

UC Santa Barbara

UC Santa Barbara Previously Published Works

Title

The SITES reserve selection system: A critical review

Permalink

<https://escholarship.org/uc/item/9hg8w8wz>

Journal

Environmental Modeling and Assessment, 10(3)

ISSN

1420-2026

Authors

Fischer, Douglas T.
Church, R L

Publication Date

2005-09-01

DOI

10.1007/s10666-005-9005-7

Peer reviewed

The SITES reserve selection system: A critical review

Douglas T. Fischer^{a,*} and Richard Church^b

^a Department of Geography, University of California, Santa Barbara, CA 93106-4060, USA
E-mail: fischer@geog.ucsb.edu

^b Department of Geography/National Center for Geographic Information and Analysis, University of California, Santa Barbara, CA 93106-4060, USA
E-mail: church@geog.ucsb.edu

Numerous models have been put forth to help with the growing demand for the establishment of biodiversity reserves. One site selection model that has been used in several recent studies is SITES [S.J. Andelman, I. Ball, F.W. Davis and D.M. Stoms, SITES V 1.0: an analytical toolbox for designing ecoregional conservation portfolios, Unpublished manual prepared for the nature conservancy, 1999, 1–43. (available at <http://www.biogeog.ucsb.edu/projects/tnc/toolbox.html>)]. SITES includes two heuristic solvers: based on Greedy and Simulated Annealing. We discuss the formulation of the SITES model, present a new formulation for that problem, and solve a number of test problems optimally using off-the-shelf software. We compared our optimal results with the SITES Simulated Annealing heuristic and found that SITES frequently returns significantly suboptimal solutions. Our results add further support to the argument, started by Underhill [L.G. Underhill, Optimal and suboptimal reserve selection algorithms, *Biol. Conserv.* 70 (1994) 85–87], continuing through Rodrigues and Gaston [A.S.L. Rodrigues and K.J. Gaston, Optimization in reserve selection procedures – why not?, *Biol. Conserv.* 107 (2002) 123–129], for greater integration of optimal methods in the reserve design/selection literature.

Keywords: reserve site selection, optimization, integer programming, heuristics, model formulation, Simulated Annealing

1. Introduction

As awareness of conservation issues has grown over the last several decades, a growing number of planners have focused on ways to conserve individual species, whole ecosystems, and other natural resources. A key strategy for conservation has been the establishment of reserves that can be managed for the benefit of the targeted conservation elements, be they endangered species, threatened vegetation communities, unique habitat types, or some other element of conservation concern. Until recently, most efforts to establish reserves have focused on areas with scenic and recreational value, resulting in ad hoc reserve networks with substantial redundancy and many gaps. As more areas experience environmental degradation and more species are threatened with extinction, greater attention has focused on designing comprehensive sets of reserves, where all conservation elements in a region are adequately represented in the reserve system [4–6].

Social and economic considerations often preclude simply conserving all land in a region; so the problem of reserve design has focused on selecting small portions of a region for conservation. Because there may be considerable flexibility in which portions are selected, the problem is far from simple. This design dilemma has fueled the development of a wide variety of computer-based reserve site selection models (e.g., [4–17]).

Reserve site selection models are different from many other applications of optimization techniques. In most cases the ecological data available to conservation planners

contain much larger uncertainties than in more traditional optimization applications in business or the military. In addition, unmodeled objectives (e.g., aesthetics, public opinion, politics, etc.) often play a much more influential role in the implementation of a reserve system than in other applications (e.g., [18]). The primary utility of reserve site selection algorithms, then, is not to produce single, prescriptive solutions. Rather, the principle utility of reserve site selection algorithms is to explore the ranges of performance possible for various modeled objectives, and the potential tradeoff curves that may exist between them. The optimal solutions thus produced then provide benchmarks against which specific on-the-ground plans can be compared.

Prendergast et al. [19] argue that there is a gap between theory and application and that current site-selection algorithms do not address many of the pressing and practical issues in reserve design, including ease of use for managers and decision makers. Although Pressey and Cowling [20] answer many of the issues raised by Prendergast et al. [19], there remains a real need to make better selection models, solution procedures, and decision support systems for reserve planning and design. Within this same spirit, there is a need to test and compare existing models and solution methods.

One model that has received significant attention and has been used in recent studies of reserve design [21–24] is the SITES model [1,25] that was developed for The Nature Conservancy. The SITES program includes various mapping and analysis functions, all built around a conceptual model for reserve selection, and a pair of heuristics for solving reserve selection problems based on this model. The model is an area-representation model similar to the

* Corresponding author.

80 BMAS model [6] with a goal-programming approach and
 81 added terms to encourage clustering. Our objective in this
 82 paper is to present a reformulation of this model, test off-
 83 the-shelf software in solving this model, and compare these
 84 results to the solvers provided with the SITES program.
 85 We will show that the performance of the existing heu-
 86 ristics for the SITES model can be improved.

87 In the next section, we describe the SITES model, based
 88 upon both the SITES user manual as well as several
 89 supporting papers. Working from this conceptual model,
 90 the SITES developers opted for the development of two
 91 heuristics. These have been tested on several problems [13]
 92 and have been compared against each other [26]. Up to
 93 now, the SITES model has been solved only heuristically;
 94 thus, the quality of the solutions determined from the
 95 heuristic solvers has never been fully evaluated. In a
 96 subsequent section, we offer an alternate formulation for
 97 the SITES model. We demonstrate that this new formula-
 98 tion can be used to solve problem instances optimally.
 99 Thus, we are able to provide an assessment of the efficacy
 100 of the heuristics already developed for solving SITES. We
 101 provide a comparison of heuristic and optimal approaches
 102 and then conclude with a summary and final assessment.

103 2. The SITES model

104 The SITES program is described as picking from among
 105 a number of feasible sites, a set that comprises a portfolio
 106 [1]. The objective is to pick sites for the portfolio in such a
 107 manner that all conservation goals are met and that the cost
 108 is minimized. That is, the stated objective is to find the
 109 minimal cost set of sites such that each conservation goal is
 110 satisfied. The conservation goals can include representation
 111 goals (coverage of species or area of habitat) and spatial
 112 configuration goals. The SITES program attempts to select
 113 a minimal cost portfolio where the portfolio cost is defined
 114 as:

$$\begin{aligned} \text{Total Portfolio Cost} = & (\text{cost of selected sites}) \quad (1) \\ & + (\text{penalty cost for not meeting the stated} \\ & \quad \text{conservation goals for each element}) \\ & + (\text{cost of spatial dispersion of the selected sites} \\ & \quad \text{as measured by the total boundary length of} \\ & \quad \text{the sites in portfolio}). \end{aligned}$$

115 This is further described as:

$$\begin{aligned} \text{Total Cost} = & \sum_i \text{Cost site } i \quad (2) \\ & + \sum_j \text{Penalty cost for element } j \\ & + w_b \sum \text{Boundary length.} \end{aligned}$$

116 This “total cost” function represents the sum of the costs of
 117 selected sites (e.g., site area, acquisition cost, opportunity
 118 cost, habitat quality), plus the sum of the penalties for not

119 meeting specific conservation targets, plus the weighted
 120 perimeter of all selected sites. The third term of the cost
 121 objective is weighted by the term, w_b (and is actually a
 122 measure of clustering and compactness rather than of spa-
 123 tial dispersion [17]). The higher the value of w_b , the more
 124 important it is to select a set of sites that are clustered with
 125 a small perimeter, even if doing so increases the other costs
 126 somewhat. Thus, the total cost function allows for tradeoffs
 127 between boundary length (i.e., an encouragement to cluster
 128 elected sites) and the costs of sites and penalty costs for not
 129 meeting specific conservation targets. In addition to the
 130 terms described above, the SITES documentation [1]
 131 includes references to the selection of spatially separated
 132 clusters of reserves. Our understanding is that this function-
 133 ality was never fully implemented in SITES (D. Stoms, per-
 134 sonal communication, 2001), and so we have omitted it.

135 The second term of the total cost function, as described
 136 by Andelman et al. [1], involves the penalty costs asso-
 137 ciated with falling short of any conservation targets. In
 138 minimizing total portfolio costs, the penalty function
 139 encourages sites to be chosen in such a manner that all
 140 conservation targets are met. When all targets are met or
 141 exceeded, then the penalty costs are zero for all conserva-
 142 tion elements. A conservation target for an element is stat-
 143 ed in terms of a minimum desired value. An element may
 144 represent a species, habitat type, or other factor of interest.
 145 It is assumed that each site contains a specific quantity of
 146 each element. The total of that element over all selected
 147 sites represents the amount protected among the sites in the
 148 selected portfolio. If the total is lower than the target for
 149 that element, then a penalty cost is incurred that is pro-
 150 portional to the shortfall. For example, consider the hy-
 151 pothetical problem and solution portfolio comprised four
 152 sites presented in table 1.

153 For this example, assume that these sites do not share
 154 any boundary in common. There are five different elements
 155 with conservation targets. The first two elements involve
 156 specific types of habitat (called habitat types 1 and 2). The
 157 remaining three elements involve the representation of
 158 three different species. For example, site 65 contains 200
 159 ha of habitat type 1, 2,500 ha of habitat type 2, contains
 160 species C, but not species A or B. (Note, for species
 161 presence data, a 1 means the species is present at the site
 162 and a 0 means that the species is absent). Together, this
 163 portfolio of four sites contains 3,500 ha of habitat type 1
 164 and 4,500 ha of habitat type 2. Because the target values
 165 for each habitat type is 4,000 ha, the portfolio falls short of
 166 habitat type 1 by 500 ha and meets the target for habitat
 167 type 2. The element target shortfall amounts are given in
 168 the penultimate column of the table. For species presence
 169 targets, it is desired to pick sites so that each species is
 170 present in at least two sites in the portfolio. Note that
 171 species B is found at sites 21 and 109. Thus, species B is
 172 found at two sites and meets the conservation target.
 173 Unfortunately, species A is present at only one site so a
 174 shortfall of representation occurs for this species. In cal-
 175 culating the total cost, we have multiplied each cost by a

t1.1 Table 1
Sample SITES problem with two habitat protection targets, and three species representation targets.

t1.2	Problem definition		Selected portfolio				Portfolio cost	
	Weight	Target value	Site 21	Site 65	Site 13	Site 109	Element shortfall	Weighted objective
t1.3								
t1.4	Boundary	0	11,267	14,321	22,456	16,728		0
t1.5	Cost	1	1,000	1,000	1,000	1,000		4,000
t1.6	Habitat type 1	3	4,000	2,000	200	0	1,300	500
t1.7	Habitat type 2	3	4,000	100	2,500	1,900	0	0
t1.8	Species A	500	2	1	0	0	0	1
t1.9	Species B	500	2	1	0	0	1	0
t1.10	Species C	500	2	0	1	1	0	0

t1.11 The selected portfolio of four sites meets the habitat protection target for type 2, but not type 1, and represent species B and C adequately, but not species A. The portfolio cost is a weighted sum of the costs of the selected sites, and the penalties for not meeting protection targets.

176 weight listed on the left side of the table. While not included in formula (2), SITES provides for separate weights
 177 for each element shortfall, as described below. (While
 178 SITES does not provide for weights for site costs, such
 179 weights are easily applied before loading cost data into the
 180 program.) The total cost for this portfolio, as shown in the
 181 rightmost column of table 1 is:
 182

$$\begin{aligned}
 \text{Total cost} &= 1 \times 4,000(\text{site cost}) & (3) \\
 &+ 3 \times 500(\text{penalty cost for habitat 1}) \\
 &+ 3 \times 0(\text{penalty cost for habitat 2}) \\
 &+ 500 \times 1(\text{penalty cost for Species A}) \\
 &+ 500 \times 0(\text{penalty cost for Species B}) \\
 &+ 500 \times 0(\text{penalty cost for Species C}) \\
 &+ 0 \times 64,772(\text{penalty for boundary length}) \\
 &= 4,000 + 1,500 + 0 + 500 + 0 + 0 + 0 \\
 &= 6,000
 \end{aligned}$$

183 It is important to note that the overall penalty cost for a
 184 given shortfall is proportional to the amount of shortfall.
 185 That is, the penalty cost is a linear function with respect
 186 to shortfall. In McDonnell et al. [26], the penalty cost for
 187 a given target is normalized by the amount of the target
 188 (and multiplied by an additional, heuristically determined
 189 weight), so that a balance can be struck between targets
 190 involving small acreage and those involving larger acreage.
 191 Each possible portfolio has a calculated total cost. The
 192 objective is to identify the portfolio with the smallest “total
 193 cost.” If the units of site cost and element penalty are very
 194 different in magnitude, lowest-cost solutions may involve
 195 selecting all the area to eliminate any penalties, or ac-
 196 cepting all penalties to avoid the cost of selecting any sites,
 197 or somewhere in between.

198 McDonnell et al. [26] provide additional description of
 199 the underpinning model of SITES. The following notation
 200 is necessary to describe their formalism:

- 201 c_i – Total area or cost of site i ;
- 202 a_{ik} – Area or other measure of conservation value k on
 203 site i ;
- 204 b_i – Total boundary length of site i ;

- b_k – Required area or amount of conservation value k
 205 needed in portfolio; 206
- b_{ij} – Length of shared boundary between sites i and j ; 207
- K – Set of all conservation elements k ; 208
- I – Set of all sites i ; 209
- m – Total number of sites available for selection 210

The decision to select a site for the portfolio can be represented by the following 0–1 decision variable: 211 212

$$x_i = \begin{cases} 1, & \text{if site } i \text{ is selected for the portfolio} \\ 0, & \text{otherwise} \end{cases}$$

Using this notation, McDonnell et al. [26] describe the following “crisp” optimization problem: 213 214

$$\begin{aligned}
 &\text{Minimize } C(x) \\
 &= \sum_i c_i x_i + w_b \left(\sum_i b_i x_i - 2 \sum_{i=1}^{m-1} \sum_{j=i+1}^m b_{ij} x_i x_j \right) \quad (4)
 \end{aligned}$$

But ensure that sites selected for portfolio contain a minimum quantity of element k 215 216

$$\sum_i a_{ik} x_i \geq L_k \text{ for each element } k \in K \quad (5)$$

Subject to: 217
 Enforce integer restrictions on site decision variables 218

$$x_i = 0, 1 \text{ for each site } i \in I \quad (6)$$

This formulation is an integer nonlinear programming model. The objective (4) involves the minimization of site costs and weighted total boundary length. The boundary length is calculated as the sum of all boundaries of each of the selected units minus twice the distances of the shared edges (since each shared edge is counted twice in the sum). If a pair of sites i and j are both selected for the portfolio, then the term $x_i x_j$ will equal 1, and two times the shared boundary of b_{ij} will be subtracted from the total sum of the individual site boundary lengths. If the term $x_i x_j$ is zero, then at least one of the two sites i and j has not been 219 220 221 222 223 224 225 226 227 228 229

230 selected and no shared boundary is subtracted. If $b_{ij} = 0$,
 231 then the two sites i and j units are not adjacent and
 232 selecting them will not alter the total boundary length. The
 233 first type of constraint (5) ensures at least a prescribed
 234 minimum area (or some other measure) of each conserva-
 235 tion element is achieved by the selected sites. The second
 236 type of constraint (6) refers to the integer restrictions on
 237 the decision variables. We refer to this as a “crisp” model
 238 as each conservation element must be protected by select-
 239 ing a set of sites that contain at least a minimum
 240 amount of area (or representation) for the element.

241 Recognizing that it is easier to develop a heuristic for an
 242 unconstrained optimization problem, McDonnell et al. [26]
 243 present a reformulated version of the above model where
 244 shortfalls in each conservation target are allowed. Consider
 245 the following additional notation:

- 246 u_k – Amount of protection shortfall, if any, for
 247 conservation element k ;
- 248 w_k – Penalty weight per unit of shortfall for
 249 conservation element k ;
- 250 SPF_k – User-specified weight for each element.

252 The value of w_k is determined heuristically as described in
 253 McDonnell et al. [26].

254 With these additional terms they formulated the follow-
 255 ing expanded model, which is the mathematically explicit
 256 version of equations (1) and (2):

$$\begin{aligned} \text{Minimize } C(x) = & \sum_i c_i x_i + \sum_k \left(\frac{w_k SPF_k}{L_k} \right) u_k \\ & + w_b \left(\sum_i b_i x_i - 2 \sum_{i=1}^{m-1} \sum_{j=i+1}^m b_{ij} x_i x_j \right) \end{aligned} \quad (7)$$

257 Subject to:
 258 Define the amount of shortfall in conservation element k
 259 associated with sites selected for portfolio

$$\sum_i a_{ik} x_i + u_k \geq L_k \text{ for each element } k \in K \quad (8)$$

260 Enforce integer restrictions on site decision variables

$$x_i = 0, 1 \text{ for each site } i \in I \quad (9)$$

261 Enforce nonnegativity on Shortfall variables

$$u_k \geq 0 \text{ for each element } k \in K \quad (10)$$

262 This second formulation contains the penalty terms de-
 263 scribed by Andelman et al. [1] in the SITES manual. That
 264 is, this model solves for the optimal portfolio set of sites
 265 that together minimize the total cost function described
 266 above in equations (1) and (2).

267 Although both integer nonlinear programming models
 268 convey exactly what is modeled in the SITES program, the
 269 formulations are not amenable to direct, optimal solution
 270 by commercially available software (except for relatively
 271 small problem instances), because they are nonlinear.

3. Description of the two SITES heuristics

272

273 There are two possible approaches to dealing with the
 274 difficulties of solving the above integer quadratic prog-
 275 ramming problem: (1) rely on heuristics, or (2) attempt a
 276 reformulation, to create a similar, linear problem that would
 277 be solvable. Taking the first route (i.e., rely on a heuristic
 278 approach) is pragmatic where the second approach proves
 279 unsatisfactory. To rely solely on the development of a
 280 heuristic approach, however, means that it may be impos-
 281 sible to truly assess the quality of the solutions generated.

282 Where the second approach is feasible, Rodrigues and
 283 Gaston [3] effectively show that it is often advantageous.
 284 Moreover, the skill and effort required to program an
 285 efficient heuristic from scratch is significantly greater than
 286 that required to format a problem for solving by an off-the-
 287 shelf optimization code.

288 The SITES program provides two solvers designed to
 289 select a portfolio [1]. Details of these approaches can also
 290 be found in [1, 13, 25, 26]. The first solver is a Greedy
 291 heuristic. Starting with no sites in the portfolio, Greedy
 292 selects the site that yields the lowest value of the total cost.
 293 For the second site Greedy picks the site that reduces the
 294 total cost the most and adds it to the portfolio. At each step,
 295 Greedy adds one more site to the portfolio [1]. The heur-
 296 istic stops when no site that can be added to the portfolio
 297 would lower total cost. Thus, Greedy may stop short of
 298 meeting all conservation targets because the reduction of
 299 penalty costs by selecting additional sites may be over-
 300 whelmed by additional site costs or weighted boundary
 301 length. At this point, according to the objective function, it
 302 is not “cost effective” to add any more planning units to
 303 the portfolio, even though it is possible that not all goals
 304 have been satisfied for all elements [1].

305 It should be understood that the greedy heuristic has
 306 been of interest in the Operations Research and Computer
 307 Science literature because it is relatively easy to prove
 308 worst-case bounds for complex problems. It is rarely used
 309 in practice because other techniques have proved to be
 310 considerably better. The Greedy heuristic in the SITES
 311 program should be used with considerable caution, espe-
 312 cially if the Simulated Annealing (SA) process described
 313 below is not used. Church and Reville [27] described how
 314 the greedy process can perform poorly in location and
 315 siting problems. Essentially, as sites are added to the port-
 316 folio, newly added sites tend to marginalize sites that are
 317 already members of the portfolio. Without the ability to
 318 remove sites from the portfolio as it is being constructed,
 319 Greedy suffers in performance because some sites added
 320 early on may not, in the end, be needed, or be justified in
 321 terms of net cost minus target penalties. This means that
 322 Greedy tends to construct solutions that are “bloated,” with
 323 more sites than are necessary.

324 One important note is that SITES Greedy is not
 325 deterministic. Most greedy heuristics have defined tie-
 326 breaking rules, so that multiple runs will always produce
 327 the same solution. SITES Greedy may produce different

328 solutions (with different objective values) in different runs,
329 because any ties during the solution process are broken by
330 random selection [26].

331 We conducted a small test of SITES Greedy using the
332 five zero-perimeter Santa Barbara datasets (described be-
333 low). The heuristic produced different solutions every time
334 for each of 50 restarts on all five problems. In contrast, when
335 using the Small Sierra dataset (SPF of 10, also described
336 below), SITES Greedy produced only one solution for each
337 of the five problems with 50 restarts each. Whatever the spec-
338 ifics of SITES Greedy, it appears more likely to be con-
339 sistent on sparser datasets because of the lower likelihood
340 of ties. In all cases, solution quality was inferior to solu-
341 tions generated by the second solver included with SITES.

342 The second solver that is provided with the SITES
343 program is based upon a solution technique called Sim-
344 ulated Annealing (SA). SA is based upon a statistical anal-
345 ogy between solution quality and energy states of particles
346 in the process of tempering glass and metals by systematic
347 heating and cooling [28]. The SA procedure in SITES
348 starts with a random set of sites. At each iteration, the
349 procedure identifies a single site at random, and then ex-
350 amines the possibility of either adding that site to the
351 portfolio, or, if currently selected, of discarding it from the
352 portfolio. If the change (dropping a site or adding a site)
353 produces an improved solution, the change is automat-
354 ically accepted. If it does not produce an improvement,
355 the change *may still be* accepted (based on comparing a
356 random number to a probability distribution). The prob-
357 ability of accepting a change that degrades solution quality
358 is taken from the Boltzmann distribution (which describes
359 the number of particles that will have a higher energy state
360 than a specified state, at a given temperature). Statisti-
361 cally, when the simulated temperature is high, and the
362 proposed change is not substantially worse, the probabil-
363 ity of accepting a change is relatively high. But, as the sim-
364 ulated temperature is lowered (systematically as the process
365 runs), the probability of accepting changes that worsen a
366 portfolio (by even a small amount) decreases. This pro-
367 cess has the capability of converging to a local optimum and
368 backing out of the local optimum (making a portfolio
369 worse) and then finding even better local optima. SA has
370 been applied to other site selection problems with varying
371 degrees of success [29,30]. As with any SA heuristic, the
372 success in application is somewhat dependent on the
373 problem being solved and the parameter settings used
374 (e.g., cooling rate and the number of iterations). The only
375 parameters in SITES SA that can be set by the user are the
376 number of restarts and the number of iterations per restart.

377 Most applications of SA require multiple restarts of the
378 process, where only the best solution or solutions found
379 among the different restarts is/are considered. SITES in-
380 cludes an option for examining the sum of results from
381 multiple restarts, showing how many times each site was
382 selected. Andelman et al. [1] suggest that this analysis pro-
383 vides a measure of the “robustness” of a solution, as though
384 the number of times a site was chosen by SA was an in-

385 dicator of its importance to an optimal solution. This
386 conclusion is not supportable. Fischer [31] described SITES
387 problems where the median solution quality was more than
388 50% worse than the best solution. With most solutions
389 being very inferior, any site that was selected a majority of
390 the time was necessarily a part of many very inferior
391 solutions. In each case examined, Fischer [31] also found
392 numerous “popular” sites (selected in more than 50% of
393 the solutions) that were not part of an optimal solution, and
394 numerous “unpopular” sites (selected in fewer than 20% of
395 the solutions) that were. The “summed solution” approach
396 described in SITES is a haphazard approach to modeling
397 robustness (a field reviewed by Owen and Daskin [32])
398 that appears to be uninformative at best.

399 McDonnell et al. [26] present a comparison of the
400 Greedy approach and the SA approach in solving the
401 SITES model applied to a vegetation dataset of Northern
402 Territory, Australia. Their comparison demonstrates that at
403 times the Greedy approach outperforms SA, although they
404 conclude that the SA process is probably better suited to
405 solving the SITES problem. As with all heuristics, there is
406 no guarantee that the Greedy or SA processes will find
407 optimal solutions. It should also be understood that the
408 quality of the solutions cannot be ascertained without ac-
409 tual optimal solutions with which to make a comparison.
410 That is, the results of a heuristic, used by itself, should be
411 interpreted with caution.

412 To evaluate the SITES solution process further, an as-
413 sessment is needed in terms of how close to optimal either
414 technique solves the SITES problem. In the next section we
415 present a reformulated model for SITES and then describe
416 how this model can be solved in practice. With this model
417 we will provide an assessment of the SITES heuristics later
418 in this paper.

4. A reformulation of the SITES model 419

420 The major obstacle to using an optimal solver for the for-
421 mulation of the SITES model (as described in McDonnell
422 et al. [26]) is the set of quadratic terms that are used to
423 define the boundary length of the sites selected for the
424 portfolio:

$$\sum_{i=1}^{m-1} \sum_{j=i+1}^m b_{ij}x_i x_j \quad (11)$$

425 With the exception of these terms, the model is an integer-
426 linear programming problem. These terms make the
427 previous two formulations difficult or impossible to solve
428 optimally. It is, however, possible to reformulate the model
429 and circumvent the need for the quadratic terms. We can
430 do this by introducing the following new variable:

$$z_{ij} = \begin{cases} 1, & \text{if sites } i \text{ and } j \text{ have been selected for} \\ & \text{the portfolio} \\ 0, & \text{otherwise} \end{cases}$$

431 We need to define such a variable for each pair of sites that
432 share an edge. Therefore, consider:

$$Z = \text{set of site pairs}(i,j) \text{ which share boundaries,} \\ \text{where } i < j$$

433 Each pair of adjacent sites will be addressed where the
434 smaller of the two site indices is given first in the site pair.
435 This distinction allows us to represent each possible edge
436 with one decision variable. Using the two discrete decision
437 variables x_i and z_{ij} , it is now possible to construct a model
438 that represents the SITES problem and eliminates the
439 quadratic terms:

$$\text{Minimize } Obj = w_c \sum_i c_i x_i \\ + \sum_{k \in K} \left(\frac{w_k \text{spf}_k}{L_k} \right) u_k + w_b \left(\sum_i b_i x_i - 2 \sum_{(i,j) \in Z} b_{ij} z_{ij} \right) \quad (12)$$

440 Subject to:
441 Define amount of shortfall for target involving conserva-
442 tion element k

$$\sum_i a_{ik} x_i + u_k \geq L_k \quad \text{for each element } k \in K \quad (13)$$

443 Ensure z_{ij} is only allowed to be 1 if adjacent sites i and j are
444 both selected

$$(a) \ x_i - z_{ij} \geq 0 \quad \text{for each shared edge where } (i,j) \in Z \\ (b) \ x_j - z_{ij} \geq 0 \quad \text{for each shared edge where } (i,j) \in Z \quad (14)$$

445 Enforce integer requirements on site decision variables

$$x_i = 0, 1 \quad \text{for each site } i \in I \quad (15)$$

446 Enforce nonnegativity on Shortfall variables

$$u_k \geq 0 \quad \text{for each element } k \in K \quad (16)$$

447 The first term in the objective function sums the costs
448 of all of the selected sites. The second term represents the
449 weighted penalty costs of incurring any shortfall in meet-
450 ing conservation targets and the third term calculates the
451 total boundary of the sites selected, accounting for any
452 shared edges. The first type of constraint (13), the same as
453 (8) used in the previous model, defines any shortfall in
454 conservation targets that may exist in the set of sites cho-
455 sen. The second type of constraint (14) is used to define
456 the values of the z_{ij} variables. Each z_{ij} variable must equal
457 zero unless both sites i and j are selected. If both i and j
458 are selected, z_{ij} is allowed to have any value between zero
459 and one. Since the objective function encourages z_{ij} to be
460 as large as possible (to reduce total boundary length), z_{ij} is
461 effectively 0 or 1. Defining z_{ij} this way accurately captures
462 the boundary length of the selected sites, without the com-
463 putational burden of defining each z_{ij} as an integer var-

464 ible. This type of model construct was recently introduced
465 for a related reserve design model by Fischer and Church
466 [17].

467 The above model represents a reformulation of the
468 SITES problem. This formulation is an integer-linear
469 programming problem. Since commercially available soft-
470 ware exist for solving this type of programming problem,
471 it makes sense to test such software first, instead of in-
472 vesting in development of a special solver for this specific
473 type of problem. As the heuristic solvers now exist, it
474 makes sense to test their performance, and in the next sec-
475 tions, we present information for three different datasets
476 and compare results from SITES heuristics with results
477 from a commercial code for solving the reformulated
478 SITES model.

5. Comparing SITES solvers with an IP/LP approach 479

480 In the following sections we provide a comparison of
481 the SA heuristic of the SITES model with solutions gen-
482 erated for the reformulated SITES model. With this new
483 model formulation, it is straightforward to solve the SITES
484 model using the techniques of linear and integer prog-
485 ramming. We used an off-the-shelf optimization pack-
486 age called CPLEX (ILOG Corporation), which is a widely
487 used, general-purpose, linear-integer programming solv-
488 er. We do not give further details of SITES Greedy since,
489 overall, we found that the SA heuristic performed consid-
490 erably better than Greedy on the problems that we analyzed.

491 As with any SA heuristic, the SITES solver has a num-
492 ber of parameters (other than weights for the different
493 objectives) that affect its performance [30]. In the inter-
494 ests of operational simplicity, the designers of SITES have
495 hard-coded a number of parameters, such as initial temper-
496 ature, cooling rate, etc. Remaining variables that must be
497 set to solve a SITES problem are the number of iterations,
498 the number of restarts, and the species penalty factor.

499 The SITES manual recommends using at least 10^6 iter-
500 ations. In the interest of speed, we tried a few exploratory
501 runs using 10^5 iterations and found solutions significantly
502 inferior. All of the results presented here used 10^6 iter-
503 ations. For each problem that we solved, we tested using
504 50 restarts and 500 restarts of the SA heuristic. We also
505 tested the sensitivity of model results to varying levels of
506 the species penalty factor.

507 As mentioned earlier, the value of one of the weights
508 applied to protection shortfalls, w_k (7), is determined heu-
509 ristically in SITES during each run [26]. The method of
510 calculation of this term is not adequately defined to allow
511 independent replication of its values. As a result, we were
512 unable to explain or replicate the penalty values derived
513 by SITES for solutions that did not meet all protection
514 targets. We conducted a number of tests with CPLEX
515 and found that unconstrained problems where shortfalls
516 were allowed, with proportional penalties, were solved
517 quite a bit faster than problems where shortfalls were

518 not allowed. Only solutions that met all protection tar-
519 getts are compared below.

520 SITES is compiled for a Windows/Intel platform. Our
521 testing was conducted on a PC with an 800-MHz AMD
522 Athlon processor and 256 MB of RAM, running Windows
523 98SE. For each problem instance, we ran the SA heuristic
524 twice, first with 50 restarts, then with 500 restarts. We
525 wrote a short code to then reformat the SITES input files
526 into an MPS file (a standard text file format for Linear and
527 Integer Programming codes), to allow side-by-side testing.
528 Each MPS file was loaded into CPLEX version 6.6 on a
529 Sun Ultra SPARC 10 Station with a 467-MHz processor.
530 To confirm that SITES was not running on a slower
531 machine, we benchmarked the computers with a billion
532 iterations of the main elements of the SA solver (generate a
533 random number, compare to an array, and store results to
534 an array). The SITES machine performed the benchmark
535 1.7 times faster than the CPLEX machine.

536 CPLEX utilizes a branch-and-bound process to solve
537 mixed-integer problems like these [33]. At each step in the
538 process, CPLEX looks for an improved solution, and
539 establishes a lower bound below which it has proven that
540 no feasible solutions exist. As the process continues, the
541 lower bound rises and (hopefully) the objective value of
542 the best-known solution decreases. We used three stopping
543 rules. The first was a gap of 0.01% (0.0001). That is, when
544 the lower bound is within 0.01% of the best known so-
545 lution, we declare the current solution optimal (while a
546 better solution is possible, it could be no more than 0.01%
547 better than the current solution). We also terminated the
548 CPLEX run after branching to 500,000 nodes or if the
549 branch-and-bound tree exceeded 500 Mb.

550 Data from three reserve selection problems were used,
551 each providing a different test environment. The datasets
552 and the results from each are described in the next two
553 sections.

554 6. Sierra datasets

555 The first two datasets were prepared as part of the Sierra
556 Nevada Ecosystem Project [34]. Both datasets are from
557 primarily forested areas of the northern Sierra Nevada that
558 roughly correspond to the Jepson Northern Sierra Ecor-
559 egion [35]. They consist of distribution and conservation
560 information for different plant communities. Distribution is
561 given as an areal extent of the vegetation formation within
562 each of the several hundred watersheds making up the
563 study area.

564 The first dataset is Okin's [36] *north.M.35*. This dataset
565 uses 663 watersheds as the sites. These sites have 1,834
566 edges that are shared between adjacent sites. Thirty-six
567 plant community types are used as conservation elements.
568 The minimum protection target for each element is 10% of
569 the existing range. This dataset is hereafter referred to as
570 the Small Sierra dataset.

571 The second dataset is distributed as sample data by
572 Andelman et al. [1]. This dataset uses an expanded area
573 of the northern Sierra Nevada, encompassing 776 wa-
574 tersheds. These sites have 2,148 edges that are shared be-
575 tween adjacent sites. It includes 55 plant community types
576 as conservation elements, and requires 25% of their exist-
577 ing ranges as minimum protection targets. As a result,
578 the problem requires a substantially greater number of sites
579 to be selected as reserves, and serves as a significantly dif-
580 ferent test environment for the model. This dataset is
581 hereafter referred to as the Large Sierra dataset. Maps of
582 three solutions from this dataset are included as figure 1.

583 In both datasets, area is measured in hectares. Site costs
584 are represented in part by a site suitability index defined by
585 Davis et al. [34], where small numbers indicate planning
586 units that are very suitable for conservation, and high
587 numbers indicate areas that are unsuitable due to road den-
588 sity, fragmented ownership, degraded habitat, etc. Because




Figure 1. Maps of the three optimal reserve networks for the Large Sierra dataset. Map A shows the network requiring minimum land area. The network that optimizes land suitability scores is shown in map B. Map C shows an optimal solution trading off area and suitability with perimeter to encourage clustering of selected land. The Small Sierra dataset covers all but the northwest corner of this area.

589 SITES utilizes only a single cost term for each site, we
 590 summed weighted area and weighted suitability index val-
 591 ues for each site to generate a single cost term before
 592 preparing each set of data files for SITES. Perimeter is mea-
 593 sured in meters.

594 For the purposes of this study, we chose five sets of
 595 weights based on methods originally outlined by Cohon
 596 et al. [37]. Figure 2 gives an example of their method for
 597 efficiently estimating two-dimensional tradeoff curves
 598 using Suitability and Perimeter for the Large Sierra dataset.
 599 Solanki et al. [38] extended this method for systematically
 600 varying objective weights to estimate n -dimensional trade-
 601 off surfaces, and we used their method for calculating the
 602 three-dimensional weights. Table 2 shows the weights used
 603 and provides a quick comparison between the three data-
 604 sets used in this study.

605 Previous experience with SITES showed that solution
 606 quality can vary a great deal depending on the magnitude
 607 of SPF values compared to cost and perimeter values [31].
 608 With that in mind, we solved each problem for the Small
 609 and Large Sierra datasets with three different SPF values.

610 The remainder of this section will compare the solvers
 611 in terms of solution speed and solution quality. For each
 612 problem instance, table 3 lists the lower bound derived by
 613 CPLEX, the solution time required by CPLEX (to reach
 614 one of its stopping rules), and solution quality for each solv-
 615 er. Solution quality is expressed as “gap” defined as the

616 difference between the best known objective and the lower
 617 bound, expressed as a percentage of the bound. The first line
 618 of table 3 lists the result from the Small Sierra dataset
 619 minimum area problem. CPLEX solved the problem in 6
 620 min or 0.10 h. The best solution found by CPLEX was
 621 0.01% above the proven lower bound of 78308. The first
 622 row under SA shows results using an SPF of 10. The best
 623 solution found by SA in 50 restarts was 10.82% above the
 624 bound. After 500 restarts, the best solution found by SA
 625 was better, at 6.61% above the bound. Finally, 49.4% of
 626 the SA500 solutions met all protection targets.

627 For all five problems using the Small Sierra dataset,
 628 CPLEX solved to within 0.01% of optimality with a me-
 629 dian solve time of 39 h (i.e., each final solution was within
 630 0.01% of its respective final lower bound, and was thus
 631 declared optimal). The two problems with perimeter
 632 weights of zero was solved in less than 7 min, whereas
 633 those with perimeter considerations took much longer to
 634 solve optimally. After running CPLEX for 1 h, both mul-
 635 tiobjective problems had solutions that eventually proved
 636 to be within 3% of optimality (at the time, the gap was only
 637 known to be less than 16%). The minimum perimeter so-
 638 lution proved difficult to solve, with the best solution after
 639 1 h being 18% above optimal, and still 8% above optimal
 640 after 12 h.

641 For the large Sierra dataset, solution times were gen-
 642 erally shorter, and CPLEX solved all but one problem to

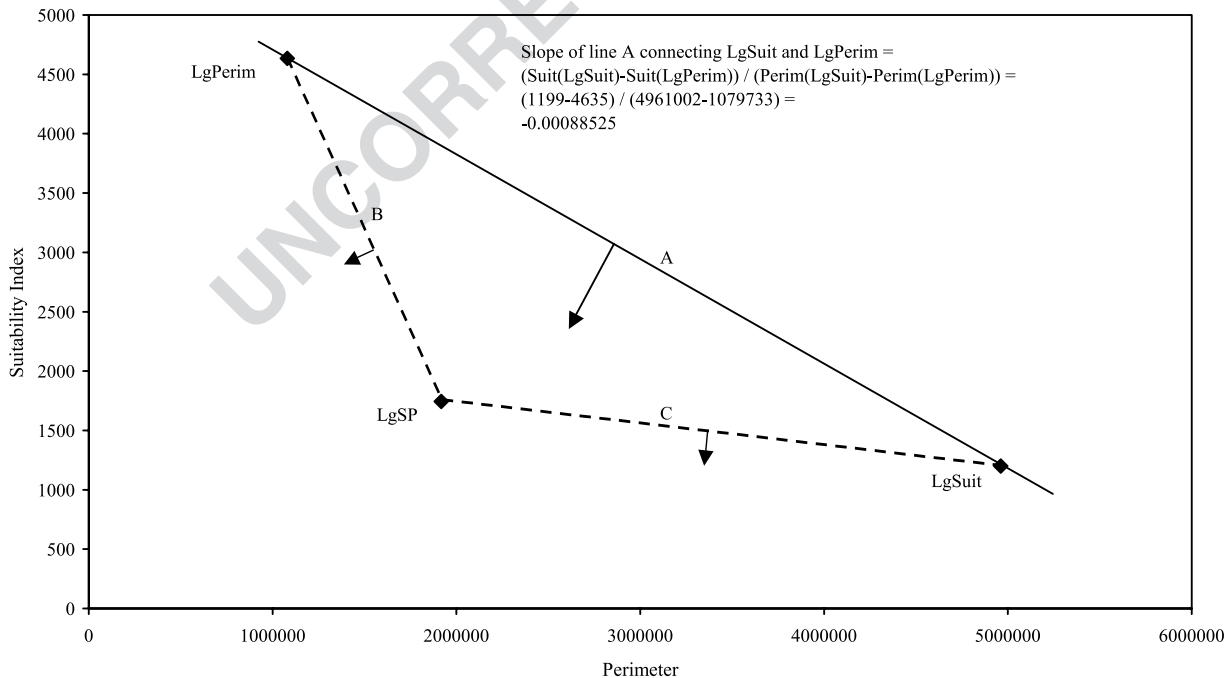


Figure 2. Estimation of a two-dimensional trade-off curve using the method of Cohon et al. [38] using the Large Sierra dataset. Steps are as follows. Optimize each single objective. Calculate the slope of line A connecting those two solutions in objective space. This is the estimated trade-off curve with two points. Apply the absolute value of the slope as the weight for the objective on the x-axis and a weight of 1 to the y-axis objective (the total weighted objectives for both solutions will be equal under the new weights). Weighting Suitability with 1, the weight for Perimeter would be 0.00088525. For scaling, we multiplied both weights by 1,000 to find solution LgSP. With three solutions, the trade-off curve is estimated as dashed lines B and C. To find additional solutions on the trade-off curve, the process is repeated for line B and C.

t2.1

Table 2
Comparison of the three datasets, showing objective weights used for each problem.

t2.2	Dataset		Problem name	Weights		
t2.3	Name	Description		Area	Suitability	Perimeter
t2.4	Small Sierra dataset		SmArea	1	0	0
t2.5		Sites = 663	SmSuit	0	1	0
t2.6		Spp. = 36	SmPerim	0	0	1
t2.7		Edges = 1834	SmSP	0	1,000	0.5052
t2.8			SmASP	1	733	0.4732
t2.9	Large Sierra dataset		LgArea	1	0	0
t2.10		Sites = 776	LgSuit	0	1	0
t2.11		Spp. = 55	LgPerim	0	0	1
t2.12		Edges = 2148	LgSP	0	1,000	0.8853
t2.13			LgASP	1	2,004	0.4194
t2.14	Santa Barbara dataset		SB25/P	1	0	0/0.01
t2.15		Sites = 1906	SB40/P	1	0	0/0.01
t2.16		Spp. = 57	SB50/P	1	0	0/0.01
t2.17		Edges = 5709	SB60/P	1	0	0/0.01
t2.18				SB75/P	1	0

t2.19 The description of each dataset includes the number of sites, the number of conservation elements (species), and the number of shared edges. Five problems were solved for each Sierra dataset. The weights used for minimizing each of the three objectives (the sums of area, suitability index, and perimeter) are given in the right three columns. Zero weights indicate the objective was not considered in that problem. For the Santa Barbara dataset, we solved ten problems, five just minimizing area (area weight = 1, suitability and perimeter weights = 0), and then five more with a small additional weight on minimizing perimeter (0.01).

643 within 0.01% of optimality before stopping. The remaining
644 problem was proven to be within 0.62% of optimal. Me-
645 dian solution time was slightly over 10 h. The Optimal
646 Suitability and Minimum Perimeter problems solved in 3
647 and 1.7 h, respectively, while the multiobjective problems
648 took longest. It is not immediately apparent why the larger
649 dataset should have had shorter solve times. Solution speed
650 of the branch-and-bound process (the process employed by
651 CPLEX) is dependent on problem characteristics and just
652 because a problem might be smaller or a subset of a larger
653 problem does not guarantee that solution times will be less
654 than that needed for the larger problem.

655 The quality of the best SITES solution for each of the
656 Sierra problems ranged from optimal (determined using
657 CPLEX) to 11% worse than optimal, depending on the
658 problem and SPF value used. The average gap for the best
659 solutions after 50 restarts was 4.5% for those problems
660 where SITES produced feasible solutions (solutions that
661 met all protection targets). For perspective, however, just
662 as we examine only the best solution in reporting SA
663 performance for a given problem, we must look at the
664 worst performance within a set of problems to give some
665 idea of how well SITES can be expected to perform on a
666 similar set of problems. Thus, when faced with problems
667 similar to those in the Sierra datasets, the best solutions
668 generated by the SITES solver after 50 restarts may vary
669 considerably from optimal.

670 One important question is whether it makes sense to run
671 a great many iterations of the heuristic. For the Sierra
672 datasets, running a problem for 500 restarts instead of 50
673 resulted in a mean decrease in the gap of 1.5 percentage
674 points. That reflects improvements in 21 of the 30 prob-
675 lems tested, and no feasible solutions returned in six cases.

In the remaining three problems, the gap after 500 restarts
was greater than the gap after a separate run of 50 restarts,
illustrating the important point that SA works by probabili-
ty, and that increasing the number of restarts only incre-
ases the probability of finding a near-optimal solution.
Even using a large number of restarts provides no guar-
antee of “getting lucky.” In the minimum perimeter prob-
lem, using the Small Sierra dataset, SITES found the
optimal solution once in 50 restarts. Subsequently, in a run
of 500 restarts, SA found the optimal solution again, only
once. However, in a subsequent run of 1,000 restarts, the
best solution that SA found was 1.7% worse than optimal.

688 For both Sierra datasets, SITES performed very differ-
689 ently depending on the SPF values used. With the recom-
690 mended SPF value of 1, SITES failed to return any feasible
691 solutions (i.e., solutions meeting all protection targets) in
692 six of the ten problems. In all cases, the percentage of
693 feasible solutions increased with increases in the SPF value
694 used. The setting of SPF values is important to the solution
695 process in SITES, with lower values allowing the heuristic
696 more flexibility to explore infeasible solutions on the way
697 to finding a good solution. Higher SPF values limit the
698 heuristic’s flexibility, but ensure a greater number of fea-
699 sible solutions from which to choose. Because the number
700 of iterations is so important to SA, it is expected that the
701 increased number of solutions from which to choose may
702 be more important to solution quality than the reduced
703 flexibility. We looked for a consistent pattern in solution
704 quality as a function of SPF, but found none.

705 An obvious point of comparison between the SITES
706 heuristic and an optimal solver such as CPLEX is solution
707 speed. Results presented above indicate that CPLEX takes
708 longer to solve these problems (optimally) than SITES

t3.1 Table 3
Comparison of results from CPLEX and SITES Simulated Annealing (SA) heuristic for solutions actually meeting all element protection goals.

t3.2	Problem name	CPLEX			SA50		SA500	
t3.3		Bound	Time	Gap (%)	SPF	Gap (%)	Gap (%)	%Feas
t3.4	A	B	C	D	E	F	G	H
t3.5	SmArea	78308	0.10	0.01	10	10.82	6.61	49.4
t3.6					5	7.80	7.60	19.6
t3.7					1	–	–	0.0
t3.8	SmSuit	95.40	0.00	0.00	10	7.97	5.03	45.8
t3.9					5	8.18	2.31	24.0
t3.10					1	0.00	0.00	1.0
t3.11	SmPerim	444255	94.79	0.01	10	0.01	0.01	82.0
t3.12					5	8.79	3.72	67.8
t3.13					1	1.67	4.24	4.2
t3.14	SmSP	428477	98.87	0.01	10	4.07	2.83	67.4
t3.15					5	2.69	2.00	48.4
t3.16					1	–	0.01	1.0
t3.17	SmASP	467085	38.69	0.01	10	7.57	5.65	62.4
t3.18					5	10.41	2.92	39.2
t3.19					1	–	–	0.0
t3.20	LgArea	1040988	10.44	0.62	10	3.29	3.55	14.0
t3.21					5	3.81	2.92	2.8
t3.22					1	–	–	0.0
t3.23	LgSuit	1199.20	3.02	0.01	10	7.18	5.51	11.6
t3.24					5	6.10	7.36	0.4
t3.25					1	–	–	0.0
t3.26	LgPerim	1079734	1.69	0.01	10	0.56	0.25	96.0
t3.27					5	1.28	0.25	86.8
t3.28					1	1.05	0.05	53.4
t3.29	LgSP	3442077	26.52	0.01	10	1.52	0.73	52.6
t3.30					5	1.44	0.83	26.8
t3.31					1	–	–	0.0
t3.32	LgASP	5066718	12.31	0.01	10	3.58	2.88	34.6
t3.33					5	4.39	2.71	8.4
t3.34					1	–	–	0.0
t3.35	SB25	118552	13.81	0.02	*	9.07	8.19	54.0
t3.36	SB40	203165	2.40	0.01	*	5.39	5.22	40.0
t3.37	SB50	261672	1.42	0.01	*	4.12	3.90	49.4
t3.38	SB60	321340	0.29	0.01	*	3.12	3.16	40.8
t3.39	SB75	412543	0.09	0.01	*	2.08	1.86	32.6
t3.40	SB25P	149321	12.00	1.75	*	25.34	20.06	18.2
t3.41	SB40P	240383	12.00	1.55	*	15.17	14.06	25.6
t3.42	SB50P	301493	12.00	1.39	*	11.91	11.19	39.6
t3.43	SB60P	363019	12.00	1.12	*	8.56	7.85	44.4
t3.44	SB75P	456677	12.00	0.85	*	5.86	5.34	17.4

t3.45 Problem name (A) is the same as in table 2. Column B is the lower bound below which CPLEX has proven that no solutions exist. Time (C) is the number of hours before CPLEX reached a stopping rule. Gap (D) is the difference between the best-known solution and the bound (B), expressed as a percentage of the bound. SPF (E) is a weight used in SITES SA. Columns F and G show the gap of the best known SITES solution after 50 and 500 restarts, respectively (expressed as a percentage of the bound in Column B). Column H is the percentage of 500 restarts that met all element protection goals. For a time comparison, the SA heuristic required 0.6–0.8 h to run 500 restarts on the Small Sierra dataset, 1.1–1.2 h for the Large Sierra dataset, and 1.8–2.0 h for the Santa Barbara dataset.

709 takes to find a solution of unknown quality. In many cases, 719
 710 even the longest of these run times would be acceptable 720
 711 given the scope of large planning efforts. However, in 721
 712 cases where rapidity of results is paramount (e.g., an 722
 713 interactive planning exercise), a key question is the quality 723
 714 of solutions that can be produced without taking much 724
 715 computer time. Table 4 shows the quality of CPLEX 725
 716 solutions after the first 2 min. The first column shows the 726
 717 gap at 2 min; this is the percentage difference between the 727
 718 best solution (after 2 min) and the bound known at that 728
 time (i.e., how good was the solution known to be at 2 min?). The second column shows the difference between that solution and the final bound proven by CPLEX at termination (i.e., how good did the 2-min solution turn out to be?). The third column lists the final gap CPLEX produced at termination. The fourth column lists the gap for the best SITES solution after running 50 restarts at each of the three SPF values (compared to the final bound proven by CPLEX). The SITES solutions represent 5–18 times more processing time on the faster SITES computer

t4.1

Table 4
Solution quality comparison between CPLEX after 2 min, final CPLEX solution, and best SA solution from a run of 50 restarts.

t4.2	Problem name	2 min Gap		Final CPLEX Gap	Best SA50
t4.3		Known	Actual		
t4.4	A	B (%)	C (%)	D (%)	E (%)
t4.5	SmArea	1.05	0.08	0.01	7.80
t4.6	SmSuit	0.00	0.00	0.00	0.00
t4.7	SmPerim	138.87	46.53	0.01	0.01
t4.8	SmSP	35.63	10.46	0.01	2.69
t4.9	SmASP	21.77	6.77	0.01	7.57
t4.10	LgArea	1.23	1.18	0.62	3.29
t4.11	LgSuit	1.20	0.45	0.01	6.10
t4.12	LgPerim	16.76	4.65	0.00	0.56
t4.13	LgSP	9.36	2.68	0.01	1.44
t4.14	LgASP	3.22	1.20	0.01	3.58
t4.15	SB25	0.21	0.20	0.02	9.07
t4.16	SB40	0.07	0.06	0.01	5.39
t4.17	SB50	0.06	0.05	0.01	4.12
t4.18	SB60	0.10	0.10	0.01	3.12
t4.19	SB75	0.03	0.03	0.01	2.08
t4.20	SB25/P	13.17	12.36	1.75	25.34
t4.21	SB40/P	12.75	12.23	1.55	15.17
t4.22	SB50/P	13.91	13.49	1.39	11.91
t4.23	SB60/P	8.28	7.98	1.12	8.56
t4.24	SB75/P	14.02	13.83	0.85	5.86

t4.25 Problem name (A) is the same as in table 2. Column B lists the difference between the best solution known after 2 min and the bound known at 2 min (expressed as a percentage of the 2-min bound). Column C shows the gap between the best 2-min solution and the final bound (expressed as a percentage of the final bound). Column D is the difference between the best CPLEX solution and the final bound, and replicates table 3 Column D. Column E shows the gap between the best SITES SA solution (of the three SPF levels used) after 50 restarts, and the final CPLEX bound. Note that for SmSP the 2-min solution was known to be within 36% of optimal. As CPLEX proceeded, the lower bound rose 25 percentage points, so this solution eventually proved to be within 10% of optimal. Meanwhile the best known solution improved by 10 percentage points. For the SB problems, most of the improvement after 2 min was due to finding improved solutions, with relatively small changes in the lower bound.

729 (or 8–31 times more computational effort). After only 2
730 min of CPLEX processing time, the known gaps were still
731 reasonably large for many of the problems, but, impor-
732 tantly, the actual solutions were generally quite good.

733 For nine of the ten Sierra problems, CPLEX identified at
734 least one solution in the first 2 min that eventually proved
735 to be within 11% (median 1.9%) of optimal. [Only the
736 minimum perimeter problem for the Small Sierra dataset
737 had a worse solution after the first 2 min (47% above
738 optimal), but, with CPLEX reporting a gap of 138%, there
739 was no question that more solution effort was indicated.]
740 These 2-min solutions compare favorably to the best so-
741 lutions returned by SITES after significantly longer times
742 (with 50 restarts for each SPF value), which had median
743 solution quality of 3.0% above optimal.

744 7. Santa Barbara dataset

745 The third dataset was originally prepared for Santa
746 Barbara County in connection with a land-use planning
747 project [39]. County-administered land was divided into
748 1,906 sites, based largely on watershed boundaries. These
749 sites have 5,709 edges that are shared between adjacent
750 sites. Habitat ranges for 57 species of regulatory concern
751 were mapped based on known ranges and wildlife– habitat

relationship models. Each species was assigned a species
penalty factor between zero and one based on regulatory
status, degree of endemicity, and threats to local habitat.

In addition to having significantly more sites than either
Sierra dataset (see table 2), the Santa Barbara dataset has
elements that are more widely distributed, with each el-
ement occupying an average of 639 sites, compared to 135
for the Large Sierra dataset and 57 for the Small Sierra
dataset. For this dataset, we explored a trade-off curve of
how much total land would be required to protect increas-
ingly large percentages of each species range. We estab-
lished five different conservation targets (25, 40, 50, 60,
and 75% of each species range), and attempted to find the
minimum amount of land that might be required to meet
these protection targets for all 57 species. For the current
study, we used the existing species-specific SPF values; we
opted against performing additional SPF sensitivity testing.

After solving these five problems, we solved a second prob-
lem for each of the five protection levels where we included
a small weight for minimizing perimeter (to increase clus-
tering and compactness of the reserve system).
For the perimeter problems, we also decided in advance
to limit the amount of computing time to 12 h per run. That
is, we used the solver to find some solutions of known
quality (whether they qualify as “good” depends on stan-
dards) in a reasonable amount of time. With the relatively

752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777

778 large areas being selected, a certain amount of clustering
779 was inevitable. Analyzing a tradeoff curve for perimeter
780 weights was therefore deemed unnecessary; we simply
781 chose an arbitrarily small, nonzero weight for perimeter to
782 somewhat further encourage compactness in our test data.

783 Solution times for the zero-perimeter problems were
784 generally shorter for the Santa Barbara dataset than for the
785 smaller, but relatively sparser Sierra datasets (see table 3).
786 In all but one case, CPLEX proved the optimality of the
787 solution – that problem terminated on reaching the tree size
788 limit, with a remaining gap of 0.02%. While some of the
789 solution times may still appear long, most of that time was
790 spent improving solutions that were already very close to
791 optimal, or in proving their optimality.

792 For each of the first five Santa Barbara problems,
793 CPLEX found at least two solutions, with a proven gap of
794 less than 0.5%, within the first 8 s of solution time. The
795 remaining hours often revealed a number of slightly better
796 solutions, but had we used a stopping rule of 0.25% instead
797 of 0.01%, only one problem would have run longer than 9
798 s. This illustrates the utility of having a lower bound, and
799 the importance of looking at solution quality as well as
800 total solution time.

801 For the first five Santa Barbara problems, CPLEX re-
802 turned solutions within 0.2% of optimal (median 0.06%) in
803 the first 2 min (table 4), comparing favorably to solutions
804 ranging 2–9% (median 4.1%) for SITES SA after 50 re-
805 starts. The perimeter problems had CPLEX solutions
806 ranging 8–14% above the final bound (median 12.4%)
807 compared to SITES SA, which ranged 6–25% (median
808 11.9%). For the perimeter problems, CPLEX returned so-
809 lutions within 1.75% of optimality after 12 h, with decreas-
810 ing gaps as the targets increased. By contrast, gaps of the
811 SITES solutions varied from 5% to 25%.

812 8. Conclusion

813 Optimal solvers offer the significant advantage of
814 revealing the quality of the present solution. At each step
815 in a branch-and-bound process, the solver can display the
816 lower bound, below which it has proven no integer so-
817 lutions exist, as well as displaying the value of the current
818 solution, and the gap between the two. With this informa-
819 tion, a planner can watch the rate of closure between lower
820 bound and current solution to forecast the longest the
821 solver is likely to run (this is a worst-case estimate, as-
822 suming no new, better solutions are found). The planner
823 can also choose to abandon the analysis at any point with a
824 clear understanding that no new solutions exist that can
825 beat the current solution by more than the reported gap. By
826 contrast, most heuristics provide little or no information on
827 the quality of the solutions they generate. Due to the ran-
828 dom seed used in SA, it is very difficult to answer the
829 question of how many restarts are needed to offer a high
830 probability of finding a near-optimal solution. As seen with
831 the SmPerim problem, 50 restarts may be sufficient to find

an optimal solution, while the next 1,000 restarts may fail
to find as good a solution again. Based on our analysis, the
user’s best approach with SITES is to use the highest
number of restarts possible, explore a range of SPF values,
apply weights efficiently (see figure 2), and carefully screen
out solutions based on objective values and success at
meeting protection targets. The “summed solution” feature,
erroneously reported to reflect a measure of robustness, is
not informative.

Performance of SA heuristics has been discussed by nu-
merous authors, as reviewed by Murray and Church [30].
The particular SA process used by SITES has a very nar-
row search neighborhood, and, echoing Underhill’s [2] call
for greater interaction between the mathematical and bio-
logical research communities, might be readily improved
[e.g., by allowing the heuristic to randomly swap two (or
more) sites at each iteration as well as randomly adding or
dropping sites [30,36]]. SA processes in general do
statistically converge on optimal solutions if the cooling
rate is very slow and the number of restarts is very high.
Statistically speaking, computer-based SA processes are
more like “simulated quenching” since the computation re-
quired for true “annealing” is well beyond realistic design [30],
and far greater than for other, superior solvers. Because of
this limitation, SA heuristics may not find optimal
solutions with any regularity for complex problems like
SITES, in any reasonable amount of computer time. In
essence, a heuristic such as SITES SA provides predict-
able solution times, but inconsistent and unpredictable so-
lution quality, whereas an optimal solver such as CPLEX
provides solution times that can only be predicted once
started [17], but precise information about solution quality.

Rodrigues and Gaston [3] offer compelling evidence
that the hardware and software limitations that in the past
have prevented optimal solvers from being used on rea-
listically large reserve selection problems have now been
largely overcome. Our results support that argument, with
CPLEX outperforming the SITES SA heuristic by larger
margins on larger problems. More research could profitably
be directed at increasing the use of optimal solvers for
existing and future reserve models.

Andelman et al. [1] claim as an advantage the multiple
solutions returned by SA. Rather than relying on chance to
produce a diversity of solutions of variable quality, another
area ripe for further research is to incorporate into reserve
selection models the existing literature in “modeling to
generate alternatives” (e.g., [40]). Having used existing
models to develop trade-off curves for a given application,
planners might then use related reserve selection models to
generate a number of optimally different alternate solutions
that all lie on or near the optimal trade-off curve. These
deliberately different alternatives are likely to offer rad-
ically different performance toward the sorts of unmodeled
objectives that are often important in reserve selection
problems. This is a systematic approach toward addressing
those concerns raised by Pressey et al. [18] that have not
already been answered by Rodrigues and Gaston [3].

889 **9. Data policy**

890 The Large Sierra datasets are available with the SITES
891 software distribution at [http://www.biogeog.ucsb.edu/
892 projects/tnc/toolbox.html](http://www.biogeog.ucsb.edu/projects/tnc/toolbox.html). The Small Sierra dataset is avail-
893 able at <http://www.geog.ucsb.edu/~fischer>. Use of the
894 small dataset should reference Davis et al. [34]. The Santa
895 Barbara dataset may be available in the future, depending
896 on county Policy (contact fischer@geog.ucsb.edu).

897 **Acknowledgements**

898 During this work, Douglas Fischer was supported in part
899 by grants from the U.S. Forest Service and the U.S. Geo-
900 logical Survey. We would like to acknowledge the support
901 of the Western Ecological Research Center of the U.S.
902 Geological Survey and thank Judd Howell and Deborah
903 Maxwell there. The U.S. Fish and Wildlife Service has also
904 provided support. We would also like to thank Frank
905 Davis, David Stoms, and Chris Pyke for providing data
906 and comments. Finally, thanks are due to both Pete Stine
907 and Klaus Barber of the U.S. Forest Service for their help
908 and support. The comments of Justin Williams and three
909 anonymous reviewers were very helpful in improving this
910 manuscript.
911

References

- 913 [1] S.J. Andelman, I. Ball, F.W. Davis and D.M. Stoms, SITES V 1.0:
914 an analytical toolbox for designing ecoregional conservation
915 portfolios, Unpublished manual prepared for the nature conservan-
916 cy, 1999, 1–43. (available at [http://www.biogeog.ucsb.edu/projects/
918 tnc/toolbox.html](http://www.biogeog.ucsb.edu/projects/
917 tnc/toolbox.html)).
- 918 [2] L.G. Underhill, Optimal and suboptimal reserve selection algo-
919 rithms, *Biol. Conserv.* 70 (1994) 85–87.
- 920 [3] A.S.L. Rodrigues and K.J. Gaston, Optimization in reserve
921 selection procedures – why not?, *Biol. Conserv.* 107 (2002)
922 123–129.
- 923 [4] B. Csuti, S. Polasky, P. Williams, R. Pressey, J. Camm, M.
924 Kershaw, A. Kiester, B. Downs, R. Hamilton, M. Huso and K.
925 Sahr, A comparison of reserve selection algorithms using data on
926 terrestrial vertebrates in Oregon, *Biol. Conserv.* 80 (1997) 83–97.
- 927 [5] R. Gerrard, R.L. Church, D.M. Stoms and F.W. Davis, Selecting
928 conservation reserves using species-covering models: adapting the
929 Arc/Info GIS, *Trans. Geol. Inf. Sci.* 2 (1997) 45–60.
- 930 [6] R. Church, D. Stoms, F. Davis and B.J. Okin, Planning manage-
931 ment activities to protect biodiversity with a GIS and an integrated
932 optimization model, in: *Proceedings, Third International Confer-
933 ence/Workshop on Integrating GIS and Environmental Modeling,
934 Santa Fe, NM, January 21–26, 1996* (National Center for
935 Geographic Information and Analysis, Santa Barbara, CA, 1996)
936 [http://www.ncgia.ucsb.edu/conf/SANTA_FE_CD-ROM/
938 main.html](http://www.ncgia.ucsb.edu/conf/SANTA_FE_CD-ROM/
937 main.html).
- 938 [7] R.L. Pressey and A.O. Nicholls, Efficiency in conservation
939 evaluation: scoring versus iterative approaches, *Biol. Conserv.* 50
940 (1989) 199–218.
- 941 [8] K.D. Cocks and I.A. Baird, Using mathematical programming to
942 address the multiple reserve selection problem: an example
943 from the Eyre Peninsula, South Australia, *Biol. Conserv.* 49
944 (1989) 113–130.
- [9] J.G. Hof and L.A. Joyce, A mixed integer linear programming
approach for spatially optimizing wildlife and timber in managed
forest ecosystems, *For. Sci.* 39 (1993) 816–834.
- [10] D. Faith, P. Walker, J. Ive and L. Belbin, Integrating conservation
and forestry production: exploring trade-offs between biodiversity
and production in regional land-use assessment, *For. Ecol. Manage.*
85 (1996) 251–260.
- [11] M. Bevers, J. Hof, D.W. Uresk and G.L. Schenbeck, Spatial
optimization of prairie dog colonies for black-footed ferret
recovery, *Oper. Res.* 45 (1997) 495–507.
- [12] C. Loehle, Optimizing wildlife habitat mitigation with a habitat
defragmentation algorithm, *For. Ecol. Manage.* 120 (1999)
245–251.
- [13] H. Possingham, I. Ball and S. Andelman, Mathematical methods
for identifying representative reserve networks, in: *Quantitative
Methods for Conservation Biology*, eds. S. Ferson and M. Burgman
(Springer-Verlag, Berlin Heidelberg New York, 2000) pp.
291–305.
- [14] R.G. Haight, C.S. ReVelle and S.A. Snyder, An integer optimiza-
tion approach to a probabilistic reserve site selection problem,
Oper. Res. 48 (2000) 697–708.
- [15] R.L. Church, R. Gerrard, A. Hollander and D.M. Stoms, Under-
standing the tradeoffs between site quality and species presence in
reserve site selection, *For. Sci.* 46 (2000) 157–167.
- [16] D.J. Nalle, J.L. Arthur and J. Sessions, Designing compact and
contiguous reserve networks with a hybrid heuristic algorithm, *For.
Sci.* 48 (2002) 59–68.
- [17] D.T. Fischer and R.L. Church, Clustering and compactness in
reserve site selection: an extension of the biodiversity management
area selection model, *For. Sci.* 49 (2003) 555–565.
- [18] R.L. Pressey, H.P. Possingham and C.R. Margules, Optimality in
reserve selection algorithms: when does it matter and how much?,
Biol. Conserv. 76 (1996) 259–267.
- [19] J.R. Prendergast, R.M. Quinn and J.H. Lawton, The gaps between
theory and practice in selecting nature reserves, *Conserv. Biol.* 13
(1999) 484–492.
- [20] R.L. Pressey and R.M. Cowling, Reserve selection algorithms and
the real world, *Conserv. Biol.* 15 (2000) 275–277.
- [21] C. Groves, L. Valutis, D. Vosick, B. Neely, K. Wheaton, J. Touval
and B. Runnels, *Designing a Geography of Hope: A Practitioners’
Handbook for Ecoregional Conservation Planning* (The Nature
Conservancy, Arlington, VA, 2000).
- [22] M.W. Beck and M. Odaya, Ecoregional planning in marine
environments: identifying priority sites for conservation in the
northern Gulf of Mexico, *Aquat. Conserv.* 11 (2001) 235–242.
- [23] S.J. Andelman and M.R. Willig, Alternative configurations of
conservation reserves for Paraguayan bats: considerations of spatial
scale, *J. Soc. Conserv. Biol.* 16 (2002) 1352–1363.
- [24] H. Leslie, M. Ruckelshaus, I.R. Ball, S. Andelman and H.P.
Possingham, Using siting algorithms in the design of marine
reserve networks, *Ecol. Appl.* 13 (2003) S185–S198 Suppl.
- [25] I.R. Ball, Mathematical applications for conservation ecology: the
dynamics of tree hollows and the design of nature reserves, Ph.D.
thesis, University of Adelaide, Adelaide, Australia, 2000.
- [26] M.D. McDonnell, H.P. Possingham, I.R. Ball and E.A. Cousins,
Mathematical methods for spatially cohesive reserve design,
Environ. Model. Assess. 7 (2002) 107–114.
- [27] R.L. Church and C.S. ReVelle, The maximal covering location
problem, *Pap. Reg. Sci. Assoc.* 32 (1974) 101–118.
- [28] N.A. Metropolis, M. Rosenbluth, A. Rosenbluth and E. Teller,
Equation of state calculations by fast computing machines, *J.
Chem. Phys.* 21 (1953) 1087–1092.
- [29] B.L. Golden and C.S. Skiscim, Using simulated annealing to solve
routing and location problems, *Naval Res. Logist. Q.* 33 (1986)
261–279.
- [30] A. Murray and R. Church, Applying simulated annealing to
location-planning models, *J. Heuristics* 2 (1996) 31–53.
- [31] D.T. Fischer, Clustering and compactness in reserve site

- 1013 selection: an extension of the Biodiversity Management Area Se- 1029
1014 lection Model, MA thesis, Univ. of Calif. Santa Barbara, 2001, 1030
1015 p. 48. 1031
- 1016 [32] S.H. Owen, M.S. Daskin, Strategic facility location: a review, Eur. 1032
1017 J. Oper. Res. 111 (1998) 423–447. 1033
- 1018 [33] ILOG. ILOG CPLEX 6.5 User’s Manual, Ilog Corp. Incline 1034
1019 Village, NV, 1999. 1035
- 1020 [34] F.W. Davis, D.M. Stoms, R.L. Church, W.J. Okin and K.N. 1036
1021 Johnson, Selecting biodiversity management areas, in: *Sierra* 1037
1022 *Nevada Ecosystem Project: Final Report to Congress, vol. II,* 1038
1023 *Assessments and Scientific Basis for Management Options* (Centers 1039
1024 for Water and Wildlands Resources, University of California, 1040
1025 Davis, 1996) pp. 1503–1527. 1041
- 1026 [35] Willis Linn Jepson, *The Jepson Manual: Higher Plants of* 1042
1027 *California*, ed. James C. Hickman (University of California Press, 1043
1028 Berkeley, 1993) p. 1400. 1044
- [36] W. Okin, The biodiversity management area selection model: 1029
constructing a solution approach, MA Thesis, Univ. of Calif. Santa 1030
Barbara, 1997, p. 67. 1031
- [37] J. Cohon, R.L. Church and D. Sheer, Generating multiobjective 1032
trade-offs: an algorithm for bicriterion problems, *Water Resour.* 1033
Res. 15 (1979) 1001–1010. 1034
- [38] R. Solanki, P.A. Appino and J.L. Cohon, Approximating the 1035
noninferior set in multiobjective linear programming problems. 1036
Eur. J. Oper. Res. 68 (1993) 356–373. 1037
- [39] Watershed Environmental, Rural resource protection project sensi- 1038
tive biological resources study, prepared for County of Santa 1039
Barbara Planning and Development Department, 2003. 1040
- [40] E.D. Brill, S.Y. Chang and L.D. Hopkins, Modelling to gener- 1041
ate alternatives – the HSJ approach and an illustration using a 1042
problem in land-use planning, *Manage. Sci.* 28 (1982) 221– 1043
235. 1044

UNCORRECTED PROOF

AUTHOR QUERY

AUTHOR PLEASE ANSWER QUERY.

Q1. Please provide a legend for the asterisks found in Table 3.

UNCORRECTED PROOF