A High-Performance Graph Model for Near-Optimal Payments for Ecosystem Services

Florian Berlinger,*[*] Lily Xu,*[*] Yiling Chen
Harvard University
fberlinger@seas.harvard.edu, lily_xu@g.harvard.edu, yiling@seas.harvard.edu

Abstract

Payments for ecosystem services (PES) offer compensation to landholders to preserve instead of develop their land, such that everyone benefits from the positive externalities such as clean air and fresh water. Ideally, payment models should be designed to maximize the total societal benefit of PES, i.e., the positive externalities minus the foregone development turnovers for landholders. Optimization is complicated by complementarity effects, which create additional value when land is preserved together. The combinatorial problem size render existing approaches intractable. Assuming knowledge of landholders' costs from an auction mechanism, we introduce a novel graph representation for PES, in which nodes hold landholder costs and preservation values, and edges represent complementarities. We offer a graph-search algorithm to find near-optimal preservation schemes in polynomial time, recuperating up to 85% of optimal preservation. We expect such algorithms to become a reliable and flexible tool for dynamic decision-making for PES.

1 Introduction

When a landholder owns a plot of forest and cuts down the trees for timber, she benefits from the financial gain of the cost of timber. If she instead decides to leave the forest intact, she then foregoes money the timber would have fetched, but instead benefits from improved air and water quality, increased biodiversity, and reduced climate change (Stein et al. 2009). Everyone else in society receives these benefits, too. These are the positive externalities of preserved forests, which unfortunately come with no economic gain for the landholder. Further ecosystem services include water and biodiversity, which, together with forestry, are examples of renewable natural capital. The cumulative value of the world's ecosystem services are enormous, providing an estimated US$33 trillion per year (Costanza et al. 1997).

Payments for ecosystem services (PES) remedy this economic incentive issue by offering payments to landholders for taking environmentally-friendly actions such as preserving forests, reducing use of pesticides, or eliminating toxic emissions. Such payments offer economic compensation for the positive externalities provided by these natural capital stocks, thereby altering potentially environmentally harmful behavior of landholders. Without external intervention, the situation is similar to the tragedy of the commons (Hardin 1968), for which well-known regulatory mechanisms exist (Ostrom 1990). However, in shared commons, the members share the positive and negative (economic) consequences of sustainable or unsustainable usage. By contrast, in PES, a multitude of private landholders individually affect non-monetary positive externalities for everyone, while optimizing their own economic situation. As a consequence, PES uses payments as behavioral incentives rather than regulation.

Figure 1: A forest in Costa Rica, where payments for ecosystem services have been implemented since the nation-wide Forest Law of 1996 (Wallbott, Siciliano, and Lederer 2019)

One particular challenge with PES allocation is the presence of complementarities, in which natural capital (such as uncut forest cover) has greater value when left untouched in large, contiguous blocks of land. For instance, contiguous land provides more territory for the movement of wildlife (Crooks and Sanjayan 2006; Dilkina et al. 2017). It may be the case that the cost of preservation for a set of individual parcels of land is higher than the societal benefits, but when combined with other parcels, complementarity effects render the combination profitable. As a consequence, a greedy budget allocation mechanism, which picks parcels of land ordered by their individual net values, does not necessarily

[*] The authors contributed equally to the work. Their names are listed in alphabetical order.

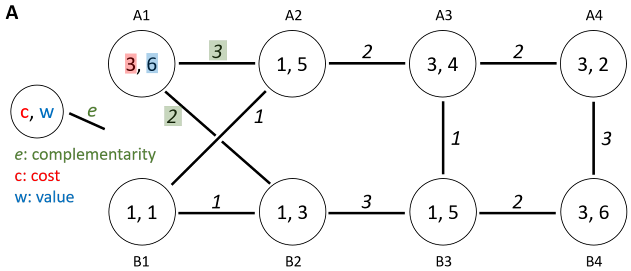

B

	1	2	3	4
A	3 6 9 11	1 5 8 10 11	3 4 5 7 9	3 2 5 7
B	1 1 2 3	1 3 6 8 9	1 5 8 10 11	3 6 9 11

Figure 2: A novel framework to represent and optimize payment for ecosystem services programs for given parcels of land and complementarities among them. (A) Our graph representation, in which nodes represent parcels of land of preservation cost *c* and value *w*, and edges represent complementarity gains *e* for parcels that are preserved together. (B) A previous tabular representation that cannot assign complementarities to specific duos of parcels that are preserved together (Polasky et al. 2014). Cost is represented by the top number, value by the leftmost bottom number, and cumulative value plus complementarities for 1, 2, 3, ... connected parcels by the following bottom numbers. (A) and (B) represent the same parcels of land, whereby our approach (A) is more specific in the way it can describe complementarities.

yield the highest overall value. Instead, more complex optimization schemes must be applied across the landscape of all possible combinations of parcels.

To address these limitations, we introduce a novel graph-based model to represent and optimize PES programs (Figure 2A). A key advantage of our new graph representation is that we are now able to leverage a suite of graph-traversal algorithms, adapted to this domain, to better optimize various payment schemes. Many graph-traversal algorithms are computationally efficient, and some also provide provable bounds on the quality of achievable solutions.

Assuming knowledge of the landholders' cost and societal benefit of preserving the land, we aim to reduce the computational complexity of finding near-optimal sets of land parcels while incurring only bounded losses in total societal value from PES. Our results include a performance evaluation of several hand-picked and customized graph-traversal algorithms against the exhaustive search, which we ran on our new graph representation of PES. We see our work as an important contribution toward high-performance near-optimal PES that can be employed flexibly and dynamically, as budgetary constraints or changes in landholders and their evaluations require. For example, the merging of two plots of land into one would simply combine those nodes, and additional plots of land can be easily added or removed. Distinguishing ourselves further from past work (Jayachandran et al. 2017), we hope that future studies on and implementations of PES take landholder costs and complementarity effects into account, and assign budgets as optimally as computational constraints allow.

2 Background

PES programs are widely implemented around the world by governments and non-governmental organizations (NGOs) alike. There are over 550 PES programs globally with an estimated US$36–42 billion in annual transactions (Salzman et al. 2018). For example, the United Nations REDD program works in over 65 countries, aiming to reduce carbon emissions by providing payments to reduce deforestation while promoting sustainable development (Cerbu, Swallow, and Thompson 2011). There has been significant literature from the perspective of environmental management to demonstrate the effectiveness of PES programs. For instance, a randomized controlled trial conducted in Uganda demonstrated that villages where a PES program was implemented had only 4.2% decline in tree cover after two years compared to 9.1% in control villages (Jayachandran et al. 2017). In addition to providing more sustainable ecosystems, PES programs have also been shown to provide broader societal benefits by strengthening social relationships and a sense of community (Alix-Garcia et al. 2018).

Ecosystem services encompass water, biodiversity, and forestry, which are examples of renewable natural capital. The challenge in renewable natural capital markets is the presence of externalities, both positive and negative. Market design can be used to address these externalities. Specific challenges with renewable natural capital markets (compared to other capital markets) are (a) heterogeneity, (b) high cost of monitoring and transactions, and (c) complementarities (Teytelboym 2019). With complementarities, the challenge is that natural capital (such as uncut forest cover) has greater value when left untouched in large, contiguous blocks of land, for instance by providing more territory and movement for wildlife (Crooks and Sanjayan 2006). Similar principles have been applied to design wildlife corridors, offering continuous paths between animal habitats (Dilkina et al. 2017).

There are significant ecological benefits from strategically selecting parcels of land to conserve. The cost of conservation may vary drastically for a single square kilometer of land, up to seven orders of magnitude (Balmford et al. 2003), enabling an agency to preserve much greater ecological value for the same amount of cost. They compute conservation benefit in terms of wilderness value, areal extent, and wildlife density. Naidoo et al. (2006) show that considering the economic cost and ecological benefits in conservation planning yields greater ecological gain under a limited budget.

Among recent efforts, one seminal piece of work addressed such implementation of optimal PES as a two-stage procedure (Polasky et al. 2014), focusing on the problem of information asymmetry to use a strategy-proof auction to learn the cost of conservation for landholders. The benefits of preservation and complementarities are assumed to be known. In the first step, landholders bid a price for conservation of their parcels in a Vickrey–Clarke–Groves (VCG) auction. The second step uses those true prices for conservation and performs an exhaustive search over all possible combinations of land parcels to find the optimal set of land parcels, whose preservation yields the highest societal benefit. A key limitation of this approach is computational tractability, since the number of combinations scales as 2^N with the number of parcels N.

We illustrate this limitation as follows. Compared to a previous study (Jayachandran et al. 2017) which studied 121 villages, 60 of which received payments for conserving land, let's assume we have $N = 80$ landholders and a supercomputer capable of 10^{15} floating-point operations per second (FLOPS). Let's further assume that the number of FLOPS to evaluate a single combination is a constant c and generously assign $c = 1$. The computer would require more than $(2^{80}/10^{15})/(3600 \cdot 24 \cdot 365) \equiv 38y$ to find the optimal solution, and this time would double with each additional landholder. Consequently, an exhaustive search may not be tractable, especially if PES is considered a continuous long-term effort. Over time, new optimal budget allocations may have to be recomputed: for instance, after changes to PES budgets, landholder costs, and societal values, or additions and removals of parcels of land. In addition, complementarities might depend on the ecosystem goal, e.g., improved air may have lower complementarities between neighboring parcels of land than movement of wildlife. In the current approach, a cost comparison between such goals would require a costly recomputation of the entire optimal budget allocation.

Other related work have modeled the problem of selecting reserve sites to preserve suitable habitat for the greatest number of species as a maximal covering problem (Church, Stoms, and Davis 1996). However, computational tractability is again a problem, as their integer linear programming solution formulation is NP-hard, and they do not consider complementarity.

Overall, past work has indicated that PES is an effective strategy to preserve the environment by combating deforestation. The challenge is to determine a payment scheme across many landholders that maximizes social welfare, while accounting for complementarities. Existing approaches are computationally intractable, and often neglect complementarities. In this work, we remedy both challenges.

3 Graph-based model

Due to the computational intractability of computing an optimal solution directly, we propose formulating the problem as a graph. A graph-based model enables us to neatly account for complementarity as the value added from the inclusion of mutually beneficial parcels. It also enables us to incorporate for spatial relationships, which are an essential characteristic of land plots.

Suppose we have N parcels of land and a budget of B. Let $c_i \geq 0$ be the value to the landholder of parcel i. This corresponds with, for example, the opportunity cost they incur according to the potential profit from developing the land. Let $w_i \geq 0$ be the societal benefit, determined by the value of the ecosystem services provided by that parcel of land, perhaps measured in terms of carbon offsets provided by the forest cover. Finally, let $e_{ij} \geq 0$ be the complementarity gain of conserving both parcels i and j.

Each parcel of land i is represented as one node, which stores the cost c_i to the landholder and the benefit to society w_i of preserving that parcel. We include edges between parcels of land whose mutual preservation would offer complementarity benefits. The weight of the edge e_{ij} represents the complementarity gain if both parcels i and j are preserved. Our graph representation is well-suited for the spatial nature of land preservation. For example, edges may be added between all adjacent land parcels. Note that this graph representation is flexible to miscellaneously shaped parcels; the land does not have to be divided into a grid. It also allows for simpler visualizations and insertion/deletion of new parcels compared to the tabular format (Figure 2B) in previous work (Polasky et al. 2014).

While previous work (Polasky et al. 2014) modelled complementarities for a given parcel as additive gains per number of adjacent conserved parcels (Figure 2B), our PES representation is more comprehensive since it allows to attribute such gains to specific connected parcels. Our representation is also more realistic, as neighboring parcels should not be treated equally. For example, conserving a neighboring parcel containing a water source will offer greater complementarity benefits than conserving a neighboring parcel that has recently experienced a wildfire. The graph in Figure 2A represents one exemplary realization recovered from the more general PES representation in Figure 2B (reroute B2 to B4 and A4 to B3 to see another valid realization).

4 Graph-based algorithms

We introduce the following five algorithms to determine a payment scheme for PES. We compare them to the optimal solution, which maximizes the net social value (benefit minus cost) for the entire set of land parcels. We also discuss properties of the algorithms, such as their time complexities, and real world-situations in which they might be especially suitable. Finally, we present our near-optimal algorithm, GRAPES.

4.1 Baseline algorithms

We use the following naive and graph-based greedy algorithms to compare against.

Brute-force optimal Our algorithmic implementation follows a bottom-up recursive approach to compute the optimal solution as in (Polasky et al. 2014) by comparing all possible combinations of land preservations. The optimal solution serves as a benchmark and theoretical upper bound for all other algorithms we discuss subsequently. Under budgetary constraints, we order all solutions according to optimality, and select the highest ranked one that fits within the budget.

Recall that although this brute-force approach provides the optimal solution, it is not scalable to larger graphs due to the 2^N possible combinations.

Naive flat rate In the absence of any optimization for PES, and without knowledge of landholders' individual costs, a decision-maker might choose to offer the same flat-rate payment to every landholder. Any flat rate cannot, however, lead to robust PES performances. A flat rate that is too low would miss out on beneficial land parcels, while a flat rate that is too high would waste financial resources by overpaying, also limiting the total number of land parcels that can be preserved.

In our algorithmic implementation, we set the flat rate to the mean of landholders' costs. Any individual landholder accepts a flat rate payment if it matches at least their cost c, and their land parcel is preserved. Note that a naive flat rate does not require any compute power.

This naive flat rate serves as a performance lower bound, which promising near-optimal algorithms should consistently beat.

Greedy node This greedy algorithm sorts nodes (i.e., land parcels) by their potential, which we define as the societal benefit w of the node plus the sum of all edges e adjacent to that node (i.e., complementarities). The algorithm then greedily adds nodes in decreasing order of potential until the budget is exhausted.

Such node-potential-based optimization works well, if the potential is in fact realized, i.e., if the neighboring nodes are added to realize the complementarities. On the other hand, this optimization might also result in dispersed individually selected nodes, particularly in low-budget scenarios. However, in many real-world landscape scenarios, nodes of high potential might indeed be neighboring, since connectivity among parcels of land follows a gradient from low to high due to physical constraints.

The greedy node algorithm has a time complexity of $O(N \log N)$, where N is the number of nodes.

Greedy maximum spanning tree The greedy maximum spanning tree prioritizes connectivity between nodes (i.e., contiguous stretches of land as would be useful for wildlife habitats). We make four modifications to Prim's algorithm for minimum spanning trees. First, edge weights are negated to transform from a minimum to a maximum spanning tree. Second, negated node values are added to the edge weight for consideration of candidate edges, i.e., a good edge to a good next node includes both, good complementarity and a valuable land parcel. Third, the final value of the maximum spanning tree is not its length, but the sum of all node values plus the sum of the values of edges between those nodes, including edges that are additional to the spanning tree. Fourth, our algorithm re-initializes spanning trees in non-tree nodes as long as budget is available, rather than terminating once the first tree is built.

The greedy maximum spanning tree has a time complexity of $O(E \log N)$, with E edges and N nodes.

Connected components A connected component of a graph is a subgraph in which any two vertices are connected by a path. For example, a fully connected graph has one connected component, and a graph with V vertices and zero edges has V connected components.

Algorithm 1: GRAPES: GRAph-based PES

Inputs: budget B, parcel costs $\mathbf{c}$, benefits $\mathbf{w}$, complementarities $\mathbf{e}$, and adjacency list *adj*
initialize $dist(i) = \infty$ for all i
initialize $pred(i) =$ `None` for all i
$s = \arg\min_i c_i / w_i$ // start node
$dist(s) = c_s / w_s$
for $k = 1, \ldots, N-1$ **do**
  for $i = 1, \ldots, N$ **do**
    for j *in* $adj(i)$ **do**
      $ratio = c_j / (w_j + e_{ij})$
      if $dist(i) + ratio < dist(j)$ **then**
        $dist(j) = dist(i) + ratio$
        $pred(j) = i$
let *selected* be an empty set
let *sortedDist* be the keys i of *dist*, sorted in ascending order of *dist(i)*
$totalCost = 0$
for i *in sortedDist* **do**
  if $cost + c_i \leq B$ **then**
    add i to *selected*
    $totalCost = totalCost + c_i$
return *selected*

Discovering connected components in a graph helps us maximize the benefit of complementarities. When we do not include the entirety of a connected component, there exists some edge (u, v) such that we pay the cost c_u (w.l.o.g.) but not the cost c_v, and therefore get the benefit of w_u but not of the edge (u, v). We have already paid for a portion of the cost of the edge, therefore paying the additional cost of c_v gives us the benefit not just of w_v but also of e_{uv}.

We compute the connected components of a graph by running breadth-first search through the graph, noting a new connected component whenever our queue is empty but there are still unvisited nodes in the tree. From there, we rank the connected components by highest social benefit per currency spent. Iterating through the components in ranked order, we then pay for each fully connected component we have budget for. When we no longer have sufficient budget for the entire component, we then compute the maximum spanning tree on the component, picking the tree that maximizes our societal benefit.

The time complexity of DFS is $O(N + E)$, and the maximum spanning tree subroutine has a worst-case complexity of $O(E \log N)$, yielding $O(N + E \log N)$ total for max connected components.

4.2 Near-optimal PES scheme

We provide our near-optimal algorithm, GRAph-based PES (GRAPES), shown in Algorithm 1. We implement the Bellman–Ford algorithm to compute the shortest path from a

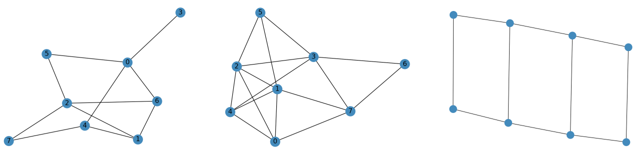

Figure 3: Examples of random graphs. (Left) $G_{N,p}$ graph with $N = 8$, $p = 0.3$. (Center) $G_{N,m}$ graph with $N = 8$, $m = 24$. (Right) Grid graph with $N = 8$.

single source. The problem of finding a longest-weight simple path in a graph is NP-hard. Thus, we cannot directly address the question of computing the path with the best value. Instead, we seek to minimize the cost–benefit ratio c_i/w_i. That is, we select the parcels that offer the most social benefit for every dollar we spend, therefore finding the most efficient way to allocate our spending.

We initialize Bellman–Ford with the node that has the lowest cost–benefit ratio, and then build the paths from there, initializing all distances except the source to 0 and then progressing through, iteratively selecting the best subsequent node to add. Suppose we are evaluating whether to relax the edge (u, v) given that node u is already in our set, and we are considering the addition of v. The cost–benefit ratio would be $c_v/(w_v + e_{uv})$. That is to say, when we add v, we get the complementarity benefit of e_{uv} while only having to pay the cost of v, given that u has already been accounted for.

The time complexity of Bellman–Ford is worst-case $O(NE)$.

5 Experiments

5.1 Graph simulation

We generate three types of random graphs. First is an Erdős–Rényi graph $G_{N,p}$ that produces a graph with N nodes with p probability of edge creation between any two nodes. Second is $G_{N,m}$ that produces a graph with N nodes and m total edges, placed randomly. Third is a grid-structured planar graph, with the nodes arranged in a grid and each node with degree up to 4. Note that this representation is flexible to various shapes of parcels. Examples of each graph type are shown in Figure 3.

5.2 Results

We show results of each algorithm in Figure 4. The y-axis measures the value achieved by each algorithm as a percentage of the optimal value. The brute-force optimal approach, clearly, is at 100%. Bellman–Ford achieves the otherwise best performance, with a mean of 89.0% achieved of the optimal value and standard deviation of 22.7 across 30 trials on a $G_{N,p}$ graph. It is followed by greedy node selection, with a mean of 83.0% and stdev 19.5. Naive flat rate lags behind with a mean of 52.0 and stdev 14.58. The large size of the error bars indicates volatility in the solutions, suggesting that the worst-case performance may be considerably lower.

Figure 4 captures algorithm performance for a single fixed budget of $B = 30$. Figure 5 shows the performances with

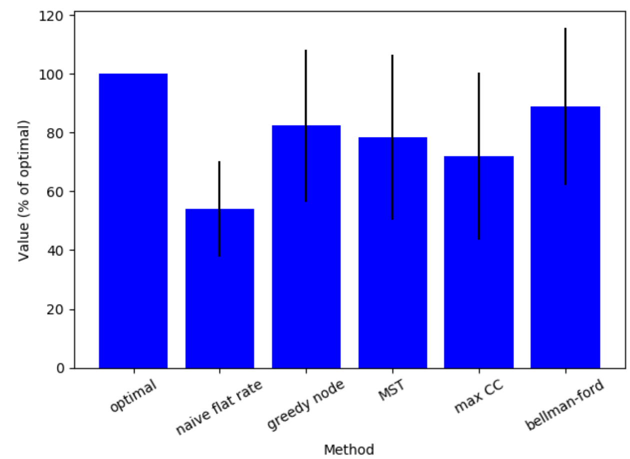

Figure 4: Experimental results, assessed as percentage achieved of optimal performance averaged over 30 trials. With graph structure $G_{N,p}$ with $N = 15$ nodes, probability $p = 0.3$, and budget $B = 30$.

varying budget. All algorithms approach optimal as the budget increases, with the exception of naive flat rate. This behavior is expected, because with a sufficiently high budget the optimal solution is trivial, which is to simply purchase every parcel of land. Observe that naive flat rate does not achieve this optimal solution because it selects a fixed payment to offer to each landowner and pays as many landowners as the budget permits. However, the fixed payment may be less than the cost to some landowners, so the landowners will never accept.

Bellman–Ford achieves the closest-to-optimal solution, most notably in low-budget scenarios. Greedy node lags during lower budgets, but catches up to Bellman–Ford with higher budgets. The max connected component approach, which relies on the maximal spanning tree algorithm as a subroutine, reveals a performance that nearly matches that of the maximal spanning tree. Although naive flat rate beats three out of four of the more sophisticated algorithms with a low budget, its performance is stagnant as budget rises.

Figure 6 shows the optimal value attainable at each budget. In this case, full payment to all landowners is achieved with a budget of around $B = 70$, with the increase in potential value rising nearly linearly before plateauing at the maximal value.

5.3 Discussion

The strong performance of Bellman–Ford can be attributed to two key insights. First, the cost–benefit ratio must be considered, not simply the value of a node. Consider the following simple example. We have three nodes, A, B, and C with benefit 90, 60, 70 and cost 10, 1, 2. If we were to greedily maximize value, we would select A to attain a value of $90 - 10 = 80$ for a cost 10. However, if we instead minimize the cost–benefit ratio, we would prioritize B and C with a ratio $1/60$ and $1/35$, which are lower than A with ratio $1/9$. Thus, we attain a higher total value of 127 for a lower

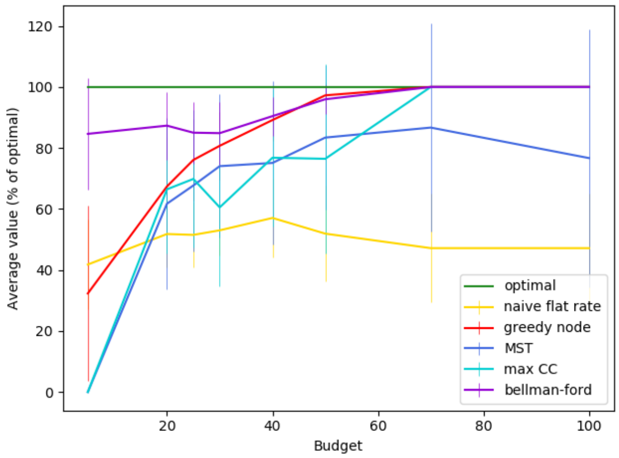

Figure 5: Performance when varying the budget, assessed as value of the mechanism averaged over 30 trials. With $N = 15$ nodes and $p = 0.3$ probability of adding a new edge.

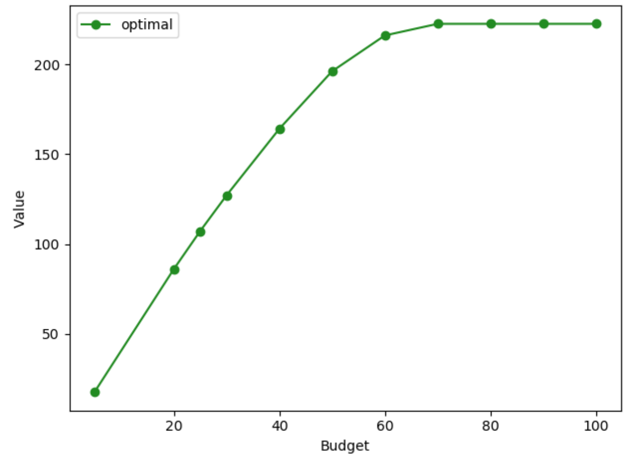

Figure 6: Value of optimal solution as the budget changes. If the budget surpasses the total cost to preserve all land parcels, the total value cannot improve anymore.

cost of only 3. Second, connectivity must be prioritized, especially in low-budget scenarios to capitalize on the benefits of complementarity. Bellman–Ford enforces connectivity, which is particularly valuable under very constrained budgets. Other algorithms such as greedy node selection will consider the added value of complementarity when greedily selecting nodes, but do not require the paired node to be selected as well, thus potentially losing out on benefits from complementarity.

6 Conclusion

We introduced a novel graph-based framework to optimize payments for ecosystem services. This framework opens the door to investigating the suitability of a range of well-known graph-traversal algorithms for optimization of PES. While the optimal solution under budgets smaller than the cost to preserve every land parcel can be ensured only by checking every single combination of land parcels, we were already able to show in preliminary experiments that standard graph-traversal algorithms—with key problem-driven modifications—can achieve near-optimal payment schemes. Importantly, our algorithms scale well with the number of land parcels and allow for efficient re-computation of PES schemes in dynamic real-world scenarios, where preservation goals, budgets, landholders, and cost–benefit evaluations may frequently change.

In a next step, the presented graph algorithms will be analyzed more thoroughly, both in theory and simulation, with the goal to characterize their worst-case, best-case, and average-case performances. Given that, they can become a reliable and predictable tool for PES. Our current results already indicate two key ingredients for robust algorithms: (1) taking into account both costs and benefits of nodes, e.g., by using a cost–benefit ratio; and (2) enforcing connectivity to capitalize on complementarity, especially for low budgets (see single-source shortest path algorithm).

We wish to emphasize two subtle differences between our current framework and previous optimal PES (Polasky et al. 2014), which can be addressed in future work. First, previous work compensated landholders for their marginal contributions to preservation value, including complementarity effects, even if that compensation exceeded their preservation costs at large (Polasky et al. 2014). This was necessary due to the nature of their strategy-proof auction mechanism design, but hurts the available budget by potentially spending more money on land parcels than required. We would be interested in investigating alternative mechanism designs to reveal landholder costs, for which at least complementarity effects do not have to be compensated. In this work, we compensated landholders based on their true costs for simplicity.

Second, our current framework is based on undirected graphs. If node A and B are preserved together, a single complementarity of value c may occur if A and B are connected by an edge. An alternative framework might consider directed graphs, offering a finer breakdown of complementarities. For instance, node A might get a complementarity of x if node B is preserved, and node B might get a complementarity of y if node A is preserved as well. In fact, this alternative directed graph more closely resembles the tabular data structure from previous work (Polasky et al. 2014). However, which scenario is more relevant to complementarities in the real world needs to be further assessed.

Overall, we find that a graph-based approach to model PES offers substantial promise, particularly due to its computational tractability and ability to flexibly account for complementarities. We hope this work lays a foundation to design future algorithms to maximize the preservation of our scarce natural resources under budget constraints.

References

[Alix-Garcia et al. 2018] Alix-Garcia, J. M.; Sims, K. R.; Orozco-Olvera, V. H.; Costica, L. E.; Medina, J. D. F.; and Monroy, S. R. 2018. Payments for environmental services supported social capital while increasing land management. *Proceedings of the National Academy of Sciences* 115(27):7016–7021.

[Balmford et al. 2003] Balmford, A.; Gaston, K. J.; Blyth, S.; James, A.; and Kapos, V. 2003. Global variation in terrestrial conservation costs, conservation benefits, and unmet conservation needs. *Proceedings of the National Academy of Sciences* 100(3):1046–1050.

[Cerbu, Swallow, and Thompson 2011] Cerbu, G. A.; Swallow, B. M.; and Thompson, D. Y. 2011. Locating redd: A global survey and analysis of redd readiness and demonstration activities. *Environmental Science & Policy* 14(2):168–180.

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

[Costanza et al. 1997] Costanza, R.; d'Arge, R.; De Groot, R.; Farber, S.; Grasso, M.; Hannon, B.; Limburg, K.; Naeem, S.; O'neill, R. V.; Paruelo, J.; et al. 1997. The value of the world's ecosystem services and natural capital. *Nature* 387(6630):253–260.

[Crooks and Sanjayan 2006] Crooks, K. R., and Sanjayan, M. 2006. *Connectivity conservation*, volume 14. Cambridge University Press.

[Dilkina et al. 2017] Dilkina, B.; Houtman, R.; Gomes, C. P.; Montgomery, C. A.; McKelvey, K. S.; Kendall, K.; Graves, T. A.; Bernstein, R.; and Schwartz, M. K. 2017. Trade-offs and efficiencies in optimal budget-constrained multispecies corridor networks. *Conservation Biology* 31(1):192–202.

[Hardin 1968] Hardin, G. 1968. The tragedy of the commons. *science* 162(3859):1243–1248.

[Jayachandran et al. 2017] Jayachandran, S.; De Laat, J.; Lambin, E. F.; Stanton, C. Y.; Audy, R.; and Thomas, N. E. 2017. Cash for carbon: A randomized trial of payments for ecosystem services to reduce deforestation. *Science* 357(6348):267–273.

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

[Ostrom 1990] Ostrom, E. 1990. *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.

[Polasky et al. 2014] Polasky, S.; Lewis, D. J.; Plantinga, A. J.; and Nelson, E. 2014. Implementing the optimal provision of ecosystem services. *Proceedings of the National Academy of Sciences* 111(17):6248–6253.

[Salzman et al. 2018] Salzman, J.; Bennett, G.; Carroll, N.; Goldstein, A.; and Jenkins, M. 2018. The global status and trends of payments for ecosystem services. *Nature Sustainability* 1(3):136–144.

[Stein et al. 2009] Stein, S. M.; McRoberts, R. E.; Mahal, L. G.; Carr, M. A.; Alig, R. J.; Comas, S. J.; Theobald, D. M.; and Cundiff, A. 2009. Private forests, public benefits. *USDA*.

[Teytelboym 2019] Teytelboym, A. 2019. Natural capital market design. *Oxford Review of Economic Policy* 35(1):138–161.

[Wallbott, Siciliano, and Lederer 2019] Wallbott, L.; Siciliano, G.; and Lederer, M. 2019. Beyond pes and redd+: Costa rica on the way to climate-smart landscape management? *Ecology & Society* 24(1):24.