Sustainable Decision Making for Chemical Process Systems via Dimensionality Reduction of Many Objective Problems

Justin M. Russell, Andrew Allman*

Department of Chemical Engineering, University of Michigan, Ann Arbor, MI 48109, USA

*Corresponding Author: allmanaa@umich.edu

Abstract

Recent global events and the rise of sustainable investing has made clear that the chemical and energy industry must consider sustainability goals beyond profit maximization to remain competitive. Multi-objective optimization provides an ideal framework for analyzing sustainability tradeoffs, but when four or more objectives are considered, the ability to rigorously solve problems and interpret results is lost. This necessitates an approach to systematically reduce the dimensionality of many objective problems to three or fewer objectives. In this work, an algorithm to group objectives based on their correlating nature *a priori* to solving the full space problem is proposed. It utilizes community detection on a novel weighted objective correlation graph to identify two or three groups of correlated objectives. Results from three representative case studies demonstrate that objective groupings obtained from this algorithm minimize the amount of tradeoff information lost and outperform intuitive groupings by economics or the environment.

Introduction

In recent times, the idea of sustainable investing has come into the mainstream[1]. In this scheme, shareholders and board members judge industrial performance not just by the traditional economic bottom line of profits, but also emphasizing environmental, social, and governance (ESG) issues. Indeed, looking at the events of the last 5 years illustrates the pitfalls of making decisions solely based on economics, as a pandemic has exposed vulnera-

bilities in global supply chains[2], social unrest and lawsuits have occurred due to inherently inequitable decisions that disproportionately harm minority communities[3], and the effects of climate change, largely driven by the use of economically preferred fossil fuels, are becoming more apparent[4]. The realities of today make it clear that a modern chemical industry cannot remain competitive by solely maximizing profits as is traditionally done, and that decisions made at all levels in the chemical enterprise, including process design, strategic planning, and real time operation and control, must weigh the tradeoffs of a large number of different objectives within the scope of both economics and ESG.

One approach to sustainable decision making is to attempt to monetize ESG outcomes to arrive at a single, profit-based objective[5;6]. This approach is commonly applied for carbon emissions via either a carbon tax[7;8;9] or through a social cost of carbon[10]. A few approaches within the life cycle assessment framework attempt to be more systematic about how monetization is performed, usually by attempting to define a marginal or opportunity cost for each sustainability objective[11]. An analogous approach to monetization is the use of multi-attribute decision making methods[12], which attempt to quantify the objective preferences of a decision maker by generating weights that correspond to these preferences. These approaches are beneficial in that the optimization algorithm returns a single solution that can be directly implemented. However, a key limitation of any monetization or multi-attribute approach is that the solution obtained is typically highly sensitive to how different objectives are weighted: for example, considering social costs of carbon which may vary over orders of magnitude depending on the source[13], or using different LCA-based approaches, can result in very different decisions being made. As such, the question of what makes a solution sustainable can vary wildly between decision makers who may hold different values, including between different organizations as well as individuals at the same organization.

An alternative approach to monetization is to solve a multi-objective optimization problem. Here, the solution is not a single decision but is a manifold of possible decisions representing tradeoffs between objectives (the Pareto frontier), where each point along the Pareto

frontier represents the best one can do for one objective without making any of the others worse. While this approach does not give a single implementable solution, obtaining and visualizing the full Pareto frontier of objective tradeoffs enables making fully informed sustainable decisions, the responsibility of which ultimately lies with human actors and can be supported by preference ranking algorithms[14]. Multi-objective approaches have been applied quite broadly in chemical process systems research[15]. With respect to sustainability, it is most common to analyze the tradeoffs between a single environmental and a single economic objective, with example applications in plant design[16;17;18], supply chain management[19;20], and process operations[21;22], to name a few.

A limitation of multi-objective approaches for sustainable decision making is that for problems of four or more objectives (many-objective problems, or MaOPs), visualization of the objective tradeoffs becomes unintuitive and rigorously generating a complete set of solution points becomes computationally prohibitive. A common approach to addressing this challenge is to lump different sustainability goals into intuitive groupings, such as economic, environmental, and social groupings based on the three pillars of sustainability[23]. Examples of this include work by Santibanez-Aguilar et al.[24], who combine metrics for damages to human health, the ecosystem, and resource extraction using a life cycle assessment derived tool in the planning and site selection of biorefineries. Garcia et al.[25] propose combining the effects of agricultural wastes, land use, and ecosystem services into a single environmental green GDP objective when designing food-energy-water-waste nexus systems. Wheeler et al.[26] propose an environmental objective that takes into account the potential for a process or supply chain to cause key outputs, such as ozone depletion or ocean acidification, to exceed "planetary boundaries" or upper bounds beyond which deleterious impacts on the planet are observed. Mota et al.[27] propose a social objective that combines the goals of creating jobs and limiting inequity by weighting job creation activity in lower GDP regions, which they used in analyzing possible locations for expansion of a electronics manufacturer. Alternatively, one can also try to aggregate all outcomes into a single objective: a popular

3

method for doing this is the eco-efficiency concept which normalizes the various economic and environmental objectives and assigns weights based on social relevance and a process' impact on a specific outcome relative to known global parameters [28].

As an example that demonstrates the shortcomings of the aforementioned approach, consider the design of an ammonia production system in a water scarce region where one could choose either to produce requisite hydrogen from fossil fuels (resulting in inherent carbon emissions) or from electrolysis of water (using large quantities of water). For such a system, there is a clear tradeoff between the two objectives of carbon emissions and water usage, but information about this tradeoff would be lost if both quantities were aggregated into a single environmental objective. This example demonstrates the importance of choosing objective groupings more systematically based on their correlating (i.e. both objectives point to similar solutions) vs. competing (i.e. a large tradeoff exists between objectives) nature to preserve tradeoff information. While there exist methods to achieve this which utilize principal component analysis [29], aggregation trees [30] or dominance preservation strategies [31], these methods require the generation of at least part of the high-dimensional Pareto frontier for the original MaOP. As such, they may be susceptible to bias based on which solution points are generated and are not particularly helpful in reducing the computational burden of solving the problem.

In this work, we assert that systematic objective dimensionality reduction for (mixed-integer) linear MaOPs can be performed *a priori* to obtaining any part of the solution of the MaOP on the basis of problem structure. We propose a graph structure to represent variable-constraint-objective connectivity, which we use to develop an objective correlation graph with edge weights corresponding to the competing vs. correlating nature of the two objectives. From the objective correlation graph, we apply a community detection approach to identify 2-3 groups of objectives, such that objectives in the same group are correlated and those in different groups are competing. We also present an information loss metric and demonstrate that our approach is able to choose groupings that preserve as much information

4

about objective tradeoffs as possible. The remainder of this paper is structured as follows: section* 2 will provide background on solving multi-objective optimization problems, as well as on identifying and exploiting optimization problem structure using graph theory. section* 3 will provide the details of the proposed algorithm and how it utilizes the problem structure of MaOPs to determine the strength of links between objective functions. section* 4 will examine three case studies adapted from the sustainable process systems literature and utilize the proposed algorithm to analyze the structure of MaOPs adapted from the original formulations. Finally, section* 5 will include some concluding remarks and discussion of areas for future work.

Background

Multi-Objective Optimization

A general multi-objective optimization problem can be written as follows:

$$\min_{x} \quad \{f_1(x), ..., f_I(x)\}$$
$$\text{s.t.} \quad x \in \mathcal{X}, \tag{1}$$

Where f_i represents the function for objective i, x are the decision variables, and $\mathcal{X}$ is the set of all values x that are feasible to the problem. Typically, it is impossible to optimize each objective f_i simultaneously, such that no single solution exists for problem (1). Instead, the solution to this problem is a set of points, known as the Pareto frontier, where each point represents the best one can do in one objective without sacrificing another. A critical concept for understanding the Pareto frontier is the idea of dominance. Consider two feasible points to problem (1), $\hat{x}$ and $\bar{x}$. Solution $\hat{x}$ is said to dominate solution $\bar{x}$ if the following

two relationships hold:

$$f_i(\hat{x}) \leq f_i(\bar{x}) \quad \forall\, i \in \{1, ..., I\} \tag{2}$$

$$\exists\, i \in \{1, ..., I\} : f_i(\hat{x}) < f_i(\bar{x}) \tag{3}$$

In words, this means that solution $\hat{x}$ performs no worse than $\bar{x}$ in all objectives, and strictly better in at least one objective. Using this concept, the Pareto frontier can be defined as the set of all feasible solutions that are not dominated by any other feasible solution.

Solution approaches to multi-objective problems seek to find the Pareto frontier by identifying non-dominated solutions. The most common approaches can be broken into two categories. The first are scalarization approaches, which generate Pareto optimal points through solving a set of single objective optimizaiton problems. Examples of methods in this classification include weighted sum[32], epsilon constraint[33] and Chebyshev scalarization approaches[34]. While rigorous deterministic global optimization can be used to solve each single objective problem and guarantee that each point found is Pareto optimal, the number of single objective problems to solve scales exponentially with number of objectives, such that these approaches are impractical to implement for MaOPs. The second class of approaches are evolutionary algorithms, which use biological principles of natural selection such as mutation and recombination to drive a population of feasible points towards optimal solutions[35]. While these methods tend to work reasonably well in practice and are more scalable than scalarization approaches, they are ultimately stochastic and heuristic approaches which do not provide any guarantees of solution quality. As such, an important goal of this work is to systematically reduce the dimensionality of MaOPs to 3 objectives or fewer, in order to apply a rigorous scalarization method for the determination of the Pareto frontier.

Graph Representation of Optimization Problems

For complex optimization problems that cannot be readily solved by off-the-shelf solvers in relevant amounts of time, such as the MaOPs of interest in this work, it is often useful to identify and exploit problem structure to derive an approach that makes solving the problem easier. A natural way to achieve this is by representing the optimization problem as a graph, or a set of nodes and edges that capture the connectivity of different objects (i.e. variables and constraints) within the optimization problem[36]. Once such a graph is developed, an effective approach for structure identification is community detection, which identifies subgroups within a graph on the basis of maximizing a quantity called modularity, effectively generating subgroups such that nodes within the same subgroup interact strongly, while minimal interaction occurs between nodes in different subgroups[37]. While modularity maximization is a known NP-hard problem, several well-known greedy algorithms give good heuristic solutions such as spectral partitioning[38], the Louvain algorithm (or fast unfolding)[39], and the Leiden algorithm[40], the latter of which is used in this work.

Community detection has been shown to be a powerful tool for identifying structure in optimization problems amenable to decomposition. Early work in this area looked at identifying distributed optimization structures for augmented Lagrangian solution approaches[41], as well as structure within a model predictive control problem for obtaining distributed controller architectures[42]. This approach was later generalized to identify communities that correspond to optimization subproblems with minimal complicating variables or constraints, and thus amenable to various decomposition approaches[43]. This approach was extended using a stochastic block modeling approach to identify both community and core-periphery structure in optimization problems that can be exploited using various decomposition solution approaches[44]. Other recent work in this area developed a new overlapping Schwarz type decomposition rooted in a problem's graph structure and the exponential decay of sensitivity propagation through graphs[45]. Beyond identifying structures for decomposition, graph-theoretic methods can also be used to identify symmetry in optimization problems[46],

which can degrade performance of global nonconvex solvers. In this work, we seek to build a graph structure relating multiple objectives with weights corresponding to their correlating vs. competing nature. From this, a community detection approach can be applied to determine subgroups of objectives such that objectives in the same subgroup are correlated, while those in different subgroups are competing.

Proposed Algorithm

This section* provides the framework and details of the mathematical algorithm to reduce objective space dimensionality in many objective optimization problems. The algorithm assumes an optimization problem with I objectives, V variables, M inequalities, and N equality constraints. In particular, we consider the linear many objective optimization problem formulated as follows:

$$
\begin{aligned}
\min_{x} \quad & \{c_1^T x, ..., c_I^T x\} \\
\text{s.t.} \quad & Ax \le b \\
& Dx = e,
\end{aligned}
\tag{4}
$$

Here, c_i are known cost vectors of size V for the different objectives, A is a $M \times V$ matrix of known constraint coefficients, b is a known M dimensional vector, D is a $N \times V$ matrix of known coefficients, and e is a known N dimensional vector. The decision variables x can be either continuous or integer. Without loss of generality, objective functions are formulated as all minimization problems. The only assumption made about the problem is that the inequality constraint matrix A fully bounds the problem, such that no decision variables or objectives can feasibly diverge to $\pm\infty$; we note that for most problems of practical interest, this assumption is not restrictive. The algorithm takes the three coefficient matrices, $C = [c_1|...|c_I]^T$, A, and D, as input and identifies groups of objectives with expected correlating
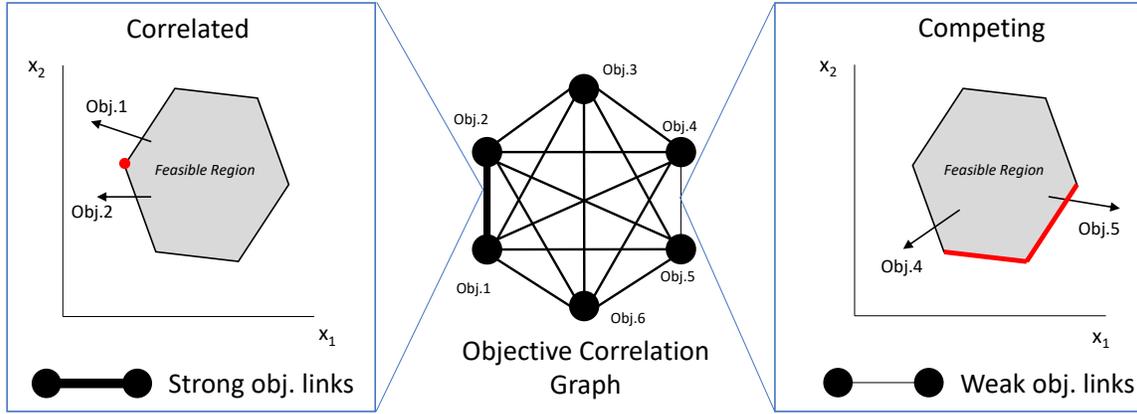
Figure 1: General structure of the weighted objective correlation graph. Strong weights (near one) correspond to objective pairs which are correlated such that minimizing one inherently minimizes the other (left), while weak weights (near zero) correspond to objective pairs that have inherent conflict (right).

behaviors with respect to the optimization/decision variables. To do this, we propose a weighted objective correlation graph, whereby nodes correspond to different objectives in the MaOP and edges connecting the different objectives are weighted between 0 and 1, corresponding to strongly conflicting to strongly correlating objectives. The general weighted objective correlation graph structure and proposed weighting scheme is depicted in Figure 1. The challenge of constructing this graph is intelligently and systematically determining the edge weights, as once these are obtained, community detection can be used to identify groups of objectives such that objectives in the same group are strongly correlated, while those in different groups conflict with each other.

Edge Weight Determination

To determine the edge weights of the objective correlation graph, we recall that for a linear programs like the MaOP (4) we are considering in this work, the optimal solution is guaranteed to lie on the boundary of the feasible region. Thus, to determine if two objectives are likely to be competing or correlating, it is useful to consider the projections of their cost vectors onto active constraint hyperplanes of the optimization problem. If these projections
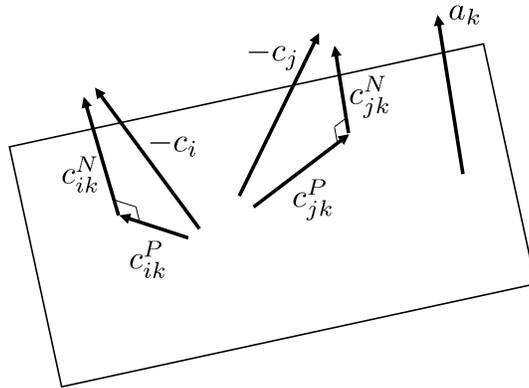
9

Figure 2: Graphical representation of vectors c_i, c_j, a_k, c_{ik}^P, c_{ik}^N, c_{jk}^P, and c_{jk}^N. The rectangle represents the constraint surface $a_k x = b_k$.

point in (nearly) the same direction along all surfaces, it is likely the objectives will be correlated, while if the projections point in different directions on at least some surfaces, conflict and tradeoffs are likely to occur. We formalize the approach mathematically as follows: consider two objective cost vectors c_i and c_j interacting along inequality constraint with normal vector a_k (a row of the matrix A). For each cost vector, we obtain the vector components normal (c_{ik}^N) and projected (c_{ik}^P) onto the constraint surface, using the following equations:

$$c_{ik}^N = \frac{-c_i^T a_k}{\|a_k\|_2^2} a_k \tag{5}$$

$$c_{ik}^P = -c_i - c_{ik}^N \tag{6}$$

A graphical representation of these various vectors is displayed in Figure ??. For normal and projected components, the vectors are normalized to $\hat{c}_{ik}^N$ and $\hat{c}_{ik}^P$, respectively, such that they point in the same direction as the original vectors but are of length 1. In the case where there is no normal or projected component, the normalized vector is left as the zero vector. The normal components are used to determine if the constraint is likely to be active in determining a tradeoff between objectives i and k. If both normal components $\hat{c}_{ik}^N$ and $\hat{c}_{jk}^N$ point inwards from the constraint surface, then both cost vectors pull away from the

10

constraint, rendering the constraint inactive for the objective pair and allowing us to neglect interaction along this surface by giving it a weight W_{ijk} of zero. Otherwise, the normalized projected component vectors are used to determine a correlation strength of two objectives along the constraint hyperplane, S_{ijk}:

$$S_{ijk} = (\hat{c}_{ik}^P)^T \hat{c}_{jk}^P \tag{7}$$

Since conflict along any constraint surface can cause a tradeoff between objectives, while correlation requires overlap on all constraint surfaces, it makes sense to weight findings of conflict more heavily when combining the interactions along different constraints. As such, we weight each interaction (W_{ijk}) using a logistic function that provides high weights when conflicts are found, and lower weights when correlation is found:

$$W_{i,j,k} = 1 - \alpha \left(\frac{1}{1 + \exp(-\beta \times S_{i,j,k})} \right) \tag{8}$$

In this equation, α and β represent hyperparameters to the algorithm. The hyperparameter α should vary between 0 and 1 and represents a maximum "discount rate" for correlated objective-constraint-objective triplets to ensure that many correlated constraints do not overwhelm a smaller number of more informative competing constraints. The hyperparameter β should be positive, and governs the smoothness of the logistic curve, with larger values making the curve more step-like at $S_{ijk} = 0$. Empirically, we have determined that values of $\alpha = 0.9$ and $\beta = 100$ tend to work well in practice.

Equality constraints will always be active, so the component of the cost vector normal to the constraint surface is unimportant. Equations (5)-(6) are used with d_k (a row of the matrix D) in place of a_k to determine the component of the cost vector projected onto the constraint surface. Strengths along equality constraint are then found again using (7). Since the Pareto solution will always lie along the equality constraint surface, weights W_{ijk} are set to 1 for all equality constraints.

11

To determine the total correlation strengths, which we denote S_{ij}^A, we calculate a weighted average of strengths along all constraints determined to be possibly active, and rescale values such that they are between 0 and 1:

$$S_{ij}^A = 0.5 \left(1 + \frac{\sum_{j=1}^{M+N} W_{i,j,k} S_{i,j,k}}{\sum_{j=1}^{M+N} W_{i,j,k}} \right) \tag{9}$$

Note that the matrix of S_{ij}^A values is the adjacency matrix of the objective correlation graph, that this matrix will always be symmetric (i.e. $S_{ij}^A = S_{ji}^A$), and that by convention, diagonal elements of this adjacency matrix are always set to zero (i.e. $S_{ii}^A = 0$). Values of this adjacency matrix near zero correspond to an objective pair with a large expected amount of conflict, while values close to one imply that the two objectives are expected to be correlating.

Note that for large scale problems, it can be inefficient to compute the effect of every constraint on every objective pair. In such cases, we provide in this algorithm the option of being more systematic about which values of S_{ijk} are obtained by using a variable-constraint-objective graph. This graph is a tripartite graph where nodes correspond to variables, constraints, or objectives in the original optimization formulation, and edges exist between variable and constraint/objective nodes if the variable appears within the constraint/objective. From this, we can identify primary linking constraints, which contain one or more variable shared by two objectives, and secondary linking constraints which contain at least one variable unique to both of the two objectives considered. An example of the variable-constraint-objective graph of a simple problem is shown in Figure 3, displaying examples of both primary and secondary linking constraints. As primary and secondary linking constraints comprise the shortest paths from objective to constraint to objective in the graph, it expected that they capture the most important objective interactions. This argument aligns with recent findings that for large scale, structured optimization problems, the sensitivity of the optimal solutions at one node with respect to perturbations at another decays with respect to the distance between nodes[47], although this approach can be

$$\min \begin{cases} o_1 = 5x_1 + 3x_2, \\ o_2 = 2x_1 + 3x_2 + 5x_3 + 2x_4 \\ o_3 = 2x_3 + 3x_4 \end{cases}$$

$$\text{s.t.} \quad x_1 + x_3 \geq 1 \quad (f_1)$$
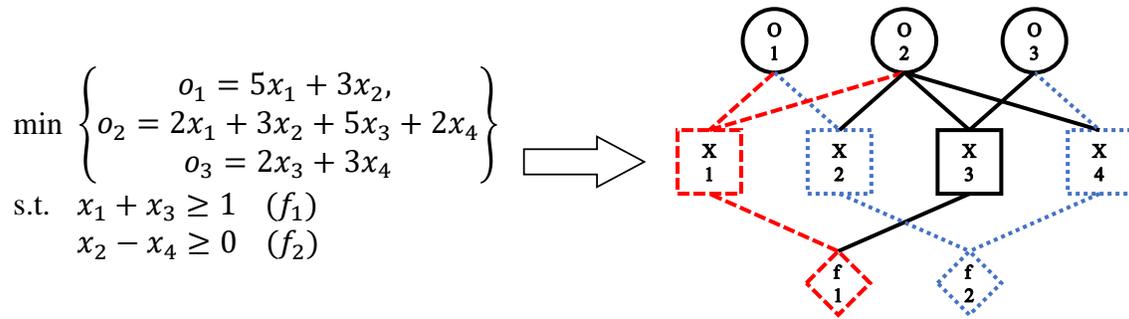$$x_2 - x_4 \geq 0 \quad (f_2)$$

Figure 3: Example variable-constraint-objective graph. In red, f_1 is a primary constraint link between objectives 1 and 2, as it contains x_1, which appears in both objectives. In blue, f_2 is a secondary constraint link between objectives 1 and 3, as it contains x_2 from objective 1 and x_4 from objective 3.

obfuscated by formulations with a large number of "auxiliary" variables or constraints.

Determination of Objective Groups

Now that a way to systematically determine edge weights in the objective correlation graph has been presented, groupings of correlated objectives are obtained by performing community detection using the Leiden algorithm[40]. This algorithm takes as input the adjacency matrix of the objective correlation graph, as well as a hyperparameter which governs the resolution of communities. In most cases, it is desirable to identify two or three communities of objectives, since two or three-objective optimization problems are not computationally prohibitive to solve and give interpretable Pareto frontiers. To achieve this, we begin by setting the resolution hyperparameter so high that each objective is in its own community, and then gradually decrease this value until the desired number of communities are achieved. Once the community structure is obtained, the dimensionality of the original MaOP is reduced by combining objectives within the same community into a single objective. There are a variety of ways that this can occur: two popular approaches are to simply add all of the objectives in the same group together, or to neglect all but one objective from the group[29]. Different approaches can be more beneficial depending on the application and how the Pareto frontier

will be used; however, the remainder of this paper will consider grouping by adding together objectives, and comparison of different grouping approaches will be considered beyond the scope of this work.

Information Loss Metric

Assuming grouped objectives are combined either by neglecting all but one objective, or using a weighted sum, Pareto frontiers obtained using the reduced-space formulation will be a subset of the full-space Pareto frontier. A good performing objective reduction will retain a larger proportion of full-space Pareto optimal solutions in its Pareto frontier, or equivalently, each point on the reduced space Pareto frontier will have embed only a small amount of tradeoff information from grouped variables. Here, we propose an information loss metric to quantify this performance. First, a solution set that is a representative sample of the full-dimensional Pareto frontier is required. Two sets of objectives are identified: the set $\mathcal{C}$ is all objectives that are being combined and the set $\mathcal{R}$ is all objectives that are kept as individuals. From the Pareto frontier solution set, unique solution values are recorded for the objectives in set $\mathcal{R}$, this gives a set of $\mathcal{P}$ unique points that are each a combination of values for objectives in $\mathcal{R}$. At each unique point p, one can often find multiple different Pareto optimal solutions that vary the values of objectives in $\mathcal{C}$ that are grouped togeter, giving another set of points $\mathcal{J}_p$. Matrix K_p is created that contains all Pareto optimal values for objectives in $\mathcal{C}$ at all points $j \in \mathcal{J}_p$ where the objective values for objectives in $\mathcal{R}$ are held constant. In each K_p the values are scaled using:;

$$\hat{K}_{p,i,j} = \frac{K_{p,i,j} - l_i}{u_i - l_i} \tag{10}$$

where $\hat{K}_{p,i,j}$ is the scaled value from objective i at point j, $K_{p,i,j}$ is the original value, l_i is the lower bound or minimum value found for objective i through the entire full space Pareto, and u_i is the upper bound or maximum value found for objective i through the entire full

14

space Pareto. Equation 10 gives us scaled values for each objective in C, such that they range between 0 and 1. With these scaled values, we can determine the total information lost at each point p. Total information loss for each point p is determined by:

$$B_p = \sum_{i \in \mathcal{C}} (\hat{K}_{p,i}^{\max} - \hat{K}_{p,i}^{\min}) \tag{11}$$

where B_p is the information lost at unique point p, $\hat{K}_{p,i}^{\max}$ is the maximum scaled value for objective i, and $\hat{K}_{p,i}^{\min}$ is the minimum scaled value for objective i. A B_p value of 0 tells us that no tradeoff exists between grouped objectives at point p, while a larger value of B_p, up to a maximum of $|\mathcal{C}|$ indicates that a wider range of tradeoffs between combined objectives is being neglected by combining them. . Finally, the average of all B_p values is taken to obtain the total average information loss. This is a single value that describes the information lost by grouping the chosen set of objectives together. Its utility is in comparing different choices of grouping objectives together to determine which group results in a lower average information loss and a more valuable and informative Pareto frontier.

Case Studies

The proposed algorithm will be demonstrated on three representative studies adapted from the sustainable process systems literature.. Each of the studies have been chosen to represent MaOP formulations which include varying numbers of variables and constraints as well as integer and binary variables. These cases will demonstrate the ability of this method to identify and group objectives which are most strongly correlated in a range of problem formulations. All optimization and calculating was completed using an Intel i9-10900 CPU with 64 GB of RAM using CPLEX 20.1[48], JuMP v0.21.10[49], and Julia v1.5[50].

UK Energy Mix

This case study is included as it is a relatively simple problem formulation that has only three objectives with a full three dimensional Pareto frontier presented in the original work, Limleamthong and Guillen-Gosalbez[51]. This work studies the energy sector in the UK, analyzing total energy mix including conventional sources and a variety of alternatives. Energy technologies including nuclear, wind, natural gas, coal, and biomass are studied. The objectives within the formulation of this study are cost, global warming potential, and worker injuries. Variables model the total electricity generated by each of the studied technologies. This leads to a 6-variable, 3-objective study with the only constraints being that the total generation is equivalent to the demand of the nation and bounds on the possible values for each variable. The full optimization problem formulation is:

$$
\min_{G_j} \quad \left\{ \sum_{j \in J}(c_j G_j), \sum_{j \in J}(w_j G_j), \sum_{j \in J}(i_j G_j) \right\}
$$
$$
\text{s.t.} \quad G_j^{min} \leq G_j \leq G_j^{max} \quad \forall \quad j \in \mathcal{J}
$$
$$
\sum_{j \in J} G_j = D \tag{12}
$$

where G_j is the amount of energy generated using technology j, c_j is the cost in £per kWh, w_j is the global warming potential in kg CO2 equivalent per kWh, i_j is worker injuries in number of injuries per kWh, G_j^{min} and G_j^{max} are the lower and upper bounds for energy generated with each technology, and D is the total energy demand which must be satisfied. Pareto frontiers presented in the paper and shown in Figure 4, with each objective rescaled to lie between 0 (the single objective minimum) and 1 (the worst case for each objective observed on the Pareto frontier), indicate that there is no strong agreement among the three objectives. However, our goal for this case study was to see if our algorithm could successfully identify the objective pair resulting in the least tradeoff information lost. After running this problem formulation through the proposed algorithm, the correlation strength
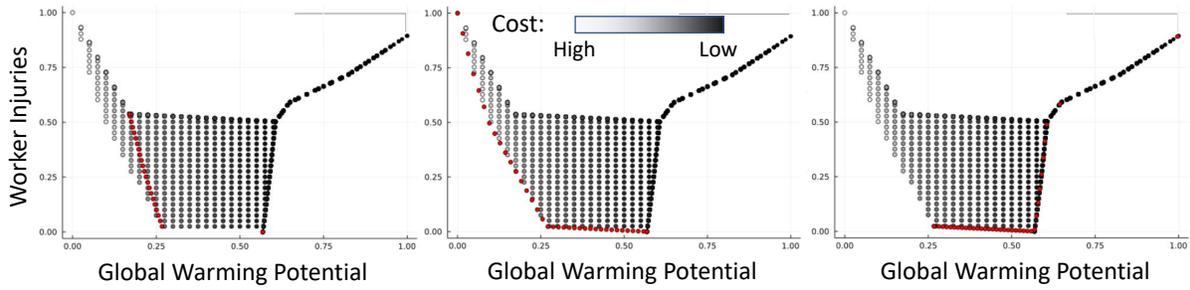
Figure 4: Three dimensional Pareto frontier for UK energy case study. Points highlighted in red indicate solutions retained in the two-dimensional Pareto frontier when (left) cost and global warming potential, (center) cost and worker injuries, or (right) global warming potential and worker injuries, are grouped together.

weights shown in Figure 5 were found between each pair of objectives.   Running the
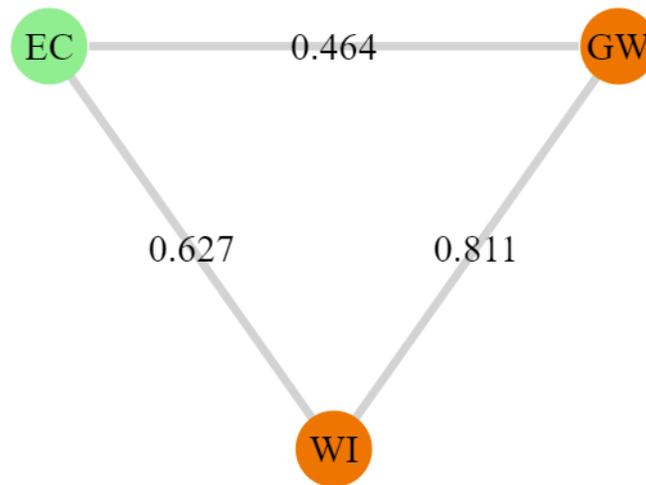

Figure 5: Objective correlation graph with correlation strengths as edge weights and identified communities in each color. Objectives are cost (EC), global warming potential (GW), and worker injuries (WI).

Leiden algorithm to detect the best grouping into two communities trivially groups the two objectives with the largest edge weight, giving one community with only the cost objective and one community with both the global warming potential and worker injury objectives. The complete algorithm of determining edge weights and objective groupings takes 0.36 seconds for this case study. Physically, the results suggest that there is some agreement between how much an energy technology emits and how many worker injuries it typically

incurs. Looking at the input data, we can confirm this finding: the lowest emitting technology (nuclear) also has the second fewest worker injuries, while the highest emitting technology (coal) also has the highest worker injuries. Similarly, the second strongest pair also makes sense to have high correlation, as the natural gas is the technology with both the lowest cost and worker injuries. However, the constraints limit the ability of natural gas to meet all of the demand, and the second lowest cost technology (coal) is the worst with worker injuries. Finally, it makes sense that cost and emissions are the most competing constraints, as the two lowest cost technologies are the highest emitting, and vice versa. The results of the algorithm are aligned with the calculated average information lost for each of the possible pairs from the full space Pareto frontier. If cost and global warming potential are grouped, the average information loss is 0.736. Combining cost and worker injury objective results in an average information loss of 0.213. The algorithm's identified group of global warming potential and worker injury objectives together results in an average information loss of 0.105. The results of our algorithm are also supported by Figure 4, where the red points are the parts of the full space Pareto frontier obtained for each of the three possible groupings of two objectives. It is evident that the selected grouping of global warming potential and worker injuries retains points that capture the largest range of values of the three objectives. Based on these results, the algorithm has successfully identified the best objective grouping if one were to be chosen. Furthermore, the identification of the worker injury-global warming potential grouping as best also matches results of using the principal component analysis method from Saxena et al.[29], which identified the emissions objective for removal, by, for example, grouping it with worker injuries.

Sustainable Ammonia Supply Chain in Minnesota

To demonstrate the application of the proposed algorithm to a MaOP with many variables and few objective functions, the optimization of the ammonia supply chain as presented in Palys et al.[9] will be used. Specifically, we adapt the formulation for modular production

18

units by considering the additional objectives of carbon emissions and water usage, while separating capital and operating costs into their own objectives. For a full listing of the notation for this problem, we refer the reader to the original work. We review the problem formulation below , with the operating cost objective function given by:

$$\text{cost} = \zeta + \sum_{r=1}^{R} \sigma x_r + \sum_{r=1}^{R}\sum_{f=1}^{F} \tau_{r,f} y_{r,f} + \sum_{p=1}^{P}\sum_{d=1}^{D} (\tau_{p,d} + \alpha_p) y_{p,d} + \sum_{d=1}^{D}\sum_{f=1}^{F} \tau_{d,f} y_{d,f} \tag{13}$$

where ζ is the capital cost of constructed renewable plants, σ is operating cost factor for renewable plants, x_r is installed renewable capacity at site r, $\tau_{i,j}$ is transportation cost and $y_{i,j}$ is amount transported from site i (renewable production, r, conventional production, p, distribution facility, d) to j (distribution facility, d, consumption site, f). The capital cost for modular, renewable-powered production is given by the following:

$$\zeta = \sum_{m=1}^{M} \rho_m \left(\sum_{n=1}^{N} n^{\gamma_m} z_{n,m} \right) \tag{14}$$

where ρ_m is the cost of one module size m, $z_{n,m}$ is the a binary variable which is one when n is modules of size m are built and zero otherwise, and γ_m is the mass production factor. An emission model is also included as an objective and is drawn from a study using a related optimization study[20]. Emissions of the supply chain are given by:

$$\text{emissions} = \sum_{r=1}^{R}\sum_{f=1}^{F} \epsilon_{r,f} y_{r,f} + \sum_{p=1}^{P}\sum_{d=1}^{D} (\epsilon_{p,d} + \eta_p) y_{p,d} + \sum_{d=1}^{D}\sum_{f=1}^{F} \epsilon_{d,f} y_{d,f} \tag{15}$$

where $\epsilon_{i,j}$ are emissions resulting from transporting ammonia from i to j, and η_p are the emissions from conventional production of ammonia at site p. Creating the bounds for this problem are the following constraints. Linearizing the capital cost objective yields the two constraints:

$$\sum_{n=1}^{N} z_{n,m} \leq 1 \ \forall m \tag{16}$$

$$\sum_{n=1}^{N} n z_{n,m} = \sum_{r=1}^{R} w_{m,r} \ \forall m \tag{17}$$

where $w_{m,r}$ is the number of modules size m at renewable production site r. The production at each site is defined from this variable:

$$x_r = \sum_{m=1}^{M} \pi_m w_{m,r} \ \forall r \tag{18}$$

where π_m is the amount of ammonia produced by a module of size m. A requirement that the supply chain meets the demand at each site:

$$\sum_{d=1}^{D} y_{d,f} + \sum_{r=1}^{R} y_{r,f} \geq \delta_f \ \forall f \tag{19}$$

where δ_f is the ammonia demand at site f. At each conventional production site, there is an upper bound on the production capacity:

$$\sum_{d=1}^{D} y_{p,d} \leq \xi_p \ \forall p \tag{20}$$

where ξ_p is the capacity at each production site. A mass balance is imposed on distribution facilities:

$$\sum_{p=1}^{P} y_{p,d} \geq \sum_{f=1}^{F} y_{d,f} \ \forall d \tag{21}$$

Additionally, there is an upper bound placed on the amount of renewable ammonia production at each site which arises from available wind-power:

$$x_r \leq \xi_r \ \forall r \tag{22}$$

20

where ξ_r is the maximum ammonia production at site r. An additional objective of water usage for each production site was modeled:

$$\text{water} = \sum_{p=1}^{P} \sum_{d=1}^{D} \omega_p y_{p,d} + \sum_{r=1}^{R} \sum_{m=1}^{M} \omega_r y_{r,m} \tag{23}$$

where ω_i is the water consumed in ammonia production at site i. A simplifying assumption was made in the modeling of this objective that the only significant difference in water consumption among production methods was the stoichiometric amount of water needed for hydrogen production with either an electrochemical technology for renewable production or a fossil fuel-based technology. Additionally, capital cost and operating costs were split into separate objectives. The resulting problem formulation is a MILP with four objectives: capital cost, operating cost, carbon emissions, and water consumption. Implementing the proposed algorithm on the linear reformulation of this case study yields the objective correlation graph shown in Figure 6: The algorithm identifies operating cost and carbon emissions
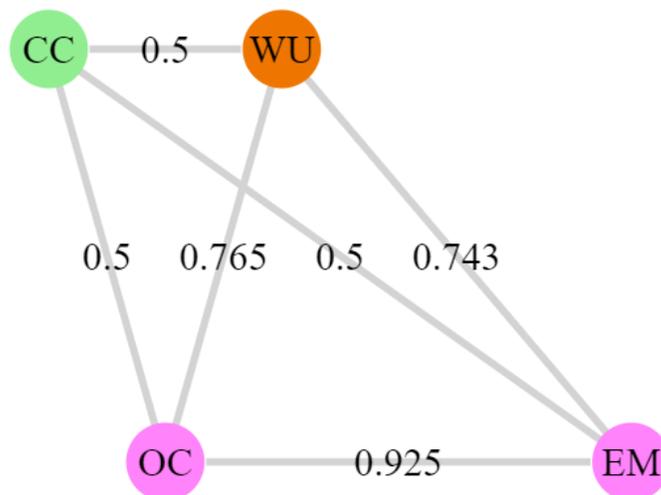

Figure 6: Objective correlation graph with correlation strengths as edge weights and identified communities in each color. Objectives are capital cost (CC), water use (WU), operating cost (OC), and carbon emissions (EM).

objectives as the best pair to combine, with the total algorithm requiring 3.34 seconds to determine graph edge weights and best objective groupings. While this finding suggests to

combine of an economic and an environmental objective which traditionally is not done, we note that this grouping intuitively makes sense for this particular problem: when deciding to reduce carbon emissions by building new distributed wind-powered ammonia facilities, we are inherently reducing operation costs as the feedstocks of onsite wind power, water, and air are essentially free, and transportation costs are significantly lower as these facilities are built closer to the farms which use the ammonia produced. As combining these objectives is not a traditional grouping, the information lost was compared with the conventional grouping that combines both economic objectives using the classical notion of net present value. Using a single dimension resolution of 40 steps, a total of $64,000$ optimization iterations, the information lost by combining economic objectives was found to be an average of 0.0883. Following the combination suggested by community detection on the correlation strength graph, the average information lost was found to be 0.0159. This is an 82% decrease in the information lost on the Pareto frontier. As such, using the algorithm's identified groups in place of the conventional groups results in additional information about the tradeoff between emissions and capital cost by grouping operating cost with emissions, which it is correlated with, rather than capital cost, which it competes with. This additional information can provide decision makers with a more complete view of tradeoff options when determining sustainability preferences and making a final design decision.

Energy Technology Selection Optimization

The multi-objective energy storage system selection optimization study presented in Li et al.[52] contains three objectives. The objectives are to minimize both levelized cost of energy storage (EC) and an environmental impact metric, ReCiPe, (RE), while the technology deployment number (TE) objective is maximized. The problem is formulated to have technologies' feasibility for different use cases determined by discharge duration and rated power parameters. These parameters are varied in order to explore what technologies are best for large-scale energy management, transmission and distribution support, customer energy

management, and distributed energy systems applications. The full optimization formulation is:

$$\min_{z_{ij}} \quad \{-\sum_{i=1}^{p} TE_i z_{ij}, \sum_{i=1}^{p} EC_{ij} z_{ij}, \sum_{i=1}^{p} RE_i z_{ij}\}$$

$$\text{s.t} \quad \sum_{i=1}^{p} z_{ij} = 1$$

$$0 \leq \left(\frac{(D_i^{max} - D_i^{min}) + (D_j^{max} - D_j^{min})}{2} - \left| \frac{D_i^{min} + D_i^{max}}{2} \right| - \left| \frac{D_j^{min} + D_j^{max}}{2} \right| \right) z_{ij}$$

$$0 \leq \left(\frac{(P_i^{max} - P_i^{min}) + (P_j^{max} - P_j^{min})}{2} - \left| \frac{P_i^{min} + P_i^{max}}{2} \right| - \left| \frac{P_j^{min} + P_j^{max}}{2} \right| \right) z_{ij} \quad (24)$$

where z_{ij} is a binary variable representing the selection of technology i in use case j, TE_i is the number of deployments of the technology i, EC_{ij} is the levelized cost of storage, RE_i is the combined environmental impact of each technology, D_i is the discharge duration, and P_i is the rated power. Practically, this problem can be simplified by identifying infeasible solutions by calculating the values of coefficients on the inequalities and eliminating any technology that is infeasible prior to optimization. This calculation and removal of infeasible technologies will be performed prior to implementation of the proposed algorithm.

Within the original environmental objective, fossil-fuel depletion (FD), particulate matter formation (PM), human toxicity (HT), and climate change impact (CC) are the individual factors that are added together. Combining these metrics gives an understanding of the overall behavior of the specific solution with respect to the environment, but separating them allows a more detailed understanding of how the different energy technologies have varied impacts on certain environmental metrics. In Oliveira et al.[53] the underlying data from these four metrics is presented for several of the studied technologies. Extracting this data and expanding the objective formulation of the optimization problem in Li et al. gives a 6-objective optimization with the original technology and cost objectives and the newly separated four environmental measures. This case study neglects technologies whose ReCiPe component values were not given by Oliveira et al. but were assumed in Li et al. giving

23

energy storage technology options of pumped hydro storage (PHS), compressed air energy storage (CAES), lead-acid batteries, sodium sulphur batteries (NaS), sodium nickel chloride batteries (NaNiCl), and lithium-ion batteries.

Two use cases from the original study are examined. The first use case is in transmission and distribution support, use case A3 in the original work. Running the algorithm takes 0.38 seconds and gives correlation strength weights and objective groups shown on the objective correlation graph in Figure 7. The groups that result from community detection are one with
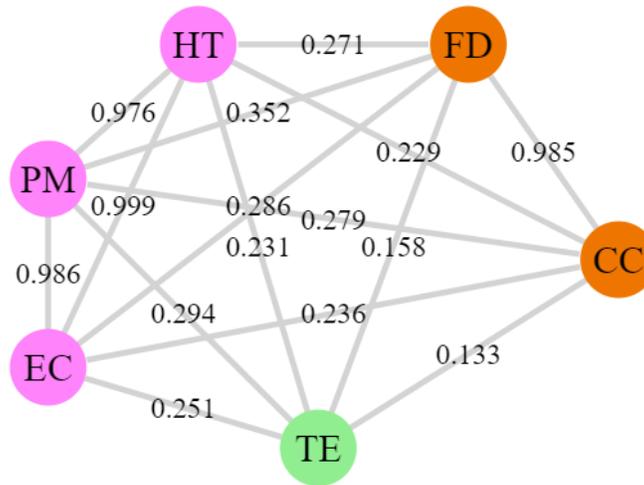

Figure 7: Objective correlation graph with correlation strengths as edge weights and colors representing detected communities for transmission and distribution support use case.

only the TE objective, the next with the FD and CC objectives, and the final group contains EC, PM, and HT objectives. This grouping indicates that combining all environmental objectives into one gives misleading results that a single technology, pumped hydro, is superior in this use case when instead there is an embedded tradeoff such that other technologies perform better for specific environmental metrics.

Next, the case where the energy storage technology is being used for customer energy management is explored; this is use case A8 in the original work. Figure 8 shows the weighted objective correlation graph and the communities identified by the algorithm, which takes 0.41 seconds to run. Groups that result from community detection are one with only the TE objective, one with only the CC objective, and the final containing EC, PM, FD,
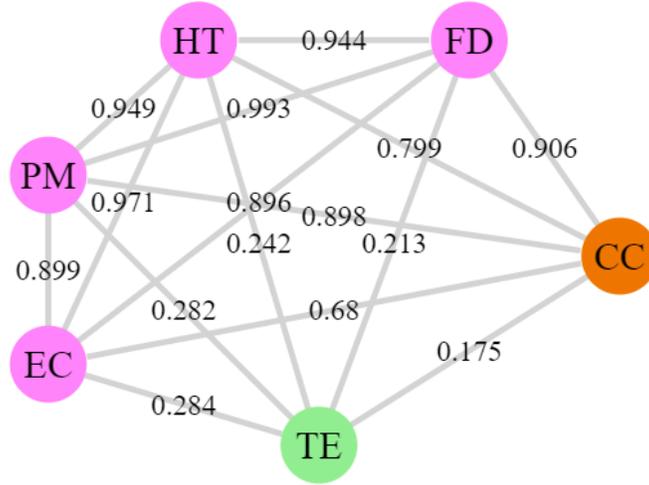
Figure 8: Objective correlation graph with correlation strengths as edge weights and colors representing detected communities for customer energy management use case.

and HT objectives. Aside from correlation strengths that link the technology objective to any other, the lowest correlation strength is 0.68. This indicates that the CC objective could be reasonably switched with the economic objective, returning the original RE objective.

Since only a small selection of technologies are available for each of these use cases, the entire Pareto frontier can be examined to understand the resulting objective groups and demonstrate the physical insights generated by our algorithm . Figure 9 shows the optimal technology choices using the objective groups determined by the algorithm, each individual objective, and the original objective functions for both use cases. In the transmission and distribution support case, the original objectives all result in PHS as the optimal choice. Looking at the individual objectives shows that in order to minimize both the FD and CC objectives, NaNiCl is chosen. The algorithm's resulting objective groups capture this behavior by combining the EC objective with the PM and HT objectives, leaving FD and CC in their own combined objective. The CAES technology also gives optimal performance for the individual HT objective; however, this technology is only weakly Pareto optimal as it is dominated by PHS, which performs equally well in HT and better in the other objectives. Overall, it is apparent that our algorithm is able to preserve tradeoff information that is lost when all environmental objectives are grouped together using the ReCiPe formulation, and

Technology	Transmission and Distribution Support			Customer Energy Management		
	Algorithm	Individual	Original	Algorithm	Individual	Original
PHS	O1,O3	EC,TE,PM,HT	EC,TE,RE	O1,O3	EC,TE,FD,PM,HT	EC,TE
CAES		HT			HT	
Lead						
NaS				O2	CC,FD,PM,HT	RE
NaNiCl	O2	FD,CC				
Li-ion						

Figure 9: Results of many-objective energy technology selection. Technologies that optimize a combined or individual objective are highlighted and those deemed infeasible due to inadequate discharge duration and rated power are grayed out. Lighter blue highlighting indicates that the solution is weakly Pareto optimal. Individual objectives listed multiple times are optimal for multiple technologies. In both cases, O1, O2, and O3 represent the green, orange, and pink objective groups, respectively, on each objective correlation graph.

gives results that make physical sense based on the environmental impacts of the various technologies considered.

In the customer energy management use case, original EC and TE objectives result in the choice of PHS while the RE objective results in choosing NaS. Examining individual objectives, it is apparent that the NaS technology optimizes all four environmental metrics, but FD, PM, and HT perform equally well when choosing PHS. HT is also optimal using CAES. The algorithm assigned FD, PM, and HT to a group with the EC objective. The overlap in technologies optimizing many objectives explains the high correlation strengths in the objective-strength graph for this use case excluding the technology objective.

Conclusions and Future Work

In this work, we presented a novel algorithm for reducing the dimensionality of linear MaOPs *a priori* to generating points which are Pareto-optimal in the original full space problem. This was achieved on the basis of a weighted objective correlation graph, where weights were determined based on the overlap of cost vector projections on constraint surfaces. From this graph, community detection using the Leiden algorithm was performed to identify two to three groups of objectives whereby objectives in the same group are correlated, while

objectives in different groups are competing. The algorithm provides a method to reduce an intractable, uninterpretable high dimensional MaOP to a tractable two or three objective problem that generates Pareto frontiers which are interpretable and retain as much information as possible from the original full-space problem. Through the analysis of three representative case studies relevant to sustainable chemical and energy production, we demonstrated the efficacy of our algorithm to systematically generate objective groups which preserve more information about tradeoffs than "intuitive" groupings of all economic or environmental objectives. As such, decision makers analyzing the Pareto solutions will be able to make more informed decisions based on full knowledge of the scope of tradeoffs among various sustainability outcomes.

Future work will attempt to extend this algorithm to nonlinear MaOPs, which provide the additional challenges that single-objective optimal points are no longer guaranteed to lie on the boundary of the feasible space, and constraint normals and objective gradients are no longer constant. We also intend to apply this framework to distributed and decomposed optimization problems in order to identify correlation between subproblem objectives, which can be useful in building minimal communication architectures which still give convergence of the decomposition solution algorithm. Finally, we plan to examine more deeply the parametric sensitivity of objective correlations, which will be important for moving horizon scheduling and control problems which must be solved repeatedly with different initial conditions over time.

Notation

Objective Reduction Algorithm

A Coefficient matrix for inequality constraints

a_k The k^{th} row of matrix A

b Vector of upper bounds for inequality constraints

C	Matrix of all objective cost vectors
c_i	Cost vector of i^{th} objective
c_{ik}^N	Component of cost vector c_i normal to constraint surface defined by a_k
c_{ik}^P	Component of cost vector c_i along the plane of the constraint surface defined by a_k
$\hat{c}_{ik}^N$	Component of cost vector c_i normal to constraint surface defined by a_k, normalized to length 1
$\hat{c}_{ik}^P$	Component of cost vector c_i along the plane of the constraint surface defined by a_k, normalized to length 1
D	Coefficient matrix for equality constraints
d_k	The k^{th} row of matrix D
e	Vector of constants for right hand side of equality constraints
S_{ijk}	Strength of correlation between objectives i and j along constraint surface k
S_{ij}^A	Total correlation strength between objectives i and j
W_{ijk}	Weight of contribution of S_{ijk} for calculating total correlation strength S_{ij}^A
α	Hyperparameter for discounting correlating constraints in weighting
β	Hyperparameter governing smoothness of weighting logistic curve

Information Loss Metric

B_p	Information loss at point p
$\mathcal{C}$	Set of all objectives combined into a new grouped objective
$\mathcal{J}_p$	Set of Pareto optimal points found when keeping objectives values for objectives in $\mathcal{R}$ constant at point p value.
K_p	Matrix of objective values at point p
K_{pij}	Value of objective i at point p in the reduced space Pareto, and point j in the full space Pareto with constant objectives values in $\mathcal{R}$
$\hat{K}_{pij}$	Objective value K_{pij} rescaled to be between 0 and 1.

l_i Minimum value of objective i in entire full space Pareto frontier.

$\mathcal{P}$ Set of points on the reduced space Pareto frontier

$\mathcal{R}$ Set of all objectives retained as individual objectives

u_i Maximum value of objective i in entire full space Pareto frontier.

UK Energy Mix Case Study

c_j Cost of generating energy by technology j

D Total energy demand

G_j Amount of energy generated by technology j

G_j^{min} Minimum amount of energy generated by technology j

G_j^{max} Maximum amount of energy generated by technology j

i_j Worker injuries caused by technology j

$\mathcal{J}$ Set of energy technologies

w_j Global warming potential of emissions from technology j

Energy Storage Selection Case Study

D_i^{max} Maximum discharge duration for technology i

D_j^{max} Maximum discharge duration for use case j

D_i^{min} Minimum discharge duration for technology i

D_j^{min} Minimum discharge duration for use case j

EC_{ij} Levelized cost of storage for technology i in use case j

P_i^{max} Maximum rated power for technology i

P_j^{max} Maximum rated power for use case j

P_i^{min} Minimum rated power for technology i

P_j^{min} Minimum rated power for use case j

RE_i ReCiPe-determined environmental impact for technology i

TE_i Number of deployments of technology i

z_{ij} Binary variable, 1 if technology i is selected in use case j

Literature Cited

1 Pastor L, Stambaugh RF, Taylor LA. Sustainable investing in equilibrium. *Journal of Financial Economics*. 2021;142:550–571.

2 Yu Z, Razzaq A, Rehman A, Shah A, Jameel K, Mor RS. Disruption in global supply chain and socio-economic shocks: a lesson from COVID-19 for sustainable production and consumption. *Operations Management Research*. 2021;.

3 Neavling S. Marathon ordered to pay fines, improve air quality in southwest Detroit. *Detroit Metro Times*. 2020;.
URL https://www.metrotimes.com/news/marathon-ordered-to-pay-fines-improve-air-quality-in-southwest-detroit-25018119

4 World Meteorological Organization. State of Climate in 2021: Extreme events and major impacts. 2021.
URL https://public.wmo.int/en/media/press-release/state-of-climate-2021-extreme-events-and-major-impacts

5 da Silva C, Barbosa-Pavoa AP, Carvalho A. Towards sustainable development: Green supply chain design and planning using monetization methods. *Business Strategy and the Environment*. 2021;.

6 Al-Qahtani A, Parkinson B, Hellgardt K, Shah N, Guillen-Gosalbez G. Uncovering the true cost of hydrogen production using life cycle monetisation. *Applied Energy*. 2021; 281:115958.

7 Zirngast K, Kravanja Z, Pintaric ZN. The influence of variable CO2 emission tax rate on flexible chemical process synthesis. *Processes*. 2021;9:1720.

8 Leonzio G, Foscolo PU, Zondervan E. Sustainable utilization and storage of carbon

dioxide: Analysis and design of an innovative supply chain. *Computers and Chemical Engineering.* 2019;131:106569.

9 Palys MJ, Allman A, Daoutidis P. Exploring the benefits of modular rebewable-powered ammonia production: a supply chain optimization study. *Industrial and Engineering Chemistry Research.* 2019;58:5898–5908.

10 Darom NA, Hishamudding H, Ramli R, Nopiah ZM. An inventory model of supply chain disruption recovery with safety stock and carbon emission consideration. *Journal of cleaner production.* 2018;197:1011–1021.

11 Pizzol M, Weidema B, Brandao M, Osset P. Monetary valuation in life cycle assessment: a review. *Journal of Cleaner Production.* 2015;86:170–179.

12 Wheeler J, Paez MA, Guillen-Gosalbez G, Mele FD. Combining multi-attribute decision-making methods with multi-objective optimization in the design of biomass supply chains. *Computers and Chemical Engineering.* 2018;113:11–31.

13 Wang P, Deng X, Zhao H, Yu S. Estimates of the social cost of carbon: A review based on meta-analysis. *Journal of Cleaner Production.* 2019;209:1494–1507.

14 Wang Z, Rangaia GP, Wang X. Preference ranking on the basis of ideal-average distance method for multi-criteria decision making. *Industrial and Engineering Chemistry Research.* 2021;60:11216–11230.

15 Rangaiah GP, Feng Z, Hoadley AF. Multi-Objective Optimization Applications in Chemical Process Engineering: Tutorial and Review. *Processes.* 2020;8.

16 Roh K, Al-Hunaidy AS, Imran H, Lee JH. Optimization-based identification of CO2 capture and utilization processing paths for life cycle greenhouse gas reduction and economic benefits. *AIChE Journal.* 2019;65:e16580.

17 Zhao X, You F. Waste high-density polyethylene recycling process systems for mitigating plastic pollution through a sustainable design and synthesis paradigm. *AIChE Journal.* 2021;67:e17127.

18 Gopalakrishnan V, Bakshi B. Ecosystems as unit operations for local techno-ecological synergy: Integrated process design with treatment wetlands. *AIChE Journal.* 2018; 64:2390–2407.

19 You F, Tao L, Graziano DJ, Snyder SW. Optimal design of sustainable cellulosic biofuel supply chains: Multiobjective optimization coupled with life cycle assessment and input–output analysis. *AIChE Journal.* 2012;58:1157–1180.

20 Allman A, Tiffany D, Kelley S, Daoutidis P. A framework for ammonia supply chain optimization incorporating conventional and renewable generation. *AIChE Journal.* 2017; 63:4390–4402.

21 Liu Z, Lim MQ, Kraft M, Wang X. Simultaneous design and operation optimization of renewable combined cooling heating and power systems. *AIChE Journal.* 2020;66:e17039.

22 Saif Y, Almansoori A, Elkamel A. Sustainable optimization of waste management network over extended planning time horizon. *AIChE Journal.* 2021;67:e17256.

23 Garbie I. Sustainability Optimization in Manufacturing Enterprises. *Procedia CIRP.* 2015;26:504–509.

24 Santibanez-Aguilar JE, Gonzalez-Campos JB, Ponce-Ortega JM, Serba-Gonzalez M, El-Halwagi MM. Optimal planning and site selection for distributed multiproduct biorefineries involving economic, environmental and social objectives. *Journal of Cleaner Production.* 2014;65:270–294.

25 Garcia DJ, Lovett BM, You F. Considering agricultural wastes and ecosystem services in

Food-Energy-Water-Waste Nexus system design. *Journal of Cleaner Production.* 2019; 28:941–955.

26 Wheeler J, Galan-Martin A, Mele FD, Guillen-Gosalbez G. Designing biomass supply chains within planetary boundaries. *AIChE Journal.* 2021;67:e17131.

27 Mota B, Gomes MI, Carvalho A, Barbosa-Pavoa AP. Sustainable supply chains: An integrated modeling approach under uncertainty. *Omega.* 2018;77:32–57.

28 Saling P, Kicherer A, Dietrich-Kramer B, Wittlinger R, Zombik W, Schmidt I, Schrott W, Schmidt S. Eco-efficiency analysis by BASF: the method. *The International Journal of Life Cycle Assessment.* 2002;7:203–218.

29 Saxena DK, Duro JA, Tiwari A, Deb K, Zhang Q. Objective Reduction in Many-Objective Optimization: Linear and Nonlinear Algorithms. *IEEE Transactions on Evolutionary Computing.* 2013;17:77–99.

30 de Freitas ARR, Fleming PJ, Guimares FG. Aggregation Trees for visualization and dimension reduction in many-objective optimization. *Information Sciences.* 2015;298:288–314.

31 Brockhoff D, Zitzler E. Objective Reduction in Evolutionary Multiobjective Optimization: Theory and Applications. *Evolutionary Computation.* 2009;17:135–166.

32 Marler RT, Arora JS. The weighted sum method for multi-objective optimization: new insights. *Structural and Multidisciplinary Optimization.* 2009;41:853–862.

33 Giagkozia I, Fleming PJ. P.J. Methods for multi-objective optimization: An analysis. *Information Sciences.* 2015;293:338–350.

34 Dutta S. *Optimization in Chemical Engineering.* Cambridge University Press. 2016.

35 Emmerich A MTM; Deutz. A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Nat Comput.* 2018;17.

36 Jalving J, Cao Y, Zavala VM. Graph-based modeling and simulation of complex systems. *Computers and Chemical Engineering.* 2019;125:134–154.

37 Fortunato S, Hric D. Community detection in networks: a user guide. *Physics Reports.* 2016;659:1–44.

38 Newman MEJ. Spectral methods for community detection and graph partitioning. *Physical Review E.* 2013;88:042822.

39 Blondel VD, Guilaume JL, Lambiotte R, Lefebver E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment.* 2008; 2008:P10008.

40 van Eck NJ TVWL. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific Reports.* 2019;9(5233).

41 del Rio-Chanona EA, Fiorelli F, Vassiliadis VS. Automated structure detection for distributed process optimizaiton. *Computers and Chemical Engineering.* 2016;89:135–148.

42 Tang W, Allman A, Pourkargar DB, Daoutidis P. Optimal decomposition for distributed optimization in nonlinear model predictive control through community detection. *Computers and Chemical Engineering.* 2018;111:43–54.

43 Allman A, Tang W, Daoutidis P. DeCODe: a community-based algorithm for generating high-quality decompositions of optimization problems. *Optimization and Engineering.* 2019;20:1067–1084.

44 Mitrai I, Tang W, Daoutidis P. Stochastic Blockmodeling for learning the structure of optimization problems. *AIChE Journal.* 2021;p. e17415.

45 Shin S, Zavala VM, Anitescu M. Decentralized Schemes With Overlap for Solving Graph-Structured Optimization Problems. *IEEE Transactions on Control of Network Systems.* 2020;7:1225–1236.

46 Bodi R, Herr K, Joswig M. Algorithms for highly symmetric linear and integer programs. *Mathematical Programming.* 2013;137:65–90.

47 Shin S, Anitescu M, Zavala VM. Exponential Decay of Sensitivity in Graph-Structured Nonlinear Programs. *SIAM Journal on Optimization.* 2022;32.

48 Cplex II. V20. 1: User's Manual for CPLEX. *International Business Machines Corporation.* 2020;.

49 Dunning I, Huchette J, Lubin M. JuMP: A Modeling Language for Mathematical Optimization. *SIAM Review.* 2017;59(2):295–320.

50 Bezanson J, Edelman A, Karpinski S, Shah VB. Julia: A Fresh Approach to Numerical Computing. *SIAM Review.* 2017;59:65–98.

51 Limleamthong P, Guillén-Gosálbez G. Rigorous analysis of Pareto fronts in sustainability studies based on bilevel optimization: Application to the redesign of the UK electricity mix. *Journal of Cleaner Production.* 2017;164:1602–1613.

52 Li L, Liu P, Li Z, Wang X. A multi-objective optimization approach for selection of energy storage systems. *Computers and Chemical Engineering.* 2018;115:213–225.

53 Oliveira L, Messagie M, Mertens J, Laget H, Coosemans T, Mierlo JV. Environmental performance of electricity storage systems for grid applications, a life cycle approach. *Energy Conversion and Management.* 2015;101:326–335.