited with any text editor. This allows it to be used under a number of operating systems. The program is designed to be used both by land managers and by those interested research questions.

### Using the Software in a reserve design process.

The methods in this thesis are automatic reserve design methods but it is important to see them in the context of actual reserve design. The role of the reserve designer is enhanced rather than replaced by the use of such a tool. The most important aspect, perhaps, is in setting up the problem. It must be decided which conservation values should be included, what type of information to incorporate and what the target levels of representation for the conservation

value should be. Should the region be used as a whole or split into sub-regional problems? What other information needs to be incorporated. The cost value for each of the sites needs to be determined and that could be used to include a wide variety of information, such as site by site adjustments for social factors which are hard to quantify. The relative importance of each conservation value has to be set. In this thesis they have been assumed to have equal importance, even though their targets might be different. In a real reserve design exercise it could make sense to lower the importance of conservation values which could potentially be preserved in other regions.

Deciding what should constitute a site is also important. In the New South Wales data set these were based around the concept of a pre-existing parcel of land, the Northern Territory data set had a regular square grid placed on it to define sites. Another alternative is to use something akin to hydrological units. Even when dealing with terrestrial species the hydrological relationship is a simple measure of topography in as much as neighbouring valleys will be placed into separate units.

Once all of the data is set up an iterative process begins where the methods used in this thesis are applied to the problem and the problem definition is refined in light of the resulting reserve system. The methods used in this thesis allow the exploration of a number of related issues, such as which sites are generally most important for preservation, the spatial configuration of the reserve system and the relationship between the reserve system in the region with other neighbouring regions which are not included in the data set.

Reserve design algorithms were originally based around heuristic algorithms which tried to capture the process of design. With the application of simulated annealing and, more importantly, the formal definition of the objective function the methods have moved away from that paradigm. In the new paradigm the definition of what would make a good reserve system is set into the objective function and powerful mathematical tools, such as simulated annealing are then used to find an optimal solution. This is a move away from the iterative heuristic approach which is based on trying to define the process of reserve design, and a move toward algorithms based on defining the objective or end point of the reserve design process. In doing this, the role of the reserve designer is still very important and revolves around the problem

definition stage. It is hoped that this thesis has aided in shifting this perspective and changing both the nature of reserve design and the design of nature reserves.

## Bibliography

## HOLSIM References

- Acevedo, M. F., Urban, D. L. and Shugart, H. H., (1996) "Models of forest dynamics based on roles of tree species.", *Ecological Modelling*, **87**, 267-284
- Ambrose, G. J., (1982) *An ecological and behavioural study of vertebrates using hollows in eucalypt branches*. Ph. D thesis, La Trobe University, Melbourne.
- Ashton, D. H., (1975a) "The root and shoot development of *Eucalyptus regnans* F. Muell.", *Australian Journal of Botany*. **23**, 867-887.
- Ashton, D. H., (1975b) "The flowering behaviour of *Eucalyptus regnans* F. Muell.", *Australian Journal of Botany*", **23**, 399-411.
- Ashton, D. H., (1975c) "Studies of the litter in *Eucalyptus regnans* F. Muell. forests.", *Australian Journal of Botany*, **23**, 413-433.
- Ashton, D. H., (1976) "The development of even aged stands of *Eucalyptus regnans* F. Muell. in Central Victoria.", *Australian Journal of Botany*, **24**, 397-414.
- Ashton, D. H., (1981) "Tall Open Forests." in R. H. Groves Ed. *Australian vegetation*. Cambridge University press, Cambridge. pp. 121-151.
- Attiwill, P. M., (1994) "Ecological disturbance and the conservative management of eucalypt forests in Australia.", *Forest Ecology and Management*, **63**, 301-346
- Banks, J. C., (1993) *Tree-ring analysis of two Mountain Ash trees *Eucalyptus regnans* F. Muell from the Watts and O'Shannassy Catchments, Central Highlands, Victoria. A Report to the Central Highlands Old Growth Forest Project. August 1993*. Department of Conservation and Natural Resources, Melbourne, Australia.
- Botkin, D. B., Janak, J. F. and Wallis, J. R., (1972) "Some ecological consequences of a computer model of forest growth.", *Journal of Ecology*, **60**, 849-873
- Bossel, H., (1996) "TREEDYN3 forest simulation model.", *Ecological Modelling*, **90**, 187-227.
- Burgman, M.A., Incoll, W., Ades, P., Ferguson, I., Fletcher, T. and Wohlers, A., (1995) "Mortality models for Mountain Ash and Alpine Ash.", *Forest Ecology and Management*, **67**, 319-327.
- The Commonwealth of Australia, (1992) *National Forest Policy Statement*. Advance Press: Perth.
- Cunningham, T.M., (1960) "The natural regeneration of *Eucalyptus regnans*.", *Bulletin School Forestry*, University of Melbourne **1**: 1-158.
- Dahl, B., (1940) "Normal stocking of regrowth forests of Mountain Ash (*Eucalyptus regnans*).", *Victorian Forester*, **2**, 20-23.
- Dale, V. H., Doyle, T. W. and Shugart, H. H., (1985) "A comparison of tree growth models.", *Ecological Modelling*, **29**, 145-169.
- Day, J. and Possingham, H.P., (1995) "A Stochastic Metapopulation Model with Variability in Patch Size and Position.", *Theoretical Population Biology*, **48**, 333-360.
- Department of the Environment, Sports and Territories, (1995) *National Strategy for the Conservation of Australia's biological diversity*. Australian Government Publishing Service: Canberra.

- Gibbons, P., (1994) "Sustaining old growth characteristics in native forests used for wood production: the case of habitat tree retention." pp. 59-84 In T.W. Norton and S.R. Dovers, (Editors). *Ecology and the sustainability of southern temperate ecosystems*. CSIRO, Melbourne, Australia.
- Gibbons, P., and Lindenmayer, D.B., (1996) "A review of issues associated with the retention of trees with hollows in wood production forests in New South Wales.", *Forest Ecology and Management*, **83**, 245-279.
- Government of Victoria, (1988) *Victoria Timber Industry Strategy*. 9. Government Statement. F.D. Atkinson, Melbourne, 104 pp.
- Horne, F., Watts, G. and Robinson, G., (1991) "Current forms and extent of retention areas within a selectively logged Blackbutt forest in N.S.W., a case study.", *Australian Forestry*, **54**, 148-153.
- Kavanagh, R.P., and Turner, R.J., (1994) "Birds in eucalypt plantations: the likely role of retained habitat trees.", *Australian Birds* **28**, 32-41.
- Lin, C. R., Buongiorno, J. and Vasievich, M., (1996) "A multi-species, density-dependent matrix growth model to predict tree diversity and income in northern hardwood stands.", *Ecological Modelling*, **91**, 193-211.
- Lindenmayer, D.B., (1992a) "The ecology and habitat requirements of arboreal marsupials in the montane ash forests of the Central Highlands of Victoria. A summary of studies. Value Adding and Silvicultural Systems Program.", *Report no. 6*. Native Forest Research. February 1992.
- Lindenmayer, D.B., (1992b) "Some impacts on arboreal marsupials of clearfelling on a 80-120 year rotation in Mountain Ash, Eucalyptus regnans forests in the Central Highlands of Victoria.", *Victorian Naturalist*, **109**, 181-186.
- Lindenmayer, D.B., (1994) "The impacts of timber harvesting on arboreal marsupials at different spatial scales and its implications for ecologically sustainable forest use and nature conservation.", *Australian Journal of Environmental Management*, **1**, 56-68.
- Lindenmayer, D.B., (1995) "Disturbance, forest wildlife conservation and a conservative basis for forest management in the mountain ash forests of Victoria.", *Forest Ecology and Management*, **74**, 223-231.
- Lindenmayer, D.B., Cunningham, R.B., Tanton, M.T. and Smith, A.P., (1990a) "The conservation of arboreal marsupials in the montane ash forests of the Central Highlands of Victoria, south-east Australia. II. The loss of trees with hollows and its implications for the conservation of Leadbeater's Possums *Gymnobelideus leadbeateri* McCoy (Marsupialia: Perauridae).", *Biological Conservation*, **54**, 133-145.
- Lindenmayer, D.B., Cunningham, R.B., Tanton, M.T., Smith, A.P., and Nix, H.A., (1990b) "The habitat requirements of the Mountain Brushtail Possum and the Greater Glider in the montane ash-type forests of the Central Highlands of Victoria.", *Australian Wildlife Research*, **17**, 467-478.
- Lindenmayer, D.B., Cunningham, R.B., Tanton, M.T. and Smith, A.P., (1990c) "The conservation of arboreal marsupials in the montane ash forests of the Central Highlands of Victoria, south-east Australia. I. Factors effecting the occupancy of trees with hollows." *Biological Conservation*, **54**, 111-131.
- Lindenmayer, D.B., Cunningham, R.B., Nix, H.A., Tanton, M.T. and Smith, A.P., (1991a) "Predicting the abundance of hollow-bearing trees in montane ash forests of south-eastern Australia.", *Australian Journal of Ecology*, **16**, 91-98.



- Lindenmayer, D.B., Cunningham, R.B., Tanton, M.T., Nix, H.A. and Smith, A.P., (1991b) "The conservation of arboreal marsupials in the montane ash forests of the Central Highlands of Victoria, south-east Australia. III. Models of the habitat requirements of Leadbeater's Possum, *Gymnobelideus leadbeateri* and the diversity and abundance of arboreal marsupials." *Biological Conservation*, **56**, 295-315.
- Lindenmayer, D.B., Cunningham, R.B., Tanton, M.T., Smith, A.P. and Nix, H.A. (1991c) "Characteristics of hollow-bearing trees occupied by arboreal marsupials in the montane ash forests of the Central Highlands of Victoria, south-east Australia.", *Forest Ecology and Management* **40**, 289-308
- Lindenmayer, D.B., Cunningham, R.B., Donnelly, C.F., Tanton, M.T. and Nix, H.A. (1993) "The abundance and development of cavities in montane ash-type eucalypt trees in the montane forests of the Central Highlands of Victoria, south-eastern Australia." *Forest Ecology and Management*, **60**, 77-104.
- Lindenmayer, D.B., Cunningham, R.B., and Donnelly, C.F. (1994) "The conservation of arboreal marsupials in the montane ash forests of the Central Highlands of Victoria, south-east Australia. VI. Tests of the performance of models of nest tree and habitat requirements of arboreal marsupials.", *Biological Conservation*, **70**, 143-147.
- Lindenmayer, D.B., and Possingham, H.P., (1995) *The risk of extinction: ranking management options for Leadbeater's Possum. Centre for Resource and Environmental Studies, The Australian National University and The Australian National Conservation Agency, Canberra. 204pp.*
- Lindenmayer, D.B., Ritman, K.R., Cunningham, R.B., Smith, J. and Howarth, D., (1995) "Predicting the spatial distribution of the Greater Glider, *Petauroides volans* Kerr in a timber production forest block in south-eastern Australia.", *Wildlife Research*, **22**, 445-456.
- Lindenmayer, D.B., Cunningham, R.B. and Donnelly, R.B., (1997) "Decay and collapse of trees with hollows in eastern Australian forests: Impacts on arboreal marsupials.", *Ecological Applications*, **7**, 625-641.
- Lindenmayer, D.B., and Franklin, J.F., (1998) "The importance of stand structure for the conservation of wildlife in logged forests: A case study from Victoria.", *Conservation Biology* (in press).
- Lindenmayer, D.B., Cunningham, R.B. and C.F. Donnelly., (1998) "The collapse of hollow trees in contiguous forest and wildlife corridors in forests of Central Victoria, south eastern Australia.", *Ecological Applications*. (in press).
- Loyn, R.H., Chesterfield, E.A., and Macfarlane, M.A., (1980) "Forest utilization and the flora and fauna in Boola Boola State Forest in south-eastern.", *Victoria. For. Comm. Vict. Bull.* 28.
- Mackowski, C. M., (1984) "The ontogeny of hollows in Blackbutt, *Eucalyptus pilularis* and its relevance to the management of forests for possums, gliders and timber.", *In: A.P. Smith and I.D. Hume (Editors), Possums and Gliders. Surrey Beatty and Sons, Sydney, pp. 517-525.*
- Mackowski, C.M., (1988) "Characteristics of eucalypts incised for sap by the Yellow-bellied Glider, *Petaurus australis* Shaw (Marsupialia: Petauridae) in north-eastern New South Wales.", *Australian Mammal.*, **11**: 5-13.
- McIlroy, J.A., (1978) "The effects of forestry practices on wildlife in Australia: a review.", *Australian Forestry*, **41**, 78-94.
- Newton, I., (1994) "The role of nest sites in limiting the numbers of hole-nesting birds: a review." *Biological Conservation*, **70**, 265-276.

- Oja, T, Yin, X, Arp, P, A., (1995) "The forest modelling series ForM-S: applications to the Solling spruce site.", *Ecological Modelling*, **83**, 207-217.
- Osho, J.S.A., (1996) "Modelling the tree population dynamics of the most abundant species in a Nigerian tropical rain forest.", *Ecological Modelling*, **89**, 175-181.
- Ough, K., and Ross, J., (1992) "Floristics, fire and clearfelling in wet forests of the Central Highlands of Victoria. Silvicultural Systems Project" *Technical Report No. 11*. Department of Conservation and Environment: Melbourne.
- Pausas, J.G., Austin, M.P. and Noble, I.R., (1997) "A forest simulation model for predicting eucalypt dynamics and habitat quality for arboreal marsupials", *Ecological Applications*, **7**, 921-933.
- Possingham, H. P., and Davies, I., (1995) "ALEX: A Population Viability Analysis Model for Spatially Structured Populations.", *Biological Conservation*, **73**, 143-150.
- Rayner, M.E. and Turner, B.J., (1990) "Growth and yield modelling of Australian eucalypt forests. I. Historical development.", *Australian Forestry*, **53**, 224-237.
- Recher, H.F., (1996) "Conservation and management of eucalypt forest vertebrates." In R.M. DeGraaf and R.I. Miller (Editors). *The conservation of wildlife diversity in forested landscapes around the globe*. Chapman and Hall, London.
- Recher, H.F., Rohan-Jones, W. and Smith, P., (1980) *Effects of the Eden woodchip industry on terrestrial vertebrates with recommendations for management*. Forestry Commission of New South Wales, Research Note No. 42.
- Recher, H. F., (1985a) "Synthesis: a model of forest and woodland bird communities." In: J. Keast, H.F. Recher and D. Saunders (Editors), *Birds of Eucalypt Forests and Woodlands: Ecology, Conservation and Management*. Surrey Beatty and Sons, Chipping Norton. Pp. 129-135.
- Recher, H.F., (1985b) "A diminishing resource: mature forest and its role in forest management." In: J. Kikkawa (Editor), *Wildlife Management in the Forest and the Forestry-Controlled Lands in the Tropics and the Southern Hemisphere. Proceedings of I.U.F.R.O. Workshop. July 1984*, pp 28-33.
- Recher H. F., (1995) "Conservation and management of eucalypt forest vertebrates." In: R.M. DeGraaf and R.I. Miller (Editors) *The Conservation of Wildlife Diversity in Forested Landscapes Around the Globe*. Chapman and Hall, London.
- Resource Assessment Commission, (1992) *Forest and Timber Inquiry. Final Report. Volume 1*. Australian Government printing Service, Canberra. 570 pp.
- Roberts, D. W., (1996) "Modelling forest dynamics with vital attributes and fuzzy systems theory." *Ecological Modelling*, **90**, 161-173.
- Saunders, D.A., Smith, G.T. and Rowley, I., (1982) "The availability and dimensions of tree hollows that provide nest sites for cockatoos (Psittaciformes) in Western Australia." *Australian Wildlife Research*, **6**, 205-16.
- Scotts, D.J., (1991) "Old-growth forests: their ecological characteristics and value to forest-dependent vertebrate fauna of south-east Australia." Pp. 147-159 in *Conservation of Australia's forest fauna* Ed by D. Lunney. Royal Zoological Society of N.S.W, Sydney.

Shugart, H. H., (1984) *A Theory of forest dynamics: The ecological implications of forest succession models*. Springer, New York 278pp

Smith, A.P., and Lindenmayer, D.B., (1988) "Tree hollow requirements of Leadbeater's Possum and other possums and gliders in timber production forests of the Victorian Central Highlands." *Australian Wildlife Research*, **15**, 347-362.

Smith, A.P., and Lindenmayer, D.B., (1992) "Forest Succession, timber production and conservation of Leadbeater's Possum (*Gymnobelideus leadbeateri* Marsupialia: Petauridae)." *Forest Ecology and Management*, **49**, 311-332.

Smith, A.P., (1985) "Effects of intensive logging on birds in eucalypt forest near Bega, New South Wales. I. Differences between forest types." *Emu*, **85**, 15-21.

State Forests of New South Wales, (1994) *Eden Management Area Proposed Forestry Operations. Environmental Impact Statement. Volume A. State Forests of New South Wales*, Sydney, pp 12-24.

Thomas, J. W., Anderson, R. G., Maser, C. and Bull, E.L., (1979) "Snags." In: J. W. Thomas (Editor), *Wildlife Habitats in Managed Forests the Blue Mountains of Oregon and Washington*. U.S. Department of Agriculture Forest Service, Agriculture Handbook No. 553. Wildlife Management Institute and U.S. Department of Interior, Bureau of Land Management, Washington, D. C., pp. 60-77

Tyre, A. J., (1999) *Understanding Spatial Variation in Population Dynamics: Enter the Virtual Ecologist*. PhD Thesis, University of Adelaide.

van Saane, L. and Gorden, B., (1991) "Management and planning of Tasmania's native forests for sawlog and pulpwood." In: F.H. McKinnell, E.R. Hopkins and J.E.D. Fox (Editors), *Forest Management in Australia*. Proc. 1987 Conf. Of the Inst. For. Of Australia. Surrey Beatty, Sydney, pp. 58-76.

Voit, E. O. and Sands, P. J., (1996) "Modeling forest growth I. Canonical approach", *Ecological Modelling*, **86**, 51-71.

## Reserve Design References

- Ball, I. R., Smith, A., Day, J. R., Pressey, R. L. and Possingham, H., P (In Press) "Comparison of Mathematical Algorithms for the Design of a Reserve System for Nature Conservation: An Application of Genetic Algorithms and Simulated Annealing." *Journal of Environmental Management*
- Beasley, D., Bull, D. R., and Martin, R. R. (1993) "An overview of genetic algorithms: Part 1, Fundamentals", *University Computing*, **15**, 58-69
- Bedward, M., Robert L. Pressey & David A. Keith. (1991) "A new approach for selecting fully representative reserve networks: addressing efficiency, reserve design and land suitability with an iterative analysis." *Biological Conservation* **62**, 115-125
- Belbin, L. (1993) "Environmental representativeness: Regional partitioning and reserve selection.", *Biological Conservation*, **66**, 223-230.
- Belew, R. K. and Vose, M. D., (1997) *Foundations of Genetic Algorithms 4*, Morgan Kaufman publishing, San Francisco.
- Brown, M. and Hickey, J., (1990) "Tasmanian Forest – Genes or Wilderness?", *Search*, **21(3)**, 86-87.
- Camm, J. D., Polasky, S. Solow, A., and Csuti, B., (1996) "A Note on Optimal Algorithms for Reserve Site Selection", *Biological Conservation*, 353-355
- Church, R. L., Stoms, D. M., and Davis, F. W., (1996) "Reserve Selection as a Maximal Covering Location Problem", *Biological Conservation*, **76**, 105-112
- Cocks, K. D. and Baird, I. A., (1989) "Using mathematical programming to address the multiple reserve selection problem: An example from the Eyre Peninsula, South Australia", *Biological Conservation*, **49**, 113-130.
- Cocks, K. D. and Baird, I. A., (1991) "The role of geographic information systems in the collection, extrapolation and use of survey data", page 74-80 in *Nature Conservation: The role of remnants of Native Vegetation* ed by Saunders, D. Q., Arnold, G. W., Burbidge, A. A., and Hopkins, A. J. M. (Surrey Beatty and Sons Pty Limited in association with CSIRO and CALM)
- Connolly, D. J. (1990) "an improved annealing scheme for QAP", *European Journal of Operations Research*, **46**, 93-100.
- Crozier, R. H. (1992) "Genetic diversity and the agony of choice", *Biological Conservation*, **61**, 11-15.
- Csuti, B., Polasky, S., Williams, P. H., Pressey, R. L., Camm, J. D., Kershaw, M., Kiester, A. R., Downs, B., Hamilton, R., Huso, M., and Sahr, K. (1997) "A Comparison of Reserve Selection Algorithms using Data on Terrestrial Vertebrates in Oregon", *Biological Conservation*, **80**, 83-97
- Faith, D. P. (1992) "Conservation evaluation and phylogenetic diversity", *Biological Conservation*, **61**, 1-10.

- Faith, D. P. and Walker, P. A., (1996) "How do indicator groups provide information about the relative biodiversity of different sets of areas?: on hotspots, complementarity and pattern-based approaches.", *Biodiversity Letters*, **3**, 18-25.
- Freitag, S. and Van Jaarsveld, A. S., (1998) "Sensitivity of selection procedures for priority conservation areas to survey extent, survey intensity and taxonomic knowledge", *Proceedings of the Royal Society of London B*, **265**, 1475 - 1482.
- Garey M.R. and Johnson D. S. (1979) *Computers and Intractability: A guide to the theory of NP-Completeness.*, W H Freeman and Company, San Fransisco.
- Gehlbach, F. R., (1975) "Investigation, Evaluation, and Priority ranking of natural areas", *Biological Conservation*, **8**, 79-88
- Goldman, F. B., (1975) "The Evaluation of Ecological Resources in the Countryside for Conservation Purposes", *Biological Conservation*, **8**, 89-96.
- JANIS (1997) *Nationally agreed criteria for the establishment of a comprehensive, adequate and representative reserve system for forests in Australia. A report by the Joint ANZECC/MCFFA National Forest Policy Statement Implementation Subcommittee. National forest conservation reserves: Commonwealth proposed criteria.* Commonwealth of Australia, Canberra.
- Kershaw, M., Williams, P. H., and Mace, G. B., (1994) "Conservation of Afrotropical Antelopes: Consequences and Efficiency of Using Different Site Selection Methods and Diversity Criteria", *Biodiversity and Conservation*, **3**, 354-372.
- Kirkpatrick, J. B. (1983) An Iterative Method for Establishing Priorities for the Selection of Nature Reserves: An Example from Tasmania *Biological Conservation* **25** 127-134
- Kirkpatrick, J. B. and Harwood, C. E., (1983) "Conservation of Tasmanian macrophytic wetland vegetation", *Papers and Proceedings of the royal Society of Tasmania*, **117**, 5-???
- Krajewski (1994) "Phylogenetic measures of biodiversity: A comparison and critique", *Biological Conservation*, **69**, 33-39.
- Lahti, T. and Ranta, E., (1985) "The SLOSS principle and conservation practice: an example", *OIKOS*, **44**, 369-370.
- Lombard, A. T., Cowling, R. M., Pressey, R. L., and Mustart, P. J., (1997) "Reserve Selection in a species-rich and fragmented landscape on the agulhas plain, South Africa", *Conservation Biology*, **11**, 1101-1116.
- Lomolino, M. V., (1994) "An Evaluation of Alternative Strategies for Building Networks of Nature Reserves" *Biological Conservation* **69** 243-249
- Mackey, B. G., Nix, H. A., Stein, J. A., and Cork, S. E., (1989) "Assessing the representativeness of the wet tropics of Queensland world heritage property", *Biological Conservation*, **50**, 279-303.
- McCullough, P. and Nelder, J. A., (1983) *Generalised Linear Models*, Chapman and Hall. London.
- McKenzie, N. L., Belbin, L., Margules, C. R., Keighery, G. J., (1989) "Selecting representative reserve systems in remote areas: A case study in the Nullarbor Region, Australia", *Biological Conservation*, **50**, 239-261.
- Margules, C. R., Higgs, A. J., and Rafe, R. W., (1982) "Modern biogeographic theory: are there any lessons for nature reserve design?", *Biological Conservation*, **24**, 115-128.

- Margules, C. R. and Nicholls, A. O., (1987) "Assessing the conservation value of remnant habitat 'islands': Mallee patches on the Western Eyre Peninsula, South Australia", page 89-102 in *Nature Conservation: The role of remnants of Native Vegetation* ed by Saunders, D. Q., Arnold, G. W., Burbidge, A. A., and Hopkins, A. J. M. (Surrey Beatty and Sons Pty Limited in association with CSIRO and CALM)
- Margules, C. and Stein, J. L., (1989) "Patterns in the distributions of species and the selection of nature reserves: An example from *Eucalyptus* forests in South-eastern New South Wales", *Biological Conservation*, **50**, 219-238.
- Margules, C. and Usher, M. B. (1981) "Criteria used in assessing wildlife conservation potential: A review", *Biological Conservation*, **21**, 79-109.
- Margules, C. R. and Usher, M. B., (1984) "Conservation evaluation in practice I. Sites of different habitats in North-east Yorkshire, Great Britain", *Journal of Environmental Management*, **18**, 13-168.
- Nantel, P., Bouchard, A., Brouillet, L., and Hay, S., (1998) "Selection of areas for protecting rare plants with integration of land use conflicts: a case study for the west coast of Newfoundland, Canada", *Biological Conservation*, **84**, 223-234.
- Nicholls, A. O., (1989) "How to make biological surveys go further with generalised linear models", *Biological Conservation*, **50**, 51-75.
- Nicholls, A. O. and C. R. Margules (1993) "An Upgraded Reserve Selection Algorithm" *Biological Conservation* **64** 165-169
- Possingham, H. and Andelman, S., (in press) "An equity formula for fauna persistence: What happens when conservation biologists have to provide prescriptions for real problems - quickly?"
- Possingham, H., Day, J., Goldfinch, M., and Salzborn, F., (1993) "The mathematics of designing a network of protected areas for conservation", pp 536-545 in "Decision sciences: tools for today" Proceedings of 12<sup>th</sup> National ASOR conference. Eds. D. J. Sutton, C. E. M. Pearce, and Cousins, E. A., ASOR Adelaide,
- Prendergast, J. R., Quinn, R. M., Lawton, J. H., Eversham, B. C., and Gibbons, D. W., (1993) "Rare species, the coincidence of diversity hotspots and conservation strategies", *Nature*, **365**, 335-337.
- Pressey, R. L., Humphries, C. J., Margules, C. R., Vane-Wright, R. I., and Williams, P. H., (1993) "Beyond opportunism: Key principles for systematic reserve selection", *TREE*, **8**, 124-128
- Pressey, R. L., Johnson, I. R., and Wilson, P. D., (1994) "Shades of irreplaceability: towards a measure of the contribution of sites to a reservation goal."
- Pressey, R. L. and Nicholls, A. O., (1989a) "Efficiency in Conservation evaluation: scoring versus iterative approaches", *Biological Conservation*, **50**, 199-218.
- Pressey, R. L. and Nicholls, A. O., (1989b) "Application of a numerical algorithm to the selection of reserves in semi-arid New South Wales", *Biological Conservation*, **50**, 263-278.
- Pressey, R. L., Possingham, H. P., Day, J. R., (1997) "Effectiveness of alternative heuristic algorithms for identifying indicative minimum requirements for conservation reserves", *Biological Conservation*, **80**, 207-219.
- Pressey, R. L., Possingham, H. P., Margules, C. R., (1996) "Optimality in Reserve Selection Algorithms: When Does it Matter and How Much?" *Biological Conservation* **76**, 259-267.
- Purdie, R. W., Blick, R., and Bolton, M. P. (1986) "Selection of a conservation reserve network in the mulga biogeographic region of Sout-Western Queensland, Australia", *Biological Conservation*, **38**, 369-384.

- Rabe, F. W. and Savage, N. L., (1979) "A methodology for the selection of aquatic natural areas", *Biological Conservation*, **15**, 291-301.
- Rebelo, A. B. and Siegfried, W. R., (1990) "Protection of Fynbos vegetation: Ideal and real-world options", *Biological Conservation*, **54**, 15-31.
- Rebelo, A. G. and Siegfried, W. R., (1992) "Where Should Nature Reserves be Located in the Cape Floristic Region, South Africa? Models for the Spatial Configuration of a Reserve Network Aimed at Maximizing the Protection of Floral Diversity" *Conservation Biology* **6** (2) 243-252.
- Rossi, E. and Kuitunen, M., (1996) "Ranking of habitats for the assessment of ecological impact in land use planning", *Biological Conservation*, **77**, 227-234.
- Sætersdal, M., Line, J. M., and Birks, H. J. B., (1993) "How to maximise biological diversity in nature reserve selection: Vascular plants and breeding birds in deciduous woodlands, Western Norway", *Biological Conservation*, **66**, 131-138.
- Simberloff, D., (1998) "Flagships, umbrellas, and keystones: is single species management passe in the landscape era?", *Biological Conservation*, **83**, 247-257.
- Sinden, J. A. and Windsor, G. K., (1981) "Estimating the value of wildlife for preservation: A comparison of approaches", *Journal of Environmental Management*, **12**, 111-125.
- Smith, P. G. R. and Theberge, J. B., (1986) "A review of criteria for evaluating natural areas", *Environmental Management*, **10**, 715-734.
- Thomas, C. D. and Mallorie, H. C., (1985) "Rarity, species richness and conservation: Butterflies of the Atlas Mountains in Morocco", *Biological Conservation*, **33**, 95-117.
- Underhill, L. G. (1994) "Optimal and Suboptimal Reserve Selection Algorithms", *Biological Conservation*, **70**, 85-87.
- Vane-Wright, R. I., Humphries, C. J., and Williams, P. H., (1991) "What to Protect? -- Systematics and the agony of choice", *Biological Conservation*, **55**, 235-254.
- Vidal R. V. V. (1993) *Applied Simulated Annealing*, Springer-Verlag: Berlin - New York.
- Williams, J. C. and ReVelle, C. S., (1997) "Applying mathematical programming to reserve selection", *Environmental Modeling and Assessment*,
- Willis, C. K., Lombard, A. T., Cowling, R. M., Heydenrych, B. J., and Burgers, C. J. (1996) "Reserve Systems for Limestone Endemic Flora of the Cape Lowland Fynbos: Iterative Versus Linear Programming", *Biological Conservation*, **77**, 53-62.