Solving conservation planning problems with integer linear programming

Hawthorne L. Beyer [a],[*], Yann Dujardin [b], Matthew E. Watts [a], Hugh P. Possingham [a]

[a] *ARC Centre of Excellence for Environmental Decisions, Centre for Biodiversity & Conservation Science, University of Queensland, Brisbane, QLD 4072, Australia*
[b] *CSIRO Ecosystem Sciences, Ecosciences Precinct, Dutton Park, QLD 4102, Australia*

ABSTRACT

Deciding where to implement conservation actions in order to meet conservation targets efficiently is an important component of systematic conservation planning. Mathematical optimisation is a quantitative and transparent framework for solving these problems. Despite several advantages of exact methods such as integer linear programming (ILP), most conservation planning problems to date have been solved using heuristic approaches such as simulated annealing (SA). We explain how to implement common conservation planning problems (e.g. Marxan and Marxan With Zones) in an ILP framework and how these formulations can be extended to account for spatial dependencies among planning units, such as those arising from environmental flows (e.g. rivers). Using simulated datasets, we demonstrate that ILP outperforms SA with respect to both solution quality (how close it is to optimality) and processing time over a range of problem sizes. For modestly sized quadratic problems (100,000 spatial units and 10 species), for example, a processing time of approximately 14 h was required for SA to achieve a solution within 19% of optimality, while ILP achieved solutions within 0.5% of optimality within 30 s. For the largest quadratic problems we evaluated processing time exceeding one day was required for SA to achieve a solution within 49% of optimality, while ILP achieved solutions within 0.5% of optimality in approximately one hour. Heuristics are conceptually simple and can be applied to large and non-linear objective functions but unlike ILP, produce solutions of unknown quality. We also discuss how ILP approaches also facilitate quantification of trade-off curves and sensitivity analysis. When solving linear or quadratic conservation planning problems we recommend using ILP over heuristic approaches whenever possible.

© 2016 Elsevier B.V. All rights reserved.

1. Introduction

Systematic conservation planning (SCP) describes the process of identifying and preserving areas of conservation value (Gaston et al., 2002; Moilanen et al., 2009). It's goal is to ensure the long-term persistence of a wide range of biodiversity using an explicit, objective, transparent, repeatable and efficient methodology (Pressey et al., 1993). The stages involved in this process typically include quantifying conservation value spatially, setting explicit goals, identifying actions for achieving those goals, identifying combinations of actions that efficiently meet goals in the context of operational limitations (e.g. budgets), and implementing these actions (Margules and Pressey, 2000; Margules et al., 2002; Kukkala and Moilanen, 2013). It is a flexible framework that has been applied to several types of conservation problem including,

for example, protected area design (Klein et al., 2013; Beger et al., 2015), the allocation of resources to deter illegal activity (Plumptre et al., 2014), evaluating the performance of protected areas in the context of climate change (Game et al., 2011; Loyola et al., 2013), terrestrial and marine zoning with multiple zone types (Mazor et al., 2014; Runting et al., 2015), and vegetation management (Levin et al., 2013). Here, we focus on the problem of deciding where to perform actions in order to efficiently achieve conservation goals.

In SCP conservation value is quantified within a set of discrete spatial units ("planning units") that can either be arbitrary (e.g. a regular grid) or based on existing boundaries (e.g. administrative, ecological, watershed or land ownership boundaries). The value of a planning unit can be estimated in several ways depending on the problem and how much is known about the "features" of conservation concern, such as individual species, habitats or ecosystems. In the design of protected areas, for example, a common approach is to estimate the occupancy (presence or absence) of a population or species within each planning unit (e.g. Giakoumi et al., 2013; Plumptre et al., 2014), though examples of other approaches

include using species abundance estimates (Williams et al., 2014), the area of a species habitat within a planning unit (particularly relevant to irregularly sized planning units), or combining measures of both occupancy and habitat condition (Klein et al., 2013).

The simplest problem formulation pertains to the binary decision of whether to include a planning unit in the selected set or not. Alternative formulations allow for multiple actions that can be implemented within a planning unit and the problem is to identify which actions to adopt and where. In either case, the value of the planning units under each action and with respect to each feature must also be quantified, along with some measure of the cost of implementing an action. Explicit targets must then be set to achieve conservation goals. In the case of protected area design, for example, the target could be a minimum habitat area falling within protected areas, potentially with different targets for each species of conservation concern. Although persistence of conservation features is a key principal of SCP (Margules and Pressey, 2000), targets are generally not explicitly evaluated to determine probabilities of persistence or persistence times. Instead, targets are often set subjectively, through expert opinion (Levin et al., 2013), community consensus (Game et al., 2011), or are informed by legislation or policy (Giakoumi et al., 2013; Runting et al., 2015).

Usually there are insufficient resources to manage or protect all planning units, hence the need for an approach for selecting a subset of planning units for conservation purposes. This can be naturally expressed as an optimisation problem in which the goal is to achieve all targets for the least cost. Maximising the efficiency of management is important because conservation resources are scarce and achieving a high return on investment means that other resources can be allocated to conservation problems elsewhere. Inefficient management plans may also be too large and expensive to implement, and less likely to succeed in the face of competing interests (Possingham et al., 2006, p520).

Generally, these are not trivial problems for which optimal solutions can be found using complete enumeration or heuristics. The value of a planning unit is conditional upon the set of other selected planning units (the issue of "complementarity", Margules and Pressey, 2000), so planning units cannot be independently ranked. For anything other than the smallest problems (perhaps a few tens of planning units) the number of permutations of planning units is too large to be enumerated and other strategies are required to identify solutions. Many conservation planning problems involve thousands of planning units and sometimes hundreds of conservation features (species, habitats, ecosystem services, etc.).

There are two main approaches to solving optimisation problems of this type. First, integer linear programming (ILP), which minimises or maximises an objective function (a mathematical equation describing the relationship between actions and outcomes) subject to a set of constraints and conditional on the decision variables (the variables corresponding to the selection of actions to implement) being integers. Second, solutions can be found using heuristic methods such as simulated annealing (SA; Kirkpatrick et al., 1983), which iteratively, stochastically explore the state-space of the decision variables. There are numerous other heuristics (e.g. ranking procedures, genetic algorithms, and mixtures of these approaches) that could also be used. Here, we focus on SA because it is the most widely used heuristic in the conservation planning literature in the form of the conservation planning software Marxan (Ball et al., 2009; Watts et al., 2009) and, unlike deterministic heuristics such as ranking, it is possible that SA could find an optimal solution to any problem.

The discussion about the relative merits and disadvantages of linear programming versus heuristics in conservation planning spans more than two decades (Cocks and Baird, 1989; Underhill, 1994; Church et al., 1996; Pressey et al., 1997; Rodrigues and Gaston, 2002; Önal, 2003). The key issues in this debate include the quality of the solution (efficiency), the size or complexity of the problem that can be addressed, and the computing time required to find a solution. The main concern with heuristics is that there is no guarantee of the quality of the solutions as it is possible for these approaches to find local rather than global minima solutions, and there is no measure of how far from optimality the solution is (Underhill, 1994; Önal, 2003). In contrast, ILP is guaranteed of finding an optimal solution or a solution guaranteed to be within a specified shortfall ("gap") of the optimum.

The main concern with ILP, on the other hand, is that it cannot be used to solve highly non-linear or complex objective functions and it is sometimes impractical for solving large problems. It is often straightforward to linearise quadratic objective functions using a combination of additional state variables and constraints (e.g. Billionnet, 2011, 2013), thereby facilitating optimisation using ILP. But it is often impractical to linearise objective functions that include more complex components, such as ecological dynamics. Perhaps the greatest advantage of SA is that it can be used to find feasible solutions to these more complex, non-linear objective functions (e.g. Westphal et al., 2007). Further, tests carried out 20 years ago showed the limitations of ILP even on non-linear problems (Pressey et al., 1997).

Here, we describe how common conservation planning problems can be linearised so that efficient solutions can be found expediently using ILP. We describe the benefits that ILP methods provide with regard to quantifying trade-offs, flexibility in problem formulation and sensitivity analysis. We demonstrate that ILP methods consistently outperform SA with respect to both processing time and solution quality across a wide range of problem sizes. Our work is an improvement over commonly used heuristic approaches as it provides a performance guarantee and finds higher quality solutions considerably faster. Given the manifest benefits of an ILP framework for solving conservation planning problems we recommend using ILP when possible and heuristics only when necessary because of recent advances in algorithms and computing power.

2. Theory

2.1. ILP formulations of conservation planning problems

Although there are numerous variations in the way that conservation planning optimisation problems have been posed (reviewed in Rodrigues et al., 2000; Williams et al., 2004, 2005), many of these problems are derived from the "reserve selection problem" (RSP), which attempts to represent each of K features to a specified threshold while minimising a measure of cost:

$$\begin{aligned} \text{min} \quad & \sum_{i=1}^{N} c_i x_i \\ \text{subject to} \quad & \sum_{i=1}^{N} r_{ik} x_i \geq T_k, \, k \in K \\ & x_i \in \{0, 1\}, \, i \in N \end{aligned} \tag{1}$$

where x_i is a binary decision variable determining whether planning unit i is selected (1) or not (0), and c_i represents the cost of planning unit i or, if the objective is simply to select the smallest number of planning units, then $c_i = 1$ for every i. The parameter r_{ik} is the contribution of planning unit i to feature k and T_k is the minimum target to be achieved for feature k among all planning units. Because the objective function is linear with respect to the decision variables the RSP is straightforward to solve as an ILP problem.

The RSP can be extended to accommodate further objectives relating to the spatial arrangement of planning units in order to

facilitate solutions in which the selected planning units are more aggregated or connected. Such extensions involve the addition of quadratic expressions to the objective function (e.g. the interaction of two decision variables: $x_i x_j$), which are problematic for ILP because the objective function is then non-linear with respect to the decision variables. The key to solving such problems using ILP is to linearise these quadratic terms, which is straightforward in the case of binary decision variables. Specifically, the quadratic term $x_i x_j$ ($x \in \{0, 1\}$) can be linearised in an ILP framework by replacing it with a new decision variable z_{ij} and implementing the following additional constraints (Billionnet, 2007):

$$
\begin{aligned}
z_{ij} - x_i &\leq 0 \\
z_{ij} - x_j &\leq 0 \\
z_{ij} - x_i - x_j &\geq -1
\end{aligned}
\tag{2}
$$

The first two of these constraints ensure that z_{ij} cannot be 1 unless both x_i and x_j are also 1, and the third constraint ensures that z_{ij} is exactly 1 if both x_i and x_j are also 1. The process of linearisation thus involves the addition of both decision variables and constraints. In practice, only a subset of these constraints needs to be implemented explicitly depending on whether the objective function is minimised or maximised and the sign of the quadratic term because the process of minimisation (or maximisation) inherently ensures some of the constraints are achieved. For example, in the case of minimisation of a negative term the decision variable must be forced to be less than a specified value but would not need to be forced to be greater than a specified value as this is achieved by the minimisation, hence only the first two constraints would be required. The simplicity of this linearisation technique belies its profound implications for solving conservation planning problems in an ILP framework (Williams et al., 2005; Billionnet, 2007, 2013).

2.2. Linearisation of the Marxan objective function

Marxan (Ball et al., 2009; Watts et al., 2009) is commonly used conservation planning software. It solves a form of the RSP whereby planning units are selected to meet targets for a minimum total cost. It includes an optional penalty for the selection of non-adjacent planning units thereby providing a mechanism to control the degree of aggregation among selected units. This penalty can also be used to facilitate selection of non-adjacent planning units that are connected through ecological, biophysical or social processes, for example those reflecting the benefits or disadvantages driven by the dispersal of juveniles or pollutants (Hermoso et al., 2011; Klein et al., 2012; Makino et al., 2013; Beger et al., 2015).

Specifically, the problem that Marxan solves is:

$$
\begin{aligned}
\min \quad & \sum_{i=1}^{N} c_i x_i + b \sum_{i=1}^{N} \sum_{j=1; i \neq j}^{N} x_i (1 - x_j) v_{ij} \\
\text{s.t.} \quad & \sum_{i=1}^{N} r_{ik} x_i \geq T_k, k \in K \\
& x_i \in \{0, 1\}, i \in N
\end{aligned}
\tag{3}
$$

where c_i is the cost of selecting site i, N is the number of planning units, K is the number of features (e.g. species) and x_i is the binary decision variable that determines whether a site is selected ($x_i = 1$) or not ($x_i = 0$). The objective function includes a cost penalty for selecting non-adjacent planning units based on a property quantifying the spatial relationship between two units (v_{ij}), such as the length of the shared boundary between them, and a scaling parameter b that is adjusted to control the strength of the penalty thereby influencing the aggregation of planning units in the solution (Watts et al., 2009). The constraints ensure that minimum targets (T_k) are

met for each of k features of interest, where r_{ik} is the value or contribution of unit i to feature k. There is considerable flexibility in the implementation of variables c, v and r which means this formulation can be used innovatively to solve many variations of the RCP problem.

Marxan does not strictly enforce constraints. Instead, it includes the constraints in the objective function using a "shortfall penalty" function, an additional penalty that is incurred whenever a target is not met (Watts et al., 2009). The premise of this approach is that even configurations of planning units that do not meet all of the targets may still have value, thereby providing a way of finding reasonable solutions even if there are no solutions that meet all targets. When all targets are met the objective function simplifies to that in Eq. (3). The shortfall penalty is not straightforward to implement in an ILP framework. Instead, we advocate implementing the objective function above and solving the problem over a range of target values, thereby explicitly describing the trade-off between the targets and objective values.

The first term in the objective function is linear with respect to the decision variables x. The second term, however, is non-linear (quadratic) with respect to the decision variable (this is clearer if we rewrite the expression $x_i(1 - x_j)v_{ij}$ as $x_i v_{ij} - x_i x_j v_{ij}$). The term $x_i v_{ij}$ can be removed by adding $b \sum_{j}^{N} v_{ij}$ to c_i (as a pre-processing step), and $x_i x_j$ can be linearised as described above. Specifically, for a negative quadratic term in a minimisation problem the first two constraints in Eq. (2) must be implemented.

The linearisation of each quadratic term involves the addition of one decision variable (z_{ij}) and two constraints. In the worst case scenario this could result in a total of $N + (N-1)^2$ decision variables and $2(N-1)^2$ constraints in addition to the structural constraints defining the targets. However, the linearisation terms can be omitted whenever $v_{ij} = 0$, which often applies to all non-neighbouring (or otherwise disconnected) planning units. In fact, in many applications the matrix v is sparse resulting in the addition of few constraints relative to the worst-case. Nevertheless, the dimension of the problem can still increase rapidly with N, which is why ILP software may be difficult to apply to very large quadratic problems (millions of planning units).

The Marxan objective function allows for asymmetric penalties for non-neighbouring (or disconnected) planning units, i.e. if x_i is selected the penalty for not selecting x_j can be different than the penalty for not selecting x_i if x_j is selected. In the context of ILP the Marxan objective function can be expressed more efficiently as:

$$
\min \sum_{i=1}^{N} c_i x_i - b \sum_{(i,j) \in E} x_i x_j v_{ij}
\tag{4}
$$

where E defines the set of neighbouring planning units. Here, $x_i x_j$ is 1 when $x_i = x_j = 1$ and 0 otherwise. If the matrix v is symmetric the set of neighbours E is defined according to the condition $i < j$, thereby ensuring the above expression is evaluated once for each pair of neighbours. If the matrix v is asymmetric then the set E is defined for each combination of i and j.

A detailed case study of how this approach can be used to linearise the Marxan With Zones objective function (Watts et al., 2009), in which multiple actions can be implemented within a planning unit and the problem is to identify which actions to adopt and where, is provided in the Appendix B.

2.3. ILP formulations of constraints for enforcing spatial dependencies among planning units

Spatial dependencies among planning units often arise as a result of underlying ecological, social or geophysical processes that affect the features of management concern. The implication of these

effects is that it may not be permissible to select one planning unit without also selecting a neighbouring planning unit, or a set of other planning units (e.g. river or coastal planning problems). Conversely, it could also be necessary to prevent neighbouring planning units from being selected (e.g. to prevent incompatible actions from occurring in adjacent units).

In an ILP framework constraints can be used to enforce spatial dependencies in the selection of planning units, and provide exact control over these dependencies (i.e. dependencies specific to each planning unit). They can also be used to enforce both uni- and bi-directional dependencies.

2.3.1. Directional, conditional dependency between planning units

To ensure that planning unit b is selected only if unit a is also selected, the following constraint is implemented:

$$x_b - x_a \leq 0$$

This constraint is directional because it does not prevent unit a from being selected if b is unselected. Importantly, this constraint can be repeated among many planning units to enforce more complex spatial dependencies. Consider a case where planning units are arranged in sequence along a linear feature, such as a river, and the flow direction of the river determines the spatial dependencies. Implementing the above constraint for each neighbouring pair of planning units will ensure that all planning units upstream of any given unit must also be selected if that unit is selected:

$$x_b - x_a \leq 0$$
$$x_c - x_b \leq 0$$
$$x_d - x_c \leq 0$$
$$x_e - x_d \leq 0$$

A more complex problem involving directional, conditional dependencies occurs when a planning unit can have multiple upstream neighbours, such as an inland planning unit bordering multiple coastal planning units or a planning unit have m units above it in a watershed. The following constraint will ensure that planning unit a is selected only if *at least one* upstream neighbouring unit is selected:

$$\sum_{i=1}^{m} x_i - x_a \geq 0$$

where $\{1 \ldots m\}$ defines the set of all the upstream neighbours to planning unit a. In contrast, the following constraint will ensure that planning unit a is selected only if *all* upstream neighbouring units are selected:

$$\sum_{i=1}^{m} x_i - m x_a \geq 0$$

2.3.2. Non-directional, conditional dependency among neighbouring planning units

One way of preventing isolated selected planning units is to make the selection of each unit conditional on the selection of a certain number of its neighbours. The following constraint ensures that a planning unit can be selected only if *at least n* neighbouring units are also selected:

$$\sum_{i=1}^{m} x_i \geq n x_a$$

where $\{1 \ldots m\}$ defines the set of all the neighbours of planning unit a. Conversely, the following constraint ensures that a planning unit can be selected only if *at most n* neighbouring units are also selected:

$$\sum_{i=1}^{m} x_i \leq m + (n - m) x_a$$

It may be desirable to prevent selected planning units from being clumped when the actions is a service that is intended to be widely distributed. The following constraint ensures that a planning unit is selected only if none of its' m neighbouring units are selected:

$$\sum_{i=1}^{m} x_i \leq m(1 - x_a)$$

2.4. Approaches to facilitating aggregation, compactness and connectivity

Planning unit selections resulting from simple objective functions often result in solutions that are highly fragmented and widely dispersed, yet spatial aggregation of planning units may be desirable for both ecological and management reasons. The ecological justification for aggregation often relates to the 'single large or several small' (SLOSS) debate (Diamond, 1975), species-area relationships and population viability. Metapopulation theory predicts that fewer, larger reserves will maximise the metapopulation capacity while an intermediate number of reserves will maximise time to extinction (Ovaskainen, 2002), though it is not clear how well these rules hold for ecosystems or communities rather than single species. Aggregations of planning units may also reduce edge and fragmentation effects that impact the conservation value of solution, or establishment and management costs.

We distinguish between aggregation and compactness. The former refers to the frequency of selection of adjacent planning units and increasing aggregation corresponds to a decrease in the number of spatially disjoint planning units. Compactness refers to the dispersion of selected planning units (how spread out the planning units are in space) and increasing compactness corresponds to reducing the total area within which selected planning units occur. The cost penalty for selecting non-adjacent planning units in the Marxan objective function determines the degree of aggregation among planning units but has a limited effect on compactness. Conversely, minimising the maximum distance between any two selected planning units (Williams et al., 2005) increases compactness but may only weakly affect aggregation except under extreme levels of compactness.

There are several ways that aggregation and compactness can be facilitated (reviewed in Williams et al., 2005; Billionnet, 2013) and the most appropriate implementation depends on the problem being addressed. For example, if the planning units are isolated patches in space (e.g. ponds) then a method based on distances among planning units rather than shared boundary lengths would be more useful. As an alternative to the boundary modifier approach used by Marxan aggregation can also be facilitated by constraining the total area to perimeter ratio of the reserve (Ohman and Lamas, 2005). Compactness has been facilitated by minimising the sum of Euclidean distances among all selected planning units (Nalle et al., 2002), maximising the inverse sum of distances among planning units (Rothley, 1999), minimising the maximum distance between any two selected planning units (Önal and Briers, 2002), and simultaneously considering compactness and aggregation for multiple species (Wang, Önal, 2015).

One issue with these approaches is that they are non-specific in the sense that they continue to cause aggregation among

clusters of planning units that may already exceed a minimum size threshold (Smith et al., 2010) when it may only be necessary to aggregate the small and isolated planning units. This may result in a substantial loss of efficiency in the final solution. It is important to be clear on why aggregation or compactness is important and to select a method that achieves these goals in a targeted and specific way if possible. Regardless of the method selected, determining the strength of the aggregation or compactness effect is a subjective decision that can be usefully visualised by trade-off curves (example below).

The term connectivity is sometimes used in a general sense to refer to any approach that increases the frequency of selection of adjacent planning units and could, therefore, apply to both aggregation and compactness. But connectivity can also refer to the specific problem of identifying a single, contiguous, fully-connected set of planning units that meet conservation objectives (Onal and Briers, 2006; Billionnet, 2012). Aggregation, compactness and connectivity all involve quadratic objective functions that can be linearised for implementation in an ILP framework. Linearisation involves the addition of both decision variables and constraints, thereby increasing the size of the problem in proportion to the number of quadratic terms, thereby requiring more computer memory (RAM) and processing time to solve the problem.

Unlike compactness and connectivity, which typically involve quadratic terms between all pairs of planning units (e.g. see Billionnet, 2013), spatial aggregation only involves quadratic terms between adjacent planning units so results in a relatively small increase in the problem size when linearised. Aggregation among non-adjacent planning units connected through ecological or biophysical processes often also applies to a subset of all possible pairs of planning units. The practical significance of this is that aggregation can be included in ILP problems with relatively large numbers of planning units (e.g. 1E6 planning units; example below) while problem sizes are much more limited for ILP solutions to compactness and full connectivity problems. Onal and Briers (2006) solve a full connectivity problem with 391 planning units and 118 species and with the more efficient formulation of (Billionnet, 2012) this could likely be expanded to several thousand planning units. Conversely, heuristics such as SA can be used to find solutions to non-linear objective functions without incurring these costs of linearisation, though the quality of those solutions relative to the optimum is unknown.

2.5. Balancing trade-offs among multiple objectives

Objective functions can contain multiple objectives (also referred to as "criteria") that may not share the same units. The simplest approach to combining multiple criteria with different units in a single objective function is the "scalarisation technique", in which additional parameters control the relative weighting among the criteria. The weights can be adjusted by the decision maker to balance the contribution of the criteria. For example, in the Marxan objective function the cost objective might be measured as an area while the boundary penalty term has arbitrary units. The relative contribution of these two criteria is controlled by the parameter b (Eq. (3)).

In general the different criteria are at least partially conflicting, implying that not every criterion can be optimised simultaneously. These trade-offs result in a Pareto frontier describing the set of every best compromise solution in the sense that every point of this set is optimal according to a specified set of preferences (relative weights) among the criteria. The role of the decision-maker is to determine the relative importance of the criteria. A strong approach to informing this subjective decision is to evaluate the objective function across a range of weights and to plot the

trade-off. We note that this approach applies regardless of whether ILP or SA is used to solve the problem.

3. Methods

We illustrate the relative performance of ILP and SA with respect to solution quality and computational time across a range of problem sizes (1E3, 1E4, 1E5 and 1E6 planning units, 10 species targets), for one linear and one quadratic problem (Eqs. (1) and (3) respectively). The contribution of each planning unit to each species (r_{ik}) was a random variable drawn from a normal distribution (mean 0, s.d. 5), with all negative values truncated to 0. Thus, for each species, approximately half of the planning units had no conservation value for that species. Targets for each species (T_k) were set at $0.3\sum_{i=1}^{N}r_{ik}$ for every k. The cost of preserving a planning unit (c_i) was a random variable drawn from a uniform distribution in the range [100, 10,000]. For the quadratic problem the penalty for selecting non-adjacent units (v_{ij}; Eq. (3)) was set to 200.

The ILP was parameterised to stop processing when a gap ≤0.5% was achieved (i.e. when the current best solution was within 0.005 times the guaranteed lower boundary of the optimal solution). Marxan was run with three levels of replicates (the number of times the SA algorithm is independently repeated, with the best solution selected among all replicates) and total number of iterations among all temperature steps: 10 replicates and 1E6 iterations (the default), 100 replicates and 1E7 iterations, and 1000 replicates and 1E8 replicates. The quality of the SA solution (the 'gap') can be quantified when the same problem is solved using ILP. All processing occurred sequentially on a single desktop computer (4-core Intel i7 3.4 GHz processor) with 16 GB RAM that was not used for any other purpose during the trial to ensure comparable processing times.

We illustrate balancing trade-offs among two objectives and the value of evaluating a range of objective weights using an ILP implementation of the Marxan objective function. The simulated data included 1E5 planning units with value data for 10 species and land costs (Fig. A.1). The species and cost datasets were generated using the RandomFields library in R (R Development Core Team, 2015) (Appendix C). Each planning unit had up to four neighbours in the cardinal directions (edge units had less than four). Targets for each species were set at 25% of the total value all planning units contributed to that species (i.e. $\forall k \in K, T_k = 0.25\sum_{i}^{N}r_{ik}$). The relative weight of the cost and aggregation components of the objective function were varied to describe the trade-off and maps of the solutions.

We evaluate the sensitivity of these solutions to uncertainty in the species data by decreasing all the species values by 10% and repeating the analysis based on these degraded values. Clearly, if the value of each planning unit to each species decreases then many more planning units are required to meet the targets and the total solution cost will be considerably higher than the expected (mean) approach described above.

The SA and ILP problems were solved using Marxan (version 2.4.3) and Gurobi (version 5.6.2) respectively (see Appendix C for details). Marxan uses simulated annealing to stochastically explore solution space. Gurobi is proprietary software that uses several algorithms, including simplex and branch and bound algorithms, to solve linear programming problems and is guaranteed of finding optimal solutions given enough time.

4. Results

ILP found higher quality solutions in less processing time compared to SA over the full range of problem sizes and for both linear and quadratic models (Fig. 1a and b). As the problem size increased, the quality of the SA solution degraded substantially unless the number of replicates and iterations (the parameters that can be
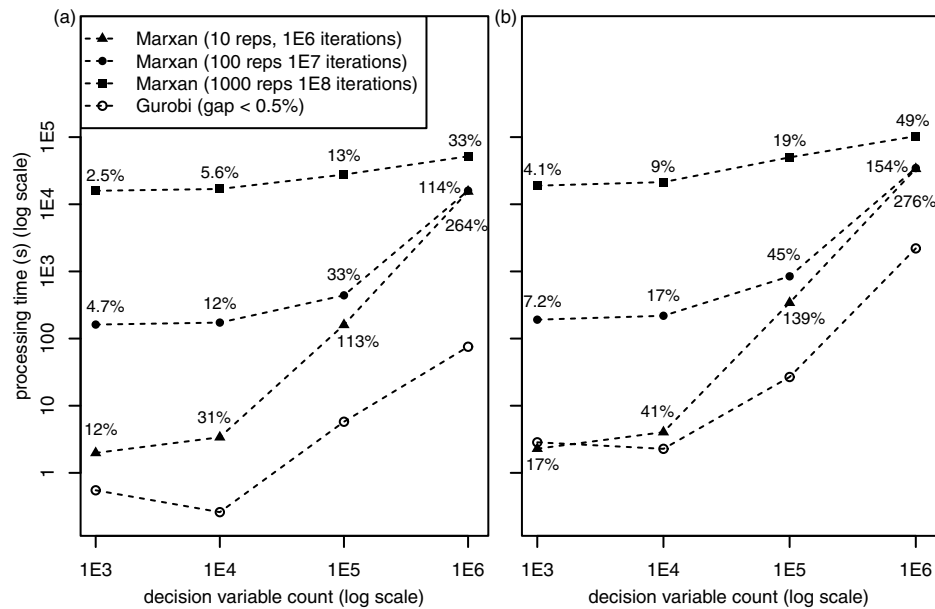
Fig. 1. Comparisons of processing times (*y*-axis) and solution quality (text labels) between integer linear programming software (Gurobi; open symbols) and simulated annealing software (Marxan; solid symbols) over a range of problem sizes (*x*-axis). Two problems were evaluated: a simple, linear reserve selection problem (a), and Marxan's quadratic objective function (b). The quality of the solution is quantified by the "gap" between the solution and the guaranteed lower bound of the optimal solution, and is expressed as a percentage of this value (thus a gap of 0% indicates the optimal solution). Simulated annealing was implemented with three levels of numbers of iterations and replicates to illustrate the trade-off between processing time and the quality of the solution. The increase in processing time for the smallest ILP solutions in each figure may result from additional automatic pre-processing that occurs within Gurobi that is omitted for larger problems.

used to tune the SA algorithm) were increased, with associated marked increases in processing time. The three implementations of the SA algorithm we evaluated never matched the solution quality that was achieved by ILP. The increase in processing time for the smallest ILP solutions may result from additional automatic pre-processing that occurs within Gurobi that is omitted for larger problems.

These results should also be interpreted in the context of what constitutes an important gap. SA consistently found "good" solutions to optimisation problems in approximately 12–24 h of processing time. The primary significance of the differences in processing times between these methods is the opportunities that fast solutions provide for quantifying trade-off curves and facilitating multi-objective optimisation in near real-time, thereby altering the way in which optimisation is used in the decision making process.

In the second analysis there is a clear trade-off between increased aggregation and the solution cost (Fig. 2, first panel). As the weighting of the boundary penalty (parameter *b*) increases more planning units are required in order to meet the targets (Fig. 2, second panel), which subsequently increases the solution cost. At the same time the fragmentation of the selected planning units decreases, quantified in Fig. 2 (third panel) as the number of contiguous groups of planning units that share a boundary with another planning unit in a cardinal direction. While the increase in solution cost is approximately linear over the range of values of *b* we have evaluated the decrease in fragmentation is non-linear, indicating that per-unit fragmentation reduction becomes increasingly expensive as aggregation increases. The distribution and fragmentation of selected planning changed little in the sensitivity analysis solutions (Fig. 2, maps) indicating that in this case study the solutions were fairly insensitive to uncertainty in the species data.

5. Discussion

With advances in algorithms and processing power integer linear programming has become a flexible and efficient framework

for identifying optimal or near-optimal solutions to conservation planning problems. Three benefits of ILP over simulated annealing are faster computational speeds, better solution qualities, and guaranteed quantification of solution quality. These advantages further facilitate the efficient and precise description of trade-offs among objectives, analysis of sensitivity of solutions to parameter uncertainty, and the exploration of multi-objective problems interactively in near real time. The primary disadvantage of heuristic methods is that the solution quality is unknown and the quality of the solution is sensitive to the way that heuristic algorithms are tuned (e.g. the number of iterations at each temperature step and the total number of replicates in SA).

Comparisons in processing time between ILP and SA are often disingenuous as they fail to account for differences in the quality of solutions found. For both approaches processing times increase as the dimension of the problem increases (e.g. as the number of planning units and actions increases) and, for a given problem size, there is a trade-off between the processing time and the quality of the solution. ILP is usually characterised by a rapid increase in the quality of the solution in the early stages (often the first few seconds) followed by a period where the rate of further improvement is much slower. If allowed to run to completion ILP will find an optimal solution, though this can take considerable time. SA can be tuned in a variety of ways in an attempt to balance processing time with the quality of solution that is likely to be found. However, the only way of definitively quantifying the quality of the solution in an absolute sense is by comparing it to an ILP solution, so it is often not clear how to tune the SA algorithm. Indeed, one of the most pertinent criticisms of SA is that there is no practical guidance on how to parameterise the algorithm for different problem sizes to ensure a consistent quality of solution.

ILP processing times may, however, increase substantially for larger or more constrained problems and a key advantage of heuristic methods is that they can find solutions to complex, non-linear problems that would be difficult or impossible to implement in an ILP framework. Adding complexity to objective functions to make them more relevant to real-world problems could have a profound
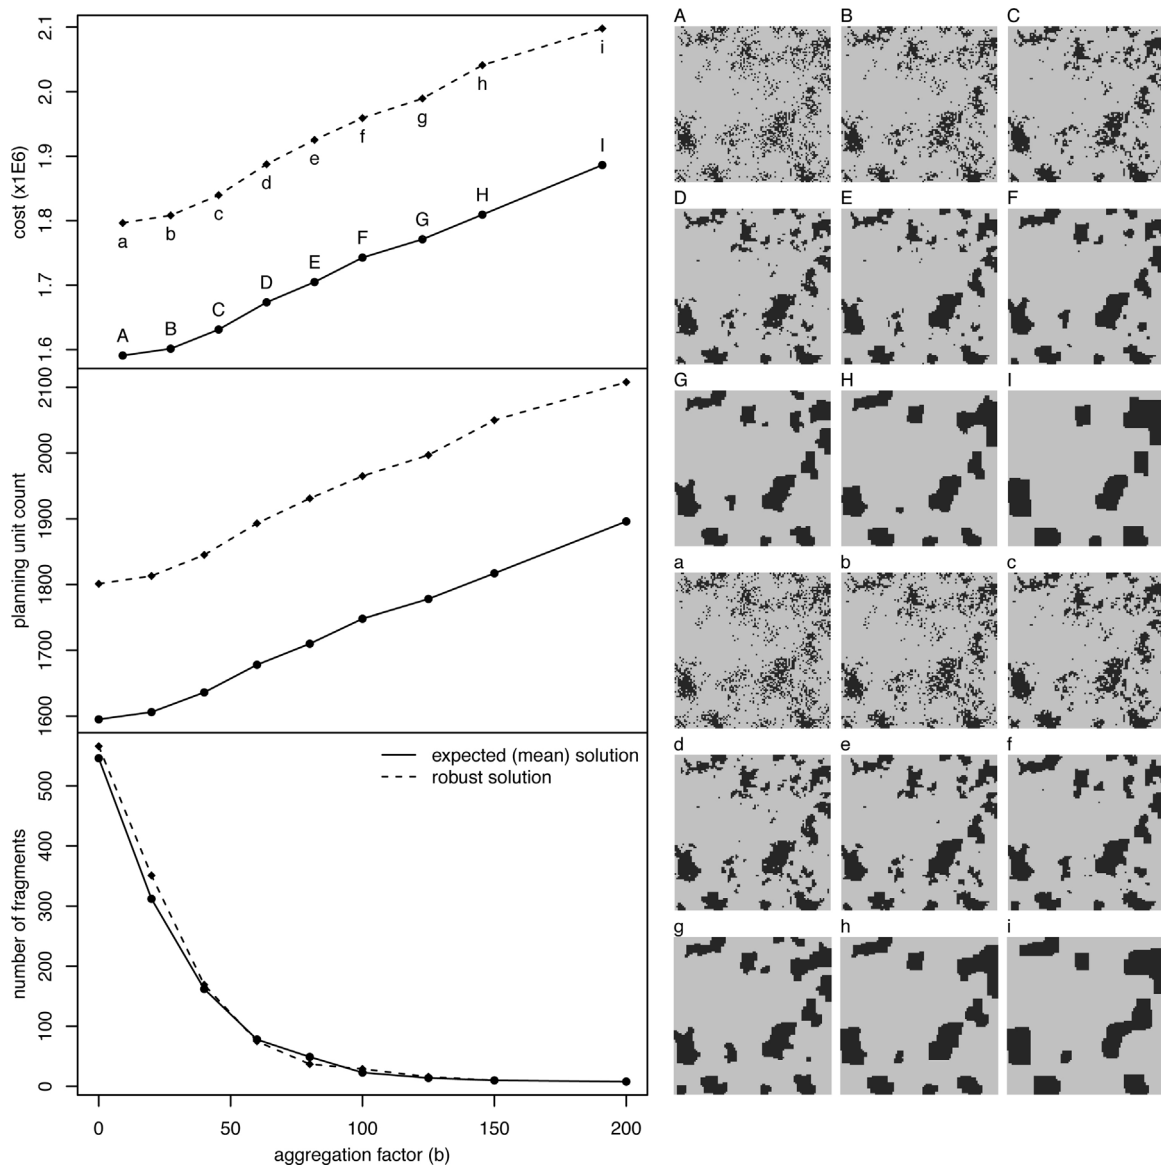
Fig. 2. Trade-offs between two objectives, cost and aggregation of planning units, in a reserve selection problem solved using integer linear programming to within 0.5% of optimality. As the weighting of the aggregation objective increases (*x*-axis), the cost of the total solution increases (top) as a result of the increased number of planning units required to meet targets (middle). The number of spatially discrete groups of planning units decreases non-linearly as aggregation increases (bottom). The sensitivity analysis solutions (dashed line) cost considerably more than the expected value optimisation (solid line) but are more robust to uncertainty in the value of planning units to each species target. Labels A–H and a–h represent solutions to the expected value and sensitivity analysis problems respectively at specific points along the trade-off curves.

affect on the solutions found. For this reason the discussion of whether SA or ILP finds 'better' solutions must also consider how well the objective function represents the problem being addressed (Moilanen, 2008). Often, problems are simplified to a linear (or linearisable) form, or the dimensions of the problem are reduced so that a solution can be identified expediently. Such structural simplifications in the formulation could dramatically alter the solution found (Langford et al., 2011). The degree to which the optimal solution to the simplified problem also represents a good solution to the complex, real-world problem is generally not known and not evaluated. This is a major shortcoming of systematic conservation planning and it is essential to explicitly validate the effectiveness of conservation plans through monitoring during and following implementation.

One reason that conservation research may fail to trigger changes to management is that it is too narrow in scope, considering only a few dimensions of a problem. In contrast, real-world management must balance numerous competing objectives and

interests. The optimal solution to a problem that considers only a few conservation objectives may provide little insight into solutions to problems that simultaneously consider multiple objectives, including social, economic and planning objectives. It is incumbent on conservation scientists to work closely with a broad range of stakeholders to bridge the research-implementation gap (Knight et al., 2008). Multi-objective optimisation is a particularly powerful framework for identifying consensus compromises among decision-makers with different priorities.

The assumptions that conservation planning problems are based on can also be obstacles to implementation. Although they are often not stated explicitly, common assumptions inherent in reserve selection problems are that all planning units are available for selection, that the costs associated with the selection of each planning unit are not dynamic, and that the benefits associated with the selection of a planning unit are guaranteed, are not dynamic, and are independent of what happens in other planning units. In reality, planning units may only become available for

purchase or management asynchronously, the costs of planning units change, and the value of planning units to conservation is both uncertain and often subject to long lags (e.g. as a result of forest restoration). Considerable progress has been made in explicitly incorporating these sorts of complexities into an ILP framework (Haight et al., 2000; Costello and Polasky, 2004; Snyder et al., 2004; Toth et al., 2011).

Targets are often formulated as constraints because it is straightforward to solve a single objective function subject to constraints. The issues with implementing targets as constraints, however, are that: (i) it may result in a problem that is insoluble; (ii) the targets are often subjectively defined and treating them as strict thresholds belies the uncertainty associated with these values; and (iii) these strict constraints constrain the solution space, potentially precluding more efficient solutions that only just miss one or more targets. Future applications could adopt multi-objective optimisation approaches (e.g. the interactive method; Ehrgott, 2005) in which a set of objective functions must be minimised (or maximised) and the targets are implemented as objectives, not constraints. Treating constraints as objectives in a multi-objective optimisation framework would allow decision-makers to more fully explore the solution space by explicitly evaluating the importance and consequences of trade-offs among objectives. Whether targets should be constraints or objectives will often depend on the social, political and economic context of the problem.

There may be considerable uncertainty in estimates of the costs of actions and values of planning units. The risk in ignoring such uncertainties is that the optimal solution identified may ultimately violate some constraints, hence becoming an infeasible solution, or may be far from optimal (Bertsimas and Sim, 2004). Robust optimisation can be used to identify solutions while explicitly accounting for this uncertainty. Solutions are described as 'robust' when they remain feasible and near-optimal regardless of how the data changes. For example, worst-case optimisation (Chinneck and Ramadan, 2000) involves solving the problem while guaranteeing that no constraint is violated whatever the realisation of the parameters and only requires that values are known within an interval. It is particularly relevant when we want to avoid the risk of failing to achieve the constraints, but solutions can be costly compared to an expected value approach (Birge and Louveaux, 2011).

An extension of worst-case optimisation involves specifying a different level of risk of violation for each constraint (Bertsimas and Sim, 2004), which is beneficial as it allows the decision-maker to assume greater risk with some constraints, thereby reducing the cost of the robust solution. Another alternative approach, min–max regret, consists of determining an optimum solution which minimises the maximum regret that could be realised in the face of parameter uncertainty (Averbakh and Lebedev, 2005), where regret is defined as the difference in the benefit between the adopted solution and any other solution. Although conceptually appealing the problem can be difficult to solve. Finally, one could also attempt a robust multi-objective approach, where each target is considered as an objective (as discussed above) with uncertain parameters (Gaspar-Cunha and Covas, 2008; Ehrgott et al., 2014; Ide et al., 2014).

6. Conclusions

This paper provides guidance on the conceptual and practical aspects of implementing a variety of conservation planning problems in an ILP framework. The three key benefits of ILP over simulated annealing are faster computational speeds, better solution qualities, and guaranteed quantification of solution quality. Despite these advantages and widespread application in operations research, the adoption of integer linear programming methods in conservation has been slow because early trials proved

unsatisfactory. When solving linear and quadratic conservation planning problems we recommend the use of exact methods (e.g. integer linear programming) when possible, and heuristics only when necessary.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.ecolmodel.2016.02.005.

References

Averbakh, I., Lebedev, V., 2005. On the complexity of minmax regret linear programming. Eur. J. Oper. Res. 160 (1), 227–231.

Ball, I.R., Possingham, H.P., Watts, M., 2009. Marxan and relatives: software for spatial conservation prioritisation. In: Moilanen, A., Wilson, K., Possingham, H. (Eds.), Spatial Conservation Prioritisation: Quantitative Methods and Computational Tools. Oxford University Press, Oxford, UK, pp. 185–195, chap. 14.

Beger, M., McGowan, J., Treml, E.A., Green, A.L., White, A.T., Wolff, N.H., Klein, C.J., Mumby, P.J., Possingham, H.P., 2015. Integrating regional conservation priorities for multiple objectives into national policy. Nat. Commun. 6, http://dx.doi.org/10.1038/ncomms9208.

Bertsimas, D., Sim, M., 2004. The price of robustness. Oper. Res. 52 (1), 35–53.

Billionnet, A., 2007. Optimisation discrète: De la modélisation à la résolution par des logiciels de programmation mathématique. Dunod.

Billionnet, A., 2011. Solving the probabilistic reserve selection problem. Ecol. Model. 222 (3), 546–554.

Billionnet, A., 2012. Designing an optimal connected nature reserve. Appl. Math. Model. 36 (5), 2213–2223.

Billionnet, A., 2013. Mathematical optimization ideas for biodiversity conservation. Eur. J. Oper. Res. 231 (3), 514–534.

Birge, J.R., Louveaux, F., 2011. Introduction to Stochastic Programming. Springer.

Chinneck, J.W., Ramadan, K., 2000. Linear programming with interval coefficients. J. Oper. Res. Soc. 51 (2), 209–220.

Church, R.L., Stoms, D.M., Davis, F.W., 1996. Reserve selection as a maximal covering location problem. Biol. Conserv. 76 (2), 105–112.

Cocks, K.D., Baird, I.A., 1989. Using mathematical-programming to address the multiple reserve selection problem – an example from the Eyre Peninsula, South-Australia. Biol. Conserv. 49 (2), 113–130.

Costello, C., Polasky, S., 2004. Dynamic reserve site selection. Resour. Energy Econ. 26 (2), 157–174.

Diamond, J., 1975. The island dilemma: lessons of modern biogeographic studies for the design of natural reserves. Biol. Conserv. 7 (2), 129–146.

Ehrgott, M., 2005. Multicriteria Optimization. Springer-Verlag.

Ehrgott, M., Ide, J., Schoebel, A., 2014. Minmax robustness for multi-objective optimization problems. Eur. J. Oper. Res. 239 (1), 17–31.

Game, E.T., Lipsett-Moore, G., Saxon, E., Peterson, N., Sheppard, S., 2011. Incorporating climate change adaptation into national conservation assessments. Global Change Biol. 17 (10), 3150–3160.

Gaspar-Cunha, A., Covas, J.A., 2008. Robustness in multi-objective optimization using evolutionary algorithms. Comput. Opt. Appl. 39 (1), 75–96.

Gaston, K.J., Pressey, R.L., Margules, C.R., 2002. Persistence and vulnerability: retaining biodiversity in the landscape and in protected areas. J. Biosci. 27 (4), 361–384.

Giakoumi, S., Sini, M., Gerovasileiou, V., Mazor, T., Beher, J., Possingham, H.P., Abdulla, A., Cinar, M.E., Dendrinos, P., Gucu, A.C., Karamanlidis, A.A., Rodic, P., Panayotidis, P., Taskin, E., Jaklin, A., Voultsiadou, E., Webster, C., Zenetos, A., Katsanevakis, S., 2013. Ecoregion-based conservation planning in the mediterranean: dealing with large-scale heterogeneity. PLOS ONE 8 (10), e76449.

Haight, R.G., Revelle, C.S., Snyder, S.A., 2000. An integer optimization approach to a probabilistic reserve site selection problem. Oper. Res. 48 (5), 697–708.

Hermoso, V., Linke, S., Prenda, J., Possingham, H.P., 2011. Addressing longitudinal connectivity in the systematic conservation planning of fresh waters. Freshw. Biol. 56 (1), 57–70.

Ide, J., Koebis, E., Kuroiwa, D., Schoebel, A., Tammer, C., 2014. The relationship between multi-objective robustness concepts and set-valued optimization. Fixed Point Theory Appl. 83, http://dx.doi.org/10.1186/1687-1812-2014-83.

Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P., 1983. Optimization by simulated annealing. Science 220 (4598), 671–680.

Klein, C.J., Jupiter, S.D., Selig, E.R., Watts, M.E., Halpern, B.S., Kamal, M., Roelfsema, C., Possingham, H.P., 2012. Forest conservation delivers highly variable coral reef conservation outcomes. Ecol. Appl. 22 (4), 1246–1256.

Klein, C.J., Tulloch, V.J., Halpern, B.S., Selkoe, K.A., Watts, M.E., Steinback, C., Scholz, A., Possingham, H.P., 2013. Tradeoffs in marine reserve design: habitat condition, representation, and socioeconomic costs. Conserv. Lett. 6 (5), 324–332.

Knight, A.T., Cowling, R.M., Rouget, M., Balmford, A., Lombard, A.T., Campbell, B.M., 2008. Knowing but not doing: selecting priority conservation areas and the research-implementation gap. Conserv. Biol. 22 (3), 610–617.

Kukkala, A.S., Moilanen, A., 2013. Core concepts of spatial prioritisation in systematic conservation planning. Biol. Rev. 88 (2), 443–464.

Langford, W.T., Gordon, A., Bastin, L., Bekessy, S.A., White, M.D., Newell, G., 2011. Raising the bar for systematic conservation planning. Trends Ecol. Evol. 26 (12), 634–640.

Levin, N., Watson, J.E.M., Joseph, L.N., Grantham, H.S., Hadar, L., Apel, N., Perevolotsky, A., DeMalach, N., Possingham, H.P., Kark, S., 2013. A framework for systematic conservation planning and management of mediterranean landscapes. Biol. Conserv. 158, 371–383.

Loyola, R.D., Lemes, P., Nabout, J.C., Trindade-Filho, J., Sagnori, M.D., Dobrovolski, R., Diniz-Filho, J.A.F., 2013. A straightforward conceptual approach for evaluating spatial conservation priorities under climate change. Biodivers. Conserv. 22 (2), 483–495.

Makino, A., Klein, C.J., Beger, M., Jupiter, S.D., Possingham, H.P., 2013. Incorporating conservation zone effectiveness for protecting biodiversity in marine planning. PLOS ONE 8 (11), e78986.

Margules, C.R., Pressey, R.L., 2000. Systematic conservation planning. Nature 405 (6783), 243–253.

Margules, C.R., Pressey, R.L., Williams, P.H., 2002. Representing biodiversity: data and procedures for identifying priority areas for conservation. J. Biosci. 27 (4), 309–326.

Mazor, T., Possingham, H.P., Edelist, D., Brokovich, E., Kark, S., 2014. The crowded sea: incorporating multiple marine activities in conservation plans can significantly alter spatial priorities. PLOS ONE 9 (8), e104489.

Moilanen, A., Wilson, K.A., Possingham, H.P., 2009. Spatial Conservation Prioritisation: Quantitative Methods and Computational Tools. Oxford University Press, Oxford.

Moilanen, A., 2008. Two paths to a suboptimal solution – once more about optimality in reserve selection. Biol. Conserv. 141 (7), 1919–1923.

Nalle, D.J., Arthur, J.L., Montgomery, C.A., Sessions, J., 2002. Economic and spatial impacts of an existing reserve network on future augmentation. Environ. Model. Assess. 7 (2), 99–105.

Ohman, K., Lamas, T., 2005. Reducing forest fragmentation in long-term forest planning by using the shape index. For. Ecol. Manage. 212 (1-3), 346–357.

Önal, H., 2003. First-best, second-best, and heuristic solutions in conservation reserve site selection. Biol. Conserv. 115 (1), 55–62.

Önal, H., Briers, R.A., 2002. Incorporating spatial criteria in optimum reserve network selection. Proc. R. Soc. B-Biol. Sci. 269 (1508), 2437–2441.

Onal, H., Briers, R.A., 2006. Optimal selection of a connected reserve network. Oper. Res. 54 (2), 379–388.

Ovaskainen, O., 2002. Long-term persistence of species and the SLOSS problem. J. Theor. Biol. 218 (4), 419–433.

Plumptre, A.J., Fuller, R.A., Rwetsiba, A., Wanyama, F., Kujirakwinja, D., Driciru, M., Nangendo, G., Watson, J.E.M., Possingham, H.P., 2014. Efficiently targeting resources to deter illegal activities in protected areas. J. Appl. Ecol. 51 (3), 714–725.

Possingham, H.P., Wilson, K.A., Andelman, S.J., Vynne, C.H., 2006. Protected areas: goals, limitations and design. In: Groom, M.J., Meffe, G.K., Carroll, C.R. (Eds.), Principles of Conservation Biology. Sinauer Associates, Sunderland, pp. 509–533.

Pressey, R.L., Humphries, C.J., Margules, C.R., Vanewright, R.I., Williams, P.H., 1993. Beyond opportunism – key principles for systematic reserve selection. Trends Ecol. Evol. 8 (4), 124–128.

Pressey, R.L., Possingham, H.P., Day, J.R., 1997. Effectiveness of alternative heuristic algorithms for identifying indicative minimum requirements for conservation reserves. Biol. Conserv. 80 (2), 207–219.

R Development Core Team, 2015. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria http://www.R-project.org.

Rodrigues, A.S., Cerdeira, J.O., Gaston, K.J., 2000. Flexibility, efficiency, and accountability: adapting reserve selection algorithms to more complex conservation problems. Ecography 23 (5), 565–574.

Rodrigues, A.S.L., Gaston, K.J., 2002. Optimisation in reserve selection procedures – why not? Biol. Conserv. 107 (1), 123–129.

Rothley, K.D., 1999. Designing bioreserve networks to satisfy multiple, conflicting demands. Ecol. Appl. 9 (3), 741–750.

Runting, R.K., Meijaard, E., Abram, N.K., Wells, J.A., Gaveau, D.L.A., Ancrenaz, M., Posssingham, H.P., Wich, S.A., Ardiansyah, F., Gumal, M.T., Ambu, L.N., Wilson, K.A., 2015. Alternative futures for Borneo show the value of integrating economic and conservation targets across borders. Nat. Commun. 6, http://dx.doi.org/10.1038/ncomms7819.

Smith, R.J., Di Minin, E., Linke, S., Segan, D.B., Possingham, H.P., 2010. An approach for ensuring minimum protected area size in systematic conservation planning. Biol. Conserv. 143 (11), 2525–2531.

Snyder, S.A., Haight, R.G., ReVelle, C.S., 2004. Scenario optimization model for dynamic reserve site selection. Environ. Model. & Assess. 9 (3), 179–187.

Toth, S.F., Haight, R.G., Rogers, L.W., 2011. Dynamic reserve selection: optimal land retention with land-price feedbacks. Oper. Res. 59 (5), 1059–1078.

Underhill, L.G., 1994. Optimal and suboptimal reserve selection algorithms. Biol. Conserv. 70 (1), 85–87.

Wang, Y., Önal, H., 2015. Optimal design of compact and connected nature reserves for multiple species. Conserv. Biol., DOI: 10.1111/cobi.12629.

Watts, M.E., Ball, I.R., Stewart, R.S., Klein, C.J., Wilson, K., Steinback, C., Lourivald, R., Kirchera, L., Possingham, H.P., 2009. Marxan with zones: Software for optimal conservation based land- and sea-use zoning. Environ. Model. Softw. 24 (12), 1513–1521.

Westphal, M.I., Field, S.A., Possingham, H.P., 2007. Optimizing landscape configuration: a case study of woodland birds in the Mount Lofty Ranges, South Australia. Landsc. Urban Plan. 81 (1–2), 56–66.

Williams, J.C., ReVelle, C.S., Levin, S.A., 2004. Using mathematical optimization models to design nature reserves. Front. Ecol. Environ. 2 (2), 98–105.

Williams, J.C., ReVelle, C.S., Levin, S.A., 2005. Spatial attributes and reserve design models: a review. Environ. Model. Assess. 10 (3), 163–181.

Williams, R., Grand, J., Hooker, S.K., Buckland, S.T., Reeves, R.R., Rojas-Bracho, L., Sandilands, D., Kaschner, K., 2014. Prioritizing global marine mammal habitats using density maps in place of range maps. Ecography 37 (3), 212–220.