Incorporating Economic and Ecological Information into the Optimal Design of Wildlife Corridors

Jon Conrad[1] Carla P. Gomes[1,2]
Willem-Jan van Hoeve[3] Ashish Sabharwal[2] Jordan F. Suter[4]

[1] Dept. of Applied Economics and Management, Cornell University, Ithaca NY
[2] Dept. of Computer Science, Cornell University, Ithaca NY
[3] Tepper School of Business, Carnegie Mellon University, Pittsburgh PA
[4] Dept. of Economics, Oberlin College, Oberlin OH
jmc16@cornell.edu, gomes@cs.cornell.edu,
vanhoeve@andrew.cmu.edu, sabhar@cs.cornell.edu, jordan.suter@oberlin.edu

July 20, 2010

Abstract

In an attempt to address the negative ecological impacts of habitat fragmentation, wildlife corridors have been proposed as a way to connect areas of biological significance. In this article we introduce a model to maximize the amount of suitable habitat in a fully connected parcel network linking core habitat areas, subject to a constraint on the funds available for land acquisition. The economic framework of maximizing benefits subject to a budget constraint that we employ is a divergence from other recently proposed models that focus only on minimizing the cost of a single parcel-wide corridor. While the budget constrained optimization model that we introduce is intuitively appealing, it presents substantial computational challenges above and beyond determining the cost-minimizing corridor.

We formulate the wildlife corridor design problem formally as the so-called connection subgraph problem. This graph problem, NP-hard in terms of the worst case computational complexity, demonstrates an easy-hard-easy pattern in solution runtime. We present a solution method for this optimization problem using a network flow based Mixed Integer Programming (MIP) formulation, and introduce a hybrid technique to improve scalability. We apply our model and methods to real data collected for the optimal design of a wildlife corridor for grizzly bears in the U.S. Northern Rockies, illustrating the underlying computational complexities by varying the granularity of the parcels available for acquisition. In addition, we show that budget constrained optimization drastically increases total habitat suitability of the corridor over parcel selection based solely on cost minimization. The model and solution method developed here are general and can be applied, in addition, to conservation of other species or even to problems arising in other fields such as social networks.

1 Introduction and Overview

In many parts of the world land development has resulted in a reduction and fragmentation of natural habitat, leading to increased rates of species decline and extinction. To combat the negative consequences of anthropogenic habitat fragmentation, the procurement of biologically valuable conservation land has been promoted as a way to ensure species viability. A large number of models for optimally selecting land parcels for conservation, formally referred to as the reserve site selection problem (RSSP), have been proposed in the conservation biology literature. These models select parcels to ensure that all targeted species in a given region are protected, as in the Set Covering Problem (SCP) [e.g. 24, 39], or they select a constrained number of parcels that maximize species richness, as in the Maximal Covering Problem (MCP) [e.g. 5, 9].

A number of subsequent studies have been added to the conservation biology literature incorporating economic variables into the RSSP. These studies seek to procure conservation parcels, given a budget constraint, that maximize the number of species protected [e.g. 1, 10, 33] or maximize the environmental benefits of the sites selected [e.g. 16, 27, 29]. The results of these economic-based studies show that incorporating spatially heterogeneous financial costs into reserve site selection models leads to a substantially different set of priority parcels than standard SCP or MCP models that ignore parcel costs. In addition, the parcels selected based on budget constrained optimization obtain considerably greater environmental benefits for the same conservation budget than traditional site selection models [28].

In recent years, researchers have recognized that a parcel's spatial location relative to other protected parcels is also an essential attribute to consider in reserve site selection. Reflecting this, a variety of models that seek to increase the degree of spatial coherence in the set of parcels selected for conservation have been developed (Williams, ReVelle, and Levin [47] provide a thorough review). One primary way in which spatial attributes have been incorporated into site selection models is through the optimal selection of a connected reserve network, which we refer to as a wildlife corridor.[1] The focus on developing models for the design of optimal wildlife corridors has come as biologists have highlighted the environmental imperative of connecting core areas of biological significance [30]. Properly implemented wildlife corridors provide numerous ecological benefits by returning the landscape to its natural connected state. By allowing species the ability to migrate between core areas of biological significance, corridors increase gene flow and reduce rates of inbreeding, thereby improving species fitness and survival [35]. Corridors also allow for greater mobility [2], thus allowing species to escape predation and respond to stochastic events such as fire. In addition, corridors allow species to respond more easily to long term climatic changes [26].

Responding to the ecological benefits of connected ecosystems, a wide range of corridor projects have been proposed or are currently being implemented worldwide. Despite the increasing number of corridors being implemented around the world and several studies documenting the positive ecological benefits of existing corridors [e.g. 14, 20, 37], *designing models for the optimal selection of corridor parcels* has received comparatively little attention. Exploration along these lines started with the work of Sessions [36]. Previous models of optimal corridor design have sought to minimize the number of sites selected such that a specific number of species are preserved [7, 17, 31] or minimize the amount of unsuitable habitat in the corridor [44, 45, 46]). However, previous models of optimal corridor design, have not considered the case of budget constrained optimization. In fact, with the exception of Sessions [36] and Williams [45], spatially heterogeneous parcel cost has been ignored altogether.

We feel that the formulation of the corridor design problem that we present in this article is more relevant to a conservation planner than other alternatives available, as these planners typically have limited funds with which to secure conservation land. Specifically, we seek the optimal construction of a wildlife corridor between multiple areas of biological significance. We propose a new budget constrained optimization model, and a corresponding hybrid solution methodology, that efficiently

[1] Wildlife corridors are also referred to more or less interchangeably as conservation, habitat, and movement corridors.

incorporates both economic and ecological information in the design of optimal corridors. We apply the techniques presented to the design of a wildlife corridor for grizzly bears connecting the Yellowstone, Salmon-Selway[2] and Northern Continental Divide Ecosystems in Idaho, Wyoming and Montana.

Our approach diverges from previous corridor design studies in several important ways. First, we model the corridor design problem in a cost-efficiency framework, where the conservation planner wishes to maximize the overall amount of suitable habitat in the corridor subject to a budget constraint. This is a change from previous studies that have modeled the problem as one of minimizing some aspect of parcel cost, either in terms of number of parcels, financial cost, or cost to wildlife traversing the corridor. Our model is the first to explicitly include a budget constraint; something that we feel greatly improves the relevance of the model for conservation planners, who generally operate in an environment with limited budgets. This reformulation, while intuitively appealing, presents additional computational challenges, since the number of feasible corridors is significantly larger than for the standard minimum cost path model.

The second primary change from previous studies is that we do not require the corridor to have a "tree" structure, which is what one obtains in models that seek the minimum cost Steiner tree as the "best" wildlife corridor. This is important in the model that we present because it means that if the budget allotted for the corridor is higher than the minimum cost corridor, then the benefits of the corridor can be improved either by selecting an additional route, or by simply making the corridor wider so that it cost-effectively includes adjacent parcels.

Another contribution of this study is that we incorporate estimated parcel costs from a naturally occurring landscape, as discussed shortly. In addition, by changing the granularity of the parcels available for selection we gain a greater understanding of the relationship between computational complexity and the number of parcels in the landscape (see Figure 5). We also gain insight into the tradeoffs between parcel size and model specificity.

Our work also provides several contributions from a computational point of view. We formalize the corridor design problem as a graph theory problem that we refer to as *the connection subgraph problem*. This approach allows us to focus on the computational issues of the problem, independently of the particular domain. It also highlights the fact that other problems, with the same structure as the corridor design problem, as they occur, for example, in social network applications, can be modeled as a connection subgraph problem. We formally characterize the worst-case computational complexity of the connection subgraph problem as a so-called NP-Hard problem. In order to further understand its *typical case* complexity, beyond the standard NP-Hard worst case notion used in computer science, we developed a randomized generator of "synthetic" instances of the connection subgraph problem, using semi-structured graphs.[3] By studying the behavior of different algorithms, combined with different model formulations, on synthetic instances, we gained insights into the structure of the problem. We also discovered an interesting "easy-hard-easy" pattern in the typical computational complexity of proving optimality for instances of this problem. Such insights led to the development of a hybrid algorithm that exploits the structure of the problem.

Our hybrid algorithm for the connection subgraph problem allows us to dramatically scale up solutions: it incorporates a *provably* efficient procedure (i.e., it runs in polynomial time) for

[2]The Salmon-Selway Ecosystem is also referred to as the Bitterroot Ecosystem.

[3]An instance of a problem results from assigning concrete values to its parameters. For example, in the corridor design problem, we assign concrete values to the parcel layout, parcel utilities, parcel costs, and the budget. One can generate multiple instances by randomly assigning parameter values. The value of analyzing multiple problem instances is that it provides a much better understanding of the problem's computational complexity, as opposed to focusing on a single instance that may, or may not, be representative.

computing the optimal minimum cost corridor; the minimum cost solution can be used to initialize a general procedure, that allows the algorithm to converge to the overall optimal solution much faster; it also incorporates so-called propagation and pruning techniques that considerably speed up the solution procedure by ruling out early candidate solutions that are guaranteed to be suboptimal. The resulting hybrid algorithm, described in detail in this article, performs remarkably well, with *strong optimality guarantees*, both when considering the synthetic instances as well as instances of the real-world wildlife corridor for grizzly bears, connecting the Yellowstone, Salmon-Selway and Northern Continental Divide Ecosystems. For the real-world wildlife corridor problem, the resulting instance has over 12,000 parcels. When considering a budget of \$8M (the minimum cost of a corridor is \$7.2M) our procedure provides a solution that is guaranteed to be within 1% of the optimal solution. This is in sharp contrast with other corridor design approaches described in the literature, which in general do not provide any optimality guarantees.

In the next section we provide a detailed discussion of related work. We formally describe the corridor problem and characterize its computational complexity in section 3. Section 4 formulates the problem as a mixed integer programming problem and describes our solution procedure in detail. Section 5 provides experimental results and describes the application of corridor design for the grizzly bear in the U.S. Northern Rockies. Section 6 concludes.

2 Related Work

A wide range of wildlife corridor projects have been proposed or are currently being implemented. The projects range from local scale projects, such as the Quimper Wildlife Corridor, which provides a 3.5 mile greenbelt in Jefferson County, WA, to much wider scale projects like the 'Yellowstone to Yukon' initiative, which seeks to implement a viable corridor stretching from Yellowstone National Park to the Yukon region of western Canada. Corridor projects are currently being planned or implemented by governments and NGO's across the world, such as the Siju-Rewak Corridor in India, the proposed Selous-Niassa Wildlife Protection Corridor Project in Africa and the Amapa Biodiversity Corridor in the Amazonian Rainforest.

From a mathematical perspective, the problem of optimal corridor design was first posed by Sessions [36], who modeled the selection of a hypothetical corridor as a network Steiner tree (NST) problem. The hypothetical formulation employed by Sessions involves a landscape composed of a set of available parcels to connect a subset of critical parcels. The cost of each parcel is defined as the opportunity cost of not harvesting the parcel's timber, which is assumed to be known and the model's objective is to connect the critical parcels with the least-cost set of available parcels. Noting that arriving at a solution may not be possible in polynomial time for a large set of parcels, Sessions uses a shortest path heuristic to select parcels that minimize the cost of connecting the critical parcels.

Williams [45] also models the optimal selection of a hypothetical corridor as an NST problem, but includes the dual objectives of minimizing the cost of the corridor and minimizing the amount of unsuitable area included in the corridor. Using integer linear programming, Williams finds all of the non-inferior solutions, which allows for a comparison of the tradeoffs between corridor cost and habitat suitability. The problem Williams poses in his 1998 article is novel in that it incorporates both the financial and environmental attributes of each parcel. In subsequent work, Williams modified his original model to consider cases where there are no predefined reserves and the planner is simply trying to form a connected reserve [44, 46]. Williams [44] considers a relaxation of the contiguity requirement by incorporating a separate contiguity parameter that can be adjusted to control the overall degree of connectivity in the parcels selected. Williams and Snyder [46] take up the special case of percolating clusters, where the corridor is selected so as to connect one end of the landscape to the other (i.e., from north to south). Önal and Wang [32] extend the NST approach to the case where no contiguous reserve is feasible. They label Steiner tree nodes as reserve nodes or "gap" nodes, and use an integer linear programming model to minimize the number of gap nodes.

The works of Sessions and Williams are ground-breaking in the different formulations of the corridor problem that they introduce. Their models, however, only allow each parcel to be connected to one other parcel in the corridor. This is done to eliminate the possibility of cycles, loops, and islands and also to decrease the computational complexity of the problem. Considering only a one parcel-wide corridor, however, rules out the possibility of a corridor being 'thicker" (i.e., multiple parcels wide) for at least some portion of the path. This would be beneficial, for example, if there is an agglomeration of high quality and low cost habitat in some portion of the corridor that could be cost-effectively incorporated into the reserve system. In addition, the authors do not extend their research to the study of an applied corridor instance, making it difficult to determine how the models perform in practice.

Recent articles [7, 17, 31] also introduce alternative models of optimal corridor design and apply them to specific study areas. Cerdeira et al. [7] formulate an integer linear programming approach to solve a fully connected set covering problem and apply their model to the case of 496 uniform and contiguous parcels in the county of Hertfordshire, UK. They find that a minimum of 22 contiguous sites are needed to optimally cover the 45 species of butterflies in the study area. A heuristic method that they develop in the paper selects 23 sites for conservation, which the authors take as evidence that their heuristic performs well in comparison to exact methods. Önal and Briers [31] also formulate a fully connected set covering problem as an integer linear program. They apply their model to 121 bird species dispersed over 391 parcels in Berkshire County, UK and show that the model is too complex to efficiently compute the optimal solution. They then outline a procedure that involves solving the problem at a more aggregate scale and then selecting the optimal set of small disaggregate sites within the aggregate solution. This selection algorithm is found to perform more favorably than a heuristic procedure that is an extension of the greedy algorithm. Finally, Fuller et al. [17] apply a three stage algorithm to select a connected conservation network in central Mexico. They begin by selecting sites for conservation based on the habitat requirements of 99 species. They then define a set of paths that link the conservation areas with parcels containing suitable habitat. Finally, in the third stage, the paths that have the smallest area and impact on human population are selected to form the connected reserve network.

Several of the previous works have incorporated connectivity as a "soft requirement" rather than a hard constraint, by using an indirect measure of connectivity in the objective function. These methods are geared towards minimizing habitat fragmentation but do not necessarily guaratee the strict connectivity requirement demanded by a wildlife corridor. For example, McDonnell et al. [25] include a term corresponding to the boundary length of the reserve in the objective function, with the understanding that minimizing the boundary length will indirectly reduce the level of fragmentation in the reserve. Their problem is then formulated as an integer *non-linear* program and solved using simulated annealing and a heuristic approach. Along similar lines, Cabeza et al. [4] also consider a linear combination of reserve size and boundary length in the objective function. They do this in the context of incorporating probabilities of species occurrence in the reserve design model.

The connectivity condition has been studied in other related contexts. For example, Cerdeira and Pinto [6] consider the problem of computing a *connected* set cover in a bipartite graph where the selected vertices, coming from the first partition, are required to (a) cover (i.e., together connect to) every vertex of the other partition and (b) form a connected subgraph of a separate underlying graph defined on just this partition. They introduce and study, from a purely mathematical perspective, several properties of such connected set covers. In recent work, Cerdeira et al. [8] have considered species-specific minimum connectivity, where for each species s there is a target number t_s indicating the minimum number of contiguous sites needed for the protection of that species. Their objective is to find a least cost subset of sites such that every species s under consideration has at least one large enough connected component suitable for s, i.e., a connected component with t_s or more sites when the subset of sites is restricted to those that are suitable for s.

From an algorithmic perspective, Vanderkam et al. [43] provide a general study of the trade-offs between optimal algorithms (in particular, integer linear programs) and heuristic algorithms. Their study is not for conservation corridors but for the related problem of reserve design using site selection in order to maximize certain desirable attributes or features of the reserve such as species coverage, habitat diversity, etc. They find that for the data sets they considered, heuristic algorithms previously proposed result in solutions that are often quite far from the optimal solution, and that an integer linear programming "exact" approach is not so slow as to be quickly discarded in favor of heuristic approaches. This aligns well with our setting in this work—for us, connectivity is a strict requirement and out goal is to find provably optimal or near-optimal solutions.

3 Problem Description: Wildlife Corridors as Connection Subgraphs

We begin by mathematically defining the wildlife corridor design problem as a problem of finding a connected subgraph of a given graph with costs and utilities assiciated with its edges. We then give a brief analysis of this problem from the traditional worst-case complexity perspective, proving that the corresponding decision problem is NP-complete and the cost optimization variant of the problem is NP-hard to approximate within a certain constant factor.

3.1 The Connection Subgraph Problem

Let $\mathbb{Z}^+$ denote the set $\{0, 1, 2, \ldots\}$ of non-negative integers. The decision version of the connection subgraph problem is defined on an undirected graph as follows:

Definition 1 (Connection Subgraph Problem)**.** Given an undirected graph $G = (V, E)$ with terminal vertices $T \subseteq V$, vertex costs $c : V \to \mathbb{Z}^+$, vertex utilities $u : V \to \mathbb{Z}^+$, a cost bound $C \in \mathbb{Z}^+$, and a desired utility $U \in \mathbb{Z}^+$, does there exist a vertex-induced subgraph H of G such that

1. H is connected,
2. $T \subseteq V(H)$, i.e., H contains all terminal vertices,
3. $\sum_{v \in V(H)} c(v) \leq C$, i.e., H has cost at most C, and
4. $\sum_{v \in V(H)} u(v) \geq U$, i.e., H has utility at least U?

In this decision problem,[4] we can relax one of the last two conditions to obtain two natural optimization problems: (1) *Utility Maximization*: given a cost bound C, maximize the utility of H; (2) *Cost Minimization*: given a desired utility U, minimize the cost of H.

The connection subgraph problem captures the key mathematical aspects of the corridor design problem if we think of the available land parcels as vertices of a graph, reserves as terminal vertices,

[4] A *decision problem* is a problem with a yes-no answer. Typically, given an algorithm for the yes-no version of a problem, it is easy to produce an equally efficient algorithm that actually produces a solution if the answer is yes.

parcel cost (or utility) as the cost (or utility, resp.) associated with the corresponding vertex, and two land parcels sharing a boundary being equivalent to having an edge between the two corresponding vertices in the graph. A connected subgraph of this graph containing the designated terminal vertices corresponds to a conservation corridor connecting the given reserves.

In the context of social networks, a similar problem has been investigated by Faloutsos et al. [15]. Here, one is interested, for example, in identifying the few people most likely to have been infected with a disease, or individuals with unexpected ties to any members of a list of other individuals. This relationship is captured through links in an associated social network graph with people forming the nodes. Faloutsos et al. consider networks containing two special nodes (the "terminals") and explore practically useful utility functions that capture the connection between these two terminal nodes. Our interest, on the other hand, is in studying this problem with the sum-of-weights utility function but with several terminals. In either case, the problem has a bounded-cost aspect that competes with the utility one is trying to maximize.

3.2 Worst-Case Complexity Analysis

From a computer science perspective, the first question one typically asks is how hard the problem under consideration is, in terms of the traditional computational complexity hierarchy. Broadly speaking, computer scientists consider a problem to be "easy" or efficiently solvable if there is a polynomial time algorithm (polynomial in the size of the input) that solves the problem. A large set of real-world problems belong to the so-called NP-complete class for which only exponential time algorithms are known and for which it is believed by many that no polynomial time algorithm exists.[5]

We next discuss the computational complexity of the connection subgraph problem, before moving on to our solution methodology and experimental evaluation. In order to maintain the focus of the paper on effective solution methods, this section is kept brief and all proofs are deferred to Appendix B.

The connection subgraph problem is a generalized variant of the standard *Steiner tree problem* [cf. 34] on undirected graphs, with the difference being that the costs are on vertices rather than on edges and that we have utilities in addition to costs. The utilities add a new dimension of hardness to the problem. In fact, while the Steiner tree problem is polynomial time solvable when $|T|$ is any fixed constant [cf. 34], we will show that the connection subgraph problems remains NP-complete even when $|T| = 0$. We prove this by a reduction from the Steiner tree problem. This reduction also applies to planar graphs, for which the Steiner tree problem is still NP-complete [cf. 34].

Theorem 1 (NP-Completeness). *The decision version of the connection subgraph problem, even on planar graphs and without any terminals, is NP-complete.*

The reader is referred to Appendix B for the relatively short proof of this theorem. The theorem immediately implies the following:

Corollary 1 (NP-Hardness of Optimization). *The cost and utility optimization versions of the connection subgraph problem, even on planar graphs and without any terminals, are both NP-hard.*

[5] NP stands for Non-deterministic Polynomial time. This captures the idea that, given a candidate solution, one can *verify* its validity as a solution in polynomial time. Note that this does not mean that one can *generate* the solution in polynomial time—being able to do that would make the problem polynomial time solvable, i.e., "easy". NP-complete problems are the *hardest* problems within the class NP and all known algorithms for them take exponential time (in the input size) in the worst case. Roughly speaking, the notion of being complete for a class means that all other problems in the class can be translated to this problem in polynomial time; therefore, if one could find a polynomial time algorithm to solve any one of the complete problems in a class such as NP, then all the problems in the class would be solved in polynomial time as well.

It turns out that in the NP-hardness reduction used in the proof of Theorem 1, the graph $\widehat{G}$ in the given Steiner tree instance has a Steiner tree with cost C' iff the graph G constructed for the connection subgraph problem has a connection subgraph with cost C'. Consequently, if the cost optimization version of the connection subgraph instance (i.e., cost minimization) can be approximated within some factor $\alpha \geq 1$ (i.e., if one can find a solution of cost at most α times the optimal), then the original Steiner tree problem can also be approximated within factor α. It is, however, known that there exists a factor α_0 such that the Steiner tree problem *cannot* be approximated within factor α_0, unless P=NP. This immediately gives us a hardness of approximation result for the utility optimization version of the connection subgraph problem. Unfortunately, the best known value of α_0 is roughly only $1 + 10^{-7}$ [cf. 34].

Fortunately, we can use a different reduction—from the NP-complete Vertex Cover problem—which will enable us to derive as a corollary a much stronger approximation hardness result for the connection subgraph problem.

Lemma 1. *There is a polynomial time reduction from Vertex Cover to the connection subgraph problem, even without any terminals, such that the size of the vertex cover in a solution to the former equals the cost of the subgraph in a solution to the latter.*

We again refer the reader to Appendix B for a proof of this statement. Combining Lemma 1 with the fact that the vertex cover problem is known to be NP-hard to approximate within a factor of 1.36 [13] immediately gives us the following:

Theorem 2 (APX-Hardness of Cost Optimization). *The cost optimization version of the connection subgraph problem, even without any terminals, is NP-hard to approximate within a factor of* 1.36*.*

We refer the reader to Appendix A for a simple example that highlights some of the combinatorial issues of the connection subgraph problem that make it computationally hard.

4 Solving the Connection Subgraph Problem

Next we present the Mixed Integer Linear Programming Model (MIP model) for the connection subgraph problem, that was used in our experiments. We will then discuss a hybrid solution method for efficiently solving this MIP model.

4.1 Mixed Integer Linear Programming Formulation

Let $G = (V, E)$ be the graph under consideration, with $V = \{1, \ldots, n\}$ and budget C. The corresponding MIP formulation is given in Figure 1, which we discuss in detail below.

For each vertex $i \in V$, we introduce a binary variable x_j, representing whether or not i is in the connected subgraph. Then, the objective function is stated as $\sum_{j \in V} u_j x_j$, as in the expression (1) of Figure 1. The budget constraint is given by inequality (2).

To ensure the connectivity of the subgraph, we apply a particular *network flow model*, where the network is obtained by replacing all undirected edges $\{i, j\} \in E$ by two directed edges (i, j) and (j, i). Call the set of directed edges E'. First, we introduce a source vertex 0, with maximum total outgoing flow n. We arbitrarily choose one terminal vertex $\hat{t} \in T$ as the "root" node, and define a directed edge $(0, \hat{t})$ to insert the flow into the network, assuming that there exists at least one terminal vertex.[6] Then, by demanding that the flow reaches all terminal vertices, the edges

[6] If there are no terminal vertices specified, we add edges from the source to all vertices in the graph, and demand that at most one of these edges is used to carry flow.

$$
\begin{aligned}
\text{maximize} \quad & \sum_{j \in V} u_j x_j && && (1)\\
\text{subject to} \quad & \sum_{j \in V} c_j x_j \le C && && (2)\\
& x_j \in \{0,1\} && \forall j \in V && (3)\\
& x_t = 1 && \forall t \in T && (4)\\
& x_0 + y_{0,\hat{t}} = n && && (5)\\
& y_{i,j} \le n x_j && \forall (i,j) \in E' && (6)\\
& \sum_{i:(i,j)\in E'} y_{i,j} = x_j + \sum_{\ell:(j,\ell)\in E'} y_{j,\ell} && \forall j \in V && (7)\\
& \sum_{j \in V} x_j = y_{0,\hat{t}} && && (8)
\end{aligned}
$$

Figure 1: Network flow based MIP model for the connection subgraph problem.

carrying flow (together with the corresponding vertices) represent a connected subgraph. To this end, each of the vertices with a positive incoming flow will act as a ‘sink’, by ‘consuming’ one unit of flow. In particular, all terminal vertices will act as sinks, and any other vertex that is part of the eventual connected subgraph will also be a sink (in other words, $x_j = 1$ will correspond simultaneously to vertex j being in the connected subgraph solution and to it acting as a sink for the network flow). Finally, we will add constraints to enforce flow conservation: for every vertex the amount of incoming flow equals the amount of outgoing flow plus the amount of consumed flow.

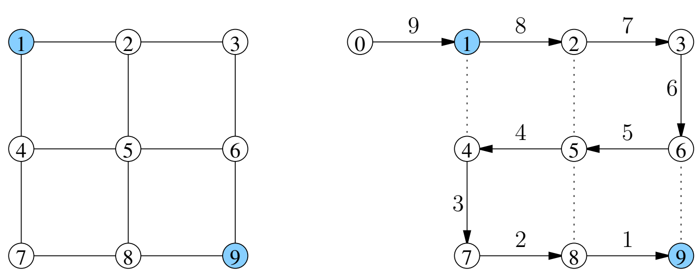

a. Original graph   b. Feasible flow

Figure 2: Flow representation of the connection subgraph problem on a graph with 9 vertices. The terminal vertices, 1 and 9, are shaded. The special terminal node, $\hat{t}$, is vertex 1.

More formally, for each (directed) edge $(i,j) \in E'$, we introduce a non-negative variable $y_{i,j}$ to indicate the amount of flow from i to j. For the source, we introduce a variable $x_0 \in [0,n]$, representing the eventual residual flow. The residual flow plus the flow injected into the network corresponds to the total system flow, as given by equation (5), where $\hat{t} \in T$ is arbitrarily chosen. Each of the vertices with a positive incoming flow retains one unit of flow, i.e., $(y_{i,j} > 0) \Rightarrow (x_j = 1), \forall (i,j) \in E'$. We convert this relation into a linear constraint, as in inequality (6). The flow conservation is modeled as equation (7). All terminal vertices are forced to retain one unit of flow and thus be in the connected subgraph constructed by this process, using equation (4). Finally, the overall flow absorbed by the network is set to equal the flow injected into the system, using equation (8).

Figure 2 depicts an example of this network flow representation, where we omit the costs for clarity. Figure 2.a presents a graph on 9 vertices with terminal vertices 1 and 9. In Figure 2.b, a feasible flow for this graph is depicted, originating from the source 0, with value 9. It visits all vertices, while each visited vertex consumes one unit of flow. The thus connected subgraph contains all vertices in this case, including all terminal vertices.

Remark 1. This network flow based MIP formulation as well as the connection subgraph problem itself allow for the possibility of cycles and loops, which we see as a favorable option given that the overall utility of the parcels selected can be increased by widening the corridor or by incorporating paths to areas of high quality habitat. It may, however, be that the conservation planner wishes to eliminate the possibility of having “peninsulas” in the network, which could represent geographic dead ends to wildlife in the corridor. While this option is not explored empirically in this article, in practice peninsulas could be reduced through the institution of an additional constraint requiring that every vertex receiving flow must output flow to at least one other vertex that is different from the input vertex. Formally, the constraint is

$$y_{i,j} \leq n \sum_{\ell \neq i}^{n} y_{j,\ell} \qquad\qquad \forall j \in V \setminus T \tag{9}$$

Note the use of the multiplier n in the right-hand-side of the above constraint, which is needed because the outgoing flow from j would, by design, be one unit less than the incoming flow (when the incoming flow is non-zero) as node j would absorb one unit of flow. While this constraint will eliminate all single parcel wide peninsulas, it is still possible for there to exist a multiple parcel wide peninsula. □

4.2 Meeting the Scalability Challenge: A Hybrid Solution Method

While the MIP formulation presented above can be solved to optimality by state-of-the-art MIP solvers, such as IBM/Ilog’s CPLEX, for relatively small size problems, scalability quickly becomes a challenge as one begins to handle real-life data. In order to address the scalability challenge, we use a two phase solution method.

In Phase I, we compute a minimum cost Steiner tree for the terminal nodes of the graph, ignoring all utilities. While there are fixed parameter tractable (FPT) algorithms for computing a minimum cost Steiner tree, we used a simpler “enumeration” method (see, e.g., [34]) based on computing all-pairs-shortest-paths (APSP) with respect to vertex costs. The APSP matrix can be computed in time $O(n^3)$ for a graph with n vertices. As we will see later in this section, this matrix also comes handy in pruning away often a large fraction ($\approx 40-60\%$) of the vertices of the graph based on their distance from the terminal vertices, significantly improving the efficiency of Phase II. The idea behind the enumeration-based Steiner tree algorithm, which runs in polynomial time for a constant number of terminal nodes, is to first compute a minimum Steiner tree $\tilde{T}$ for the “complete shortest distance graph” $\tilde{G}$ of the original graph G, where $\tilde{G}$ is a complete graph with as many vertices as G and where the weight assigned to an edge $\{u, v\}$ in $\tilde{G}$ equals the cost of a shortest path between the corresponding vertices u and v in G (provided by the APSP matrix). The algorithm uses the fact that in any complete shortest distance graph (such as $\tilde{G}$), there exists a minimum Steiner tree whose non-terminal nodes have degree at least three, thereby limiting the total number of nodes in the Steiner tree to be two fewer than the number of terminal nodes. A minimum Steiner tree $\tilde{T}$ of $\tilde{G}$ yields a minimum Steiner tree T for the original graph G by simply replacing edges $\{u, v\}$ in $\tilde{T}$ by paths in T corresponding to a shortest path between u and v in G.

The computation of the Steiner tree in Phase I typically took a few minutes to a few hours on our problem instances, which was in fact an almost negligible amount of time compared to Phase II, which we describe next. The Steiner tree computation either classifies the problem instance as infeasible for the given budget or provides a feasible (but often sub-optimal) "mincost" solution. In the latter case, Phase II of the computation translates the problem into a MIP instance using the encoding discussed in Section 4.1, and solves it using IBM/Ilog's CPLEX solver [21]. Solving the MIP formulation using CPLEX is the most computationally-intensive part of the whole process. The mincost solution obtained from Phase I is passed on to CPLEX as the starting solution.[7]

Further, the APSP matrix computed in Phase I is also passed on to Phase II. It is used to statically (i.e., at the beginning) *prune* away all nodes that are easily deduced to be too far to be part of a solution (e.g., if the minimum Steiner tree containing that node and all of the terminal vertices already exceeds the budget). This significantly reduces the search space size, often in the range of 40-60%. We also experimented with dynamic pruning, performed during the branch-and-bound search of CPLEX, but the overhead was a bit too high for our problem instances for dynamic pruning to pay off. The experimental results are reported with static pruning only.

Overall, Phase II is designed to compute an *optimal solution to the utility-maximization version* of the connection subgraph problem. In case it runs out of time, which happened on our large instances, it provides a feasible solution along with a conservative bound on how far this solution is from the optimal (i.e., the *optimality gap*).

As a comparison point, we also use a very efficient *greedy method* to improve the quality of the mincost solution provided by the Steiner tree computation. The idea is to use any residual budget to acquire additional vertices in a greedy fashion as follows. We consider those vertices in the graph that are adjacent to the current solution and have cost lower than the residual budget, and identify one whose *gain*, defined as the utility-to-cost ratio, is the highest. If there is such a vertex, we add it to the current solution, appropriately reduce the residual budget, and repeat until no more vertices can be added. This process often significantly increases the solution quality compared to the mincost solution provided by the Steiner tree computation. The resulting greedy solution is an example of what we call an *extended-mincost solution*.

We will also be interested in computing the *optimal extended-mincost solution*, obtained by "freezing" the Steiner tree vertices to be in the constructed solution and, rather than extending this partial solution greedily, solving the MIP encoding of Phase II using CPLEX with these frozen parcels "forced" to be in the solution. The solution quality (i.e., the overall utility) of the optimal extended-mincost solution lies between that of the greedy solution and the optimal solution to the problem. The computation of this solution also follows the same trend: it is slower than the greedy computation but faster than the full MIP encoding for computing the optimal solution.

5 Experimental Results and an Application to the Grizzly Bear Corridor in the U.S. Northern Rockies

While our main goal is to identify the optimal corridor for grizzly bear in the U.S. Northern Rockies, we are also interested in underestanding properties of general instances of the connection subgraph problem. To that end, we conducted a series of experiments to study the typical case complexity of the problem. In particular, we investigate the *empirical* computational hardness of the problem

[7]In reality, we actually pass on to CPLEX the *greedy solution* to be described shortly. This provides a major boost to the efficiency of CPLEX in solving the MIP encoding.

with respect to computing the optimal solution or the extended-mincost solution mentioned in Section 4.2, as we vary the the cost bound (or budget).

All our experiments were conducted on a number of 3.8 GHz Intel Xeon machines with 2 GB memory, running Linux 2.6.9-22.ELsmp. We used the CPLEX 10.1 solver [21] to solve the mixed integer programming formulation of the problem instances considered. For the larger instances, which would not fit in the 2 GB RAM of the computers used, we relied on the built-in disk use capabilities of CPLEX (rather than the computer's virtual memory mechanism) to store and manage very large search trees.[8]

5.1 Scaling Behavior: Semi-Structured Instances and Easy-Hard-Easy Pattern

For the experiments in this and later sections, our main parameter is the feasibility component of the problem, i.e., the cost bound or budget. Here, for a varying budget level, we investigate the computational hardness of the problem with respect to computing the optimal solution or the extended-mincost solution. For the study in this section, we make use of semi-structured graphs, with uniform random utility and cost functions. The graphs are composed of an $m \times m$ rectangular lattice or grid, in which we place up to 3 terminal vertices. The details of this semi-structured graph model, as well as more empirical data on them, may be found in Appendix C.

In the figures below, each data point is based on statistics from 100 to 500 randomly generated instances. The hardness curves are represented by median running times over all instances per data point, normalized for the small but non-negligible variation in the characteristics of various randomly generated instances with the same parameters. On the x-axis of the plots is the '**budget slack percentage**', rather than simply the budget, computed as follows. For every instance, we consider the *mincost*, i.e., the minimum budget needed to obtain a valid connection subgraph. The *budget slack %* with respect to mincost is defined as: $100 \times (budget - mincost)/mincost$. In other words, we consider computational hardness and other measured quantities as a function of the extra budget available for the problem beyond the minimum required to guarantee a feasible solution. The results are shown in Figures 3(a) and 3(b).

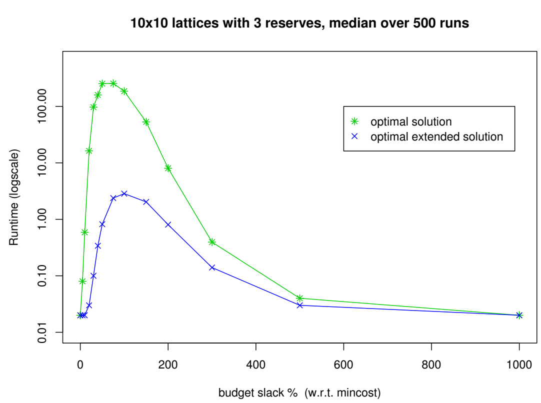

(a) Hardness profile (runtime in log-scale). Upper curve: optimal solution. Lower curve: optimal extended-mincost solution.

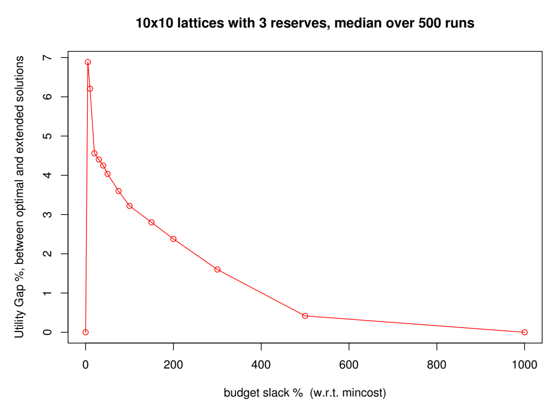

(b) Percentage gap in the utility of optimal and extended-mincost solutions.

Figure 3: Results for lattices of order 10 with 3 terminal vertices

[8]Specifically, we used the following parameter settings: `cplex.setParam(IloCplex::WorkMem, 1024)` and `cplex.setParam(IloCplex::NodeFileInd, 3)`.

In Figure 3(a), we show the hardness profile of lattices of order 10 with 3 terminals.[9] These optimization problems exhibit an easy-hard-easy pattern, the peak of which is to the right of the mincost point (shown as 0 on the relative x-scale). As one might expect, computing the optimal extended-mincost solution (the lower curve) is significantly easier than computing the true optimal solution (the upper curve; note that the y-axis is in log-scale).

This naturally raises the question: how much "better" are the true optimal solutions compared to the easier-to-find extended-mincost solutions? Figure 3(b) shows the *relative gap %* between the solution qualities (i.e., attained utilities) in the two cases, defined as $100 \times (optimal - extended)/optimal$. We see that at mincost, both optimal and extended-mincost solutions have similar quality, which is not too surprising. The gap between the qualities reaches its maximum shortly thereafter, and then starts to decrease rapidly, so that the extended-mincost solution at 100% budget slack is roughly 3.2% worse than the optimal solution, and at 500% budget slack, only around 0.4% worse. This suggests that for much larger, real-world problems representable as the connection subgraph problem, where computing the optimal solution is out of the question due to limited computational resources, it may suffice for practical purposes to only compute an optimal extended-mincost solution.

5.2 Application to Corridor Design for U.S. Northern Rockies

5.2.1 Data Collection

Study Area. The study area for our analysis is comprised of 64 counties in Idaho and western Montana, located in the Northern Continental Divide region. At the aggregate level, the parcels that we consider for inclusion in the corridor are the 64 counties themselves. While securing an entire county to be included in the reserve may seem infeasible, the county-level analysis provides an illustrative example for a case where the optimization problem is relatively simple from a computational perspective. The county level model allows us to identify general corridor areas that contain low cost, suitable habitat, similar to Ando et al. [1]. The county model also provides a means of comparing the results of an aggregate model with relatively few sites, to more granular models with greater numbers of parcels. A map of the study area is included as Figure 4.

To investigate the impact of increasing the granularity of the available parcels, we segment the study are into contiguous sets of square grid cells. The largest grid cells are 60km on each side and segment the study area into 118 parcels. The parcel size is then incrementally reduced to square grids with sides of 50km, 40km, 25km, 10km and 5km. With the most granular grid size of 5km, the study area is segmented into 12,788 cells. Given the relatively large range of an adult grizzly (the home range of an adult female grizzly bear is approximately 125 square km), grid sizes smaller than 5km are unlikely to be suitable for grizzly bear movement [23]. Increasing the granularity of the grid cells allows for much more precision in defining parcel habitat suitability and acquisition costs and it also increases the number of parcels in the landscape. Given the greater number of parcels available for the corridor, increasing the granularity also increases the complexity of the optimization problem. Thus, by comparing results across the continuum of cell sizes, we are able to investigate the tradeoffs inherent in the granularity of the model that allows for increased specificity at the cost of greater computational complexity.

In addition to square grid cells, we also consider a grid composed of 25 square km hexagonal parcels. The hexagonal grid allows parcel connections to occur diagonally and therefore generates more direct pathways between reserves that result in significantly lower costs than comparably sized square grid cells. Hexagonal grids are utilized by the Environmental Protection Agency's

[9] We obtained similar results with 10 and 20 terminals as well.

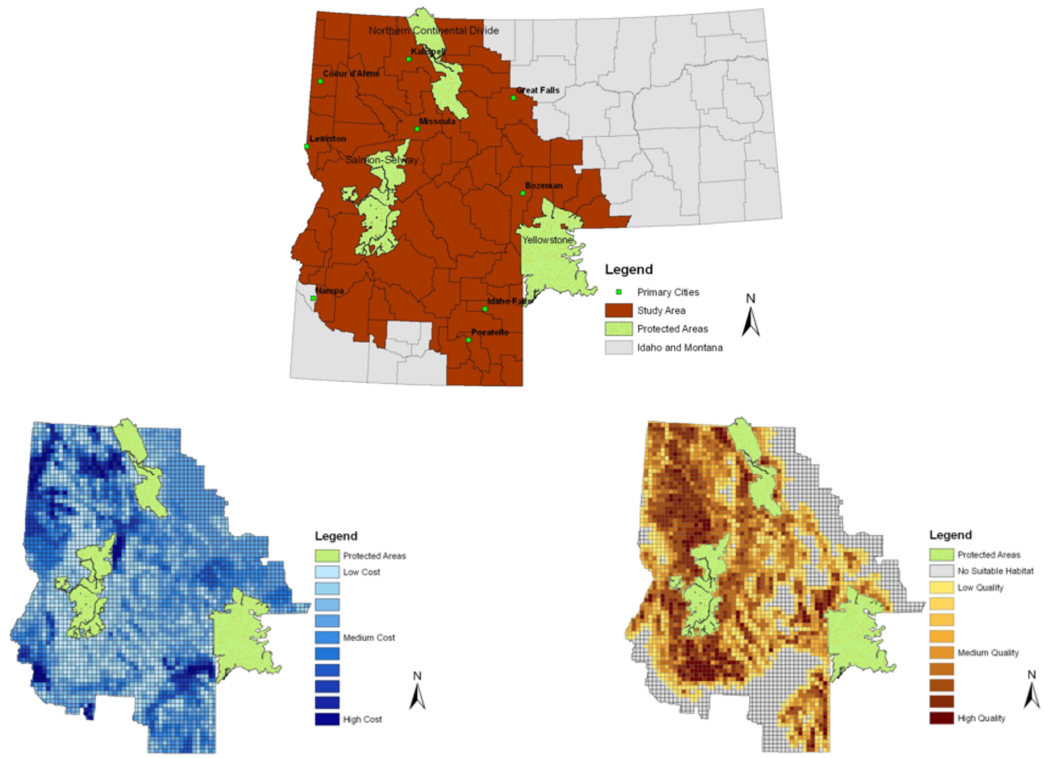

Figure 4: Top: the study area for the corridor problem. Bottom: cost and utility landscape for 10 $\times$ 10 km parcels.

Environmental Monitoring and Assessment Program (EMAP) and have been used in reserve site selection research by Polasky et al. [33] as well as Csuti et al. [11].

Utility Computation for Parcels. To measure the utility of each parcel, we utilize grizzly bear habitat suitability data developed and provided by the Craighead Environmental Research Institute (CERI). These data spatially define habitat that is considered to be suitable for grizzlies. The suitable habitat is measured on a 30 meter grid and given a score from 2 to 4, with 4 being the highest quality habitat. We then aggregate the habitat suitability data to the grid and county levels by summing the habitat scores within each parcel boundary. This method of aggregation implicitly assumes, for example, that a cell with a habitat suitability value of 4 is twice as beneficial as a cell with a habitat suitability value of 2.

Cost Computation of Parcels. We next discuss the process by which cost values are assigned to each land parcel under consideration. The estimate of parcel cost is calculated in three steps. First, spatial data on land stewardship, available for the states of Montana and Idaho from the Gap Analysis Project [42], are used to classify privately and publicly owned land in the study area. Next, the amount of private land acreage within each parcel is calculated. The private land acreage is then multiplied by the county specific average value of farm real estate per acre, available from the United States Department of Agriculture [40]. For grid cells with land acreage in multiple counties, the county specific real estate value per acre is multiplied by the amount of private acreage in each county and then summed. Using the value of farm real estate is a proxy for the cost of all private land, as it reflects the opportunity costs faced by private land owners. Ando et al. [1]

similarly use county level average farm real estate value in their reserve selection model.

In delineating the cost of each parcel, we assume that land already in the public domain is essentially freely available for inclusion in the corridor. One could, however, imagine incorporating the opportunity cost of lost timber or mining contracts as proxies for the cost of acquiring public land as done by Polasky et al. [33] and Sessions [36]. We have chosen not to incorporate costs on public land in the present analysis as there is insufficient data with which to accurately predict the heterogeneity in lost resource profitability associated with each parcel [23].

By calculating the cost of each parcel based on the real estate value of its privately owned acreage, we are essentially assuming that the parcels included in the corridor will be acquired with fee-simple purchases. For large projects, such as the corridor connecting the three large ecosystems in the Northern Rockies that we are considering, the funds necessary to purchase a viable corridor outright will be large. Yet our cost estimates should be put into perspective by comparison to the significant amount of both public and private funding currently being spent on land conservation. For example, the federal government has an annual budget of $900 million through the Land and Water Conservation Fund (LWCF) through which it can support land conservation at the local, state, and federal level. In addition, the Trust for Public Land estimates that in the past decade more than $36 billion in public land conservation funding has been approved over 1,000 separate ballot initiatives across the U.S. [38]. This funding is in addition to federal conservation programs such as the Conservation Reserve Program (CRP), which have average annual expenditures exceeding $1.6 billion [41]. It should also be noted that parcels may not necessarily need to be purchased outright in order to be included in the corridor, as easements and other *voluntary agreements* may be sufficient to maintain habitat. This voluntary type of arrangement is being used, for example, in the 'Alps to Artherton' project in Australia, where the Australian government is seeking agreements with private land owners to abstain from certain land use practices in exchange for annual payments.

While securing voluntary agreements for habitat protection may be a more viable strategy for cost-effectively targeting parcels to include in the corridor, there is insufficient data on the incentives necessary to secure such voluntary arrangements. We therefore use real estate value as an upper-bound on a parcel's cost, noting that the potential for voluntary habitat protection could significantly reduce the funds necessary to acquire the corridor. Future research on the incentives necessary for voluntary habitat protection could provide useful information for conservation planners.

One additional consideration in terms of the overall cost of the corridor is the transaction and management costs associated with securing property rights and maintaining the parcels. Researchers have identified transaction and management costs as being an important consideration in reserve design [e.g. 28, 29], yet these costs are rarely included in optimal conservation models. One notable exception is Groeneveld [19] who looks at the theoretical implications of varying transactions costs on the number of sites included in a reserve. In the present analysis, we investigate the influence of transaction costs on corridor design for the 5km grid parcels. Transaction costs are likely to play a more significant role when the cell granularity is small, as the transaction cost represents a greater proportion of the overall cost of the parcel and the number of potential paths is large. We include a fixed $5,000 transaction cost for each parcel that is included, which would cover legal fees, signage and other fees associated with defining a particular land area as part of the corridor. The actual transaction and maintenance cost of a particular parcel is likely to be variable, but we have chosen $5,000 simply as an approximation that is in line with reported transaction costs for conservation lands in New York State [22].

Parcel Adjancencies. Beyond defining the costs and utilities of parcel acquisition, it is also necessary to define the parcel adjacencies for all of the parcels in the study area. The adjacencies for both the aggregate county and square grid parcels are defined based on shared borders/edges. For the grid parcels this implies that interior parcels are adjacent to exactly four other parcels. This is referred to as a *rook pattern* of adjacency, which differentiates itself from a *queen adjacency pattern* where adjacency is defined based on shared edges and corners.

5.2.2 Results for U.S. Northern Rockies

Influence of Parcel Granularity on the Cost and Shape of the Corridor. We begin with a study of the effect of parcel granularity on the cost and shape of the resulting wildlife corridor. To maintain the focus on corridor cost and shape, we include for simplicity the results for *minimum cost corridors*, ignoring the utility maximization aspect for now. The experiments were conducted for granularity as large as County level parcels down to 5 km $\times$ 5 km square grid parcels (henceforth referred to simply as 5 km grid parcels) and the 25 square km *hexagonal grid*. Each parcel in a hexagonal grid has 6 neighboring parcels. Overall, the better connectivity properties of the hexagonal grid allowed us to obtain feasible solutions to the corridor problem with the lowest costs.

The minimum cost corridor (ignoring utilities) for each of these grid levels was computed to optimality using Phase I of our solution methodology based on Steiner Tree computation (cf. Section 4.2). Depending on the parcel granularity, the computation required anywhere from a few seconds to a few hours of CPU time.

Figure 5 visually depicts the maps of the minimum cost wildlife corridors at various granularities. Overall, with the increase in the granularity of the parcels available for acquisition, the minimum cost of a corridor that connects the three ecosystems decreases considerably. For example, the cost of the cheapest corridor is \$1.9 billion for the County level and drops to as low as \$11.8 million for the 5 km $\times$ 5 km grid, and further down to \$7.3 million for the 25 square km hexagonal grid. It is, of course, not surprising that purchasing all of the private land in five counties is extremely expensive and having the option of buying smaller parcels results in significant cost savings as the corridor is able to better incorporate low cost areas, which are composed primarily, and in some cases exclusively, of zero cost national forest land.

The corridor cost could be further reduced by evaluating parcels with area less than 25 sq km. However, there are critical tradeoffs in terms of the minimum corridor width necessary for wide-ranging species such as the grizzly bear. In addition, there are tradeoffs in terms of the computational complexity of solving the minimum cost corridor problem. For grid parcels that are 25 km or larger, the computation time necessary to prove optimality is less than one second. At the 10km and 5km grid sizes, the solution time is no longer trivial, increasing to close to thirty minutes in the case of the 5 km grid and to a couple of hours for the 25 sq km hexagonal grid. Thus, we begin to see the tradeoffs inherent in corridor design in terms of the model granularity, or alternatively the size of the study area, and the computational complexity of the problem.

Changing the parcel granularity not only influences the cost of the parcels selected, but it also influences the general path or *shape* that the corridor follows. For the county level, 60 km, and 50 km parcel maps, the minimum cost corridor essentially forms the shape of an upside-down T, where the parcels selected are concentrated in the area in the middle of the three ecosystems. When the parcel size is reduced to 40km and below, the minimum cost corridor traces a path connecting the three reserves that resembles the shape of a C, with the Salmon-Selway Ecosystem connecting directly to the Northern Continental Divide Ecosystem via a parcel path in the northwestern portion of the study area. By increasing the parcel granularity, the model avoids higher priced

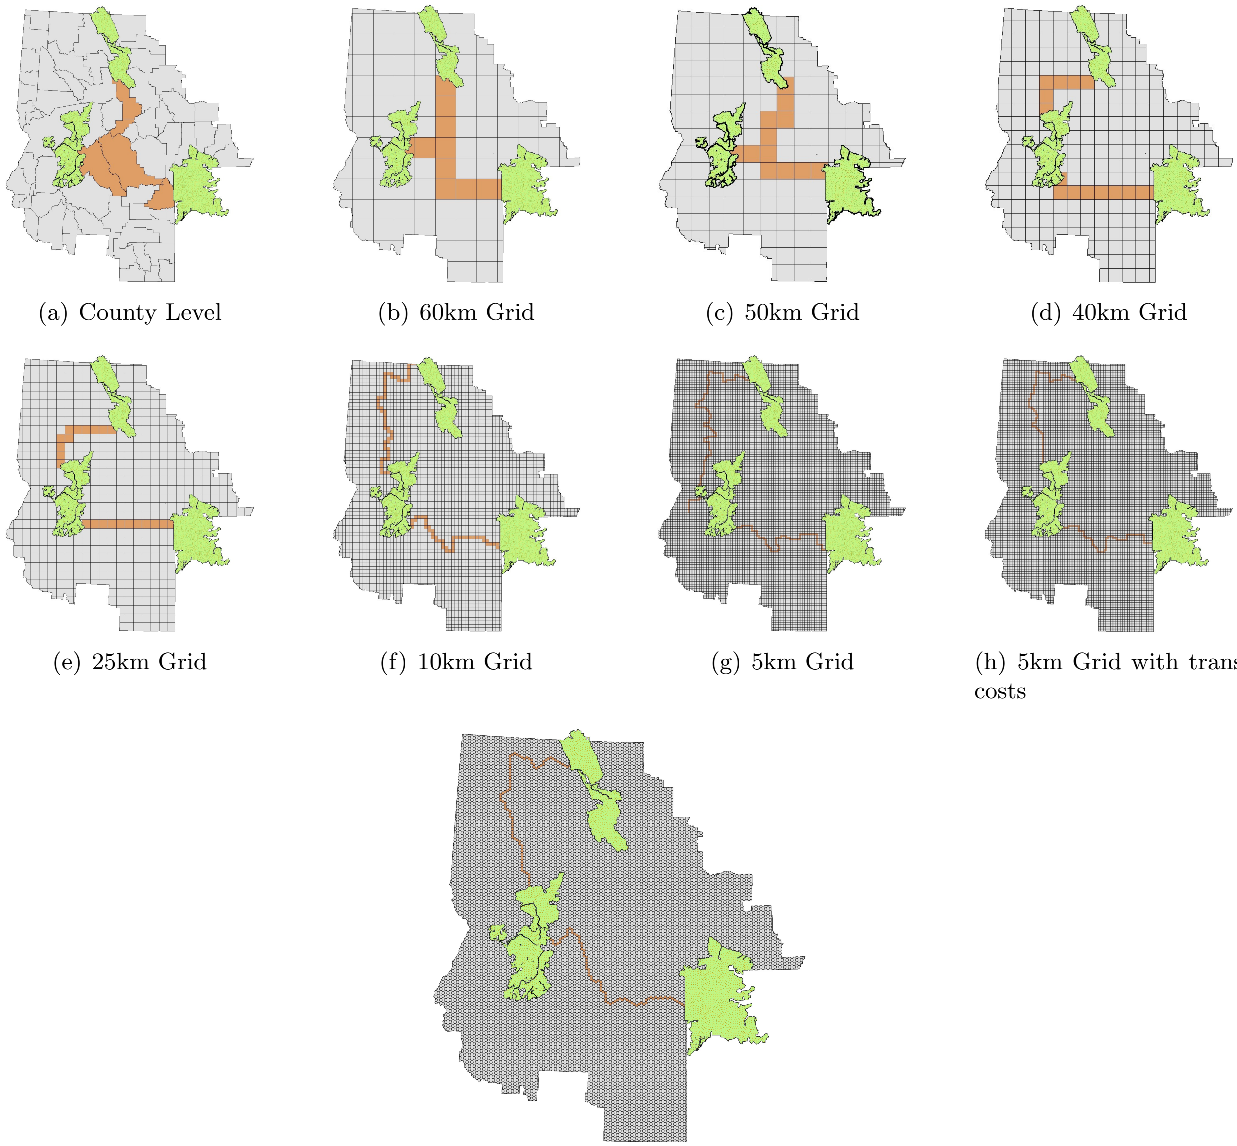

(a) County Level (b) 60km Grid (c) 50km Grid (d) 40km Grid

(e) 25km Grid (f) 10km Grid (g) 5km Grid (h) 5km Grid with trans costs

(i) 25 sq km Hex Grid

Figure 5: Minimum cost solutions for the corridor problem at various granularities: County level, 60 km grid, 50 km grid, 40 km grid, 25 km grid, 10 km grid, 5 km grid, 5 km grid with transaction costs, and 25 sq km hexagonal grid with transaction costs.

areas in southwestern Montana and instead chooses a slightly longer corridor that incorporates more national forest land. Thus, influencing the parcel granularity not only influences the estimated cost of the cheapest corridor, but it also has a significant influence on the general path that the corridor follows across the landscape.

In terms of *computational hardness*, as the granularity of the parcels is refined, the problem size and the corresponding search space grows rapidly. For example, while the County level abstraction of the problem has only 67 parcels, the 40 km square grid already has 242 parcels, and the 25 square km hexagonal grid—the one that yielded the most cost efficient solution—has 12,889 parcels. As a result, solving the connection subgraph model in a naïve manner using the CPLEX solver quickly becomes infeasible, especially with various budget levels beyond the basic minimum. Hence, we used the *two phase* solution procedure discussed in Section 4.2 in order to scale up our method. As mentioned earlier, the minimum cost solution obtained from a polynomial time Steiner tree implementation for three reserves was first extended greedily up to the available budget, and this solution was passed on to CPLEX as the starting solution. This turned out to be critical for CPLEX to even find *any* feasible solution to the corridor problem in a reasonable amount of time, once the granularity was increased to a 40 km square grid or finer.

The Impact of Transaction Costs. The addition of transaction costs to the model also alters the structure of optimal wildlife corridor. We again use as an example the minimum cost corridors considered in the discussion above and depicted in Figure 5. The most noticeable difference of the inclusion of a \$5,000 transaction cost per parcel at the 5 km level is that the number of parcels selected is reduced from 265 to 196. Given that each additional parcel adds incrementally to the overall cost of the corridor, even if there is no private land on the parcel, the minimum cost corridor selects parcels that provide more of a direct link between the reserve sites, rather than following a slightly longer path that includes more zero cost, national forest parcels. This difference is illustrated in panels (g) and (h) of Figure 5, which show the chosen 5 km corridor both with and without transaction costs. The most noticeable difference between the two corridors is the portion of the corridor connecting the Salmon-Selway to the Northern Continental Divide Ecosystem. With the inclusion of transaction costs, the parcels selected link directly to the northern portion of the Salmon-Selway, rather than the longer path selected without transaction costs that connects to the western edge of the ecosystem. With transaction costs the model also does not select the zero cost parcels that form a peninsula starting from the western edge of the Salmon-Selway. Thus, incorporating transactions costs has a significant influence on both the number and shape of the resulting corridor that is selected and represents an important consideration for land use planners.

Budgets Larger than the Minimum Cost. It is reasonable to expect land use planners to have budgets that are somewhat larger than the precise minimum cost needed for the cheapest wildlife corridor. In such cases a natural question to ask is: *Can neighboring land parcels be acquired which will significantly increase the net utility of the corridor?* Using results from the 25 sq km hexagonal grid, we show that a relatively small increase in the budget beyond the minimum cost can often lead to solutions of much higher utility.

Table 1 shows the impact of larger budgets on the utility of the resulting wildlife corridor. Specifically, it reports the utility levels of our *best found solutions* for various budgets.[10] In particular, while the minimum cost corridor (costing \$7.289 million) results in 169 parcels with an overall utility of 1,362,000 units, increasing the budget slightly—to only \$8 million—results in acquiring 311 parcels with an overall utility of 2,756,000, a more than two-fold increase. By doubling the budget, i.e., going for a budget of \$15 million rather than \$7.289 million, we have an over 10-fold increase in both the number of parcels purchased and the overall utility, which increases to 14,371,000.

These results suggest an interesting tradeoff between additional budget resources and expanding the wildlife corridor to beyond the basic minimum required for achieving connectivity. Instead of

[10]As mentioned earlier, we consider two related solution strategies—aiming for the full optimal solution or for the best possible solution that includes all parcels belonging to the minimum cost solution as part of the corridor. The table reports the best solution we found (not necessarily optimal but often provably close to optimal) with either approach for each budget level. A relative comparison of the two approaches is included later in this section.

Table 1: Wildlife corridors with budgets beyond the minimum cost, in the context of 25 sq km hexagonal grid with minimum cost = $7.289 million.

budget value (unit: $1M)	25 sq km hex grid corridor				
	cost (unit: $1M)	parcels number	parcels fractional increase	utility (1000×)	utility fractional increase
min	7.289	169	—	1,362	—
8.000	7.999	311	1.84x	2,756	2.02x
9.000	9.000	511	3.02x	4,599	3.38x
10.000	10.000	711	4.21x	6,498	4.77x
11.000	11.000	911	5.39x	8,270	6.07x
12.000	12.000	1,111	6.57x	9,973	7.32x
15.000	15.000	1,708	10.11x	14,371	10.55x
20.000	20.000	2,205	13.05x	17,477	12.83x
25.000	25.000	2,421	14.33x	19,068	14.00x
50.000	50.000	2,837	16.79x	22,229	16.32x

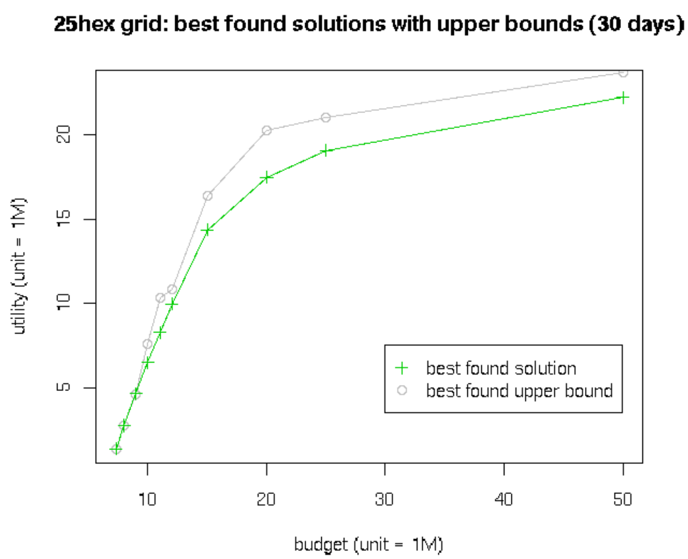

Figure 6: The best found solutions, along with provable optimality upper bounds, for the 25 square km hexagonal grid, with a 30 day computation cutoff. The upper bound in each case is computed based on the optimality gap reported by CPLEX. The plot shows utility values on y-axis for various budgets on x-axis.

focusing all resources on constructing a cost minimizing corridor, conservation planners may be better served by generating a feasible corridor of good utility (i.e., adequate width, suitable habitat, etc.) and then using the remaining budget to acquire nearby land with high net benefits.

Figure 6 shows the same data in two pictorial forms. The left pane of the figure highlights how the resulting utility rapidly increases as the budget level is initially increased, and both panels now also depict that our best found solutions for this challenging problem have utility levels *provably* very close to the respective optimal solutions. Specifically, the general trend in the left pane of the figure illustrates that the benefit-to-cost ratio of the best found corridor slowly flattens out as the budget is increased beyond $20 million. In other words, while there is a near-linear increase in

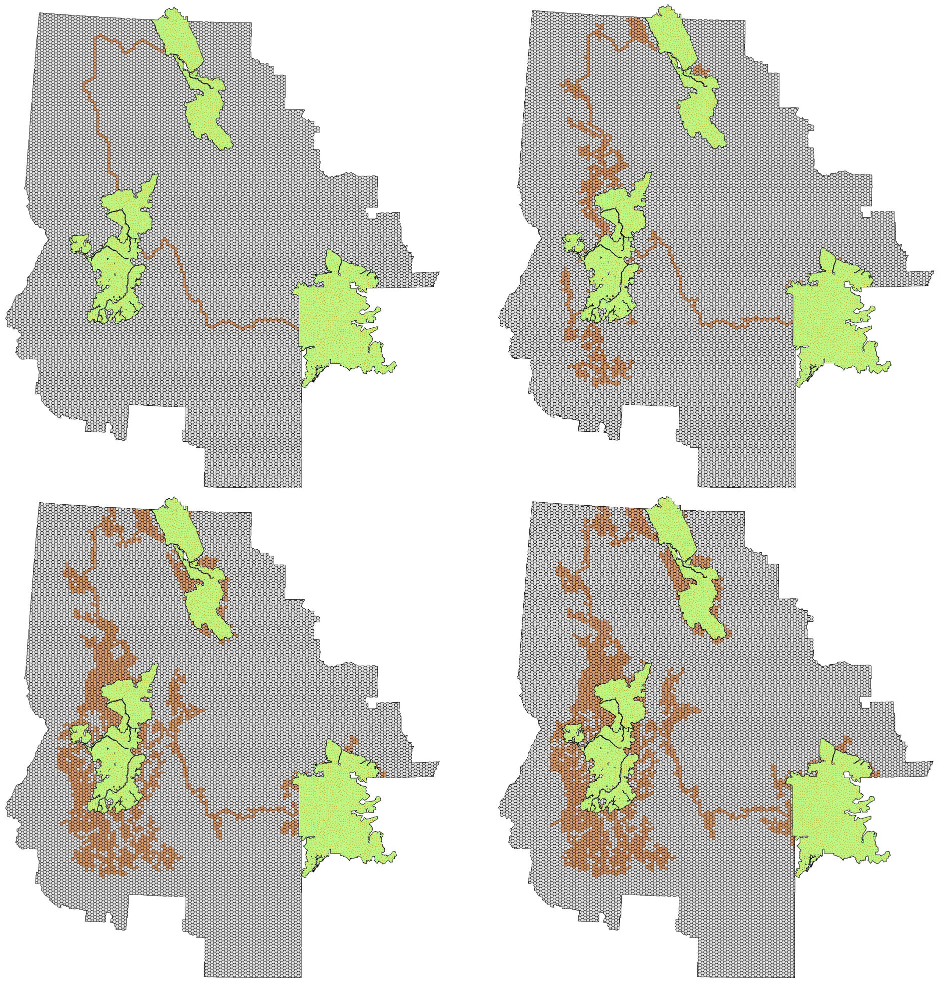

Figure 7: Map of corridors for the 25 sq km hexagonal grid with varying budget levels: minimum cost (\$7.289 million), \$10 million, \$15 million, \$20 million.

the attainable utility when increasing the budget from nearly \$7.3 million to \$15 million, there is a significant decline in the rate at which the utility increases when further increasing the budget level. For example, it could perhaps be argued that it is justifiable to invest \$15 million in order to obtain a 10-fold increase in utility and land over the minimum cost solution. At the same time, we also see that even with a budget as large as \$50 million, the best solutions found had a utility of "only" 16-fold that of the minimum cost corridor. Thus, depending on the available budget, land use planners can gain substantial insights from similar cost-benefit tradeoffs going beyond the cost of the cheapest corridor.

Figure 7 shows maps for a few different budget levels for the 25 sq km hex grid. Consider, for example, the solution for budget \$15 million and compare it with the minimum cost solution. The \$15 million solution suggests a general trend that by investing roughly twice the money needed for a minimally thin corridor, conservation agencies can in addition expand the usable natural area to a much larger general area. In this particular case, the map suggests that a substantial amount of land near Salmon-Selway (the reserve depicted towards the bottom-left) can be acquired to significantly enhance the overall value of the investment. In addition, the path leading north from this reserve also happens to be very "thick", increasing the value of the corridor itself.

Streamlining as an Aid for Very Large Instances. As noted earlier, we employed both a full MIP formulation as well as a restricted or "streamlined" version of it—where we restrict the search to only those corridors that include all of the parcels that are in the minimum cost solution found in Phase I. The motivation behind using such "extended mincost solutions" is computational feasibility—by restricting the search space, we hope to be able to attack larger problem instances than can be solved using the full MIP encoding. Figure 9 shows that such techniques do often pay off once instances become very large.

For smaller size instances that we can still solve completely to optimality, Figure 8 shows the relative gap between the utilities obtained with the optimal solution and the extended solution—for the 50 km and 40 km square grid abstractions of the corridor problem, corresponding to Figure 3(b) discussed earlier for the connection subgraph problem on semi-structured lattice instances. These plots show that the relative gap in this case is of the order of 2-5% when it is at its peak, and is usually within 2% of the optimal. This suggests that for this problem, one does not lose too much by solving only for the extended solution, which freezes the parcels included in the mincost solution and extends this solution optimally based on the available budget.

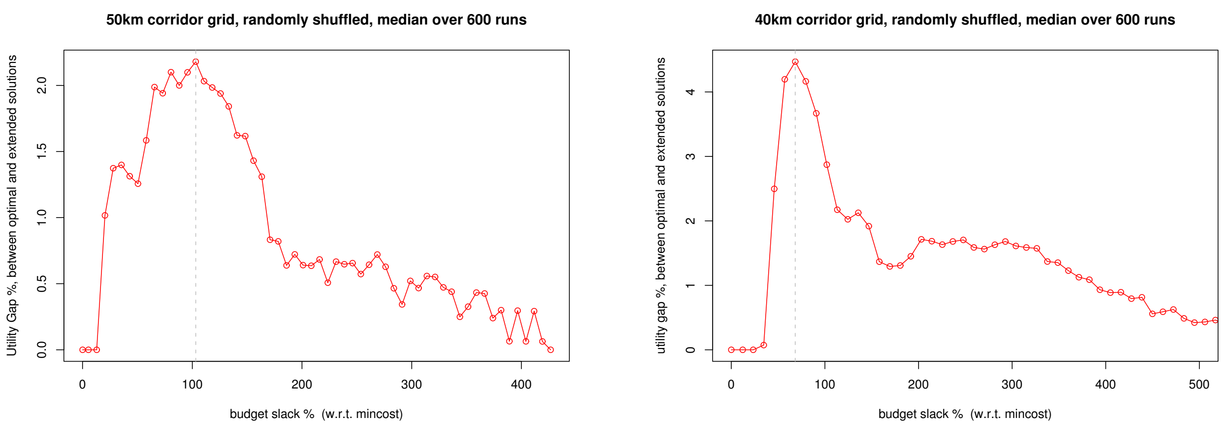

Figure 8: Percentage gap in the utility of optimal and extended-mincost solutions for the 50 km corridor grid (left) and the 40 km grid (right).

The 25 sq km hexagonal abstraction is, as one may expect, very challenging to solve, for both optimal and extended solutions. For instance, while the County level and the 50 km $\times$ 50 km square grid abstractions were solved to optimality within seconds with our two phase solution process, and the 40 km $\times$ 40 km square grid took only a few minutes to half an hour depending on the budget, the extended solution itself for the hexagonal grid required several days of computation, and for many budgets, could not be solved optimally in over 10 days. Fortunately, the eventual optimality gap for the best extended solutions found was quite low for budgets up to \$15 million, between 0 and 0.07%, meaning that the extended solutions were found to near optimality (the "best found" curve for extended solutions in Figure 9 is visually right on top of the corresponding "upper bound" curve up to a budget of \$15 million). The best found solutions for the true optimality runs, on the other hand, had a similarly low optimality gap for budgets under \$10 million but up to a 26.9% gap for higher budgets,[11] as seen from the top grey curve and the blue curve of Figure 9.

[11]For budget \$8 million, our full optimality runs reported an optimality gap of 57% for the best found solution

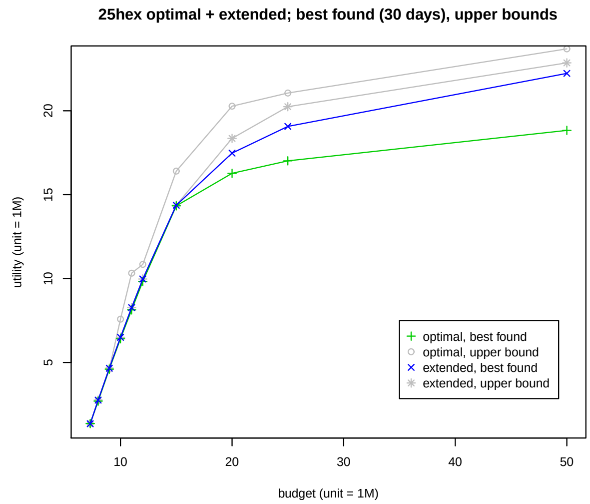

Figure 9: A comparison between best found solutions for the optimality run vs. the extended mincost solution runs, for the 25 square km hexagonal grid, with a 30 day cutoff. The upper bound, as before, is computed based on the optimality gap reported by CPLEX. The plot shows utility values on y-axis for various budgets on x-axis.

Interestingly, for higher budgets (especially \$20 million and higher, but also for \$11, \$12, and \$15 million), the best extended minimum cost solutions found for this challenging grid size were in fact of better quality than the best optimal solutions found, as seen from Figure 9. This aligns well with the concept of streamlining—if done carefully, restricting the search space to a small but promising part of the full space can result in much better solutions for computationally challenging problems. In this case, restricting the problem to only extended minimum cost corridors allowed CPLEX to focus the search and compute better quality solutions than usual within the limited amount of computation time.

Note also in the figure that the best found upper bound was also lower for the extended minimum cost solutions as compared to the full optimality runs. While this shows that the extended solutions were found to near optimality, it of course does not mean that for the full problem there is necessarily no solution of better quality than this upper bound.

6 Conclusion

Designing effective wildlife corridors is one of the key resource and environmental management problems in the newly emerging field of *Computational Sustainability*, which aims to apply state-of-the-art computational models and methods in order to solve challenging problems in the area of sustainability [18]. Real-world conservation problems are computationally highly demanding. Designing effective conservation corridors is particularly complex due to the intricate combinatorial structure of the problem induced by the spatial connectivity requirement in addition to a strict budget constraint. Solving this problem challenges the scalability limits of current computational optimization methods. Unlike many other conservation problems for which a marginal change in the available budget affects the resulting solution only marginally, the corridor design problem is

within 30 days of computation. However, recent ongoing work of Dilkina and Gomes [12] has shown that this solution is actually within 1% of the optimal solution.

quite unique in the sense that a *marginal change in the available budget can result in the selection of very different, potentially mutually exclusive, sets of parcels.*

In this work, we have developed algorithms to solve large scale corridor design problems, considering the minimum cost corridor but also going beyond the minimum cost solution to the best use of resources when a conservation planner has at his or her disposal a somewhat larger budget. Our empirical investigation into the general properties of the problem revealed the difficulty of solving it to optimality and an interesting easy-hard-easy pattern as a key problem parameter—the amount of extra budget beyond the minimum cost—is slowly increased. We presented a case study for a real-world instance exploring whether a corridor for grizzly bears in the U.S. Northern Rockies is financially feasible. The corridor being explored would provide a link between the Yellowstone, Salmon-Selway, and Northern Continental Divide ecosystems in Idaho, Wyoming, and Montana. Our study explores the implementation of such a corridor at various budget levels, showing, for example, that with a budget level of roughly twice the minimum cost, we can achieve over a 10-fold increase in both the number of parcels purchased and the overall utility.

Despite the evaluation of our method on a particular data set for the grizzly bear mentioned above, the methods and models developed here are novel and general. They could be applied to other cost and utility models beyond the ones used in our experiments, or even to designing effective corridors for other endangered species such as jaguar in Mexico and Central America. More generally, the techniques introduced here are applicable to any problem domain that can be re-formulated as the connection subgraph problem, such as in the context of social networks.

Acknowledgments

This work was supported by NSF Expeditions in Computing award for Computational Sustainability (grant 0832782), NSF IIS award (grant 0514429), and the Intelligent Information Systems Institute, Cornell University (AFOSR grant FA9550-04-1-0151). Part of this work was performed while the third and fifth authors were at Cornell University.

References

[1] A. Ando, J. Camm, S. Polasky, and A. Solow. Special distributions, land values, and efficient conservation. *Science*, 279(5359):2126–2128, 1998.

[2] H. P. Andreassen, S. Halle, and R. A. Ims. Optimal width of movement corridors for root voles: Not too narrow and not too wide. *Journal of Applied Ecology*, 33:63–70, 1996.

[3] M. W. Bern and P. E. Plassmann. The Steiner tree problem with edge lengths 1 and 2. *Information Processing Letters*, 32(4):171–176, 1989.

[4] M. Cabeza, M. B. Araújo, R. J. Wilson, C. D. Thomas, M. J. R. Cowley, and A. Moilanen. Combining probabilities of occurrence with spatial reserve design. *Journal of Applied Ecology*, 41(2):252–262, 2004.

[5] J. D. Camm, S. Polasky, A. R. Solow, and B. Csuti. A note on optimal algorithms for reserve site selection. *Biological Conservation*, 78(3):353–355, Dec. 1996.

[6] J. O. Cerdeira and L. S. Pinto. Requiring connectivity in the set covering problem. *Journal of Combinatorial Optimization*, 9:35–47, 2005.

[7] J. O. Cerdeira, K. J. Gaston, and L. S. Pinto. Connectivity in priority area selection for conservation. *Environmental Modeling and Assessment*, 10(3):183–192, 2005.

[8] J. O. Cerdeira, L. S. Pinto, M. Cabeza, and K. J. Gaston. Species specific connectivity in reserve-network design using graphs, 2009. In preparation.

[9] R. L. Church, D. M. Stoms, and F. W. Davis. Reserve selection as a maximal covering location problem. *Biological Conservation*, 76(2):105–112, 1996.

[10] C. Costello and S. Polasky. Dynamic reserve site selection. *Resource and Energy Economics*, 26(2):157–174, June 2004.

[11] B. Csuti, S. Polasky, P. Williams, R. Pressey, J. Camm, M. Kershaw, R. Kiester, B. Downs, R. Hamilton, M. Huso, and K. Sahr. A comparison of reserve selection algorithms using data on terrestrial vertebrates in Oregon. *Biological Conservation*, 80(1):83–97, 1997.

[12] B. Dilkina and C. P. Gomes. Solving connected subgraph problems in wildlife conservation. In *CPAIOR-10: 7th International Conference on Integration of AI and OR Techniques in Constraint Programming*, volume 6140 of *Lecture Notes in Computer Science*, pages 102–116, Bologna, Italy, June 2010.

[13] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–486, 2005.

[14] J. D. Dixon, M. K. Oli, M. C. Wooten, T. H. Eason, J. W. McCown, and D. Paetkau. Effectiveness of a regional corridor in connecting two Florida black bear populations. *Conservation Biology*, 20(1):155–162, Feb. 2006.

[15] C. Faloutsos, K. McCurley, and A. Tomkins. Fast discovery of connection subgraphs. In *Proceedings of the 2004 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 118–127. ACM Press, 2004.

[16] P. J. Ferraro. Conservation contracting in heterogeneous landscapes: An application to watershed protection with threshold constraints. *Agricultural and Resource Economics Review*, 32(1):53–64, 2003.

[17] T. Fuller, M. Munguía, M. Mayfield, V. Sánchez-Cordero, and S. Sarkar. Incorporating connectivity into conservation planning: A multi-criteria case study from central mexico. *Biological Conservation*, 133(2):131–142, 2006.

[18] C. P. Gomes. Computational Sustainability: Computational methods for a sustainable environment, economy, and society. *The Bridge, National Academy of Engineering*, 39(4), Winter 2009.

[19] R. Groeneveld. Economic considerations in the optimal size and number of reserve sites. *Ecological Economics*, 52:219–228, 2005.

[20] N. M. Haddad, D. R. Bowne, A. Cunningham, B. J. Danielson, D. J. Levey, S. Sargent, and T. Spira. Corridor use by diverse taxa. *Ecology*, 84(3):609–615, 2003.

[21] ILOG, SA. CPLEX 10.1 Reference Manual, 2006.

[22] Land Trust Alliance. New York State Conservation Partnership Program 2003 awards, 2004. Available at http://www.lta.org/yourregion/ny/nyscpp_2003_awards.htm.

[23] R. K. Mace and J. Waller. Final report: Grizzly bear ecology in the swan mountains. Technical report, Montana Department of Fish, Wildlife and Parks, 1997.

[24] C. R. Margules, A. O. Nicholls, and R. L. Pressey. Selecting networks of reserves to maximise biological diversity. *Biological Conservation*, 43(1):63–76, 1988.

[25] M. D. McDonnell, H. P. Possingham, I. R. Ball, and E. A. Cousins. Mathematical methods for spatially cohesive reserve design. *Environmental Modeling and Assessment*, 7:107–114, 2002.

[26] A. McEuen. The wildlife corridor controversy: a review. *Endangered Species Update*, pages 1–6, 1993.

[27] K. D. Messer. The conservation benefits of cost-effective land acquisition: A case study in Maryland. *Journal of Environmental Management*, 79(3):305–315, 2006.

[28] R. Naidoo, A. Balmford, P. J. Ferraro, S. Polasky, T. H. Ricketts, and M. Rouget. Integrating economic costs into conservation planning. *Trends in Ecology & Evolution*, 21(12):681–687, 2006.

[29] D. A. Newburn, P. Berck, and A. M. Merenlender. Habitat and open space at risk of land-use conversion: Targeting strategies for land conservation. *American Journal of Agricultural Economics*, 88(1):28–42, 2006.

[30] R. F. Noss. Corridors in real landscapes: a reply to Simberloff and Cox. *Conservation Biology*, 1(2):159–164, 1987.

[31] H. Önal and R. A. Briers. Optimal selection of a connected reserve network. *Operations Research*, 54(2):379–388, 2006.

[32] H. Önal and Y. Wang. A graph theory approach for designing conservation reserve networks with minimal fragmentation. *Networks*, 51(2):142–152, 2007.

[33] S. Polasky, J. D. Camm, and B. Garber-Yonts. Selecting biological reserves cost-effectively: An application to terrestrial vertebrate conservation in Oregon. *Land Economics*, 77(1):68–78, 2001.

[34] H. J. Prömel and A. Steger. *The Steiner Tree Problem: A Tour Through Graphs, Algorithms, and Complexity*. Vieweg, 2002.

[35] T. Schmitt and A. Seitz. Influence of habitat fragmentation on the genetic structure of *Polyommatus Coridon* (Lepidoptera: Lycaenidae): Implications for conservation. *Biological Conservation*, 107(3):291–297, Oct. 2002.

[36] J. Sessions. Solving for habitat connections as a Steiner network problem. *Forest Science*, 38 (1):203–207, 1992.

[37] J. J. Tewksbury, D. J. Levey, N. M. Haddad, S. Sargent, J. L. Orrock, A. Weldon, B. J. Danielson, J. Brinkerhoff, E. I. Damschen, and P. Townsend. Corridors affect plants, animals, and their interactions in fragmented landscapes. *Proceedings of the National Academy of Science*, 99(2):12923–12926, Oct. 2002.

[38] Trust for Public Land. TPL landvote database, 2010. Available at https://www.quickbase.com/db/bbqna2qct?a=dbpage&pageID=10.

[39] L. G. Underhill. Optimal and suboptimal reserve selection algorithms. *Biological Conservation*, 70(1):85–87, 1994.

[40] United States Department of Agriculture. Census of agriculture, 2002.

[41] United States Department of Agriculture. Conservation Reserve Program: Annual summary, 2002.

[42] United States Geologic Survey. Gap analysis – land stewardship coverage, 1999. Available at http://gapanalysis.nbii.gov.

[43] R. P. D. Vanderkam, Y. F. Wiersma, and D. J. King. Heuristic algorithms vs. linear programs for designing efficient conservation reserve networks: Evaluation of solution optimality and processing time. *Biological Conservation*, 137:349–358, 2007.

[44] J. C. Williams. A zero-one programming model for contiguous land acquisition. *Geographical Analysis*, 34(4):330–349, 2002.

[45] J. C. Williams. Delineating protected wildlife corridors with multi-objective programming. *Environmental Modeling and Assessment*, 3(1-2):77–86, 1998.

[46] J. C. Williams and S. A. Snyder. Restoring habitat corridors in fragmented landscapes using optimization and percolation models. *Environmental Modeling and Assessment*, 10(3):239–250, 2005.

[47] J. C. Williams, C. S. ReVelle, and S. A. Levin. Spatial attributes and reserve design models: A review. *Environmental Modeling and Assessment*, 10(3):163–181, 2005.

[48] W. Zhang and R. E. Korf. A study of complexity transitions on the asymmetric traveling salesman problem. *Artificial Intelligence*, 81:223–239, 1996.

A Illustrative example

Example 1. We consider a simple examle to illustrate some of the combinatorial issues of the connection subgraph problem that make it computationally hard. Consider the hypothetical 3×3 parcel map presented in Figure 10. We use this map to illustrate, among other aspects, how both the optimal choice of parcels and the complexity of finding these parcels can vary dramatically when we formulate the problem as a cost constrained utility optimization problem rather than an unconstrained cost minimization problem. In this simple example, when ignoring utilities, the cost of the corridor is minimized with the selection of parcels B, E, and H, as shown in panel I. With this selection, the cost is 7 units and the utility of the parcels selected is 5. Now suppose that the conservation planner has available a budget of 10 units. Rather than simply selecting the least cost path consisting of parcels B, E, and H, the planner would now be interested in finding the corridor that yields the highest possible utility, with a cost of no more than 10 units. In panel II, we show that for a budget of 10 units, the planner maximizes utility by selecting parcels E, F, H, and I, for a total utility of 9 units. If the conservation planner's budget is increased to 11 units, as in panel III, the optimal selection of parcels is A, B, D, with a corresponding aggregate utility of 10 units.

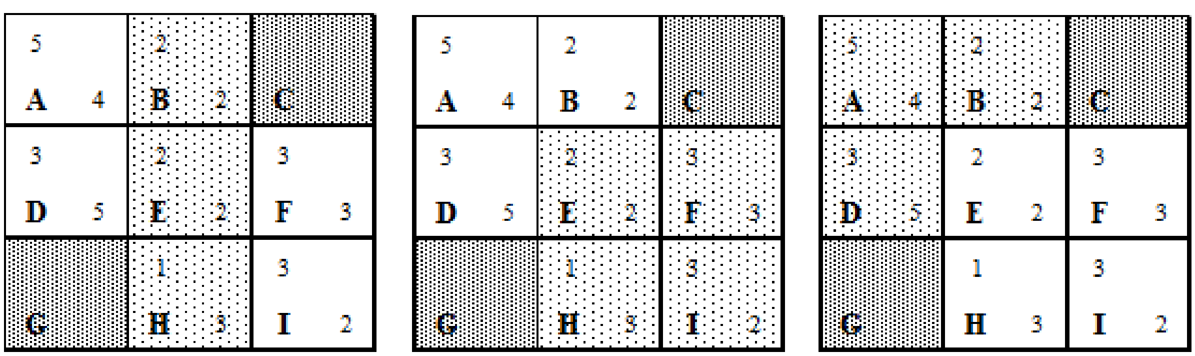

I. Minimum Cost corridor
cost = 7, utility = 5

II. Budget = 10
cost = 10, utility = 9

III. Budget = 11
cost = 11, utility = 10

Figure 10: Hypothetical corridor optimization. Parcel labels are provided in the lower left corner of each parcel, costs are in the lower right corner, and utilities are in the upper left corner. The reserves, A and C, are marked as dark gray. The optimal corridor in each case is shaded.

It is not surprising that considering only parcel costs as in panel I results in a very different set of selected parcels from that in panels II and III, where both parcel cost and utility are considered. What is unique about the constrained corridor optimization problem is that a *marginal change in the available budget can result in the selection of very different, potentially mutually exclusive, sets of parcels*, as illustrated in panels II and III. Given the constraint that all of the selected parcels must be connected, the model outcomes can change drastically as budget levels are varied, which is different from typical reserve site selection models where additional budget levels generally only influence the selection of a small subset of the available parcels.

Figure 10 also illustrates the computational challenges of the budget constrained utility maximization problem. If the goal is to find a least cost path, as has been done in all previous studies, only six possible paths in the 3×3 parcel grid need to be considered. The optimal selection will never include paths that are more than one parcel wide, as this can only add to the cost of the corridor. For the case of constrained utility maximization, however, the set of feasible corridors jumps from six to thirty. Thus, even in this small hypothetical case, the challenge of maximizing utility given a budget constraint is considerably greater than simply finding the single-parcel-wide least cost path. The computation complexity of the problem is analyzed more rigorously in in the following sections and we deal with the challenges of reaching an optimal solution for the Northern Rockies corridor later in the paper. □

B Proof Details

Proof of Theorem 1. The problem is clearly in NP, because a certificate subgraph H can be easily verified to have the desired properties, namely, connectedness, low enough cost, and high enough utility. For NP-hardness, consider the Steiner tree problem on a graph $\widehat{G} = (\widehat{V}, \widehat{E})$ with terminal set $\widehat{T} \subseteq \widehat{V}$, edge cost function $\widehat{c} : \widehat{E} \to \mathbb{Z}^+$, and cost bound $\widehat{C}$.

An instance of the connection subgraph problem can be constructed from this as follows. Construct a graph $G = (V, E)$ with $V = \widehat{V} \cup \widehat{E}$ and edges defined as follows. For every edge $e = \{v, w\} \in \widehat{E}$, create edges $\{v, e\}, \{w, e\} \in E$. The terminal set remains the same: $T = \widehat{T}$. Overall, $|V| = |\widehat{V}| + |\widehat{E}|$, $|E| = 2|\widehat{E}|$, and $|T| = |\widehat{T}|$. For costs, set $c(v) = 0$ for $v \in \widehat{V}$ and $c(e) = \widehat{c}(e)$. For utilities, set $u(v) = 1$ for $v \in T$ and $u(v) = 0$ for $v \notin T$. Finally, the cost bound for the connection subgraph is $C = \widehat{C}$ and the utility bound is $U = |E|$.

It is easy to verify that the Steiner tree problem on $\widehat{G}$ and $\widehat{T}$ has a solution with cost at most C iff the connection subgraph problem on G and T has a solution with cost at most C and utility at least U. This completes the reduction.

Note that if $\widehat{G}$ is planar, then so is G. Further, the reduction is oblivious to the number of terminals in G. Hence, NP-completeness holds even on planar graphs and without any terminals. □

Proof of Lemma 1. We give a reduction along the lines of the one given by Bern and Plassmann [3] for the Steiner tree problem. The reduction is oblivious to the number of terminals, and holds in particular even when there are no terminals.

Recall that a vertex cover of a graph $\widehat{G} = (\widehat{V}, \widehat{E})$ is a set of vertices $V' \subseteq \widehat{V}$ such that for every edge $\{v, w\} \in \widehat{E}$, at least one of v and w is in V'. The vertex cover problem is to determine whether, given $\widehat{G}$ and $C \geq 0$, there exists a vertex cover V' of $\widehat{G}$ with $|V'| \leq C$. We convert this into an instance of the connection subgraph problem. An example of such a graph is depicted in Fig. B.

Create a graph $G = (V, E)$ with $V = \widehat{V} \cup \widehat{E}$ and edges defined as follows. For every $v, w \in \widehat{V}, v \neq w$, create edge $\{v, w\} \in E$; for every $e = \{v, w\} \in \widehat{E}$, create edges $\{v, e\}, \{w, e\} \in E$.

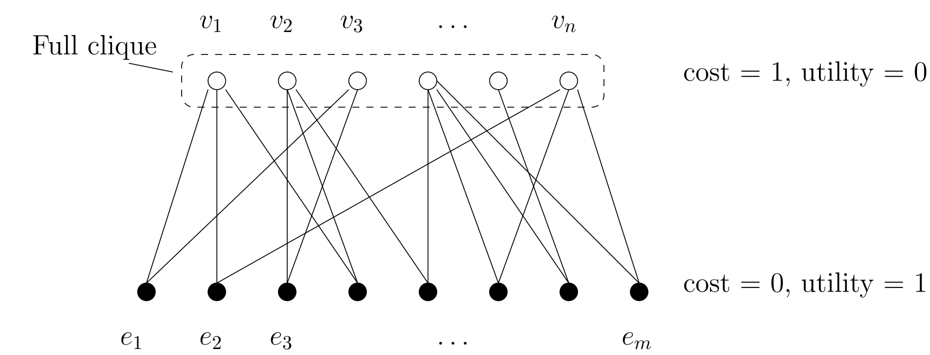

Edges in the original graph $\widehat{G}$:
$e_1 = (v_1, v_3), e_2 = (v_1, v_n), e_3 = (v_2, v_3), \ldots, e_m = (v_{n-2}, v_n)$

Figure 11: Reduction from Vertex Cover

Overall, G has $|\widehat{V}| + |\widehat{E}|$ vertices and $\binom{\widehat{V}}{2} + 2\widehat{E}$ edges. For costs, set $c(v)$ to be 1 if $v \in \widehat{V}$, and 0 otherwise. For utilities, set $u(e)$ to be 1 if $e \in \widehat{E}$, and 0 otherwise. Finally, fix the set of terminals to be an arbitrary subset of $\widehat{E}$.

We prove that solutions to the connection subgraph problem on G with costs and utilities as above, cost bound C, and desired utility $U = |\widehat{E}|$ are in one-to-one correspondence with vertex covers of $\widehat{G}$ of size at most C.

First, let vertex-induced subgraph H of G be a solution to the connection subgraph instance. Let $V' = V(H) \cap \widehat{V}$. We claim that V' is a vertex cover of $\widehat{G}$ of size at most C. Clearly, $|V'| \leq C$ because of the cost constraint on H. To see that V' is indeed a vertex cover of $\widehat{G}$, note that (A) because of the utility constraint, V' must contain *all* of the vertices from $\widehat{E}$, and (B) because of the connectedness constraint, every such vertex must have at least one edge in $E(H)$, i.e., for each $e = \{v, w\} \in \widehat{E}$, V' must include at least one of v and w.

Conversely, let V' be a vertex cover of $\widehat{G}$ with at most C vertices. This directly yields a solution H of the connection subgraph problem: let H be the subgraph of G induced by vertices $V' \cup \widehat{E}$. By construction, H has the same cost as V' (in particular, at most C) and has utility exactly U. Since V' is a vertex cover, for every edge $e = \{v, w\} \in \widehat{E}$, at least one of v and w must be in V', which implies that H must have at least one edge involving e and a vertex in V'. From this, and the fact that all vertices of V' already form a clique in H, it follows that H itself is connected.

This settles our claim that solutions to the two problem instances are in one-to-one correspondence, and finishes the proof. □

C Computational Hardness Profies for Synthetic Data

We make use of semi-structured graphs, with uniform random utility and cost functions. The graphs are composed of an $m \times m$ rectangular lattice or grid, where the order m is either 6, 8, or 10. This lattice graph is motivated by the structure of the original conservation corridors problem. In this lattice, we place k terminal vertices; in the results reported here, k is 0 or 3. When $k \geq 2$, we place two terminal vertices in the 'upper left' and 'lower right' corners of the lattice, so as to maximize the distance between them and "cover" most of the graph. This is done to avoid the occurrence of too many pathological cases, where most of the graph does not play any role in constructing an optimal connection subgraph. The remaining $k-2$ terminal vertices are placed uniformly at random in the graph. To define the utility and cost functions, we assign uniformly and independently at random a utility and a cost from the set $\{1, 2, \ldots, 10\}$ to each vertex in the graph. The cost and utility functions are uncorrelated.

In the figures below, each data point is based on statistics over 100 to 500 randomly generated instances. Figures 12 shows computational hardness results for the case of zero reserves. Note that these instances are always feasible, even with zero budget. Using the budget slack percentage relative to mincost as in Figures 3(a) and 3(b) earlier, therefore, does not make sense in this case. We simply use use for the x-axis the fraction *budget/total-budget*, where *total-budget* is the total cost of all vertices.

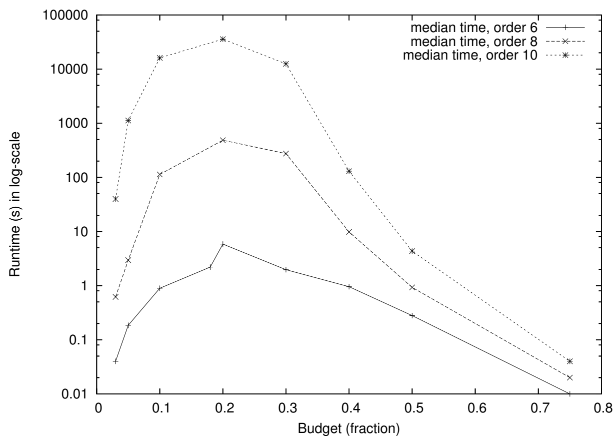

Figure 12: Hardness profile for lattices of order 6, 8, and 10, without terminal vertices.

These connection subgraph instances, defined on graphs without terminal vertices, are always satisfiable and can thus be seen as instances of a pure optimization problems. Figure 12 shows the hardness profile (i.e., the running time) on lattices of order 6, 8, and 10. Notice that the median runtime (y-axis) is plotted in log-scale in this figure. The plots clearly indicate an easy-hard-easy pattern for these instances, even though they are all feasible with respect to the budget. Such patterns have been observed previously in some pure optimization problems, but only under specialized random distributions. For example, Zhang and Korf [48] identify a similar pattern for the Traveling Salesperson Problem, using a log-normal distribution of the distance function. In our case, the pattern naturally emerges from the model.