Multiple Criterion Synchronisation for Conservation Area Network Design: The Use of Non-dominated Alternative Sets

Sahotra Sarkar and Justin Garson

*This article shows how a standard technique from multiple-criteria decision making, the computation of non-dominated alternative sets, can be adapted to incorporate non-biological criteria such as socio-political ones during the design of biodiversity conservation area networks. There are three chief advantages of this approach: (*a*) unlike almost all other methods for incorporating multiple criteria in decision making, this technique avoids making arbitrary utility assignments to alternatives; (*b*) it only requires comparative rankings of the alternatives under the criteria to be synchronised; and (*c*) it often results in several alternatives of equal status, leaving a further choice to be made by political decision-making bodies. This allows those bodies to bring into consideration criteria that cannot be satisfactorily formally modelled. The use of this method is demonstrated using a data set from Texas, one from Ecuador and two artificially constructed data sets.*

Introduction

The design of biodiversity conservation area networks (CANs) almost always occurs in the socio-political context of attempting to find compromises between designating a place for biodiversity conservation and allowing alternative uses. For instance, a place with high biodiversity value if left untransformed may also have great agricultural potential or recreational value if transformed. To choose any one of these potential uses involves a trade-off between what are often incompatible alternatives. To find such compromises requires the simultaneous

Acknowledgements: Thanks are due to James Justus, Chris Pappas and three anonymous referees, for comments on an earlier draft of this article. JG was supported by a Summer Research Grant from the Environmental Sciences Institute of the University of Texas at Austin.

Software availability: The software package used for this analysis, 'MultSync', can be obtained free of charge by contacting consbio@uts.cc.utexas.edu. ResNet Ver 1.2, the package mostly used to generate solutions for Ecuador and Texas can be downloaded from http://uts.cc.utexas.edu/~consbio/Cons/ResNet.html.

Sahotra Sarkar and **Justin Garson** are at the Biodiversity and Biocultural Conservation Laboratory, Section of Integrative Biology, University of Texas at Austin, 1 University Station, C3500, Austin, TX 78712 –1180, USA. E-mail: sarkar@mail.utexas.edu.

Conservation & Society, **2, 2 (2004)**
SAGE Publications New Delhi/Thousand Oaks/London

evaluation of a plan or 'alternative' using multiple criteria. This problem has traditionally been studied in economic contexts and has variously been called multiple objective (or multi-objective) decision making (Bogetoft and Pruzan 1997; Jannsen 1992; Keeney and Raiffa 1993), multi-criterion decision making (Arrow and Raynaud 1986), multi-objective programming (Cohon 1978) or multi-attribute utility evaluation (Edwards and Newman 1982), with many other variants. Here, for reasons that will be discussed later, the problem will be called multiple criterion synchronisation (MCS).

For biodiversity conservation planning, the importance of this problem cannot be stated too strongly. Biodiversity conservation is as much a socio-political enterprise as a scientific one. Consequently, socio-political criteria, which are usually incommensurable and often incomparable, are central to such planning. There are deep ethical reasons why these criteria should be explicitly incorporated into all discussion of biodiversity conservation (Sarkar 2004). However, even from a purely pragmatic or prudential perspective, the importance of incorporating these criteria is beyond controversy: it is a truism—sometimes learnt very grudgingly by Western conservation biologists—that conservation measures do not succeed without socio-political support.

There are three methods that have commonly been used to attempt to resolve problems involving the use of multiple criteria: (*a*) reduction to a single utility function; (*b*) reduction to a parameterised set of utility functions with accompanying robustness/sensitivity analyses; and (*c*) identification of a set of 'non-dominated' alternatives, sometimes followed by further refinement of that set. It will be pointed out below that methods (*a*) and (*b*) are often open to the charge of arbitrariness, whereas method (*c*) can avoid that problem in the context of CAN design. A running theme throughout this article will be the value of avoiding arbitrariness in decision-making protocols. However, because the emphasis here is on ensuring non-arbitrariness, this article is largely restricted to the most elementary step of decision-making protocols that begin with the set of non-dominated alternatives and then refine it using powerful but partly arbitrary heuristics.

Method (*a*) originates in neoclassical economics where it is assumed that all values can be assessed on a single quantitative scale of utility (or cost, sometimes called the 'common currency'). Typically, this assumes that all values associated with an alternative can be assigned a monetary cost. These values include the biodiversity value of a place, its recreational value and all other values. This utility is supposed to be measured by agents' preferences, which can be estimated by what they are willing to pay (WTP) for these values or the compensation they require to give them up (Silberberg et al. 2000).

As an example, consider the system described in Table 1. There are three CANs (CAN_1, CAN_2 and CAN_3) that are to be evaluated using three criteria: area, measured in sq. km; forgone opportunity of conservation, measured in dollars; and social cost, measured in the number of households displaced because of the implementation of a conservation plan. Suppose that land costs \$100/sq. km and that each household is willing to accept \$10,000 in compensation for displacement.

With these assumptions, each CAN may be attributed a cost in dollars according to each criterion, as indicated in Table 2. CAN_2 emerges as the optimal choice, that with the lowest cost ($532,000).

Table 1
Area (sq. km), Forgone Opportunity of Conservation ($) and Social Cost

Criteria	*Area (in sq. km)*	*Forgone opportunity (in $)*	*Social cost (in displaced households)*
CAN_1	2,700	100,000	50
CAN_2	2,200	92,000	22
CAN_3	1,400	144,000	75

Table 2
Area, Forgone Opportunity and Social Cost (all in $)

Criteria	*Area (in $)*	*Forgone opportunity (in $)*	*Social cost (in $)*
CAN_1	270,000	100,000	500,000
CAN_2	220,000	92,000	220,000
CAN_3	140,000	144,000	750,000

There are many conceptual and empirical problems associated with this approach (Norton 1994). Perhaps the most important conceptual problem is that this method assumes that all the relevant values are commensurable (that is, they can be assessed on the same scale) (Sarkar 2004). Empirical issues include the difficulty of carrying out WTP estimation and similar assessments, as well as the intransitivity and temporal instability of such preferences (Norton 1987). Whatever be the merits or problems associated with such uses of preferences, in practise in environmental contexts, systematic WTP or other assessments are almost never carried out because of these difficulties. Rather, specific utility functions are usually chosen by educated intuition.

In environmental contexts this method has most credibly been used to assess the value of ecosystem services for comparison with the value of forgone economic opportunities when an ecosystem is left undeveloped. It has occasionally been used to identify high-priority biodiversity conservation areas, for instance, by Sierra et al. (2002) in the case of continental Ecuador. However, since systematic WTP and similar assessments were not carried out, such exercises are open to charges of arbitrariness. As a result, in recent years, most attempts to develop systematic methods for environmental decision evaluation have turned to more sophisticated analyses.

Method (*b*) also involves using an utility function, but each additive component of that function represents a different criterion and has a weight associated with it. These weights are used for robustness/sensitivity analyses. Suppose that a set of alternatives is ranked using this utility function with one set of weights. Each weight can then be varied across its entire range to test the invariance of the ranking. High invariance indicates robustness, that is, a high reliability of the ranking. More importantly, the weights can be used to analyse trade-offs between the criteria. In the context of CAN design, this method has been advocated by Faith (1995) and used in several contexts, including Papua New Guinea (Faith et al. 2001).

Returning to the example discussed earlier, Table 3 shows the normalised value of each CAN using the values from Table 1. Suppose area, forgone opportunity and social cost are given weights of 0.7, 0.15, and 0.15 respectively. The weighted costs are shown in Table 4 and CAN_3 is the optimal alternative. However, if the three criteria are given weights of 0.5, 0.25 and 0.25 respectively, the weighted costs are shown in Table 5. CAN_2 is the optimal solution. There is no set of positive weights under which CAN_1 can emerge as the optimal solution. However, because of the change of ranking induced by variation in the relative weights, neither CAN_2 nor CAN_3 can be said to be robustly preferred.

Table 3
Normalised Values of Each CAN

Criteria	*Area (normalised)*	*Forgone opportunity (normalised)*	*Social cost (normalised)*
CAN_1	0.43	0.30	0.34
CAN_2	0.35	0.27	0.15
CAN_3	0.22	0.43	0.51

Table 4
Weighted Costs of Each CAN

Criteria	*Area (normalised)*	*Forgone opportunity (normalised)*	*Social cost (normalised)*	*Sum*
CAN_1	0.301	0.0450	0.0510	0.397
CAN_2	0.245	0.0405	0.0225	0.308
CAN_3	0.154	0.0645	0.0765	0.295

Table 5
Weighted Costs of Each CAN

Criteria	*Area (normalised)*	*Forgone opportunity (normalised)*	*Social cost (normalised)*	*Sum*
CAN_1	0.215	0.0750	0.0850	0.375
CAN_2	0.175	0.0675	0.0375	0.280
CAN_3	0.110	0.1075	0.1275	0.345

Method (*b*) is an advance over method (*a*). Nevertheless, it does not fully address the possible incommensurability of values. Moreover, it only provides unequivocal results in the (mathematically) trivial case. Suppose that for two alternatives, α_1 and α_2, and two criteria, κ_1 and κ_2, ω_1 and ω_2 are the associated weights. α_1 is unequivocally preferable to α_2 if and only if its utility is higher than that of α_2 for every value of ω_1 and ω_2. This is the situation in which there is no conflict between the criteria, which is both rare and uninteresting. In any other situation some choice will have to be made as to how the values of ω_1 and ω_2 will be restricted. This choice will almost always be partly arbitrary. For instance, Faith (1995) used biodiversity and economic value as the two criteria that are usually in conflict. Unless some alternative is unequivocally better than another, it is hard to see how the weights for such criteria can be restricted non-arbitrarily.

Method (*c*) will be used here. Though it is hardly new, it has very rarely been advocated in the context of CAN design. Rothley (1999) provides an important early exception that will be discussed later; (see also, Cameron 2003 and Sarkar et al. 2000). It has two advantages over the previous two methods: (*a*) it makes no attempt to compound different criteria, thereby allowing them to be treated as genuinely incommensurable; and (*b*) because it does not use a utility function, it only requires that alternatives be ranked (that is, weak linear ordered, generating an ordinal but not a metric structure) by the criteria. It does not require an attribution of definite numerical values by some metric. This is important because, for some criteria (for instance, biodiversity content) it is possible to determine non-arbitrarily whether some alternative is better than (the same as, or worse than) some other alternative without being able to determine by how much in a non-arbitrary fashion. A third feature of method (*c*) is that it often selects more than one alternative that is often desirable in the context of CAN design though it can also present problems.

The formalism used for method (*c*) will be developed in the next section. Here, the basic ideas will be sketched (for more detail see, for example, Bogetoft and Pruzan 1997). The first step in the process is to select a set of 'feasible' alternatives. These may be the entire set of alternatives initially considered or they may be only those alternatives that satisfy some global constraint (distinct from the criteria to be synchronised). For instance, in the case of CAN design, the feasible alternatives may be the sets of places that meet an imposed target of representing biodiversity surrogates or they may be those that can be acquired within a monetary budget (Rothley 1999). The next step is to rank each choice according to each criterion. It will be assumed here that a lower rank is always more desirable. (This is not a limitation: if higher ranks are desired according to some criterion as originally formulated, the inverse of that criterion can be used to make lower ranks more desirable. The data set from Ecuador analysed in the third section will illustrate this transformation.)

An alternative, α_1, 'dominates' (or 'weakly dominates') another alternative, α_2, if it is ranked lower than the latter by at least one criterion and ranked at least as low as the latter by all criteria. Thus, if α_1 dominates α_2, there is no criterion by which α_2 can be regarded as better than α_1 because α_2 is never ranked lower than α_1. The set of 'non-dominated' alternatives consists of those alternatives that are not dominated by any alternative. (Formally, the set of non-dominated alternatives is a multidimensional discrete analogue of the indifference curves of elementary economics.) The determination of the non-dominated set introduces no arbitrary assumption provided that the elements of the feasible set can be non-arbitrarily ranked by each criterion. Note that, in this process, there is no mathematical sense in which the various criteria are being optimised. They are not being compounded and no utility function is being estimated. Rather, the criteria are being brought into synchrony in the sense that, as a totality, they do not allow a conflict between non-dominated alternatives. This is the reason why the procedure is being called ‘multiple criterion synchronisation’. (However, the term ‘synchronisation’ carries with it a temporal connotation; this is a disadvantage of the terminology advocated here.)

Returning to the earlier example, Table 6 shows the ordinal rankings of the three CANs according to the three criteria using the values in Table 1 to assign rankings. The integers in the table refer to the ranking (1 is lower than 2, and so on); their quantitative magnitudes (2 is twice 1) play no role in the subsequent analysis. Both CAN_2 and CAN_3 are non-dominated, which is consistent with the results found using method (*b*).

Table 6
Ordinal Rankings of Each CAN

Criteria	*Area (ranked ordinally)*	*Forgone opportunity (ranked ordinally)*	*Social cost (ranked ordinally)*
CAN_1	3	2	2
CAN_2	2	1	1
CAN_3	1	3	3

Ordinarily, the non-dominated set will consist of more than one alternative. If a unique solution is required, this will be a problem. Fortunately, in the case of CAN design, unique solutions are not only not necessary but often undesirable. Ultimately, decisions to designate individual places for biodiversity conservation rest with political decision-making bodies. The role of conservation biologists is to present these decision-making bodies with sets of options that all satisfy the biological targets and other formal criteria (just as a preference for low economic cost) that have been set. The decision-making bodies can now bring other considerations, including criteria that are not amenable to formal modelling, to bear upon the problem of selecting one of the options. Having a set of options that are indistinguishable by the biological and other formal criteria is usually an advantage provided that the set is not intractably large. (Rothley [1999] has also emphasised this aspect of using non-dominated solutions.)

In the two non-artificial data sets that are analysed in the third section, there were eleven and four alternatives in the non-dominated sets which are not unreasonable numbers of alternatives to present to decision-making bodies. However, it is possible, as the artificial data sets also analysed demonstrate, that there may be situations in which the non-dominated set is intractably large. In general, the greater the number of criteria, the larger the size of the non-dominated set (see the second section). A wide variety of techniques have been developed during the last two decades to refine the non-dominated set. (For the earlier work, see Dyer et al. [1992] and the references therein; a more recent review is Bogetoft and Pruzan [1997]. Rothley [1999] specifically advocates the use of the simple multi-attribute rating technique [SMART] in the context of CAN design.) However, none is free from partly arbitrary assumptions that usually involve the imposition of a metric structure on the performance of the alternatives under the criteria. That imposition will be avoided here; instead, it will be assumed that the criteria can themselves be ordered on the basis of their importance. (There are well-known techniques for using pair-wise comparisons for generating such rankings, for instance, during the analytical hierarchy process [for example, Saaty 1980; for an application to CAN design, see Anselin et al. 1989]. However, none are entirely non-arbitrary.) If the non-dominated set is intractably large, criteria will be iteratively dropped in inverse order of their importance, and a new non-dominated set will be computed from the current one. The process terminates when the size of the non-dominated set becomes acceptably small. (If the criteria cannot be ordered, this process can be carried out by dropping them at random.) This process is also not fully non-arbitrary since this is not the only possible way to model how the order of criteria should enter the selection process.

Finally, MCS for CAN design starts with a set of places and lists of biodiversity surrogates (such as species) at each place. These lists capture the biodiversity 'content' of an individual place. Besides other criteria that have to be synchronised, there is always a biodiversity optimisation target that takes one of two canonical forms (Sarkar et al. 2003): (*a*) all surrogates must meet a pre-assigned target of representation in as few places as possible; or (*b*) in a specified number of places, as many surrogates must meet their target as possible. Place prioritisation algorithms are designed to solve these problems (Margules et al. 1988; Margules and Pressey 2000; Sarkar et al. 2003). These algorithms iteratively add places to a list sequentially until the target is satisfied. MCS can take place either at the iterative step or after entire sets of places that satisfy the biodiversity targets have been selected. In the former case, the feasible alternatives are all the original places taken individually. In the latter case, the feasible alternatives are sets of places that satisfy the biodiversity targets. For brevity, only the latter case will be explored here; the former case does not introduce any new conceptual or methodological issue. This choice gives biodiversity content preferential treatment compared to the other criteria that are synchronised with respect to each other. (Faith [1995] and Faith et al. [2001] provide examples, in the context of method (*b*), in which biodiversity content is treated on par with the other criteria.) The criteria that now have to be synchronised include potential threat to habitat, economic cost, social cost, etc., all of which have to be kept as low as possible.

THEORY

The theoretical basis for the procedure used here to carry out the computations described in the last section is straightforward. MCS starts with two structures, a set of criteria, $K = \{\kappa_i : i = 1, 2, \ldots, n\}$, and a set of feasible alternatives, $A = \{\alpha_j : j = 1, 2, \ldots, m\}$. It is assumed that each criterion κ_i induces a weak linear order $\leq_\iota^*$ on A, that is, for $\forall \kappa_i \in K$ and $\forall \alpha_j, \alpha_k, \alpha_l \in A$, $\alpha_j \leq_i^* \alpha_k$ or $\alpha_k \leq_i^* \alpha_j$ or both; if $\alpha_j \leq_i^* \alpha_k$ and $\alpha_k \leq_i^* \alpha_l$, then $\alpha_j \leq_i^* \alpha_l$ (that is, $\leq_i^*$ is transitive). Each $\leq_i^*$ thus imposes an ordinal structure on A giving ranks to the α_j ($j = 1, 2, \ldots, m$). However, since the order is only weak, two different alternatives may have the same rank. (Table 6 shows an example of such an order.)

Let v_{ij} be the rank of alternative α_j by the criterion κ_i. The v_{ij} form an $(n \times m)$–matrix, N. For all alternatives, α_e and α_f, $e \neq f$, α_e dominates α_e or $\alpha_f \phi \alpha_f$ if and only if:

$$(\exists i)(v_{ie} < v_{if}) \wedge (\forall k)(v_{ke} \leq v_{kf}) \tag{1}$$

Therefore, for all alternatives, α_e and α_f, $e \neq f$, α_e does not dominate α_f or $\neg(\alpha_e \phi \alpha_f)$ if and only if:

$$\begin{aligned} &\neg((\exists i)(v_{ie} < v_{if}) \wedge (\forall k)(v_{ke} \leq v_{kf})) \\ \Leftrightarrow\ &\neg(\exists i)(v_{ie} < v_{if}) \vee \neg(\forall k)(v_{ke} \leq v_{kf}) \\ \Leftrightarrow\ &(\forall i)\neg(v_{ie} < v_{if}) \vee (\exists k)\neg(v_{ke} \leq v_{kf}) \\ \Leftrightarrow\ &(\forall i)(v_{ie} \geq v_{if}) \vee (\exists k)(v_{ke} > v_{kf}) \end{aligned} \tag{2}$$

It is assumed, by convention, that $\neg(\alpha_e \phi \alpha_e)$. The last form of expression (2) provides a straightforwardly implemented algorithm for finding the set of non-dominated alternatives. Using the matrix N, *an* $(m \times m)$ dominance matrix, $\Delta = (\delta_{ef})$ can be computed as:

$$\delta_{ef} = \begin{cases} 1 & \text{if} \quad (\exists i)(v_{ie} < v_{if}) \wedge (\forall k)(v_{ke} \leq v_{kf}) \\ 0 & \text{if} \quad (\forall i)(v_{ie} \geq v_{if}) \vee (\exists k)(v_{ke} > v_{kf}) \end{cases} \tag{3}$$

that is, $\delta_{ef} = 1$ if $\alpha_e \phi \alpha_f$ (α_e dominates α_f) and $\delta_{ef} = 0$ otherwise. Alternatives for which an entire column of Δ consists only of 0s are non-dominated. Equation (3) provides an algorithm for computing the set of non-dominated alternative. Given the matrix, N, δ_{ef} can be computed by evaluating each disjunct/conjunct of equation (3) by iteratively counting how often each inequality is satisfied. (Computationally, this algorithm is $o(n^2m)$.)

Let A^* be the set of non-dominated alternatives. From expression (3) it follows that any alternative that is ranked least by even one criterion (and there may be more than one for each criterion) belongs to A^*. In the fourth section alternatives that enter A^* only because of this property for a single criterion will be called 'extrema'. In general, the cardinality of A^* will increase with n. However, this is not always be the case. Let A^*_p and A^*_{p-1} be the non-dominated sets when p and $(p-1)$ criteria are used, respectively. Let $\alpha_r \in A^*_p$ and $\alpha_r \in A^*_{p-1}$. Let κ_s be the criterion that distinguishes A^*_p and A^*_{p-1}. This means that, there exists an α_j, $j \neq r$,

$$v_{sj} < v_{sr} \tag{4}$$

and also, for all κ_l, $l \neq s$,

$$v_{lr} < v_{lj} \tag{5}$$

Conditions (5) are stringent and the situation in which the size of the non-dominated set decreases when the number of criteria increases is likely to be rare.

A numerical example will clarify the last point. Let there be three criteria and three alternatives. Let each alternative have only two possible ranks (1 or 2) under each of the criteria. Let $N = \begin{pmatrix} 2 & 2 & 2 \\ 2 & 2 & 2 \\ 1 & 2 & 2 \end{pmatrix}$; then $\Delta = \begin{pmatrix} 0 & 1 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$, and only the first alternative is non-dominated. Now suppose that the third criterion is dropped. Then $N' = \begin{pmatrix} 2 & 2 & 2 \\ 2 & 2 & 2 \end{pmatrix}$ and $\Delta' = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}$ and all three alternatives are non-dominated. However, this requires all entries of N' to be equal which is a stringent requirement.

As mentioned in the last section, if the cardinality of A^* is intractably large, A^* must be further refined for practical use. The procedure suggested there consists of starting with the current non-dominated set, dropping criteria iteratively, and computing a new non-dominated set from the old one. This procedure guarantees that the cardinality of new non-dominated set will not be higher than that of the old one. However, it should be emphasised that, unlike the computation of the full non-dominated set A^*, this procedure does not have theoretical justification and must be regarded only as a heuristic.

The method developed here should be contrasted with that of Rothley (1999). She advocated the use of multi-objective programming (MOP) to compute the set of feasible alternatives in contrast to the heuristic algorithms used in the next section. Once the feasible set is obtained, the remaining problem is that of finding the non-dominated alternatives. Rothley solved this problem by visual inspection (using the equivalent of the MCS plots described below [Figures 1 and 2]). The method developed here is designed to replace such intuitive techniques that cannot be reliably used when the number of criteria is large.

Figure 1
MCS Plot for Ecuador

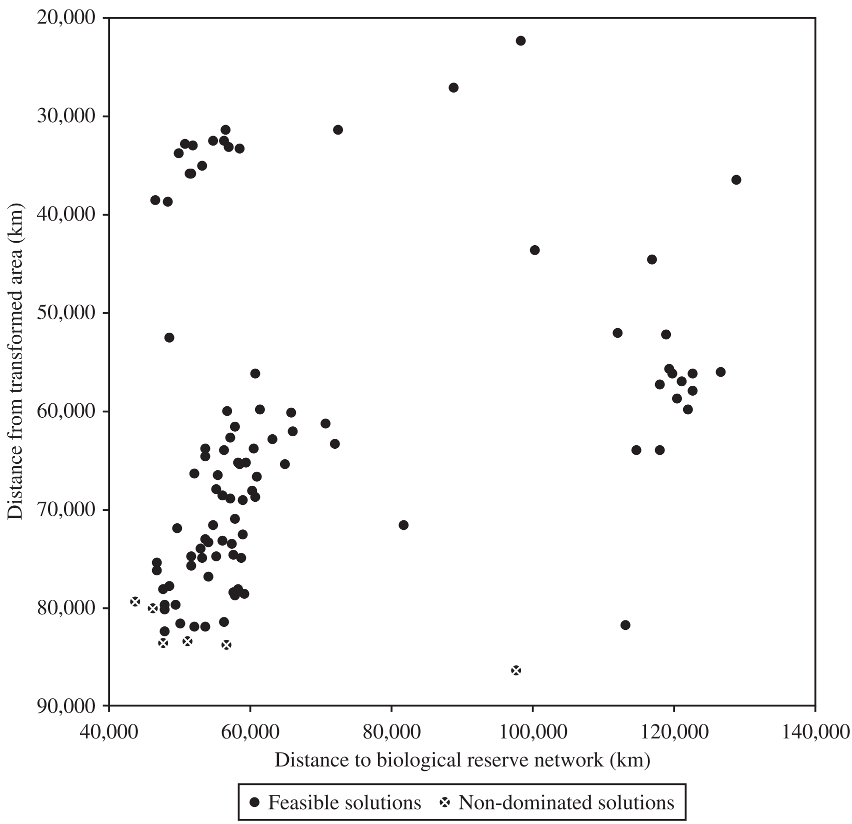

Note: *X*-axis: compounded (summed) distance from one of the units within the national biological reserve system of Ecuador; *y*-axis: compounded distance from anthropogenically transformed areas. The non-dominated solutions are those that have no other alternatives to the left or to the bottom of them.

Figure 2
MCS Plot for Texas

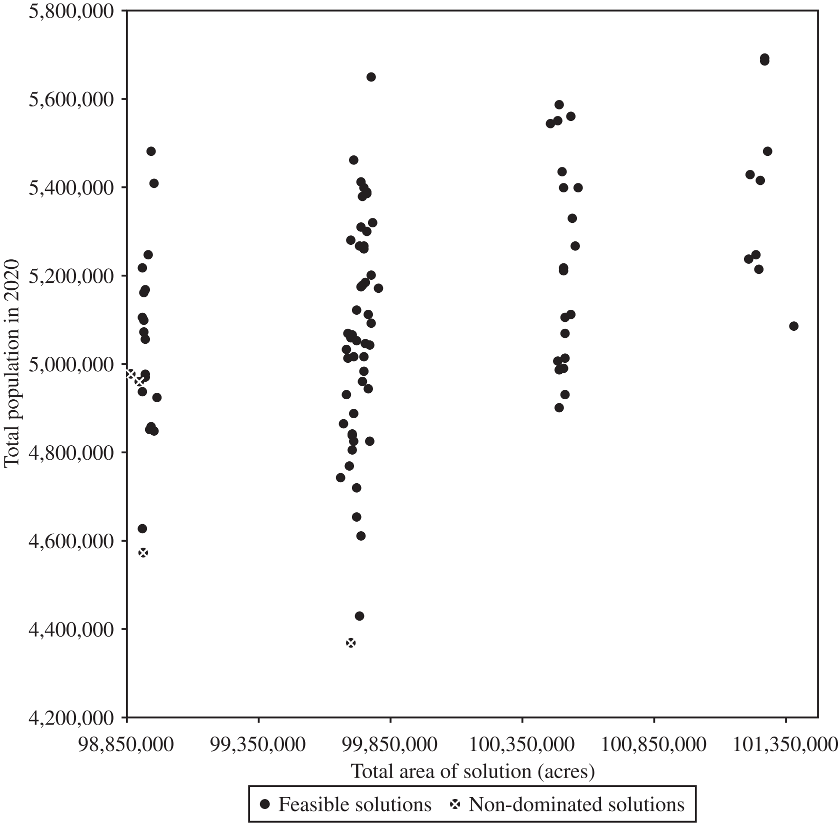

Note: *X*-axis: total area of the selected sites; *y*-axis: total projected population in 2020 in the selected sites. The non-dominated solutions are those that have no other alternatives to the left or to the bottom of them.

Examples

The MCS procedure outlined in the last section was used to analyse four data sets, two natural and two artificially constructed sets:

The first data set analysed comes from Ecuador and was used by Cameron (2003) to prioritise new places for conservation action. The original data consisted of the geographical distribution of forty-six vegetation types. This data was used to devise 100 different plans (or alternatives) that achieved a 10 per cent representation of each vegetation type within a CAN. Each alternative consists of a set of areas, each of size 4 sq. km, beyond those already within the existing biological reserve network of Ecuador. Thus, all these alternatives equally satisfy the basic biodiversity representation criterion. For each alternative, there are three other criteria of potential practical importance for conservation: (*a*) the summed distance from an anthropogenically altered area; (*b*) the summed distance to an existing reserve; and (*c*) the total area. Criterion (*a*) can be interpreted in at least two ways: (*a*) as representing the inverse risk to a place; or (*b*) as an inverse measure of the potential cost of incorporation into a conservation area network. Either way, it should be maximised or its inverse minimised. Criterion (*b*) can also be interpreted in two ways: (*a*) as inversely representing the ease with which a place can be incorporated into a network; or, (*b*) again inversely, as the probability of persistence of biota. Thus, it should be minimised. Criterion (*c*) is straightforwardly interpreted as providing an estimate of the cost of acquisition and forgone opportunity that should be minimised. Numerically, the 100 alternatives showed high variation under criteria (*a*) and (*b*), but were almost equivalent under criterion (*c*).

When all three criteria are used, there are eleven non-dominated alternatives. Because of the near-equivalence of the alternatives under criterion (*e*), it makes sense to refine the non-dominated set further by dropping that criterion. When this is done sequentially, there remain six non-dominated solutions. When there are only two criteria being considered, a helpful visualisation of the non-dominated alternatives is an MCS plot (see Figure 1). Non-dominated solutions are those that have no other solution to their left or below. If criterion (*e*) is not used from the beginning, the same six non-dominated solutions are obtained in this case. This means that removing the third criterion iteratively is functionally equivalent to ignoring it altogether.

The second data set consists of the modelled distribution of 655 animal species in Texas produced by the Texas GAP Analysis Project and previously analysed by Sarkar et al. (2000). These distributions were known for 1,183 hexagons with an average area of 649 sq. km. Using these data, purely as an exercise because such hexagons are far too large in area to be effective units of conservation planning, 100 different plans or alternatives were generated, each representing at least 10 per cent of the distribution of each species. For each alternative, two additional criteria were available: (*a*) the projected human population in 2020 in that set of hexagons; and (*b*) the total area of the set. Criterion (*a*) can be interpreted as a measure of the social or political cost of conservation that should be minimised. Alternatively, and not inconsistently, it can be interpreted as a measure of future threat. Similarly, criterion (*b*) can be taken as a measure of the economic cost of conservation, which should also be minimised. There were four non-dominated solutions that are shown in the MCS plot of Figure 2. Any further refinement of this non-dominated set would result in the use of only one criterion and the lowest alternative on that criterion would be selected. This would be equivalent to using only that criterion from the beginning.

In the artificial data set, there were ten criteria and 100 alternatives. Thus, N is a (10×100) matrix with each row corresponding to a criterion and column corresponding to an alternative. Each entry of the matrix was a random integer drawn from 1 to 100 and naturally interpreted as the rank of an alternative under a criterion. (This method of assigning ranks to entries of N allows the possibility of ties in rank for two alternatives under some criterion though this is unlikely.) There were ninety-three non-dominated alternatives, which is intractably large in most planning contexts.

Next, the refinement procedure was executed on this data set. It was assumed that the order of the criteria reflected their importance; that is, criterion (10) is the most expendable one and criterion (1) is the least expendable one. It was removed from consideration and the non-dominated set alone was recomputed for synchronisation. There were now eighty-six non-dominated alternatives; if criterion (10) is left out originally, the same non-dominated alternatives are obtained. Next criterion (9) was removed and the result was eighty non-dominated alternatives; the solution is the same as when the two criteria were removed simultaneously. Table 7 lists the performance of the MCS procedures as criteria are iteratively removed. Very little progress is made until only three criteria are left. In every case, iterative and simultaneous removal gave the same results. Once again, iteratively removing criteria was functionally equivalent to ignoring them altogether.

Table 7
Performance of MCS Procedures

Number of criteria	$1 \le v_{ij} \le 100$	$1 \le v_{ij} \le 5$ *(iterative removal)*	$1 \le v_{ij} \le 5$ *(simultaneous removal)*
10	93	65	65
9	86	59	59
8	80	48	48
7	76	36	36
6	67	18	18
5	48	7	8
4	29	1	2
3	19	1	3
2	6	1	10
1	1	1	32

Note: The first column shows the number of criteria that were being synchronised in the artificial data set. The second shows the size of the non-dominated set if the criteria were removed from the first artificial data set. (It does not matter whether the criteria are removed iteratively or simultaneously). The third shows the size if the criteria are removed iteratively from the second artificial data set; the fourth if they are removed simultaneously.

The second artificial data set differs from the first only insofar as the ranks in N were randomly selected between one and five, thus encouraging ties between alternatives. There initially were sixty-five non-dominated alternatives. Table 7 also shows the results of removing criteria iteratively. After three removals, a tractable set of eighteen non-dominated alternatives is found. It also shows the results of simultaneous removal. For the first three removals, the results are the same, after which they diverge radically. In this case the ranking of the criteria makes a difference.

Discussion

The main advantages of the MCS procedure described here are the two features mentioned in the first section: First, so long as the non-dominated set is sufficiently small as not to require any further refinement, unlike almost all other multiple criteria decision-making methods, this procedure involves no arbitrary assumption about how criteria should be compounded, utilities of alternatives assessed and so on. It is encouraging that both natural data sets and associated criteria returned tractably small non-dominated sets. This suggests that the MCS procedure can be valuable in practise. Thus, there is a non-arbitrary decision-making protocol in CAN design that can accommodate non-biological criteria.

Second, the method only requires that the alternatives be in a weak linear order by each of the criteria. It does not even require strong ordering, let alone the assignment of definite quantitative values to alternatives under the criteria, an assignment that is almost always partly arbitrary but required by the first two methods discussed in the first section. For a weak linear order to be imposed on a set of alternatives by a criterion, all that is required is a qualitative comparative evaluation of their performance under that criterion. Moreover, there are many systematic methods to incorporate the lack of evaluative results for all comparisons (Triantaphyllou 2000). (Mathematically, this amounts to a conversion of a partial order to a linear order.)

Further, because this method does not compound criteria, it does not matter whether the different criteria are independent of each other. (In two dimensions, on an MCS plot, a correlation between criteria can be modelled by having the two axes not be perpendicular to each other. What was to the left or below a point remains invariant under such a coordinate transformation.) Nevertheless, there are two aspects of this procedure that are not fully satisfactory:

1. Extrema automatically get included in the solution set because they are non-dominated alternatives. For instance, in Figure 1, one of the non-dominated alternatives is an outlier (at the extreme right) because it is the sole extremum with respect to the *x*-axis. It is at least arguable that alternatives such as these do not truly offer a synchronisation of the different criteria. One option worth considering is the exclusion of those non-dominated alternatives that have that status only because of being extrema. This may be a relatively non-arbitrary way of refining a non-dominated set. The possible justification for taking this option, and exploring other possibilities for dealing with extrema remains a task for the future.

2. There is as yet no fully non-arbitrary method to incorporate a ranking of the criteria. The computation of the full non-dominated set makes no use of any such ranking even if available. Thus, there may be relevant information about synchronisation that is not being used in the MCS procedure. It would be ironic if such information cannot be non-arbitrarily used in decision-making protocols.

The last point is particularly relevant when the cardinality of the non-dominated set is high and it has to be further refined for practical use. Given the amount of effort expended on refinement methods during the last few decades (Bogetoft and Pruzan 1997; Dyer et al. 1992), no general non-arbitrary method for refinement seems forthcoming. However, it remains an open question whether the biological aspects of the problem of CAN design can be used to develop a theoretically justified procedure. As the analysis of the first artificial data set shows, the method for further refinement of the non-dominated set used here is less than satisfactory, sometimes amounting to ignoring some of the criteria altogether. Given that the second artificial data set does not exhibit this property to the same extent, it may be an artefact because of the particular random construction of the first data set, though the same result is also obtained with the data set from Ecuador. The development of a theoretical framework for incorporating the ranking of criteria in the synchronisation process without making arbitrary choices is another important task for the future, and the use of specific biological knowledge may be necessary for further progress.

References

Anselin, A., P.M. Meire and L. Anselin (1989), 'Multicriteria Techniques in Ecological Evaluation: An Example Using the Analytical Hierarchy Process', *Biological Conservation*, 49: 215–29.

Arrow, K.J. and H. Raynaud (1986), *Social Choice and Multicriterion Decision-making*. Cambridge, MA: MIT Press.

Bogetoft, P. and P. Pruzan (1997), *Planning with Multiple Criteria: Investigation, Communication and Choice*. Copenhagen: Copenhagen Business School Press.

Cameron, S. (2003), 'Place Prioritization and Irreplaceability Analysis for Biodiversity Conservation in Ecuador Using Modeled Vegetation Classes as Surrogates', senior honours thesis, Austin TX: Department of Biology, University of Texas at Austin.

Cohon, J.L. (1978), *Multiobjective Programming and Planning*. New York: Academic Press.

Dyer, J.S., P.C. Fishburn, R.E. Steuer, J. Wallenius and S. Zionts (1992), 'Multiple Criteria Decision Making, Multiattribute Utility Theory: The Next Ten Years', *Management Science*, 38: 645–54.

Edwards, W. and R.J. Newman (1982), *Multiattribute Evaluation*. London: Sage Publications.

Faith, D.P. (1995), *Biodiversity and Regional Sustainability Analysis*. Lyneham: CISRO Division of Wildlife and Ecology.

Faith, D.P., C.R. Margules and P.A. Walker (2001), 'A Biodiversity Conservation Plan for Papua New Guinea Based on Biodiversity Trade-Offs Analysis', *Pacific Conservation Biology*, 6: 304–24.

Janssen, R. (1992), *Multiobjective Decision Support for Environmental Management*. Dordrecht: Kluwer.

Keeney, R.L. and H. Raiffa (1993), *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge: Cambridge University Press.

Margules, C.R., A.O. Nicholls and R.L. Pressey (1988), 'Selecting Networks of Reserves to Maximize Biological Diversity', *Biological Conservation*, 43: 63–76.

Margules, C.R. and R.L. Pressey (2000), 'Systematic Conservation Planning', *Nature*, 405: 243–53.

Norton, B.G. (1987), *Why Preserve Natural Variety?* Princeton: Princeton University Press.

——— (1994), 'Economists' Preferences and the Preferences of Economists', *Environmental Values*, 3: 313–32.

Rothley, K.D. (1999), 'Designing Bioreserve Networks to Satisfy Multiple, Conflicting Demands', *Ecological Applications*, 9: 741–50.

Saaty, T. L. (1980), *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. New York: McGraw-Hill.

Sarkar, S. (2004), *Biodiversity and Environmental Philosophy*. New York: Cambridge University Press.

Sarkar, S., N.C. Parker, J. Garson, A. Aggarwal and S. Haskell (2000), 'Place Prioritization for Texas Using GAP Data: The Use of Biodiversity and Environmental Surrogates within Socioeconomic Constraints', *Gap Analysis Program Bulletin*, 9: 48–50.

Sarkar, S., C. Pappas, J. Garson, A. Aggarwal and S. Cameron (2003), 'Place Prioritization for Biodiversity Conservation Using Probabilistic Surrogate Distribution Data', *Diversity and Distributions*, 10: 125–33.

Sierra, R., F. Campos and J. Chamberlin (2002), 'Assessing Biodiversity Conservation Priorities: Ecosystem Risk and Representativeness in Continental Ecuador', *Landscape and Urban Planning*, 59: 95–110.

Silberberg, E., W. Suen and W.C. Suen (2000), *The Structure of Economics: A Mathematical Approach*. New York: McGraw-Hill.

Triantaphyllou, E. (2000), *Multi-criteria Decision Methods: A Comparative Study*. Dordrecht: Kluwer.