Exploring natural resource management tradeoffs in an agricultural landscape - an application of the MOSAIC model

Kenton Lawson[a1], *Christine Hill*[b], *Andrew Hodges*[a] *and Brent Jacobs*[b2]
a. Australian Bureau of Agricultural and Resource Economics and
b. New South Wales Department of Natural Resources

51st Annual Conference of the Australian Agricultural and Resource Economics Society, Rydges Lakeland, Queenstown, New Zealand, 13 – 16 February 2007

Abstract

We describe a landscape scale non-linear discrete choice spatial optimisation model for identifying cost-effective strategies for achieving environmental goals. Spatial heterogeneity and configuration issues such as fencing costs, patch sizes and network linkages are explicitly accounted for and quasi-optimal allocations are determined using simulated annealing. Applications of the model being developed with New South Wales Catchment Management Authorities are discussed. These focus on targeting investments in revegetation to control dryland salinity and erosion and provide biodiversity benefits whilst minimising direct and opportunity costs. We compare our approach with alternate investment approaches.

Key words: natural resource management, cost effectiveness, land use change, multicriteria, spatial optimisation

[1] klawson@abare.gov.au, ABARE, GPO Box 1563, Canberra ACT 2601. First time AARES presenter.

[2] The views presented in this paper, drawn from preliminary work in progress, are those of the authors and do not represent the official view of ABARE, the New South Wales Government or the Australian Government.

Introduction

Public policy makers face the complex problem of how best to obtain the highest value to society, over time, from the use (including conservation) of Australia's natural resources. The use of natural resources provides Australia with enormous economic benefits, contributing significantly to the country's income and standard of living. Past resource use, which may have been considered appropriate at the time (given available information and prevailing social values), has also left a legacy of resource degradation such as salinity and biodiversity loss.

The potential cost of addressing resource degradation is large and the impacts on regional communities are likely to be significant. Given the large cost involved, it may not be in society's best interests to repair all of the degradation. Instead, it will be important to determine a framework for prioritising and focusing efforts in areas where the benefits of repairing damage exceed the costs. This will ensure that the benefits from the public investments being undertaken are maximised.

In this paper we present a spatially explicit model, the MOSAIC model, for analysing the potential of changed land management and targeted remediation measures to deliver improved environmental outcomes – primarily water quality and biodiversity – in particular rural landscapes. Although we do not discuss the question of how best to achieve these changes in land management we see the information outputs generated by this research as being a necessary ingredient in designing efficient policy interventions.

The major approach used by governments to address these problems has been the use of regulations to govern resource use. However, other approaches to addressing natural resource degradation such as those that fund activities that contribute to the provision of these public goods are increasingly used. Choosing the best mix of regulations, funding public good provision and other policies requires that a range of options be compared against a range of criteria including:

- the effectiveness of the policy in delivering explicitly stated objectives,
- the efficiency of the policy in minimising the full costs of achieving the objectives and in achieving an appropriate balance between the objectives achieved and the costs,
- the ability of the policy to adjust to changed conditions and information over time and in different locations,
- the security provided to resource managers for future investment, and
- the distributional effects of the policy.

In practice, realising these criteria is difficult because of the complexities of real landscapes. Not only are landscapes heterogeneous but actions and effects are often linked by a number of biophysical and socioeconomic processes operating at different scales, over different timelines and are often imperfectly understood.

The MOSAIC model is a decision support tool that:

1. embeds our current understanding of the linkages between on ground actions and environmental outcomes,
2. uses optimisation techniques to find least cost combinations of actions to deliver the targeted environmental benefits, and
3. provides information on the scope for trade offs involved in delivering varying levels of environmental benefits.

The explicit spatial representation of the MOSAIC model allows for consideration of spatial heterogeneity of benefits and costs of different land management options on different land units. Threshold effects in environmental systems and diminishing marginal returns from additional land units devoted to a given management option are all able to be included in the calculation of landscape values and constraints. These characteristics help to ensure that the greatest benefits of landuse change are achieved at the lowest cost.

Since the majority of the impacts of changed land management, whether on rural communities or on environmental benefits, occur at a regional scale it is appropriate that efforts to improve the delivery of these benefits be facilitated through processes based at the regional level. However, such processes need to be designed to deliver on a set of clearly specified objectives and need to be appropriately resourced. They also need to be supported with the provision of data and scientific and economic expertise. This regional planning problem provides the context for the application of the MOSAIC model which aims to assist with the integration of existing scientific and economic information in order to support regional environmental planning.

Modelling approach

The basic approach is to search for land use mosaics that deliver specified combinations of environmental outputs at least cost. Here cost includes the foregone returns from commodity production as well as any direct costs associated with the changed land management. This search is repeated for a range of combinations of environmental outputs in order to identify an economic-environmental efficiency frontier for the landscape being investigated (as in Polasky et al. 2005). Landscapes lying on the frontier trace out the potential for trade offs between commodity production values and the delivery of the different environmental benefits. For landscapes lying inside the frontier it is possible to achieve an overall improvement in both commodity production values and environmental benefits.

If there were agreed monetary values or shadow prices associated with each environmental benefit metric then we could search for mosaics that maximise the overall landscape value:

$$Max \quad V^m + \sum_e p^e V^e \tag{1}$$

Where the maximisation is over the choice of land management option for each land unit, V^m is the monetary landscape value (commodity production value less direct costs) and p^e and V^e are the shadow price and landscape value of environmental benefit *e* respectively. All values are long term values since the changed landscape mosaic is assumed to be maintained indefinitely.

In fact, the approach taken in the prototype implementation is to solve this dual specification and to systematically alter the shadow prices in order to trace out the efficiency frontier for the landscape. For this multi-criteria analysis approach to be workable the number of environmental benefits included needs to be kept quite modest.

In order to define a landscape mosaic we partition the region being studied into discrete polygons or land units and identify a range of management options for each land unit. How these decision units are defined, for example, their shape and size, constrains the analysis and will affect how accurately any biophysical relationships are represented. There is also a trade off between the use of finer scales and computational limits. In the reserve design literature the size of land units used has been shown to influence which land units are selected for reservation (Pressey and Logan 1998). The approach adopted here is to aim for reasonably homogeneous management units that respect biophysical boundaries. Where the spatial configuration of land uses is important, small to medium scales are called for. This allows patch sizes and fencing and other boundary costs to be explicitly considered. These aspects would be difficult to represent using larger and more heterogeneous land units.

A key component to the modelling approach is the representation of the biophysical linkages between a particular configuration of land management options and the environmental benefits generated by the landscape. These are discussed further below. In general the strategy is to embed existing biophysical models (or simplified versions) where these are available.

Once the land units, management options, environmental metrics and objective function have been defined the next task is to search for optimal configurations of management options. If there are no spatial interactions then the problem falls within the class of mixed integer problems. For small problem sizes these can be solved exactly using branch and bound algorithms[3]. However, in practice, problems with at least thousands of land units are more likely. This is required to realistically represent the true heterogeneity of rural landscapes, to allow the flexibility to manage areas with different characteristics differently and to include a large enough area to address landscape level environmental objectives and for different parts of the landscape to compete with each other in supplying environmental benefits. In this case heuristic optimisation procedures are a more practical alternative. Whilst they do not guarantee optimality effective heuristic algorithms impose only a small inefficiency cost – especially when compared with the likely much greater uncertainty in the data and modelling of biophysical linkages.

A number of different heuristics are available and range from simple greedy algorithms where at each step the best alternative in a neighbourhood of the current solution is chosen. These are generally fast and are easy to implement (Hajkowicz et al. 2005). They fail when, firstly, there are effective linkages between the value of a land use option in that neighbourhood and what is being done in the rest of the landscape and, secondly, when they are trapped by local optima. The former issue is likely when there is interdependence between the land units, such as where downstream changes affect the impact of upstream sediment yields, or where different areas are competing to supply a particular environmental benefit. A simpler example is where the land use in an adjoining land unit determines whether the boundary needs to be fenced. In these cases more sophisticated heuristics, such as genetic algorithms or simulated annealing, have both been successfully applied (Possingham et al. 2000; Aerts et al. 2005)[4].

Following on from past research we have adopted simulated annealing as the main search method. Simulated annealing (Kirkpatrick et al. 1983) is essentially a randomized version

---

[3] Non-linearities make the exact solution of these problems significantly harder. If the only non-linearity is associated with needing to account for boundary lengths then the problem can be linearised by adding an extra variable and constraint for each possible boundary. This significantly increases the number of variables that need to be solved but may assist in determining exact solutions for some small problems.

[4] Another approach has been used for solving linear problems where the relaxed form (where discrete variables are allowed to take intermediate non-integer values) is solved and any non-integer variables in the solution are rounded to the nearest integer. For such problems this method can also perform quite well.

of the greedy algorithm. In this case, when a particular land management option is being considered for a given land unit, whether or not it is accepted at this step of the search depends on the magnitude of the change in the objective function resulting from acceptance and on a random factor. If the objective function increases then the alternative is accepted. If the objective function decreases – the change, ΔV, is negative – then the change is accepted with probability given by the Metropolis Criterion:

$$\exp\left(\Delta V / T\right) \tag{2}$$

Where the parameter T is large at the beginning of the search – so that many changes that decrease the objective are accepted – and is gradually reduced until, at the end of the search, essentially only objective improving changes are accepted. The magnitude of these parameters are determined through a purely random search in order to estimate the distribution of likely changes in the objective function. In a typical search the algorithm considers changing the management option for each land unit thousands of times. Some experimentation is required with each implementation in order to determine an adequate search length. Also it is recommended that for each specification of the objective function (with different shadow prices) the search be repeated a number of times to reduce the likelihood of the results of the analysis being affected by any inefficiency in the optimisation.

Landscape values

Commodity production landscape value

The commodity production landscape value is simply the present value of commodity production in each land unit summed over all land units. The value at the individual land unit level depends on its characteristics and the selected land management option. This assumes that prices are not affected by changes in supply from the region being investigated and there are no production externalities such as pollution from adjacent land units. Where these effects are significant and can be estimated they can be readily included.

For non-commodity production land uses and where the management option includes specified remediation measures, the direct costs of these activities, including any management costs, are included in the present value calculation for these management options. In specifications where fencing and other boundary costs, which depend on the length of the boundary between each pair of land uses, are explicitly included, they form an additional component of V^m.

Biological landscape value

For public investment in biodiversity conservation to be cost effective it must achieve the best biodiversity conservation outcome possible for the given level of costs. The first obstacle to achieving this is the lack of a clear statement of the biodiversity conservation objective. Once a clear objective is agreed, different policy options can be compared according to their contribution to this objective, a full assessment of their costs and other criteria as mentioned above. An example of such a biodiversity objective is the weighted sum of the viabilities of species and ecological communities in a region:

$$V^b = \Sigma w_i V_i^b \tag{3}$$

Where V^b is the biological landscape value, w_i is a weighting factor and V_i^b is the expected viability of species or community *i*. Viability for a species or community is defined as the probability that the population size or area at the end of some time interval exceeds a given threshold level. A simple case would be where all native vertebrate and vascular plant species whose biology is reasonably well known are included and all the weights are equal. Westphal and Possingham (2003) and Montgomery et al. (1999) discuss cases where only bird species are included and Possingham et al. (2002) use the number of species saved as a simple approximation of V. In this case all weights are one and the viabilities are either zero or one. However, varying weights may be appropriate for reasons of aesthetics, ecological function or utility.

This specification of the biodiversity conservation objective is consistent with the nature of the biodiversity conservation task. It allows actions that affect the viability of a number of species to be accorded a greater value than those with a much narrower effect. Similarly, including a broad range of species is essential when species have conflicting habitat requirements. Restricting attention to only a few focal species makes it unlikely that many species' habitat requirements will be met and may mean that significant potential biodiversity benefits are missed, particularly when the focal species are the most sensitive or endangered (Westphal and Possingham 2003).

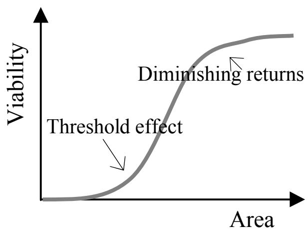

Figure 1. A species viability curve.

Figure 1 illustrates the relationship between the area of good quality habitat for a species and its viability. If actions affecting the habitat area when the species is currently on either the far right hand side or far left hand side — that is, where the species is either secure or critically endangered — are considered then the action may have little impact. However, if the species is on the steeper part of the curve, which may correspond to an endangered, vulnerable or ofconcern status, such actions would have a larger impact. Thus, while ignoring the most secure species or ecosystems may not compromise the efficiency of biodiversity investments, prioritising according to degree of endangerment will not yield the best outcome in terms of viability across species. For a given level of public investment there are likely to be species that are too costly to recover when compared with the alternative of increasing the viability of a larger number of species.

Another important factor is the spatial interdependence of biodiversity values within a region. A species' viability for a given area of habitat depends on the size and shape of habitat patches and the spatial configuration of those patches in the landscape. Thus the biodiversity value of a given patch of vegetation cannot be based solely on intrinsic characteristics of the patch. Strategic investments in biodiversity conservation need to account for these landscape and regional level interactions. Attempting to achieve biodiversity benefits at the scale of individual properties is unlikely to be cost effective.

For perhaps largely practical reasons, existing conservation investments and regulations have used a more piecemeal or ad hoc approach based on intermediate measures of biodiversity outcomes. For example, where only a fixed area of a given vegetation community is allowed to be cleared, the areas that are actually cleared are generally determined by the order in which landholders apply for permits rather than some more efficient rationing process.

The use of area targets is widespread in the implementation of biodiversity conservation policies in Australia. While the arbitrariness of many of these targets is often openly acknowledged the biases they introduce into conservation planning need to be examined. Such area targets are often determined as a given percentage of the estimated pre-European areal extent of the vegetation community. While this is a simple pragmatic approach it should only really be a temporary approach in those cases where there is no better understanding of the determinants of ecosystem viability. Such targets largely ignore critical aspects of the biodiversity values such as

– the current quality of habitat,
– the configuration of the habitat, and
– other aspects of land management such as fire, weed or feral species management.

Sattler and Creighton (2002) highlight the need to achieve an appropriate balance across a large range of actions that all affect biodiversity outcomes.

Even if these aspects were appropriately included, it is likely that there would be only a weak relationship between the areal extent of a vegetation type in say 1770 and the area required to achieve a given viability threshold. More fundamentally these area statements are only an intermediate measure of environmental benefits. Threshold effects and diminishing returns, as illustrated in figure 1, mean that maximising areas for a given budget is not generally equivalent to maximising environmental benefits. Wu and Boggess (1999) show how maximising the area can actually minimise the environmental benefits of a given investment. For example, this would be the case when limited funding is being allocated across two identical catchments where there is an increasing marginal cost of land and there are threshold effects. Maximising the area means splitting the investment equally across the two catchments. Because of threshold effects the environmental benefits are maximised by targeting the investment to one catchment first. Similarly, where the otherwise identical catchments differ in environmental quality due to past conservation actions, they demonstrate that when there are threshold effects limited investments are likely to generate greater environmental benefits if the 'cleaner' catchment is funded first. Maximising the area funded would target the lower environmental quality catchment first.

A similar spatial optimisation model using simulated annealing was developed during the Regional Forest Agreement (RFA) process. The model in Possingham et al. (2000) is based on this work. Here the problem was to decide which publicly owned forested land units would be added to the conservation reserve system. This process adopted the objective of achieving quantitative biodiversity targets at minimum cost. These include area targets for vegetation communities as discussed above as well as targets for endangered and vulnerable species that were more closely linked to the biology of the species in question. A major drawback with this process of target setting was that it gave no guidance for prioritising investments when not all the targets could be achieved. Opportunities, such as where a small reduction in the viability of one community could be traded off for a large gain in the viability of another community, were also not able to be identified. However, the model was valuable in identifying what targets could be achieved and at what cost. The model included a penalty for the boundary length of the reserve system and this could be weighted to bias the search towards more compact reserve systems.

A key problem with conservation planning relates to the accuracy of the data used. In the RFA process some areas selected to represent endangered communities in an expanded national park system did not in fact contain the communities they were selected to represent. Importantly, there were processes to revise the allocations in the cases that were identified (Davey, S. 2000). Similarly the Productivity Commission Native Vegetation and Biodiversity (PCNV&B) inquiry heard cases where the mapped vegetation underlying regulations were shown to be in error (eg S. Doust and L. Acton on pp. 57, 70 and 73 of the Brisbane transcript (Productivity Commission 2003)). While the

National Vegetation Information System is making progress in providing a national vegetation database with consistent classifications and consistent quality, there is clearly a need for regional processes to be able to incorporate more accurate data as it becomes available, and there is scope for such processes to better incorporate local knowledge. This is likely to remain problematic when information relating to public goods such as biodiversity values can lead to private costs.

Another key problem is that the actual effects of vegetation management and biodiversity conservation actions on species and ecosystem viability are poorly understood. This is particularly the case for ecological communities. At the species level, population viability analysis is relatively well advanced if only for a small set of species. For the RFA process it was possible to derive crude viability measures for hundreds of species using modelling and expert knowledge. While current processes seek to make the best use of available knowledge there is clearly a need to be able to adapt strategies and actions as new knowledge is gained. The RFA process at least nominally specified a twenty year period after which the allocation decisions could be reassessed.

The reserve selection literature usually treats non-reserved areas as having no biodiversity value with reserve areas being treated as islands. Areas outside formally protected areas do contribute towards species viabilities and this may be incorporated by allowing partial compatibility of different management options. In this case each management option has a compatibility index value which can be used in determining habitat areas. The calculation of the viability of different biodiversity components can use this index in different ways. Where this is the case the production-biodiversity efficiency frontier will show less conflict (Polasky et al. 2005).

Water quality and quantity landscape values

There are a number of water quality issues that may be important for public policy in different regions. Here we discuss stream salinity and sediment loads.

Sediment loads

Suspended sediments with associated nutrients are significant contributors to poor water quality in many waterways. This imposes costs on irrigators and graziers, urban water treatment and on the viability of natural systems.

Sediment loads derive from a number of sources and the relative significance of the different sources varies in different regions. The sediment load in a single stream link is derived from upstream links, gully erosion along the link and in the link's contributing area, bank erosion on the link and hillslope erosion. Deposition along the link reduces the sediment transported downstream. Similarly, not all the sediment generated by hillslope erosion in the link's contributing area flows into the link. This is affected by rainfall, topography, land cover, soil types and land management (Prosser et al. 2001). Each sediment source and sink can be targeted by policy interventions and because sediment is transported unevenly through the stream network targeting the areas that generate the most sediment is not generally an efficient strategy.

Lu et al. (2004) compare four scenarios for prioritising investments to reduce erosion relative to natural levels that include progressively more information about the sediment distribution networks of the Murray Darling basin:

(A) random selection of which locations and processes are treated. This scenario has low transaction costs and mimics a first come first served strategy which has been used in the past.

(B) Here investment is prioritised to areas (stream links and contributing areas) with the highest total erosion rates. This reflects the hot spot approach which assumes that the biggest source areas have the largest effect on downstream water quality.

(C) This strategy accounts for the spatially variable proportion of the sediment from hillslope erosion that does not make it to the stream network. Trapping sediment before it reaches the stream network may be more cost effective than controlling large contributing areas.

(D) The last scenario also accounts for the deposition processes that may remove sediment and thereby reduce downstream loads. For example, major storages may effectively disconnect upstream sediment sources from affecting downstream water quality. Here the targeted water quality objectives are measured at a few specified locations on the stream network. The appropriate specification should account for the different costs imposed by reduced water quality in different parts of the stream network.

Representative direct costs are included for controlling each of the major types of erosion and for reducing the hillslope delivery ratio. Opportunity costs were not included. In all the basins modelled the cost of reducing sediment loads was least, and was radically below the random scenario A, for the fully informed scenario D. Otherwise there was generally a reduction in costs as more information was used in targeting the interventions. However, there were cases where erosion hotspot scenario B was more expensive in achieving a given reduction in sediment loads than most realisations of the random prioritisation scenario. This analysis demonstrates that explicitly accounting for the spatial heterogeneity of different sediment sources and their variable linkages permits dramatic improvements in the cost effectiveness of achieving given water quality targets.

Stream salinity

The scientific and physical causes of salinisation are complex, but relatively well researched (see Eberbach 1998). As a result of more water entering the water table than is being extracted, naturally occurring salty groundwater finds its way to the surface and into streams and rivers. This causes the overall load and concentration of salt in rivers to increase, reducing water quality both for the natural environment and for consumptive uses such as irrigated agriculture.

From an economic perspective, salinisation occurs because activities that cause, affect or mitigate salinity have spatial and temporal consequences not borne by those undertaking the activity. When this occurs, decision makers often undertake actions that are privately optimal, but may not be optimal from a social perspective. Hence, many actions that have a negative impact on water quality have been undertaken, while actions that have a positive impact are not provided, or are not provided at an adequate level.

The actions that can mitigate salinity are also well documented. Some of these include revegetation, reducing water extractions from rivers, limiting the water used in consumptive land based activities, constructing drainage systems and pumping ground water.

From a public policy perspective, there are ways of both ensuring beneficial actions are undertaken, and of reducing the number of undesirable activities. However, given the site specific and often complex nature of water salinity problems, these changes in management often need to be targeted to specific parts of the landscape to achieve the greatest benefit and avoid exacerbating other resource degradation problems. For example, poorly located revegetation actions to mitigate in-stream salinity may achieve a proportionally greater reduction in surface water runoff than in salt load, thereby contributing to an increase in salt concentrations in rivers.

Another important consideration is that many of these mitigation activities can impose significant costs on agriculture and rural communities more generally. This would be the case, for example, where broad scale reforestation both reduces the surface water yield and increases the salinity concentrations in rivers in the near term to the detriment of downstream irrigators. Even in the long term, the proportional reduction in salt transported to rivers and streams may be less than the reduction in surface water runoff. That is, the cessation of tree clearing, regrowth management or broad scale tree planting such as through plantation forestry has the potential to capture water that would otherwise make its way into rivers and other watercourses. Not only could this make the problem worse, not better, it could impose significant costs on downstream communities as the volume of fresh water is reduced.

Reforestation may generate substantial net salinity mitigation benefits, however, if it is targeted to specific parts of the landscape to ensure it delivers a proportionally greater reduction in salt mobilisation than any reduction in surface water runoff. In this regard, two important landscape characteristics are soil types and the salinity of groundwater. Revegetation that reduces recharge to relatively fast responding aquifers would also be generally preferred. This suggests the level and type of actions undertaken should vary according to the biophysical characteristics of each region. A carefully designed scheme would ensure that trees were planted, or left uncleared, in areas where they would generate net salinity benefits.

The required level of information to ensure actions are correctly targeted in the landscape is generally significant and costly. As information and data improves, however, the scope for better outcomes increases. Current levels and accuracy of data are generally insufficient at the scale of individual properties although this is being facilitated as regions prioritise areas where the greatest benefits would be achieved.

When considering revegetation as an option to mitigate the effects of salinity, acknowledgment may need to be given to the opportunity costs of any water captured. One way this could be done is to draw revegetation activities within the scope and framework of the formal water market. This would require extending the coverage of water property rights to cover changes in land use. If this were done, large scale revegetation activities would be required to buy any water they need through the market, in the same way as other consumptive uses.

Whether the extension of rights to include water use by forestry will actually lead to a more efficient outcome will depend on the transaction costs associated with extending those rights. Any apparent efficiency improvements in moving to such a system may be outweighed by the costs of negotiating, administrating, monitoring and enforcing the extension of rights, as well as by any costs in gathering information on the level of water use and marginal net benefits from water use by forestry (Goesch and Hanna 2002).

Prototype specification

In order to test the feasibility of the integrated spatial optimisation approach a simple application was developed for an area of about 30 000 ha in the south west slopes region of New South Wales. The current land use is primarily livestock grazing with smaller areas of pine plantations and remnant native vegetation. The analysis considers changing the grazing land use areas to either new plantations or revegetation for nature conservation. As well as commodity production values the model was specified to include biodiversity values, water quantity and stream salinity values and a value for the expected additional carbon sequestration under the changed land management.

The biophysical specification and production values in the prototype use synthetic data and have not been calibrated. So although the prototype successfully generates allocations according to the objective function with increasing shadow prices resulting in increased supply of the associated environmental benefit the results presented here are only illustrative.

The data and model outputs are stored in a Microsoft Access database to which the user interface and optimisation functions, coded in C++, read and write. This allows a degree of data transparency so that the data and outputs can be analysed independently from the MOSAIC code.

The software defines a scenario as a set of shadow prices and other parameters, such as the time horizon and discount rate, which uniquely defines a particular optimisation problem. Once a scenario is specified a number of searches of specified lengths can then be run. The resulting allocations can be visualised in the user interface, or other GIS, and tabular reports can also be generated on the different value measures.

In order to represent the hydrological linkages, the land units were derived from the contributing area of each stream link and from the current land use. This implementation includes 622 land units with a mean size of a little less than 50 ha. At this size, as shown in figure 2, there is still a fair degree of heterogeneity of land cover within the land units. In such cases it may be worthwhile creating one or more intermediate land cover management options. However, at this resolution it is not possible to address configuration issues below the land unit level.

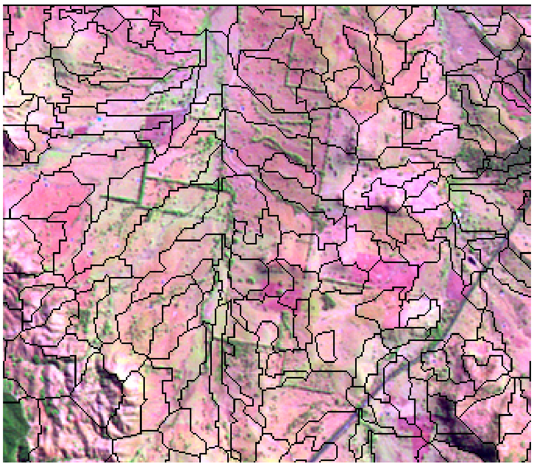

Figure 2. Land units in the prototype specification

The land units were classified into productivity classes for each land use. Each land use productivity class has a yield table which specifies the net returns, carbon flux, land cover and log yields at each year of the rotation or growth of that land management activity.

This information together with the individual land unit attributes allows the different landscape values to be calculated. For the water values additional attributes on the surface water catchments and groundwater flow systems are also required in order to assess the overall impact of the changed land management on the average flow levels and stream salinity at predefined measurement nodes on the stream network. The changed recharge associated with changing the land cover leads to changes in the discharges from the ground water system with a lagged distribution (Bell and Heaney 2000). Fencing costs on a dollars per km basis are specified for each pair of land uses.

The production, carbon sequestration and water values are calculated as present values over a user specified time horizon, for example, 100 years, and user specified discount rate. The biodiversity values are calculated as at equation (3) above where the individual viability indices are based on the area of suitable habitat relative to a threshold habitat area in the region according to a logistic curve. The amount of potential habitat for each biodiversity component in each land unit is specified in advance as is the compatibility of each land management option. A minimum patch size can be specified to discount the contribution of potentially suitable habitat in patches below the minimum patch size. When this option is specified the algorithm needs to keep track of all the habitat patches during the search. This leads to significant increases in the required search time. The other way of achieving more compact areas of habitat, than that resulting from simply including fencing costs, is to increase the cost of the boundary length of compatible management options. Thus the biodiversity value is purely pattern based and does not attempt to account for population dynamics.

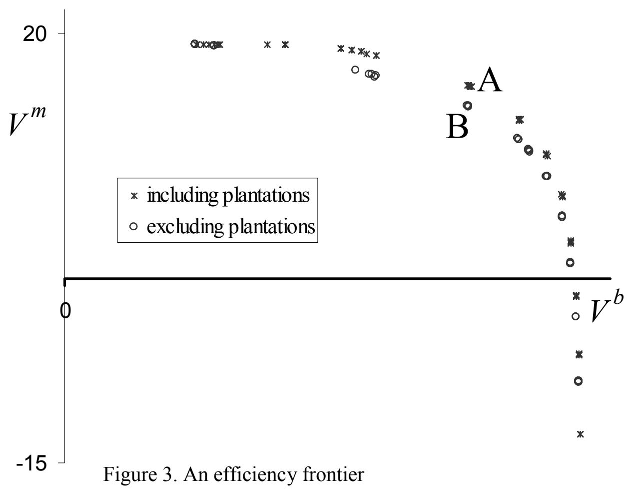

Figure 3. An efficiency frontier

Figure 3 illustrates an efficiency frontier between the commodity production and biological landscape values. To simplify the interpretation the carbon sequestration, water quantity and salinity shadow prices have been set to zero. The commodity production value is measured in millions of dollars and includes the present value of the agricultural and forestry plantation activities less the boundary and revegetation management costs. All the allocations generated for a range of shadow prices for the biological landscape value are plotted. Some of these are slightly inside the frontier although these examples of inefficient optimisation also illustrate that these inefficiencies are probably not significant. Figure 3 also shows the decreasing returns to investment in biodiversity conservation.

In this example 20 per cent of the plantation area contributes towards the habitat areas in the calculation of the biological landscape value. Thus both plantations and revegetation activities are competing to supply habitat. The circles in figure 3 represent land allocations generated when the plantation activity is excluded. In this case additional habitat is supplied by the revegetation activity. For this specification including the new plantation activity allows more efficient land allocations. Figure 4 shows the two allocations marked A and B in figure 3.

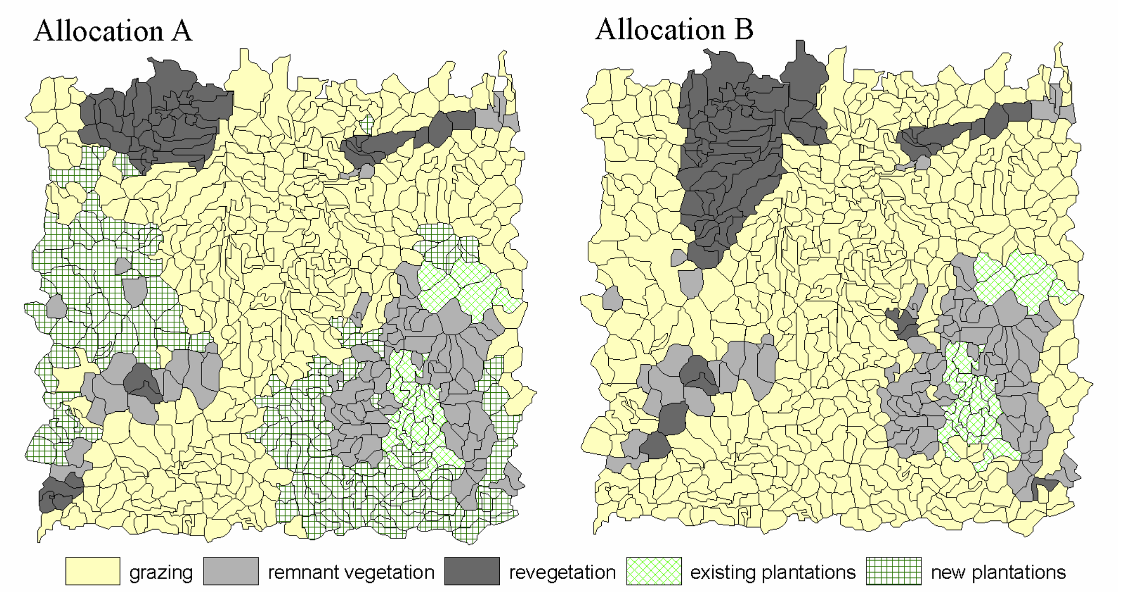

Figure 4. Two allocations

Pilot studies

Two desktop studies are being developed for the Lachlan and Hawkesbury/Nepean Catchment Management Authorities (CMAs) in New South Wales. The Lachlan study focuses on the Upper Lachlan and Lachlan Slopes sub-catchments, which are in the Eastern part of the Lachlan CMA, and may include up to one million hectares. This would give around 20 000 land units if the average size of about 50 hectares is used.

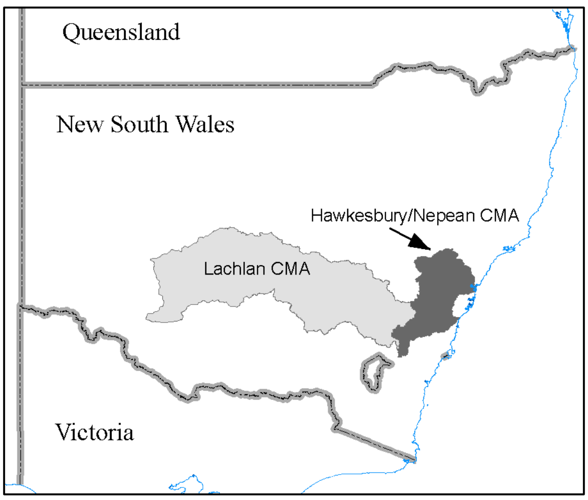

Figure 5. Pilot study CMAs

The land management changes being considered include establishing perennial pastures, changing pasture management to maintain cover and revegetation with endemic species with stock largely excluded. As well as production values and management costs, water quality in terms of stream salinity and sediment loads and biodiversity (primarily at the vegetation community level) are to be included. The main change from the prototype specification is in the addition of a model of sediment loads. This is likely to be based on the SedNet sediment budgeting approach (Lu et al. 2004). In addition to developing trade-off curves there is interest in exploring the costs of non participation of landholders in different areas.

The Hawkesbury/Nepean study focuses on a smaller area of under 100 000 ha of the Upper and Middle Cox's sub-catchments. This is in the far north-west part of the Hawkesbury/Nepean CMA. In comparison with the Lachlan study different potential erosion remediation works will also be included and stream salinity will not be modelled. A key output of this study will be advice concerning prioritising and targeting of the remediation works. The smaller area of this study should permit a comparison of the costs and benefits of using smaller land units.

These studies will be based on existing data held by or accessible by the NSW CMAs and state government departments. If the pilots prove to be useful the plan is to build the MOSAIC capabilities into tools available for use by other CMAs. The specification of the environmental metrics and the modelling of the biophysical linkages between management changes and the environmental benefits will also be developed in consultation with the NSW agencies.

References

Aerts, J., Van Herwijnen, M., Janssen, R. and Stewart, T 2005, 'Evaluating spatial design techniques for solving land-use allocation problems', *Journal of Environmental Planning and Management*, vol. 48, no. 1, pp. 121-142.

Bell, R. and Heaney, A. 2000. *A basin scale model for assessing salinity management options: Model documentation*, ABARE Technical Working Paper 2000.1, Canberra.

Davey, S. 2000, personal communication, Bureau of Rural Sciences, Canberra.

Eberbach, P. 1998, Salt affected soils: their cause, management and cost. In Agriculture and the Environmental Imperative, Edited by J. Pratley and A. Robertson, CSIRO, Melbourne, pp70-97.

Goesch, T. and Hanna, N. 2002, ‘Efficient use of water: role of secure property rights’, *Australian Commodities*, vol. 9, no. 2, pp. 372 – 84.

Hajkowicz, S., Perraud, J.M., Dawes, W. and DeRose, R. 2005, ‘The strategic landscape investment model: a tool for mapping optimal environmental expenditure’, *Environmental Modelling & Software*, vol. 20, pp. 1251–1262.

Kirkpatrick, S., C.D. Gelatt Jr., and M.P. Vecchi. 1983, “Optimization by Simulated Annealing”, *Science,* vol. 220, pp. 671-680.

Lu, H., Moran, C.J., Prosser, I., DeRose, R. 2004,’Investment prioritization based on broadscale spatial budgeting to meet downstream targets for suspended sediment loads’, *Water Resources Research*, vol. 40, W09501, doi:10.1029/2003WR002966.

Montgomery, C., Pollak, R., Freemark, K. and White, D. 1999, ‘Pricing biodiversity’, *Journal of Environmental Economics and Management*, vol. 38, pp. 1–19.

Polasky, S., Nelson, E., Lonsdorf, E., Fackler, P., Starfield, A. 2005, ‘Conserving species in a working landscape: land use with biological and economic objectives’, *Ecological Applications*, vol. 15, no. 4, pp.1387-1401.

Possingham, H., Ball, I., Andelman, S. 2000, Mathematical methods for identifying representative reserve networks. In *Quantitative Methods for Conservation Biology*, edited by S. Ferson and M. Bergman, Springer-Verlag, New York, pp. 291-335.

Possingham, H., Ryan, S., Baxter, J. and Morton, S. 2002, Setting Biodiversity Priorities, A Report for the Prime Minister’s Science, Engineering And Innovation Council, Department of the Environment and Heritage, Canberra (www.dest.gov.au/science/pmseic/documents/Setting_Biodiversity_Priorities.pdf).

Pressey, R.L. and Logan, V.S. 1998, ‘Size of selection units for future reserves and its influence on actual versus targeted representation of features: a case study in western New South Wales’, *Biological Conservation*, vol. 85, pp.305-319.

Productivity Commission 2003, Inquiry into impacts of native vegetation and biodiversity regulations, Transcript of Proceedings, Brisbane, Monday 28 July (www.pc.gov.au/inquiry/nativevegetation/trans/brisbane030728.pdf).

Prosser, I.,Young, B.,Rusomji, P.,Moran, C. and Hughes, A. 2001, 'Constructing river basin sediment budgets for the National Land and Water Resources Audit', Technical Report 15/01, CSIRO Land and Water, Canberra.

Sattler, P. and Creighton, C. 2002, Australian Terrestrial Biodiversity Assessment 2002, National Land and Water Resources Audit, Canberra, (audit.ea.gov.au/ANRA/vegetation/docs/biodiversity/bio_assess_contents.cfm).

Westphal, M. and Possingham, H. 2003, 'Applying a decision-theory framework to landscape planning for biodiversity: Follow-up to Watson et al.', *Conservation Biology*, vol. 17, no. 1, pp. 327–329.

Wu, JunJie and Boggess, W. 1999, 'The optimal allocation of conservation funds', *Journal of Environmental Economics and Management*, vol. 38, pp. 302–21.