Diversity and Distributions (2002) **8**, 297–306

BIODIVERSITY RESEARCH

Place prioritization for biodiversity reserve network design: a comparison of the SITES and ResNet software packages for coverage and efficiency

CHRIS KELLEY, JUSTIN GARSON, ANSHU AGGARWAL and SAHOTRA SARKAR*
Biodiversity and Biocultural Conservation Laboratory, Program in the History and Philosophy of Science, Waggener 316, University of Texas at Austin, Austin, TX 78712–1180, USA

Abstract. The place prioritization problem in conservation biology is that of establishing a sequentially prioritized list of places on the basis of biodiversity content. Such a list can then be used to select reserve networks that are designed to be fully representative of the biodiversity of an area as efficiently as possible (for instance, with minimum area or cost). The usual goal is the representation of all chosen biodiversity surrogates up to or beyond a required target, or to the greatest available extent. The purpose of this paper is to compare the respective performances of two place prioritization software packages, SITES and ResNet, on four datasets (distributions of termite genera in Namibia, breeding bird species in the Falkland Islands/Islas Malvinas, vertebrate species in Texas and flora and fauna species that are at risk in Québec), to determine their respective merits. The two software packages implement radically different algorithms: SITES is based on a simulated annealing procedure for finding (local) optima; ResNet uses an algorithm based on rarity and complementarity. This analysis indicates that the rarity-complementarity based algorithm of ResNet surpasses the simulated annealing approach of SITES with respect to time and completeness. SITES, however, contains other features that are useful in conservation planning. Ways in which the two packages can be used together effectively are suggested.

Key words. conservation planning, place prioritization, reserve design, ResNet, SITES, site selection.

INTRODUCTION

One of the major tasks of conservation biology is that of selecting reserve networks that are fully representative of the biodiversity of an area. In real-world scenarios, however, full representation is often impossible, because not all places of biological interest can be conserved. Consequently, the practical problem must be reduced to that of the representation of biodiversity surrogates up to a specified target as efficiently as possible, with efficiency being measured by criteria such as the minimum total area or cost. Places must therefore be prioritized before a final set is selected for reservation or other conservation action. Place prioritization algorithms are becoming a critical tool in the practice of biodiversity conservation. The purpose of this note is to compare the respective performances of two software packages implementing radically different place prioritization algorithms, SITES and ResNet, on four datasets to determine which is preferable for use in systematic conservation planning, or how they may be used together profitably. Both of these packages can be downloaded freely from the web (see Ball, 2000 for the design of SITES;

* Corresponding author: E-mail: sarkar@mail.utexas.edu

Andelman *et al.*, 1999 for downloading SITES; and Aggarwal *et al.*, 2000 and Sarkar *et al.*, 2002 for downloading ResNet). The algorithms incorporated in these packages, although not the specific packages themselves, have been compared previously by Csuti *et al.*, 1997; see also Pressey *et al.* (1997).

The algorithms implemented in SITES and ResNet are radically different. SITES implements the widely used simulated annealing algorithm (Kirkpatrick *et al.*, 1983; Otten *et al.*, 1989). Using a cost equation which computes the sum of a set of weighted components, the procedure iteratively selects new places, compares resulting changes to the cost equation, and attempts to minimize the total cost. A unique feature of SITES is the large number of optional parameters which can be incorporated into the calculation of the cost, and which require user-defined weights for the search process. These parameters include a boundary length modifier (which allows the user greater control over the level of adjacency or 'clumping' of cells), conservation categories and a penalty factor for under-representing the target. SITES has been utilized in several conservation projects in the United States, including regions in southern California and the Columbia Plateau (Andelman & Fagan, 2000).

ResNet implements a hierarchically controlled, rarity–complementarity-based selection algorithm (Margules *et al.*, 1988; Aggarwal *et al.*, 2000). The place prioritization procedure is initialized by the selection of a first cell on the basis of rarity, richness or from a user-defined set of cells. (The last option is relevant if a reserve network is to be built by extension of an existing set of reserves.) ResNet then employs an iterative procedure that selects places on the basis of rarity; then, to break ties between cells with equally rare surrogates, it uses complementarity (that is, it selects the cell that contains the most surrogates not already represented up to the specified target among the cells that have already been chosen). Should ties still remain, ResNet optionally uses adjacency (when a cell adjacent to already selected cells is preferred to one that is not); and finally selects by lexical order (that is, randomly). The user also has the option of incorporating a redundancy check, as well as area and cost constraints. ResNet has been applied to datasets from Québec, Namibia and Islas Malvinas/Falkland Islands (Sarakinos *et al.*, 2001; Sarkar *et al.*, 2001), and has been used in conjunction with the GAP Analysis Project in Texas (Sarkar *et al.*, 2001).

Although the two algorithms can be compared using a variety of different criteria, the radically disparate strategies used by the two algorithms narrow the ambit of quantitative evaluation. Only completeness in surrogate coverage (meeting the specified target levels), efficiency (doing so while selecting the minimum possible area) and computational time are parameters common to the performance of both algorithms. These were used as the most important indices of comparison in this analysis. Efficiency is particularly important because both SITES and ResNet use heuristic algorithms. While both attempt to produce a solution that meets a specified target, neither guarantees that a better (that is to say, more efficient) solution is not possible. [In other words, both algorithms guarantee that a local optimum solution is obtained; neither guarantees that the global optimum is found. For a discussion of algorithms that find the global optimum, but require much more computational time, and for a discussion of the relative merits of algorithms that are minor variants of the one implemented in ResNet, see Csuti *et al.* (1997).]

METHODS

Both packages were run on an IBM-compatible personal computer. The first dataset consists of 1250 data cells (each 0.25×0.25 degrees of longitude and latitude) in Namibia containing records of the presence of 33 termite genera (cf. Sarkar *et al.*, 2002 for a full description of the dataset). The second dataset consists of 254 cells (each 10 km × 10 km) in the Falkland Islands/Islas Malvinas, containing records of the presence of 65 breeding bird species (cf. Sarkar *et al.*, 2002 for a full description of the dataset). The third dataset consists of 1183 cells (each a hexagon of about 649 km^2 produced by the GAP Analysis Project) in Texas containing records of the presence of 655 vertebrate species (cf. Sarkar *et al.*, 2001 for a full description of the dataset). The fourth dataset consists of 16 500 data cells (each 0.2×0.2 degrees of longitude and latitude) in Québec, containing records of the presence of 387 species of flora and fauna at risk distributed over 709 of the data cells (cf. Sarakinos *et al.*, 2001 for a full description of the dataset). Each algorithm was run with the target of representation of surrogates specified at 1, 5, 10, 15 and 20. Only the target of 1 is guaranteed to be met by all surrogates because each surrogate must occur at least once in the dataset. When faced with a surrogate for which the assigned target cannot be met because it does not occur that many times in the dataset, ResNet automatically revises that surrogate's target level to its total presence. We will refer to these target levels of representation as the 'initial target' and the 'revised target'. The first two datasets are computationally simple; they are also far too sparse to be of much use in any realistic policy formulation situation; the last two are much more complex and can potentially be used in practice for reserve network design. (The Québec dataset has already been used by Sarakinos *et al.*, 2001 to make policy recommendations.)

For each target of representation, SITES produces up to six output files: (i) *projectName*_best.txt ('best file'), a prioritized list of places containing all cells selected in the single most cost-effective of the runs; (ii) *projectName*_ssoln.txt ('summed solution file'); which contains the complete list of cells selected in every run where these cells are grouped into frequency categories depending on the number of runs in which they were selected (thus for our purposes there were 50 frequency categories because there were up to 50 runs — see below); (iii) *projectName*_mvbest.txt file, which contains information on the extent to which the surrogates selected in the best run meet the targets of representation; (iv) *projectName*_sum.txt file, which records summary statistics for all runs for a given target; (v) *projectName*_sen.dat file, which provides a general performance report and information on the algorithm employed; and (vi) finally, a file in which one also has the option of recording the different solutions obtained for each of the runs. The last option was not used for this analysis. The summed solution file is a unique feature of the SITES software package, intended to support the analysis of the robustness of the corresponding 'best file'. Its potential use in this context will be discussed below. (cf. Andelman *et al.*, 1999 for a more detailed explanation of the various data reporting options offered in the SITES software package.)

ResNet produces up to five output files. Four of these output files consist of lists of prioritized places. These five files are: (i) for every application of ResNet, the first output file ('basic output file') records the solution arrived at without the use of adjacency or a check for redundancy; (ii) if the 'adjacency' option is selected, a second output file records the corresponding solution when adjacent cells are given preference as explained above; (iii) if the 'redundancy' option is selected, a third output file records the corresponding solution when redundant cells are removed; (iv) if the procedure described in (iii) is modified with a 'secondary adjacency' checking option, which results in the retention of redundant cells which are adjacent to a selected non-redundant cell, a fourth output file records the corresponding solution; and (v) finally, for every application of ResNet, a log file is produced that records, for each cell contained in the prioritized list, the rule under which it was selected, and the percentage of surrogates which have cumulatively met the revised target of representation. The 'adjacency' and 'secondary adjacency' modifications were not used in this analysis.

SITES was tested according to the recommended minimum standards (50 runs, 10^6 iterations per run) (Andelman *et al.*, 1999). To ensure testing integrity, SITES was also subjected to a longer 'control run' using 100 runs and 2×10^6 iterations at a target level of 10 to see if increasing the size of the computation improved performance. The result was negligibly different (and, in fact, slightly worse). For all targets of representation, ResNet was initialized by rarity which is known to produce the most efficient results (Sarkar *et al.*, 2002), and run with and without the removal of redundant cells.

The distance between two solutions was defined as the ratio of the number of unshared selected cells to the total number of selected cells (in both solutions). Thus, if all cells are shared, the distance is 0; if no cells are shared, the distance is 1.

In the first comparison, reported in Table 1, the relative efficiency of the two place prioritization algorithms was compared by setting an initial target of representation for each surrogate and comparing the total number of cells selected by each algorithm in attaining the revised target. ResNet automatically guarantees complete surrogate coverage by the end of a run of the

Table 1a Data from Namibia. Results of the direct comparison of SITES and ResNet. SITES was forced to achieve complete surrogate coverage by assigning a severe penalty for incomplete coverage. For more details, see text

1	2	3	4	5	6	7	8
1	33/33	6	6	0.333	3	< 0 : 00 : 01	0 : 07 : 05
5	26/33	29	29	0.349	2	< 0 : 00 : 01	0 : 07 : 26
10	23/33	54	54	0.241	6	< 0 : 00 : 01	0 : 07 : 15
15	21/33	73	73	0.068	12	0 : 00 : 01	0 : 07 : 10
20	19/33	85	85	0.141	1	0 : 00 : 02	0 : 07 : 16

Explanation of the columns:
1. The initial target of representation for each surrogate.
2. The number of surrogates for which the initial target can be met/the total number of surrogates.
3. The number of cells selected for revised target levels by ResNet's basic output file.
4. The number of cells selected for revised target levels by the SITES best file. (Parameters: surrogate absence penalty factor sufficiently weighted to ensure target completeness; all optional modifiers off; annealing schedule set at 10^6 iterations per run; 50 runs per target level. See text for more details.)
5. The distance between the solutions.
6. The ordinal number of the SITES run (of the 50 runs in the annealing schedule) that was selected by SITES as the best file.
7. The run time of ResNet.
8. The run time of SITES over the full annealing schedule.

program. In general, SITES does not guarantee such coverage; since it tries to optimize over other parameters besides the presence of surrogates, the program may terminate before complete coverage is attained because other parameters have been optimized sufficiently. For this analysis, complete surrogate coverage was achieved through a trial-and-error process of sufficiently weighting the surrogate absence penalty factor of the cost equation to prevent target incompleteness.

In the second comparison, SITES was run with site costs uniformly set equal to each other, and the penalty factor costs for surrogate absence were given equal weights to the cost of a site. In this method of comparison efficiency can no longer be measured simply by comparing the number of cells selected by the algorithms for a given target level. The comparison of efficiency of the two algorithms must be indirect.

Three different methods were used. ResNet was asked to terminate after selecting exactly the same number of cells as contained in the SITES 'best file' and the number of surrogates for which targets were met by each algorithm were compared. In an effort to establish more wide-ranging comparisons between the two algorithms, similar tests were conducted with two subsets of the SITES 'summed solution file': (i) the group of cells selected in every run for a given target (and thus falling within the highest frequency category, which will be referred to as the 'common set'); and (ii) the group of cells with the highest frequency measured across all runs for a given target, up to the number of cells selected by ResNet for that target (which will be referred to as the 'top set'). (In other words, if ResNet selects n cells for a given target, then the 'top set' consists of SITES' n most frequently selected cells.)

RESULTS

The numerical results are reported in Tables 1a,b–4a,b. In all Tables, initial targets of representation for each run are listed in Column 1. Column 2 lists the number of surrogates that met or exceeded the initial target level.

In the first (direct) comparison for efficiency (Tables 1a–4a), the number of cells selected by ResNet (column 3) was occasionally slightly lower than that selected by SITES (column 4) for revised targets of one and five representations, whereas for revised targets of 10, 15 and 20 representations the algorithms were almost equally efficient. Owing to the high penalty for surrogate absence imposed on the SITES algorithm (see above), ResNet and SITES both

Table 1b Data from Namibia. Results of the indirect comparison of SITES and ResNet. For more details, see text

1	2	3	4	5	6	7	8	9
1	33/33	6	6	31/33	31/33	2	22/22	32
5	26/33	29	27	27/33	27/31	9	15/11	31
10	23/33	54	44	26/33	26/30	19	16/17	29
15	21/33	73	52	24/33	24/29	18	8/12	29
20	19/33	85	71	26/33	26/31	34	9/20	31

10	11	12	13	14	15	16	17	18
7	4	0.333	2	5	55	56′	< 0 : 00 : 01	0 : 06 : 19
30	18	0.357	9	23	90	90	< 0 : 00 : 01	0 : 06 : 47
61	37	0.245	19	46	133	135′	< 0 : 00 : 01	0 : 06 : 53
81	44	0.296	18	61	135	137′	0 : 00 : 01	0 : 06 : 56
91	58	0.256	34	73	178	178	0 : 00 : 02	0 : 06 : 57

The SITES selections (column 15) were insufficient for complete target representation. The additional cells were selected using ResNet.
Explanation of the columns:
1. The initial target of representation for each surrogate.
2. The number of surrogates for which the initial target can be met/the total number of surrogates.
3. The number of cells selected for revised target levels by ResNet's basic output file.
4. The number of cells selected for the revised target levels by the SITES best file. (Parameters: surrogate absence penalty factor and site cost factor weighted equally; all optional modifiers off; annealing schedule set at 10^6 iterations per run; 50 runs per target level. See text for more details.)
5. The number of surrogates selected in the SITES best file/the total number of surrogates in the dataset.
6. The number of surrogates for which revised targets were met in the SITES best file/the number of surrogates for which the revised target levels were met when ResNet selected only the same number of cells as in the SITES best file (column 4).
7. The number of cells selected in the SITES 'common set' (that is, those grouped into the highest frequency category of 50).
8. The number of surrogates which meet revised targets in the SITES 'common set'/the number of surrogates for which the revised target levels were met when ResNet selected the same number of cells as in the SITES 'common set'.
9. The number of surrogates in the SITES 'top set' (the cells most frequently selected by SITES when it selects as many cells as are required by ResNet to meet the targets) that meet revised target.
10. The number of cells in a solution when: (i) from the SITES summed solution file, the same number of cells is selected as is required by ResNet to meet the target (that is, column 3); and (ii) then ResNet is used to extend the selected set until all targets are met.
11. The number of cells shared by the SITES best files (column 4) and the corresponding ResNet solutions (column 3).
12. The distance between the two solutions mentioned in column 11.
13. The number of cells shared by the SITES 'common sets' (column 7) and the corresponding ResNet solution (column 3).
14. The number of cells shared by the SITES 'top sets' (column 8) and the corresponding ResNet solution (column 3).
15. The total number of cells contained in the SITES summed solution file (that is, those selected at least once over all 50 runs).
16. The number of cells within the summed solution file required to attain completeness of surrogate coverage (as evaluated by running ResNet with the summed solution file as an input sets).
17. The run time of ResNet.
18. The run time of SITES over the full annealing schedule.

Table 2a Data from Falkland Islands/Islas Malvinas. Results of the direct comparison of SITES and ResNet. SITES was forced to achieve complete surrogate coverage by assigning a severe penalty for incomplete coverage. For more details, see text

1	2	3	4	5	6	7	8
1	61/65	7	7	0.143	1	< 0 : 00 : 01	0 : 12 : 53
5	53/65	26	24	0.200	4	0 : 00 : 01	0 : 13 : 11
10	49/65	53	53	0.132	1	0 : 00 : 01	0 : 13 : 37
15	45/65	78	78	0.064	1	0 : 00 : 01	0 : 13 : 47
20	40/65	100	100	0.050	1	0 : 00 : 02	0 : 13 : 27

Table 2b Data from Falkland Islands/Islas Malvinas. Results of the indirect comparison of SITES and ResNet. For more details, see text

1	2	3	4	5	6	7	8	9
1	61/65	7	7	65/65	65/65	3	57/52	64
5	53/65	26	24	65/65	65/64	14	46/47	65
10	49/65	53	46	61/65	61/58	30	49/49	63
15	45/65	78	67	57/65	57/59	49	50/53	62
20	40/65	100	85	53/65	53/59	61	48/51	59

10	11	12	13	14	15	16	17	18
8	5	0.286	3	5	12	12	< 0 : 00 : 01	0 : 12 : 05
26	20	0.200	12	21	40	40	0 : 00 : 01	0 : 13 : 03
56	41	0.167	30	45	72	74′	0 : 00 : 01	0 : 13 : 36
81	62	0.145	49	73	92	94′	0 : 00 : 01	0 : 13 : 53
109	83	0.103	61	91	105	111′	0 : 00 : 02	0 : 13 : 50

The SITES selections (column 15) were insufficient for complete target representation. The additional cells were selected using ResNet.

achieved completeness in surrogate coverage. For all targets, ResNet completed computation in approximately 1% of the time required by SITES (see columns 7, 8) for the datasets from Namibia, the Falklands/Malvinas and Québec. In the case of Texas, SITES took more time to make its selection over ResNet by a factor of at least 250 (Table 3a, columns 7, 8). Thus, both complex datasets presented severe problems for SITES with respect to the time required to produce a solution. Column 6 records the ordinal number of the run during the annealing procedure in which SITES found an optimal solution. These vary widely (from 1 to 38); that not more than 38 runs were needed in any case may suggest that a smaller number of runs than recommended (50) can be used. However, this increases the risk of not finding a solution during one annealing schedule. Since the number of runs does not seem to make a significant difference in the computation time (column 8), such a reduction of the annealing procedure is not recommended. The distance between the solutions (Table 1, column 5) is large when the target is 1 for the two larger datasets (Texas and Québec), but becomes small with bigger targets indicating that similar sets of cells were being selected by both programs. For the smaller datasets (Namibia and Falklands/Malvinas), the distance is consistently small.

In the second (indirect) comparison of the two software packages (Tables 1b–4b), ResNet consistently selected more cells per target level than the SITES 'best files' (columns 3, 4) but satisfied the (revised) targets of representation. However, SITES never achieved complete surrogate coverage except in two cases (Table 2b, column 5 with targets of 1 and 5); the performance became worse with higher targets (column 5). When

Table 3a Data from Texas. Results of the direct comparison of SITES and ResNet. SITES was forced to achieve complete surrogate coverage by assigning a severe penalty for incomplete coverage. For more details, see text

1	2	3	4	5	6	7	8
1	655/655	21	22	0.721	2	0 : 00 : 06	4 : 06 : 53
5	646/655	107*	109	0.435	12	0 : 00 : 24	4 : 14 : 09
10	641/655	200**	199	0.328	19	0 : 00 : 44	4 : 15 : 59
15	632/655	291	292	0.238	38	0 : 01 : 03	4 : 14 : 54
20	619/655	372	372	0.220	22	0 : 01 : 43	4 : 14 : 25

Redundancy checks removed * one cell; ** removed three cells.

Table 3b Data from Texas. Results of the indirect comparison of SITES and ResNet. For more details, see text

1	2	3	4	5	6	7	8	9
1	655/655	21	21	652/655	652/655	1	325/331	638
5	646/655	107*	98	646/655	646/653	10	322/331	641
10	641/655	200**	177	646/655	646/650	47	475/484	642
15	632/655	291	266	635/655	635/652	85	495/531	644
20	619/655	372	359	649/655	649/653	170	566/582	641

10	11	12	13	14	15	16	17	18
29	6	0.714	1	9	222	222	0 : 00 : 06	3 : 37 : 00
126	61	0.404	10	73	445	446	0 : 00 : 24	4 : 03 : 15
227	119	0.369	46	149	549	550′	0 : 00 : 44	4 : 13 : 46
314	204	0.268	85	247	573	574	0 : 01 : 03	4 : 17 : 09
401	296	0.190	170	319	699	659	0 : 01 : 43	4 : 18 : 51

Redundancy check removed * one cell; ** removed three cells.
The SITES selections (column 15) were insufficient for complete target representation. The additional cells were selected using ResNet.

ResNet was constrained to select the same number of cells as contained in the SITES 'best files' solution, the SITES best file achieved a higher target completeness than the ResNet solution for a target of 1; for all higher target levels, the SITES achieved less completeness (column 6) except in three cases (Table 2b, column 6 with targets of 1, 5 and 10). Nevertheless, the number of shared cells in the corresponding solutions (column 11) and the distance between the two solution sets (Table 2, column 12) indicate that the cell sets selected by both algorithms are highly similar. This is increasingly so for greater target levels, although the increase does not appear to be significant.

The cells in the SITES 'common set', which were common to all 50 runs (Tables 1b–4b, column 7) did not yield effective surrogate coverage. This is not surprising because these sets contained so few cells; as is to be expected, these cells were also selected by ResNet (which is why there is a great deal of agreement between columns 7 and 13). When ResNet was constrained to select the same number of cells as in the SITES 'common set' there were mixed results on completeness of surrogate coverage. The ResNet solutions also greatly under-represented the revised targets of representation (column 8).

While the SITES 'top sets' attained a higher level of surrogate coverage than the SITES 'common sets' for each target (compare Tables 1b–4b, columns 9 and 8), these sets were still less effective than ResNet solutions which obtain complete coverage. However, the difference does

Table 4a Data from Quebec. Results of the direct comparison of SITES and ResNet. SITES was forced to achieve complete surrogate coverage by assigning a severe penalty for incomplete coverage. For more details, see text

1	2	3	4	5	6	7	8
1	387/387	73	75	0.230	5	0 : 00 : 15	1 : 36 : 02
5	240/387	257*	258	0.076	29	0 : 01 : 11	1 : 39 : 14
10	151/387	376	376	0.027	22	0 : 01 : 40	1 : 44 : 49
15	95/387	458	458	0.031	25	0 : 02 : 06	1 : 40 : 33
20	56/387	523	523	0.027	1	0 : 02 : 17	1 : 42 : 52

* Redundancy checks removed one cell.

Table 4b Data from Quebec. Results of the indirect comparison of SITES and ResNet. For more details, see text

1	2	3	4	5	6	7	8	9
1	387/387	73	62	376/387	376/372	19	254/257	382
5	240/387	257*	212	352/387	352/354	117	262/265	380
10	151/387	376	295	329/387	329/338	205	283/268	384
15	95/387	458	375	328/387	328/347	273	279/272	382
20	56/387	523	434	330/387	330/359	316	286/293	380

10	11	12	13	14	15	16	17	18
78	57	0.156	19	65	200	111	0 : 00 : 15	1 : 34 : 58
265	200	0.147	117	243	346	281	0 : 01 : 11	1 : 38 : 52
380	287	0.145	205	369	451	393	0 : 01 : 40	1 : 43 : 01
468	362	0.131	273	445	525	484	0 : 02 : 06	1 : 40 : 55
533	418	0.126	316	487	576	554	0 : 02 : 17	1 : 42 : 21

* The redundancy check removed one cell from this set.

not appear to be significant (see column 2, column 9). The number of cells shared between the 'top sets' and those selected by ResNet are shown in Table 2, column 14.

As a final test, the entirety of the 'summed solutions' file was evaluated with respect to completeness of surrogate coverage for each target level. The number of cells contained in the 'summed solution file' is shown in Tables 1b–4b, column 15, whereas the number of those actually needed to obtain full surrogate coverage is listed in column 16. When cells were selected from these, using ResNet, to meet targets, the solutions were only slightly larger than those produced by ResNet (column 10). This comparison is important because the SITES 'summed solutions' file provides one way to use SITES and ResNet together. Finally, columns 17 and 18 once again underscore how computationally expensive SITES is when compared to ResNet.

DISCUSSION

The main advantages of rarity-complementarity based algorithms such as the one implemented in ResNet are computational speed and a guarantee of complete surrogate target coverage. SITES took over an hour to process the Québec dataset and more than four hours to process the Texas dataset (with one exception). Moreover, as will be detailed further below, the speed permits the exploration of a variety of solutions using different initializations or randomizations of the dataset. This is important for three reasons. Starting with a large number of solutions which all satisfy the target set for the biodiversity surrogates: (i) biological considerations that did not enter into the prioritization process, such as the prognosis for the biodiversity contents of a place (their 'viability') can be used to refine the solution set; (ii) other considerations such as social and economic cost can be used to find 'non-dominated' solutions (which are not worse than some other on every relevant criterion). These non-dominated solution sets can then be presented as alternatives to be perused by policy makers (see below). As has been recognized in the literature on reserve design (Faith, 1995; Sarkar *et al.*, 2001), most conservation projects require the production of several such alternative lists of prioritized places in order to represent the feasibility of trade-offs between multiple, and often conflicting, constraints on their implementation on the ground; and (iii) it is likely to produce solutions that are, purely by accident, close to the global optimum (see Csuti *et al.*, 1987 who show that an exact algorithm does marginally better than a rarity–complementarity-based algorithm even with a single initialization).

Compared to SITES, ResNet is so much faster because its algorithm iteratively adds cells singly, unlike the more global approach in the SITES algorithm which computes the cost equation for an entire set of cells at each step. In many planning scenarios, new plans, for instance, starting with a new initial set of selected areas may have to be devised very rapidly, perhaps during a planning meeting. In such situations, the several hours required by a reasonable annealing schedule in programs such as SITES compromises the utility of these programs. Another advantage of ResNet is that the cell selection is transparent, because the rule used for every selection is explicitly documented. Thus, decision-makers can know exactly what a new cell is adding to a plan.

However, an apparent disadvantage of using ResNet is that the differences in cost or area between cells are not incorporated explicitly into the iterative selection procedure. ResNet can be run using the option of terminating the selection procedure if some maximum permissible cost or area is exceeded. ResNet was written with the view that considerations other than biodiversity should not be reduced to a single parameter along with biodiversity. Multiple criteria, although important, are often incompatible and must be analysed separately. Sarkar (in preparation) generated 32 distinct solutions from the Texas dataset, all satisfying the criterion of meeting biodiversity targets, using ResNet. These 'feasible' solutions were then analysed further to find those two solutions which were 'non-dominated' with respect to social and economic costs. Potentially these can now be forwarded to planners for the last stage of selecting a reserve network. Such an extensive analysis was possible only because of the speed with which ResNet produces results; the program was run 100 times to find these solutions in less than 2 hours.

The main advantage of the SITES software package is the availability of the large number of optional parameters which allows a user to tailor its cost equation in many ways. Moreover, the assignment of user-defined weights to selection criteria, insofar as they allow cost considerations to play a role in the iterative procedure, may give SITES the ability to produce solutions that take many more relevant criteria for reserve design than ResNet. However, the problem is that many of these weights can only be assigned arbitrarily. This results in two potential problems: (i) the weights may be misleading; and (ii) it assumes that all the criteria that go into reserve design (biodiversity, social cost, economic cost, etc.) are commensurable on a single scale. (As noted earlier, this approach is rejected intentionally in ResNet.) At the technical level, another disadvantage to using SITES is that its approach to place prioritization requires back-propagation of error detection to correct weights in the algorithm, and this is a resource-intensive and imprecise procedure.

The set of cells in the 'summed solution file' provides one potential way in which ResNet and SITES can be used together. The initial set of cells can be screened once to remove those that are apparently undesirable under all the criteria of interest using SITES in a computationally intensive phase. Then ResNet can be used to achieve completeness of surrogate coverage and efficiency with different initializations in a computationally fast way. This may well be the most effective use of SITES and ResNet together. Since most conservation projects require the production of several alternative lists of prioritized places, we suggest that the two programs could be applied fruitfully together to yield a diversity of conservation scenarios for a given region. These scenarios could then be the basis for the selection of individual plans on socio-political grounds which can probably never be incorporated into precise algorithms.

ACKNOWLEDGMENTS

We thank Hugh Possingham and an anonymous referee for comments on an earlier draft.

REFERENCES

Aggarwal, A., Garson, J., Margules, C.R., Nicholls, A.O. & Sarkar, S. (2000) ResNet. Manual, V.1.1. Technical report, Biodiversity and Biocultural Conservation Laboratory, University of Texas at Austin. http://uts.cc.utexas.edu/~consbio/Cons/Labframeset.html

Andelman, S.J. & Fagan, W.F. (2000) Umbrellas or flagships: efficient conservation surrogates or expensive mistakes? *Proceedings of the National Academy of Sciences*, *USA* **97**, 5954–5959.

Andelman, S.J., Ball, I., Davis, F.W. & Stoms, D.M. (1999) SITES V.1.0, an analytic toolbox for eco-regional conservation portfolios. Technical report, The Nature Conservancy. http://www.biogeog.ucsb.edu/projects/tnc/toolbox.html

Ball, I.R. (2000) *Mathematical applications for conservation ecology: the dynamics of tree hollows and the design of nature reserves*. PhD Dissertation, University of Adelaide.

Csuti, B., Polasky, S., Williams, P.H., Pressey, R.L., Camm, J.D., Kershaw, M., Kiester, A.R., Downs, B., Hamilton, R., Huso, M. & Sahr, K. (1987) A comparison of reserve selection algorithms using data on terrestrial vertebrates in Oregon. *Biological Conservation* **80**, 83–97.

Faith, D.P. (1995) *Biodiversity and regional sustainability analysis*. CSIRO, Australia.

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

Margules, C.R., Nicholls, A.O. & Pressey, R.L. (1988) Selecting networks of reserves to maximize biological diversity. *Biological Diversity* **43**, 63–76.

Otten, R. & van Ginneken, L. (1989) *The Annealing Algorithm*. Dordrecht, Kluwer.

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

Sarakinos, H., Nicholls, A.O., Tubert, A., Aggarwal, A., Margules, C.R. & Sarkar, S. (2001) Area prioritization for biodiversity conservation in Québec on the basis of species distributions: a preliminary analysis. *Biodiversity and Conservation* **10**, 1419–1472.

Sarkar, S., Aggarwal, A., Garson, J., Margules, C. & Zeidler J. (2002) Place prioritization for biodiversity content. *Journal of Biosciences* in press.

Sarkar, S., Parker, N., Garson, J., Aggarwal, A. & Haskell S. (2000) Place prioritization for Texas using GAP data: the use of biodiversity and surrogates within socioeconomic constraints. *GAP Bulletin* **9**, 48–50.