

Maximizing a new quantity in sequential reserve selection

ADAM W. SCHAPAUGH* AND ANDREW J. TYRE

School of Natural Resources, Hardin Hall, 3310 Holdrege Street, University of Nebraska-Lincoln, Lincoln, Nebraska 68510, USA

Date submitted: 4 December 2012; Date accepted: 16 August 2013; First published online: 19 December 2013

THEMATIC SECTION
Spatial Simulation Models
in Planning for Resilience

SUMMARY

The fundamental goal of conservation planning is biodiversity persistence, yet most reserve selection methods prioritize sites using occurrence data. Numerous empirical studies support the notion that defining and measuring objectives in terms of species richness (where the value of a site is equal to the number of species it contains, or contributes to an existing reserve network) can be inadequate for maintaining biodiversity in the long-term. An existing site-assessment framework that implicitly maximized the persistence probability of multiple species was integrated with a dynamic optimization model. The problem of sequential reserve selection as a Markov decision process was combined with stochastic dynamic programming to find the optimal solution. The approach represents a compromise between representation-based approaches (maximizing occurrences) and more complex tools, like spatially-explicit population models. The method, the inherent problems and interesting conclusions are illustrated with a land acquisition case study on the central Platte River.

Keywords: Bayesian network, Markov decision process, reserve adequacy, reserve selection, stochastic dynamic programming

INTRODUCTION

Land acquisition is one way that conservation organizations try to cope with declines in biodiversity (Soule 1991). The expense associated with such investments means that decision-makers face a resource allocation dilemma. The challenge is analogous to portfolio optimization; management must select the subset of assets (sites) that gives the highest return (conservation value) for an acceptably-low risk (Markowitz 1952). This problem has promoted the development of systematic conservation assessment techniques, especially among organizations requiring efficient well-informed methods for spatial priority-setting (see for example Possingham *et al.* 2000).

Systematic conservation assessment techniques (hereafter referred to as ‘reserve selection methods’) generate priorities from spatial data. These priorities, complemented with an implementation strategy, can be viewed as a plan of action, or investment. The utility of such plans is often questioned (see Cowling *et al.* 2003; Faith *et al.* 2003) and conservation scientists have been criticized for not adequately considering the objectives and constraints of actual planning processes. A primary reproach (Costello & Polasky 2004) is that most methods are ‘static’ (see for example Kirkpatrick 1983; Kirkpatrick & Harwood 1983; Margules *et al.* 1988; Pressey & Nicholls 1989; Pressey & Tully 1994; Possingham *et al.* 2000), namely the assumption is that once the assessment is complete, the resulting plan can be executed immediately. Conservation organizations regularly face financial and political imperatives that render this assumption invalid. For instance, they cannot buy what is not for sale. Land tenure is just one reason why many conservation plans take time to execute, which makes reserve selection a sequential decision-making process (papers studying sequential reserve selection include Possingham *et al.* 1993; Costello & Polasky 2004; Snyder *et al.* 2004; Haight *et al.* 2005; McBride *et al.* 2005; McDonald-Madden *et al.* 2008).

A second criticism is that while the fundamental goal of conservation planning is biodiversity persistence (Pressey *et al.* 2007), most reserve selection methods (including all those sequential methods cited above) prioritize sites using occurrence data. Numerous empirical and theoretical studies (see Margules *et al.* 1994; Araujo *et al.* 2002) support the notion that defining and measuring objectives in terms of species richness (that the value of a site is equal to the number of species it contains, or contributes to an existing reserve network) can be inadequate for maintaining biodiversity in the long term.

One way to measure the impact of land use on viability is with a population model. A population model can be used to link factors such as habitat quantity and quality with a direct measure of persistence (such as extinction probability). Nicholson *et al.* (2006), for example, parameterized a set of stochastic patch-occupancy models that predicted the extinction probability of each of ten species of conservation concern. They then used simulated annealing, a relatively efficient alternative to linear programming, to find the reserve network that minimized the expected number of extinctions across all ten species. While this research represents a pragmatic step forward, simulated annealing and similar optimization algorithms assume that once the optimal reserve

*Correspondence: Dr Adam Schapaugh Tel: +1 785 317-2571
e-mail: adam.schapaugh@huskers.unl.edu

network has been identified, all sites appearing in the solution can be acquired. Consequently, the utility of this and related studies (see Calkin *et al.* 2002; Root *et al.* 2003) is still limited by the fact that they are static, assuming a one-time decision about which sites to protect. For conservation planning to be relevant, dynamic approaches are required that in some way account for population viability.

Maximizing persistence in sequential reserve selection is a non-trivial task; the challenge grows as multiple species are considered. One measure of persistence is the probability of extinction over a given time frame (Beissinger & Westphal 1998). This quantity is straightforward to estimate using a population model, and any model that can be expressed as a Markov chain can, in principle, have an objective maximized using stochastic dynamic programming (Mangel & Tier 1993). There are, however, both practical and computational limits in the context of reserve selection. First, population models require data (for example demographic data, patch colonization and extinction rates) that link land use with persistence. Because gathering such data is so costly, this criterion will only be met in a few cases (Beissinger & Westphal 1998). Second, formulating a population model as a Markov chain is computationally demanding; adding new state variables inevitably leads to large increases in the size of the state space. Combined with existing constraints on computer speed and storage capacity, Bellman's (1961) 'curse of dimensionality' can make generating exact solutions to even single-species planning problems computationally impossible. When the goal is to account for the viability of multiple species, less intensive numerical approaches are needed.

Schapaugh and Tyre (2012) described a site-assessment framework that implicitly maximizes the persistence probability of multiple species. They dodged the practical and computational limitations of population models by developing a Bayesian network to assess site quality, which assigns an expected value to a property based on conditions arrayed into a causal diagram. This represents a compromise between representation-based approaches (such as those using occurrence data) and more complex tools, like spatially-explicit population models. Here, we demonstrate how to integrate this site-assessment framework with a dynamic optimization model. We formulate the problem of sequential reserve selection as a Markov decision process and use stochastic dynamic programming to find the optimal solution. The method, problems with it, and interesting conclusions are illustrated with a land acquisition case study on the central Platte River (Nebraska, USA).

METHODS

Overview

The decision context assumes that the objective of a conservation agency is to maximize the persistence probability of multiple species. The agency affects this probability by acquiring sites through time. In doing so, the agency is

restricted to purchasing sites that have been placed on the public market voluntarily. Adding this realism reflects the possibility that site availability may be unpredictable in advance (for example, the need for a willing seller). When a site does become available, the agency faces a decision: (1) purchase the site; or (2) reject the site. Making this decision requires a way to assign value to the investment; we have discussed how it would be desirable to parameterize a set of population models (one for each species of concern) that translate site-specific characteristics into contributions to viability. To illustrate a typical data-poor scenario, however, we assumed that this is not possible. Instead, a Bayesian network was constructed that integrated correlates of persistence into a single currency, namely site quality (Schapaugh & Tyre 2012). This quantity is, in turn, an explicit measure of performance used in optimization. Our optimization framework is similar to those appearing in Costello and Polasky (2004) and, although we focused our attention on Bayesian networks, our method is applicable to any site-assessment framework that prioritizes sites based on a scoring system (as opposed to a system based on complementarity). This is one of the primary differences between our approach and related frameworks.

We modelled this problem as a Markov decision process (MDP; Bellman 1957). MDPs provide a mathematical framework for modelling sequential decision-making problems. As the name implies, MDPs are an extension of Markov chains; the difference is the addition of actions (to influence the state of the system) and rewards (giving motivation). This model for formal decision analysis is defined by the following components: an overall objective; a set of states $s \in S$; a set of actions: $d \in D$ and constraints; a state transition function; and a reward or value function: $V(\bullet)$. At each time step, the decision-maker observes the state of the system and selects an action. The state and action choice produce two results: the decision-maker receives a reward and the system transitions from one stage to the next. These transitions are not deterministic; instead, each action is represented by a transition matrix containing the probability that performing action d in state s will move the system to state s' (Putterman 1994). Using a land acquisition problem on the central Platte River as a case example, we elaborate on the components of the MDP.

Land acquisition on the central Platte River

In 1997, Nebraska, Wyoming, Colorado and the USA's Department of the Interior signed a *Cooperative Agreement for Platte River Research and Other Efforts Relating to Endangered Species Habitats along the Central Platte River, Nebraska* (Platte River Recovery Information Program 1997). This agreement was negotiated as a means to maintain and improve habitat for three threatened and endangered species: the whooping crane (*Grus americana*), interior least tern (*Sterna antillarum athalassos*) and piping plover (*Charadrius melodus*). The relevant objectives of this agreement are: (1) to improve

production (via number of nesting pairs and fledge ratios) of the two shorebird species; and (2) to increase the migratory survival of whooping cranes. These objectives serve as the desired outcomes resulting from the acquisition of 4000 hectares of habitat along a 143 kilometre reach of the central Platte River between Lexington and Chapman (Nebraska, USA).

Objective

The Cooperative Agreement was negotiated as a means to improve production of interior least terns and piping plovers and to increase the migratory survival of whooping cranes. The extent to which these objectives are met is positively related to the quality of sites that are acquired, and, as described by Schapaugh and Tyre (2012), site quality may be modelled using a Bayesian network parameterized from an inventory of site characteristics. Our objective was thus to maximize the sum of the expected values of the realized site-quality index (see Schapaugh & Tyre 2012) in the purchased sites.

States of the system

We define a state as a description of the system at a particular point in time. More specifically, a state is the minimally-dimensioned function of history relevant to the decision-making process. The term ‘minimally-dimensioned’ is included such that the state is as compact as possible, while still capturing the information needed to make a decision at time t (Boutilier *et al.* 1999). To define the states of the system, we first assumed that there were r sites to select from, each having an expected value (EV) of realized site-quality index. Then, let b be an $r \times 1$ vector with elements, $b_i = EV$ of site i , for $i = 1, \dots, r$. At any point in time, every site is unreserved and unavailable, unreserved and available, or included in the reserve network. We defined two state variables, $x(t)$ and $y(t)$, (each being an $r \times 1$ vector) that describe the state of the system at time t :

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

$$y_i(t) = \begin{cases} 1 & \text{if site } i \text{ becomes available in period } t \\ 0 & \text{otherwise.} \end{cases}$$

Note that $y_i(t)$ can be 1 if and only if $x_i(t) = 0$. The states of the system are given by the different assignments to these two vectors of state variables.

Actions and constraints

For simplicity, we restricted the feasible set of actions to include only two options: (1) purchase or (2) reject the site placed on the market. We defined the control variable z_r as an

$r \times 1$ vector where:

$$z_i(t) = \begin{cases} 1 & \text{if site } i \text{ is acquired in period } t \\ 0 & \text{otherwise.} \end{cases}$$

However, if the system enters an absorbing state (where the budget has been exhausted, see later), neither of these actions is possible, and the decision is forced to be (3) do nothing.

We introduced two constraints. First, it is only possible to purchase what is on the market. We also assumed a limited budget, that is it was only possible to acquire a limited number of sites, C . Thus, at any stage t , a site can only be acquired if $y_i(t) = 1$ and $\sum_i x_i(t) < C$. For simplicity, we did not incorporate variation in site cost.

State transitions

The state transition function constitutes a model of how the system evolves over time. We assumed that the system evolved in stages, where the occurrence of an event marks the transition from one stage to the next. The progression through stages is analogous to the passage of time; the two are identical if an action is taken at each stage and every action occupies one unit of time. The system is Markovian in that knowledge of the current state renders information about the past irrelevant to predictions of the future, that is: $\Pr(s_t | s_{t-1}, s_{t-2}, \dots, s_0) = \Pr(s_t | s_{t-1})$. We can represent a stationary Markov chain (i.e., the distribution predicting the next state is the same regardless of stage) with a single transition matrix, of size $S \times S$, where S is the number of states the system can occupy. This transition matrix, A , captures the probabilities governing the system as it moves from stage t to stage $t + 1$ (Boutilier *et al.* 1999).

Next, we focused our attention on how the system evolves given actions. At each stage and state of the process, the agency has available a feasible set of actions (namely, buy site or reject site). A transition matrix is required for each action. The transition matrices take the form $a_{ij} = \Pr(s_{t+1} = s_n | s_t = s_m, d_t = d)$. Recall that at any stage t , the state of the system s is described by two vectors, x and y . The transition matrix for each action is constructed in two parts. First, it is necessary calculate a matrix of transitions for the vector x . If the decision is made to purchase the site, the transitions are:

$$X_1(m, n) = \begin{cases} 1, & \text{if } m(x) = n(x + z) \\ 1, & \text{if } m(x) = n(x) \\ 0, & \text{otherwise} \end{cases}$$

where $m(x)$ is the vector x in state m , $n(x)$ is the vector x in state n , and $n(x + z)$ is the vector $x + z$ in state n . If the decision is made to reject the site, the transitions are:

$$X_2(m, n) = \begin{cases} 1, & \text{if } m(x) = n(x) \\ 0, & \text{otherwise.} \end{cases}$$

Second, we calculated a matrix of transitions for the vector y . To do so, we had to estimate for each site a relative likelihood, q_i , that it becomes available at stage t (q_i may be thought of as an instantaneous probability, whereupon its status in stage t does not affect its availability in subsequent stages, unless the site is purchased). For convenience, we also defined an indicator variable, I_i , where

$$I_i = \begin{cases} 0, & \text{if site } i \text{ has been acquired} \\ 1, & \text{otherwise.} \end{cases}$$

If the decision is made to purchase a site, the transitions are:

$$Y_1(m, n) = \begin{cases} \frac{q_i I_i}{\sum_{i=1}^r q_i I_i}, & \text{if } m(x) = n(x + z) \\ 0, & \text{otherwise} \end{cases}$$

If the decision is made to reject the site, the transitions are:

$$Y_2(m, n) = \begin{cases} \frac{q_i I_i}{\sum_{i=1}^r q_i I_i}, & \text{if } m(x) = n(x) \\ 0, & \text{otherwise} \end{cases}$$

The full transition matrix for each action is thus constructed from the component matrices by multiplication:

$$A_1 = X_1 \times Y_1;$$

$$A_2 = X_2 \times Y_2.$$

Rewards and solution

The problem facing the agency may be viewed as deciding which action to perform given the current state of the system. More generally, we seek a policy, π , which is defined as a mapping from the state and stage to actions, that is $\pi : s \times t \rightarrow d$. The problem formulated above is solved optimally by backward induction beginning at the end of the planning horizon (namely the beginning of stage $T + 1$). We defined a value function $V(\bullet)$ as a function mapping the state of the system into the real numbers, that is $V : S \rightarrow \mathbb{R}$. At the terminal time, the reward in each state is defined by the sum of the expected value of the realized site-quality index in the purchased sites:

$$V(T, T, s) = \sum_{i=1}^r x_i(t) b_i$$

Stepping back one period to the beginning of T , we took advantage of the fact that we know the value of endowing the future ($T + 1$) with the levels of each state variable:

$$V(t, T, s) = \max_{d \in D} \left\{ \sum_{s'} V(t + 1, T, s') A_{ds'} \right\} \quad (1)$$

where t is the current stage, s' is the current state, s' is the state at the next stage, and $A_{ds'} = \Pr(s_{t+1} = s_j | s_t = s_i, d_t = d)$. In words, maximizing actions are chosen in reverse order. At the terminal time, T , the best action in each state is selected. In $T - 1$, $V(t, T, s)$ is found by selecting the action that

Table 1 Sites, their expected values, and associated probabilities in the vector q . †EV = expected value realized site-quality index. ‡Parameterization, denoting the site-specific entries in the vector q . Each entry is a relative likelihood, which can be thought of as an instantaneous probability whereupon its status in stage t does not affect its availability in subsequent stages, unless the site is purchased.

Site	EV†	Parameterization‡	
		A	B
1	1	0.1	0.06
2	2	0.1	0.22
3	3	0.1	0.20
4	4	0.1	0.18
5	5	0.1	0.12
6	6	0.1	0.10
7	7	0.1	0.06
8	8	0.1	0.04
9	9	0.1	0.02
10	10	0.1	0.00

maximizes the expected terminal reward. These expected values are calculated by weighting all possible outcomes over the next time step by their probability of occurrence and summing the results. This process is repeated in stage $T - 2$, $T - 3$, and so on, until stage $t = 1$. This step-by-step procedure accomplishes one primary objective: it finds the set of actions that maximize the Bellman equation in Eq. (1). This set of actions is the optimal policy. For more discussion on MDPs and dynamic programming techniques, see Putterman (1994) and Mangel and Clark (2000).

Example reserve selection problem

For the purpose of demonstration, consider the following reserve selection problem: we assumed that ten sites were available for acquisition, the budget allowed for the selection of three of those sites, and the reward of each site was considered known (Table 1). We developed this example to explore two themes: (1) the importance of a finite decision period (we assumed that the agency cannot ‘hold-out’ for the best sites forever, therefore resources must be invested by the end of the decision period); and (2) to investigate how uncertainty in the distribution governing state transitions influences optimal decision-making.

We made two different assumptions about the site-specific entries in the vector q . The first (hereafter referred to as parameterization A) assumes no prior knowledge and thus, we adopted the ‘principle of indifference’ as the rule for assigning these epistemic probabilities. In this context, the principle of indifference states that if there are m sites, then each entry in the vector q should be assigned an equal probability of $1/m$. In Bayesian statistics, this would be referred to as the simplest non-informative prior. The second (hereafter referred to as parameterization B) assumes that the entries in the vector q are related to site quality. We assigned these

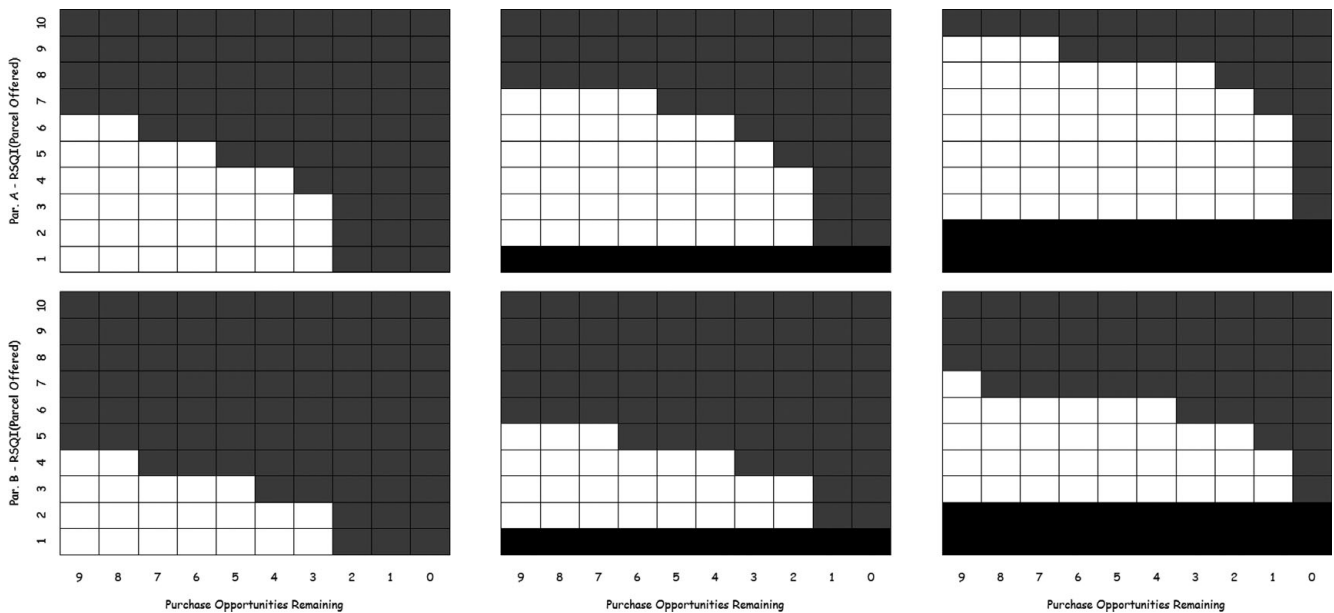


Figure 1 Optimal decision space for a portion of the state space. The optimal decision is given by the colour of the square: White = reject site; grey = purchase site; black indicates the site(s) that have already been purchased. Parameterization A (uniform probabilities) is the top row of the panel; parameterization B (weighted probabilities) is the bottom row of the panel. The left column represents a sub-portion of the state space where no sites have been purchased; the centre column represents a sub-portion of the state space where one site has been purchased; the right column represents a sub-portion of the state space where two sites have been purchased. RSQI-(parcel offered) refers to the realized site-quality index (value) of the property put up for sale.

probabilities according to the frequency distribution of the realized site-quality index on a sample of 50 properties in the Central Platte River Basin (see Schapaugh & Tyre 2012). This parameterization was selected to reflect the possibility that higher quality sites, for multiple reasons, may be harder to come by.

In the results that follow, it is cumbersome to examine the complete decision space; most of the information in the optimal policy will not be realized because, given any particular trajectory, the system will not visit much of the state space. Instead, we focused on general patterns in the results and illustrate our discussion with relevant examples.

RESULTS

We first considered the portion of the state space where no sites have been purchased. Stochastic dynamic programming simultaneously gives the optimal d_t (the decision to purchase or reject a site in stage t) and the value function one stage forward. Irrespective of which parameterization we chose, the optimal decision to purchase or reject depends on the time remaining in the decision period (Fig. 1, left column). As the number of purchase opportunities remaining decreases, the likelihood of high-quality sites becoming available decreases. The optimal decision to purchase or reject also depends on the distribution governing state transitions. The expected terminal reward for rejecting low-quality sites (namely $EV = 1, 2, 3$) was higher when we assumed uniform probabilities (Fig. 1, left column). When the best properties

are harder to come by (as in our second parameterization), the optimal policy is to become less selective.

We next considered the portion of the state space where one site has been purchased. Again, the optimal decision to purchase or reject depends on the time remaining in the decision period and on the distribution governing state transitions (Fig. 1, central column). The expected terminal reward for rejecting low-quality sites was substantially higher when we assumed uniform probabilities. The optimal policy given parameterization A associates less risk with rejecting high quality properties, especially early in the decision period, as compared to the optimal policy given parameterization B. Comparing this portion of the state space with the last (where no sites have been purchased), we found that the information in the optimal policy accounts for the number of investments that have already been made. Having already purchased one site, management can be more selective.

We finally considered the portion of the state space where two sites have been purchased. It is under these circumstances that the optimal policy is the most selective. Again, the expected terminal reward for rejecting low quality sites was substantially higher when we assumed uniform probabilities (Fig. 1, right column). With only one site left to purchase, we found that a time-independent strategy existed within the 10-purchase opportunities time horizon. With at least seven purchase opportunities remaining, the optimal strategy was generally to simply wait for the highest quality site ($EV = 10$) to become available (parameterization A). Given our second parameterization, the optimal strategy was more conservative,

nonetheless it is necessary to still hold out for a high quality site ($EV = 7, 8, 9, 10$; Fig. 1, right column).

DISCUSSION

The primary goal of this exercise has been to build upon the framework first described by Schapaugh and Tyre (2012). In this framework, a Bayesian network is used to integrate correlates of persistence for multiple species into a single currency: site quality. This quantity is, in turn, an explicit measure of performance used in optimization. In their initial presentation, Schapaugh and Tyre (2012) focused on a single acquisition; we have extended this model to the problem of acquiring multiple sites through time. We stress that this framework is not intended to be a replacement for more traditional population-level analyses when sufficient data and expertise are present. Instead, we consider it an alternative that extends the reserve selection framework to include population viability. We hope to provide a discussion of the method and results as they relate, generally speaking, to the problem of accounting for reserve adequacy in sequential reserve selection. In doing so, we discuss limitations of and alternatives to our approach and suggest directions for extending this work.

Systematic approaches to decision-making are essential, especially in this context, which involves deciding how to allocate limited resources in space and time. Reserve selection methods are simply one way of seeking the ‘biggest bang for the conservation buck’ (Moilanen *et al.* 2009). Our method illustrates how such rewards will respond to changes in key (sometimes circumstantial) factors, such as the number of purchase opportunities remaining and how the system evolves over time. The first of these, the temporal or opportunity aspect, is important regardless of how the system evolves (speaking primarily of the vector q) or the current state of the system. Our results indicate that the conservation agency should become less selective as the number of purchase opportunities remaining decreases, and accept sites with lower, but guaranteed, rewards. This strategy may result in the rejection of a property with a comparatively high realized site-quality index early in the decision period, only to later purchase one or more sites with a lower reward. The optimal policy, which is generated by explicit state-space enumeration, accounts for this possibility, and has determined that such time-dependent selectivity will result in superior expected terminal rewards (for a similar result, see McDonald-Madden *et al.* 2008).

The second key factor is the vector q ; the optimal decision to purchase or reject depends on the distribution governing state transitions. The expected terminal reward for rejecting low quality sites was higher when we assumed uniform probabilities. This is because the expected values are calculated by weighting all of the possible outcomes over the next time step by their probability of occurrence and summing the results. When high quality sites are harder to come by (as in parameterization B), the expected terminal reward for

‘holding out’ for such sites is lower because their associated probability of occurrence is also lower. The optimal policy is thus to become less selective (see Haight *et al.* 2005 for a similar result).

A simple way of including dynamics into the reserve selection problem is to assume land managers are restricted to purchasing sites that have been placed on the public market voluntarily. It should be noted, however, that every aspect of a planning problem can be a function of time (Possingham *et al.* 2009). Consider, for instance, the chance of habitat loss in sites that have not been purchased. This could be incorporated with a relatively minor modification of the existing model. The probability of being put up for sale could be reinterpreted as the probability of development in a given stage. Then a site cannot be reserved once it is developed. It would also be possible to accommodate additional complexities such as varying levels of protection (through compensation payments or conservation easements), ecological restoration, or the possibility of selling previously acquired sites. We presented a binary case, where a site was purchased or rejected, and purchases were assumed to be irreversible. ‘Un-reserving’ a site, however, is conceptually simple; the model must include an additional control variable that allows the decision-maker to sell back a previously acquired site. This idea of swapping out some areas for others is relatively new (see Fuller *et al.* 2010), even though global investments in land acquisition have slowed in recent decades (Emerton *et al.* 2006). Return-on-investment analyses should therefore receive more attention, especially considering that a ‘trade-in to trade-up’ strategy can increase the quality, and perhaps amount, of area that can be protected, with no increase in spending (Fuller *et al.* 2010).

We assumed that land values and rewards were independent of what sites had been purchased or put up for sale. It is unlikely that site availability, land values or rewards will not depend on what has happened on neighbouring or nearby sites (Costello & Polasky 2004). For example, if a site is reserved, the value of neighbouring sites may increase (*sensu* Sabbadin *et al.* 2007; Toth *et al.* 2011), which introduces spatial correlation among land prices. In this case, decisions must take into account not only the reward of the site, but also the effect that buying the site would have on land values of other potential acquisitions. Dynamic optimization models can incorporate value functions that depend on the history of decisions (the pattern of reserve selection). The complete history would be the sequence of states and actions from stage 0 to the point of interest, and would be represented by a (possibly infinite) sequence of tuples of the form $s_0, x_0, s_1, x_1, \dots, s_T, x_T$. The value function would be additive, namely the sum of the reward and/or cost function values amassed over the history of stages. Because of the probable influence of spatial correlation, making land values and/or rewards endogenous in this way would most likely increase the value of a dynamic approach (Costello & Polasky 2004).

While the framework we have presented provides a suitable conceptual foundation for sequential reserve selection, the direct implementation of dynamic programming algorithms

often proves difficult when applied to some realistically-sized problems (i.e., hundreds of sites). Our approach does not take advantage of the fact that the goal and initial states may be known; we compute the value assignments for all states at all stages. This can be wasteful from a computational perspective since optimal actions will be computed for states that cannot possibly be reached from an initial state or lead to a goal region. When the initial and goal states are known, it may be advantageous to consider the problem as a tree (or graph) search. Each state in the state space would correspond to a leaf (or node) of the tree. With the initial and goal states identified, the search proceeds forward or backward through the tree. In forward search, the initial state forms the root of the search tree. Then, each action is applied which extends the plan by one stage, generating a unique successor state (this is a new leaf node). This node can be bounded if the state it defines is already in the tree or the search may end when a state is identified as a member of the goal set (in which case a solution can be drawn from the tree). In backward search, the goal state forms the root of the search tree, and the search is expanded by adding all states that a given action would prompt the system to enter the chosen state. A state can again be pruned if it appears in the tree already. The search terminates when the initial state is added to the tree, giving a solution that can be extracted. The important point to observe is that both forward and backward searches restrict their attention to the relevant and reachable states. Both can have advantages over explicit enumeration strategies, especially if only a fraction of the state space is reachable or connected to the goal region (Boutilier *et al.* 1999).

State-based search techniques are not limited to deterministic goal-oriented problems. Knowledge of the initial state can also be exploited in stochastic settings, forming the basis of decision tree search. Each action at the initial state forms the first level of the tree. The states that result when each action is applied are placed at the second level. The third level has the actions applicable at the states at the second level, and so on. Values at the leaves of the tree are computed first and then values at successively higher levels are determined using the preceding values. This is referred to as a 'rollback' procedure and the maximizing actions form the optimal policy (Boutilier *et al.* 1999).

Unfortunately, the branching factor for stochastic problems is generally much greater than that in deterministic settings. One way around this difficulty is real-time dynamic programming (Barto *et al.* 1995). Nicol *et al.* (2010) provided the only example (to our knowledge) of this in the ecological literature. They applied an on-line sparse sampling algorithm developed by Kearns *et al.* (2002) to a hypothetical fish metapopulation, where the objective was to maximize the number of occupied patches during the management horizon. The term 'on-line' means that the policy is evaluated one step at a time based on the current state of the system. The algorithm looks ahead a defined number of steps and a rollback procedure is applied to this partially expanded search. Because the algorithm only looks at states in the vicinity of the current

state, the policy will only approximate the optimal solution. Nonetheless, the method is attractive because the running time is determined primarily by the number of look-ahead steps, which is independent of the size of the state space (Kearns *et al.* 2002).

CONCLUSION

While research on the development and refinement of reserve selection methods is accelerating, many authors have criticized conservation planners of being preoccupied with the process which has, in turn, manifested an implementation crisis (see, for example, Knight *et al.* 2006). Carefully deciding which method is best suited to the task at hand, while considering who the intended user is, is just as important to the process as evaluating decisions. Maximizing site quality represents a compromise between the use of ad hoc or generic spatial design criteria and more intensive computational tools, like spatially-explicit population models. There may be a loss in precision by using site quality as a surrogate for more direct measures of persistence. However, we believe this simplification is defensible when sufficient data, expertise, or other resources are lacking. We hope that our work will stimulate additional interest in the problem of accounting for reserve adequacy in conservation planning.

ACKNOWLEDGEMENTS

T. Buckley, T. Hefley, E. Blankenship, and two anonymous reviewers provided helpful comments on an earlier version of this manuscript. A. Schapaugh was funded by the US Army Corps of Engineers.

References

- Araujo, M.B., Williams, P.H. & Fuller, R.J. (2002) Dynamics of extinction and the selection of nature reserves. *Proceedings of the Royal Society London, Biological Series* **269**: 1971–1980.
- Barto, A.S., Bradtke, S. & Singh, S. (1995) Learning to act using real-time dynamic programming. *Artificial Intelligence* **72**: 81–138.
- Beissinger, S. & Westphal, M. (1998) On the use of demographic models of population viability in endangered species management. *Journal of Wildlife Management* **62**: 821–841.
- Bellman, R. (1957) *Dynamic Programming*. Princeton, NJ, USA: Princeton University Press.
- Bellman, R. (1961) *Adaptive Control Processes: a Guided Tour*. Princeton, NJ, USA: Princeton University Press.
- Boutilier, C., Dean, T. & Hanks, S. (1999) Decision theoretic planning: structural assumptions and computational leverage. *Journal of Artificial Intelligence Research* **11**: 1–94.
- Calkin, D., Montgomery, C., Schumaker, N., Polasky, S., Arthur, J. & Nalle, D. (2002) Developing a production possibility set of wildlife species persistence and timber harvest value. *Canadian Journal of Forest Research* **32**: 1329–1343.
- Costello, C. & Polasky, S. (2004) Dynamic reserve site selection. *Resources and Energy Economics* **26**: 157–174.

- Cowling, R.M., Pressey, R.L., Rouget, M. & Lombard, A.T. (2003) A conservation plan for a global biodiversity hotspot: the Cape Floristic Region, South Africa. *Biological Conservation* **112**: 191–216.
- Emerton, L., Bishop, J. & Thomas, L. (2006) *Sustainable Financing of Protected Areas: a Global Review of Challenges and Options*. Gland, Switzerland: The World Conservation Union (IUCN).
- Faith, D.P., Carter, G., Cassis, G., Ferrier, S. & Wilkie, L. (2003) Complementarity, biodiversity viability analysis, and policy-based algorithms for conservation. *Environmental Science and Policy* **6**: 311–328.
- Fuller, R.A., McDonald-Madden, E., Wilson, K., Carwardine, J., Grantham, H., Watson, J., Klein, C., Green, D. & Possingham, H. (2010) Replacing underperforming protected areas achieves better conservation outcomes. *Nature* **466**: 365–367.
- Haight, R.G., Snyder, S.A. & ReVelle, C.S. (2005) Metropolitan open-space protection with uncertain site availability. *Conservation Biology* **19**: 327–337.
- Kearns, M., Mansour, Y. & Ng, A.Y. (2002) A sparse sampling algorithm for near-optimal planning in large Markov decision processes. *Machine Learning* **49**: 193–208.
- Kirkpatrick, J.B. (1983) An iterative method for establishing priorities for the selection of nature reserves: an example for Tasmania. *Biological Conservation* **25**: 127–134.
- Kirkpatrick, J.B. & Harwood, C. (1983) Conservation of Tasmanian macrophytic wetland vegetation. *Proceedings of the Royal Society of Tasmania* **117**: 5–20.
- Knight, A., Cowling, R. & Campbell, B. (2006) An operational model for implementing conservation action. *Conservation Biology* **20**: 408–419.
- Mangel, M. & Clark, C.W. (2000) *Dynamic State Variable Models in Ecology: Methods and Applications*. Oxford Series in Ecology and Evolution. New York, NY, USA: Oxford University Press.
- Mangel, C.R. & Tier, C. (1993) A simple direct method for finding persistence times of populations and applications to conservation problems. *Proceedings of the National Academy of Sciences USA* **90**: 1083–1086.
- Margules, C.R., Nicholls, A.O. & Pressey, R.L. (1988) Selecting networks of reserves to maximize biological diversity. *Biological Conservation* **43**: 63–76.
- Margules, C.R., Nicholls, A.O. & Usher, M. (1994) Apparent species turnover, probability of extinction and the selection of nature reserves: a case study on the Ingelborough limestone pavements. *Conservation Biology* **8**: 398–409.
- Markowitz, H.M. (1952) Portfolio selection. *The Journal of Finance* **7**: 77–91.
- McBride, M.F., Wilson, K.A., Bode, M. & Possingham, H.P. (2005) Incorporating the effects of socioeconomic uncertainty into priority setting for conservation investment. *Conservation Biology* **21**: 1463–1474.
- McDonald-Madden, E., Bode, M., Game, E., Grantham, H. & Possingham, H. (2008) The need for speed: informed land acquisitions for conservation in a dynamic property market. *Ecology Letters* **11**: 1169–1177.
- Moilanen, A., Possingham, H. & Polasky, S. (2009) A mathematical classification of conservation prioritization problems. In: *Spatial Conservation Prioritization: Quantitative Methods and Computational Tools*, ed. A. Moilanen, K. Wilson & H. Possingham, pp. 28–42. New York, NY, USA: Oxford University Press Inc.
- Nicol, S., Chades, I., Linke, S. & Possingham, H. (2010) Conservation decision-making in large state spaces. *Ecological Modelling* **221**: 2531–2536.
- Nicholson, E., Westphal, M., Frank, K., Rochester, W., Pressey, R., Lindenmayer, D. & Possingham, H. (2006) A new method for conservation planning for the persistence of multiple species. *Ecology Letters* **9**: 1049–1069.
- Platte River Recovery Information Program (1997) Cooperative agreement for Platte River research and other efforts relating to endangered species habitats along the Central Platte River, Nebraska [www document]. URL <https://www.platteriverprogram.org/PubsAndData/Pages/ProgramLibrary.aspx>
- Possingham, H., Day, J., Goldfinch, M. & Salzborn, F. (1993) The mathematics of designing a network of protected areas for conservation, In: *Proceedings of the 12th Australian Operation Research Conference*, ed. D. Pearce, pp. 536–545. Adelaide, Australia: Adelaide University.
- Possingham, H., Ball, I. & Andelman, S. (2000) Mathematical methods for identifying representative reserve networks. In: *Quantitative Methods for Conservation Biology*, ed. S. Ferson & M. Burgman, pp. 291–309. New York, NY, USA: Springer-Verlag.
- Possingham, H., Moilanen, A. & Wilson, K. (2009) Accounting for habitat dynamics in conservation planning. In: *Spatial Conservation Prioritization: Quantitative Methods and Computational Tools*, ed. A. Moilanen, K. Wilson & H. Possingham, pp. 135–144. New York, NY, USA: Oxford University Press Inc.
- Pressey, R. & Nicholls, A. (1989) Application of a numerical algorithm to the selection of reserves in semi-arid New South Wales. *Biological Conservation* **50**: 263–278.
- Pressey, R. & Tully, S. (1994) The cost of *ad hoc* reservation: a case study in the Western Division of New South Wales. *Australian Journal of Ecology* **19**: 375–384.
- Pressey, R., Cabeza, M., Watts, M., Cowling, R. & Wilson, K. (2007) Conservation planning in a changing world. *Trends in Ecology and Evolution* **22**: 583–592.
- Putterman, M. (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. New York, NY, USA: Wiley.
- Root, K., Ackakaya, H. & Ginsberg, L. (2003) A multispecies approach to ecological valuation and conservation. *Conservation Biology* **17**: 196–206.
- Sabbadin, R., Spring, D. & Rabier, C. (2007) Dynamic reserve site selection under contagion risk of deforestation. *Ecological Modelling* **210**: 75–81.
- Schapaugh, A.W. & Tyre, A.J. (2012) Bayesian networks and the quest for reserve adequacy. *Biological Conservation* **152**: 178–186.
- Snyder, S.A., Haight, R.G. & ReVelle, C.S. (2004) A scenario optimization model for dynamic reserve site selection. *Environmental Modeling and Assessment* **9**: 179–187.
- Soule, M. (1991) Conservation: tactics for a constant crisis. *Science* **253**: 744–750.
- Toth, S., Haight, R.G. & Rogers, L. (2011) Dynamic reserve selection: optimal land retention with price feedbacks. *Operations Research* **59**: 1059–1078.