DYNAMIC OBJECTIVE CORRELATION IN MANY-OBJECTIVE OPTIMAL OPERATION: AN AMMONIA PRODUCTION CASE STUDY

Andrew Allman[1] and Justin M. Russell

Department of Chemical Engineering, University of Michigan, Ann Arbor, MI 48105

Abstract

Multi-objective optimization is an important tool for sustainable decision making that takes into account environmental, social, and governance outcomes. For problems of four or more objectives, dimensionality reduction approaches are needed to be able to apply rigorous global optimization methods for finding the Pareto frontier in a tractable manner. For optimal real-time operation, scheduling, and control problems, the correlating versus competing nature of objectives , which influences how best to reduce objective dimensionality, can vary as dynamic parameters change when solving the optimization problem repeatedly in a moving horizon. This work attempts to analyze the sensitivity of objective correlation through the use of a green ammonia production case study. It considers operating cost, carbon emissions, water usage, and electrolyzer safety objectives and examines their correlation strength as the cost and related emissions of purchasing electricity from the power grid changes. Results of the study show that correlation between cost and emissions are near perfect when the time-varying cost and emissions factor profiles are in phase, and low when they are anti-phase. Furthermore, the variance of objective correlations and the best objective grouping over time is demonstrated through the analysis of historical power cost and average grid emissions data from CAISO.

Keywords

Multi-objective Optimization, Planning and Scheduling, Green Ammonia

Introduction and Background

Recent global events and investor emphasis on environmental, social, and governance (ESG) outcomes have motivated the consideration of more than traditional economic objectives when making decisions at all levels of the chemical enterprise, including those associated with real-time operation, scheduling, and control of chemical systems. Multi-objective optimization is an ideal tool for understanding the tradeoffs inherent when considering disparate objectives. Rigorous scalarization methods such as the weighted sum or epsilon constraint approaches, combined with global optimization solvers such as CPLEX or BARON, can provide a set of trade-off solutions provably on the Pareto frontier, such that each point in the solution set represents the best one can feasibly do in one objective without harming another. Unfortunately, these methods scale very poorly to "many-objective" optimization problems of four objective functions or more, to the point where heuristic evolutionary algorithms are the current method of choice for solving problems in the many-objective optimization community (Emmerich, 2018). Historically, the process systems community has been resistant to using such heuristic approaches in decision making when they can be avoided, due to concerns with degradation in economic performance and process safety associated with suboptimal decision making and a desire to be able to calculate an optimality gap, or worst case of how far the current solution is from the best solution.

In order to make a rigorous solution approach tractable for many-objective problems, a reduction of objective dimensionality to three or fewer objectives is required. There exist several methods for systematically achieving this, the vast majority of which perform dimensionality *a posteriori* to generating at least part of or an approximation of the full space Pareto frontier. Approaches which achieve this include principal component analysis (Saxena et al., 2013), aggregation trees (de Freitas et al., 2015) or dominance preservation strategies (Brockhoff and Zitzler, 2009). The goal of these methods is to remove objectives that are uninformative in the full space Pareto solution, or equivalently, group together objectives which are inherently correlated, pointing towards the same solution. Alternatively, one can also group objectives together based on physical classification, with the "three pillars of sustainability" (economic, environmental, and social objectives) often forming the basis of such a grouping. More recently, our team has developed a dimensionality reduction method for (mixed integer) linear many-objective optimization problems that systematically achieves dimensionality reduction *a priori* to solving the many-objective problem on the basis of problem structure and the overlap of cost vector projections onto constraint surfaces (Russell and Allman, 2022). Results from this work showed the algorithm's ability to identify groupings of objectives such that correlating objectives are put in the same group, competing objectives are put in different groups, and minimal information is lost

[1] Corresponding author: A. Allman (E-mail: allmanaa@umich.edu).

in objective grouping in multiple energy system design and supply chain optimization case studies.

In optimal real-time operation, scheduling, and control problems, the standard decision making framework is such that an optimization problem is not just solved once but instead repeatedly over time in a moving-horizon fashion. At each time point, critical problem parameters such as initial conditions, feedstock costs, product demands, weather forecasts, or output set points can be different than they were in the previous solve. As these parameters change, so can the correlating vs. competing nature of objectives in a many-objective optimization problem. As such, it makes sense that the objectives we wish to group together in order to achieve an informative should not be static, but instead evolve dynamically based on the sensitivity of objective correlations with dynamically varying problem parameters.

In this work, we present an initial effort towards understanding the sensitivity of objective correlations in many-objective problems via the analysis of an ammonia production operation case study. Ammonia is a critical chemical supporting modern society, with its primary use for nitrogenous fertilizer, either through direct application to crops or as a chemical precursor to other fertilizers such as urea. It has also been proposed as a medium for long-term energy storage, or as an easily transportable liquid energy or hydrogen carrier. Traditionally, ammonia production utilizes fossil fuels for hydrogen generation, and accounts for about 2% of global greenhouse gas emissions. Recently, there has been immense academic, governmental, and industrial interest in green ammonia, whereby the ammonia production process is powered completely by renewable energy, typically generating hydrogen from water electrolysis. From a systems perspective, recent research has shown that green ammonia systems can be implemented at a distributed scale using modular units (Palys et al., 2019), and can be operated in a time-varying fashion to exploit intermittency in, for example, renewable energy production or energy costs (Allman and Daoutidis, 2018). Due to recent global events and economic inflation, natural gas and conventional ammonia prices have rapidly increased in the past year from between $400 and $850 for the previous ten years prior up to over $1450 in the last year, making green ammonia production more economically promising than ever before. This system is an ideal system to analyze for determining objective sensitivites to dynamic parameters, as there are inherently various economic, environmental, and safety objectives that should be considered in operation, and because of its reliance on electricity whose availability, cost, and related emissions may vary with time.

This work extends upon previous analysis of ammonia production scheduling in two ways: first, four different objectives are considered, including the traditional objective of operating cost, along with carbon emissions, water usage, and electrolyzer safety. Second, we consider that energy purchased from the grid has not only time varying costs, but also time varying carbon emissions, which can occur due to different modes of generation, including renewable sources such as wind and solar as well as inefficient sources such as fast-ramping fossil fuel generators, being active at different times. The remainder of this paper is organized as follows: in the next section, we briefly overview our approach for systematically grouping objectives *a priori* to generating a Pareto frontier. Then, the problem formulation for many-objective ammonia production is presented. Next, we present the details of the specific case studies considered and analyze the results of our objective reduction algorithm for the various problem instances. Finally, we conclude by summarizing important observations and proposing avenues for future work.

Objective Reduction Algorithm

Consider a linear N-objective optimization problem of the form:

$$\begin{aligned} \min_{x} \quad & \{c_1^T x, c_2^T x, ..., c_N^T x\} \\ \text{s.t.} \quad & Ax \leq b \end{aligned} \tag{1}$$

The variables x may be continuous or integer, while c_i represent the cost vectors for the various objectives and A represents the constraint matrix. To reduce the number of objectives, we need to be able to use the information embedded in the problem to understand if objectives interact in a correlated or competing manner. For linear problems, optimal solutions always lie on the boundary of the feasible region when the solution is finite; thus, it makes sense to focus the analysis to objective interactions on constraint surfaces. For each objective-objective-constraint triplet, we calculate an interaction strength by first projecting cost vectors onto the surface of individual constraints defined by normal vector a_k, a row of the constraint matrix A:

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

$$\hat{c}_{ik}^P = \frac{-c_i - c_{ik}^N}{\| - c_i - c_{ik}^N\|_2} \tag{3}$$

Where c_{ik}^N represents the component of the cost vector c_i normal to constraint surface k, and $\hat{c}_{ik}^P$ is the component of the cost vector c_i along the surface of constraint k, normalized to length 1. For two objectives i and j interacting on constraint k, we define a strength of interaction S_{ijk}:

$$S_{ijk} = (\hat{c}_{ik}^P)^T \hat{c}_{jk}^P \tag{4}$$

Since projection vectors are normalized, the strength of objective interaction along a constraint can vary between 1, when objectives are perfectly correlated and pointing in the exact same direction, and -1, when objectives are pointing in the exact opposite direction. To determine the total objective correlation strength, S_{ij}^A, we consider a weighted sum of the constraint interaction strengths. Weights W_{ijk} are zero when both normal vectors c_{ik}^N and c_{jk}^N point into the feasible region, as this implies constraint k is unlikely to be active in optimizing either objective, and are otherwise determined using the

following logistic function:

$$W_{ijk} = 1 - \alpha \left(\frac{1}{1 + \exp(-\beta \times S_{ijk})} \right) \quad (5)$$

This equation captures the fact that competition on any constraint can lead to an objective tradeoff, while correlation needs to occur on all constraints for objectives to be correlated. As such, it makes sense to weight negative values of S_{ijk} more than positive ones. The hyperparameters α and β alter the shape of the logistic curve, such that α, which should lie between zero and one, determines how much correlating constraints are discounted and β, which should be positive, determines the slope of the change from minimum to maximum weight, with larger values giving a more steplike change. The total objective interaction strength is determined as follows:

$$S_{ij}^{A} = 0.5 \left(1 + \frac{\sum_{j=1}^{M+N} W_{i,j,k} S_{i,j,k}}{\sum_{j=1}^{M+N} W_{i,j,k}} \right) \quad (6)$$

The objective correlation strength is scaled to be between 0 and 1, with low values corresponding to highly competing objectives and high values corresponding to highly correlating objectives. The strengths are used to define edge weights in an objective correlation graph, where nodes correspond to each individual objective. Community detection can then be applied to the objective correlation graph to identify groups of objectives which are highly correlated to others in the same group, but competing with those in different groups. In this work, we perform community detection using the Leiden algorithm (van Eck N.J., 2019), with a resolution hyperparameter tuned to give two groups of objectives.

System Description and Problem Formulation

The ammonia production system analyzed in this work is adapted from previous works (Allman and Daoutidis, 2018), and described below. First, hydrogen is either produced onsite via water electrolysis or purchased from a methane steam reformer. Nitrogen is also produced onsite via a pressure swing adsorption (PSA) unit. A small amount of storage capacity exists for both nitrogen and hydrogen, which are then consumed by a downstream, small-scale ammonia production reactor based on a Haber process. Energy for all process is obtained from onsite wind production and supplemented with purchases from the power grid. Both the costs of these purchases and the associated carbon emissions from energy generation are time-varying parameters. Excess energy can also be sold back to the power grid at a rate lower than the purchase rate, but no emissions credit is given for selling power.

Our goal in this problem is to meet some nominal ammonia demand over a 48-hour period considering four disparate objectives and the physical constraints of the system. Variable notation in the problem is described as introduced in the system constraints. For the objectives, we first consider minimizing the traditional objective of operating cost Z, where ζ_i represent relative cost terms for various variables in the system:

$$\begin{aligned} Z = \sum_{t=1}^{48} (\zeta_b(t) p_b(t) - \zeta_s(t) p_s(t) + \zeta_{NH3} m_{NH3}^{dev} + \zeta_{H2,buy} b_{H2}(t) \\ + \zeta_{on} s_e(t) + \zeta_{el} m_{H2}(t) + \zeta_{psa} m_{N2}(t)) + \zeta_{H2,sto} U_{H2} \\ + \zeta_{N2} U_{N2} \end{aligned} \quad (7)$$

Cost contributions come from exchanging power with the grid, purchasing hydrogen from the steam reforming plant, deviating from the target ammonia production rate, operating the electrolyzer, PSA unit, and ammonia reactor, and depleting gases from storage. Next, we consider the minimization of carbon emissions associated with operation of the facility H, where η_i represent relative emissions terms for various variables in the system:

$$H = \sum_{t=1}^{48} (\eta_{b_{H2}} b_{H2}(t) + \eta_{p_b}(t) p_b(t)) \quad (8)$$

Carbon emissions for this system are incurred from external purchases, either via the production of hydrogen through steam reforming, or via the production of power. The latter contribution can vary with time. Next, we consider the minimizing usage of water by the system Ψ, which can be of large concern if the system is being operated in water scarce regions of the world, where ψ_i represent relative water usage for various variables in the system.

$$\Psi = \psi_{H2,el} m_{H2}(t) + \psi_{H2,buy} b_{H2}(t) \quad (9)$$

Water is used by the system in the production of hydrogen, both in electrolysis in steam reforming, although the latter process has lower water requirements. Finally, we consider a safety objective, Ξ, related to the number of times an electrolyzer is started up throughout the scheduling horizon, an action which carries with it an above average safety risk:

$$\Xi = \sum_{t=1}^{48} s_e(t) \quad (10)$$

The four objectives above are constrained by physical and practical limitations in the system. First, we consider a power balance that must hold at all times:

$$\begin{aligned} p_w(t) + p_b(t) = p_s(t) + p_{N2}(t) + p_{H2}(t) + p_{NH3}(t) \\ \forall t \in \{1, ..., 48\} \end{aligned} \quad (11)$$

Here, the amount of power produced by wind p_2 and bought from the grid p_b must equal the amount of power sold to the grid p_s or consumed by the chemical units in the system (p_{H2}, p_{N2}, and p_{NH3}. The amount of power used by chemical units is related to the production rate of that chemical m_i by a linear conversion factor ρ_i:

$$p_i(t) = \rho_i m_i(t) \quad \forall i \in \{H2, N2, NH2\}, t \in \{1, ..., 48\} \quad (12)$$

Furthermore, the PSA unit and ammonia reactor have inherent upper and lower bound on their production rates:

$$m_i^{min} \leq m_i(t) \leq m_i^{max} \quad \forall t \in \{1, ..., 48\}, i \in \{N2, NH3\} \quad (13)$$

A discrete number of electrolyzers are present in the system which can be in an "on" or "off" state. The number of "on" electrolyzers n_e is limited to the number of electrolyzers installed N_e:

$$0 \leq n_e(t) \leq N_e \quad \forall t \in \{1,...,48\} \tag{14}$$

The production rate of hydrogen is bound by the number of electrolyzers turned on:

$$m_{H2}^{min} n_e(t) \leq m_{H2}(t) \leq m_{H2}^{max} n_e(t) \quad \forall t \in \{1,...,48\} \tag{15}$$

It is also important to track the number of electrolyzers starting up in a time period s_e for cost and safety purposes:

$$s_e(t) \geq n_e(t) - n_e(t-1) \quad \forall t \in \{1,...,48\} \tag{16}$$

$$s_e(t) \geq 0 \quad \forall t \in \{1,...,48\} \tag{17}$$

The amount of hydrogen that can be purchased from the steam reforming plant b_{H2} is also bounded:

$$0 \leq b_{H2}(t) \leq b_{H2}^{max} \quad \forall \quad t \in \{1,...,48\} \tag{18}$$

Intermediate storage units exist to store hydrogen and nitrogen for later use in the system, with the amount of component i stored denoted M_i. Each has a minimum and maximum capacity:

$$M_i^{min} \leq M_i(t) \leq M_i^{max} \quad \forall t \in \{1,..,48\}, i \in \{N_2, H_2\} \tag{19}$$

Amount of hydrogen and nitrogen in storage is tracked by dynamic mass balances for each component:

$$M_{H2}(t) = M_{H2}(t-1) + b_{H2}(t) + m_{H2}(t) - \frac{3}{17} m_{NH3}(t) \quad \forall t \in \{1,...,48\} \tag{20}$$

$$M_{N2}(t) = M_{N2}(t-1) + m_{N2}(t) - \frac{14}{17} m_{NH3}(t) \quad \forall t \in \{1,...,48\} \tag{21}$$

It is also important to track the amount of gases depleted from storage at the end of the scheduling horizon u_i, such that this usage is accounted for in the cost objective and gases are not depleted at the end of every 48-hour horizon unless there is a strong economic driving force to do so.

$$U_i = M_i(0) - M_i(48) \quad \forall \quad i \in \{N_2, H_2\} \tag{22}$$

As the ammonia reactor is inherently a slower responding unit than the two upstream units, we embed a set of constraints limiting its ramp rate by R^+ for ramping up and R^- for ramping down, as well as how long its production rate must remain constant D_a after a change:

$$-d_a(t)R^- \leq m_{NH3}(t) - m_{NH3}(t-1) \leq d_a(t)R^+ \quad \forall t \in \{1,...,48\} \tag{23}$$

$$\sum_{\tau=t+1-T}^{t} d_a(\tau) \leq D_a \quad \forall t \in \{1,...,48\} \tag{24}$$

In the above equations, d_a is a binary variable that is one if ammonia production rate changes in time period t, and zero otherwise. Finally, it is important to keep track of the ammonia production rate relative to demand m_{NH3}^{tar}, such that an economic penalty can be provided for underproduction m_{NH3}^{dev}:

$$m_{NH3}^{dev}(t) = m_{NH3}^{tar} - m_{NH3}(t) \quad \forall \quad t \tag{25}$$

The total four-objective optimization problem is as follows:

$$\begin{aligned} \min \quad & \{Z, H, \Psi, \Xi\} \\ \text{s.t.} \quad & (11)-(25) \end{aligned} \tag{26}$$

Such a problem can be rigorously solved by applying the weighted sum or epsilon constraint approaches. Doing so results in a set of single objective mixed integer linear programs that can be solved using CPLEX or Gurobi to obtain Pareto-optimal operating points.

Case Study Data

The majority of the parameter values used in this work are unchanged from a previous study in green ammonia production scheduling (Allman and Daoutidis, 2018). In this section, we detail any changes in parameter values or new parameters added that are considered in this work. First, the addition of the option to purchase fossil fuel-based hydrogen introduces a maximum purchase per hour of 100 kg at a cost of $ 2.40 which falls in line with data from Ramsden et al. (2009). Additional stoichiometric calculations were completed to determine the carbon emission factor from hydrogen derived from steam reforming of methane, $\eta_{b_{H2}}$ of 9.3 kg CO2/kg H2. Similarly, water usage coefficients are derived using stoichiometry to be $\psi_{H2,el} = 9$ kg water/kg H2 for electrolytic hydrogen and $\psi_{H2,buy} = 3$ kg water/kg H2 for steam reformed hydrogen. For the first 2 instances considered, power costs follow time varying pricing structure that peaks in the daytime at 25 ¢/kWh and bottoms out with negative prices at -2 ¢/kWh. For these studies, grid emissions are artificially generated to show the extreme cases of when emissions over time are completely in-phase or anti-phase with respect to power costs, using high and low emissions factors of 1.08 and 0 kg CO2/kWh, respectively. The time variation of costs and emissions for these two studies are displayed in Figure 1. Instances using real historic power cost and average emissions data obtained from CAISO are also analyzed. For each of these instances, the overlap between the cost and emissions profile is quantified using an approach described as follows. First, we shift and normalize the power ($\hat{\zeta}_b$) and emissions ($\hat{\eta}_{p_b}$) profiles such that the average value of the parameter over time is zero, and the maximum deviation from this value is ± 1:

$$\hat{\zeta}_b(t) = \frac{\zeta_b(t) - \sum_{t'=1}^{48} \zeta_b(t')/48}{\max_{t'} |\zeta_b(t') - \sum_{t''=1}^{48} \zeta_b(t'')/48|} \tag{27}$$

$$\hat{\eta}_{p_b}(t) = \frac{\eta_{p_b}(t) - \sum_{t'=1}^{48} \eta_{p_b}(t')/48}{\max_{t'} |\eta_{p_b}(t') - \sum_{t''=1}^{48} \eta_{p_b}(t'')/48|} \tag{28}$$

An overlap factor O_{ZH} is then defined as the time average of the sum of the product of the normalized cost and emis-

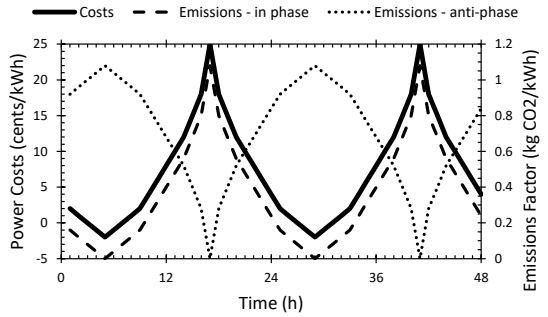

Figure 1: Plots of power costs and emissions factors over time for instances where costs and emissions are in-phase (dashed line) and anti-phase (dotted line).

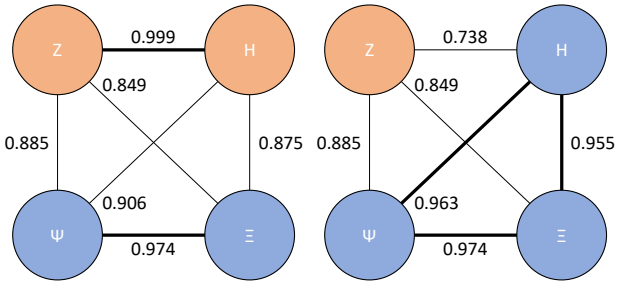

Figure 2: Objective correlation graph with color-coded objective groupings for in-phase (left) and anti-phase (right) instances.

sions profiles, divided by the maximum of the square of the normalized cost or emissions:

$$O_{ZH} = \frac{1}{48}\sum_{t=1}^{48}\frac{\hat{\zeta}_b(t)\hat{\eta}_{p_b}(t)}{\max(\hat{\zeta}_b(t)^2, \hat{\eta}_{p_b}(t)^2)} \tag{29}$$

This overlap metric can vary between 1, where costs and emissions are completely in phase as in the dotted example in Figure 1, and -1, where costs and emissions are completely out of phase as in the dashed example in Figure 1.

Finally, wind speeds used to calculate wind power generation in each time period are pulled from the National Centers for Environmental Information climate normals for the past 15 years in Sioux Falls, North Dakota (NCEI, 2020). This location was chosen as it is located in the upper Midwest and is a representative candidate location for a distributed ammonia production facility, in a location of both high wind availability and high corn production.

Results and Discussion

To understand the range over which objective correlations (namely, those between cost and carbon emissions) can vary as the power cost and emissions factor parameters vary, it is useful to first focus on the completely in-phase and anti-phase instances described in the previous section. In both instances, the objective reduction algorithm is performed to generate objective correlation strengths and determine which two objectives should be grouped together when reducing the system from four to two objectives. The resulting objective correlation graphs and community detection results in each instance are depicted in Figure 2. It is apparent that, as expected, when power costs and emissions are in phase, the two corresponding objectives have very strong correlation extremely close to one, and as such, the algorithm groups these two objectives together. Conversely, when power costs and emissions factor are anti-phase, the correlation strength becomes the weakest of all objective pairs at 0.74, indicating the presence of strong tradeoffs when considering optimal operation with respect to the two objectives. We further note that physically, it makes a lot of sense that water usage and electrolyzer safety are paired together in both cases, as both objectives promote avoiding electrolyzer usage, either due to startups being unsafe or its higher water usage than steam reforming. Somewhat surprisingly, the anti-phase instance groups carbon emissions with water usage and electrolyzer safety. This is likely because electrolysis is the most energy-intensive process in the system, thus contributing strongly to carbon emissions from grid-purchased power.

To make physical sense of the objective groupings, we show the single objective optimal schedules in Figure 3. Here, many of the trends intuited in the previous paragraph are supported visually. In particular, we see near-perfect agreement between the optimal operating schedules for operating cost and carbon emissions when the two are in phase, with most of the hydrogen and nitrogen production shifted to the times of lower power cost and emissions factor, respectively. Similarly, we see very good agreement between the optimal operating schedules for water usage and electrolyzer safety, with both producing as little hydrogen from electrolysis as possible, making as little ammonia as possible, and purchasing as much hydrogen from steam reforming as possible in order to produce ammonia at the minimum allowed production rate. Finally, we note that the anti-phase emissions optimal schedule looks the most different than all the rest, as it tries to put the bulk of its production in the midday hours when the grid's emission factor is lower.

The two instances previously analyzed are illustrative but extreme examples that demonstrate the two edge cases of how carbon emissions and operating cost can give correlating or competing solutions for optimal scheduling. Most realistic cases will have correlation strengths that lie somewhere in the middle of the two. In Table 1, we summarize the results of looking at five different 48 hour periods using real historical power price and average grid emissions data from CAISO. For each instance, the overlap factor (as defined in the previous section), correlation strength between operation cost and carbon emissions objectives, and determined objective grouping are reported. We note that in most cases, power costs and average grid emissions tend to be closer to completely in-phase than not. This is likely a consequence of using data from the California grid, where there is a large proportion of solar generation and times of low prices tend to correspond to times where solar generation is high, and thus average grid emissions are also low. An exception to this trend is during times of extremely high demand, such as due to a heat wave, as seen in the September 1-2 instance. Here, the overlap factor between cost and emissions is near

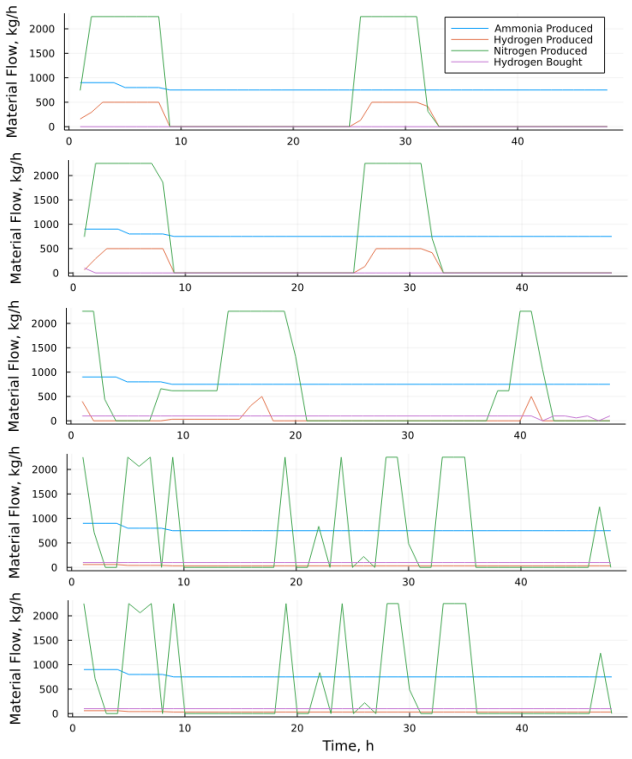

Figure 3: Optimal ammonia system production schedules when optimizing (from top to bottom) operating cost, carbon emissions (in-phase instance), carbon emissions (anti-phase instance), water usage, and electrolyzer safety.

Table 1: Results of analyzing 5 instances of historical power price and average grid emissions.

Dates	O_{ZH}	S^A_{ZH}	**Obj. Grouping**
Jan. 10-11	0.553	0.987	$\{Z,H\},\{\Psi,\Xi\}$
March 9-10	0.600	0.986	$\{Z,H\},\{\Psi,\Xi\}$
June 18-19	0.601	0.946	$\{Z,H\},\{\Psi,\Xi\}$
July 14-15	0.335	0.992	$\{Z,H,\Psi\},\{\Xi\}$
Sept. 1-2	0.007	0.832	$\{Z\},\{H,\Psi,\Xi\}$

zero, and the objective correlation strength is similarly low. In this instance, despite high solar generation in the midday lowering average grid emissions, power costs were still quite high due to the extreme demand of power. Like in the anti-phase instance tested earlier, the disconnect between times of high power costs and high average emissions caused the algorithm to group emissions with water usage and electrolyzer safety instead of with cost. Interestingly, the July 14-15 instance gave a unique objective grouping of operating cost, emissions, and water usage. In this instance, we saw that all objective correlation strengths were quite high (above 0.9), indicating that although we asked the Leiden algorithm to return two objective groups, we may have only needed one. Overall, the results show that in general, the most informative objective groupings that give maximum tradeoff information can change in time in response to changes in problem parameters.

Conclusion

In this work, we demonstrated the variability of objective correlation strength as dynamic parameters, namely time-varying power cost and grid emissions factor profiles, varied. Correlation results were shown to make practical sense, with in-phase profiles leading to high correlation strength between operating cost and carbon emissions, and anti-phase profiles leading to low correlation strength. Moreover, the *a priori* objective reduction results aligned with the optimal schedules found considering the different single objectives of operating cost, carbon emissions, water usage, and electrolyzer safety. In our future work, we will attempt to build a high-throughput machine learning approach to build classifier models that can predict when different objective groupings are more informative based on the optimization problem parameters. We will also look to build a tool to identify a single objective grouping that performs robustly well over the full range of possible dynamic parameter values, such that we do not need to worry about changing objective groupings in time.

References

Allman, A. and P. Daoutidis (2018). Optimal scheduling for wind-powered ammonia generation: Effects of key design parameters. *Chemical Engineering Research and Design*, 5–15.

Brockhoff, D. and E. Zitzler (2009). Objective reduction in evolutionary multiobjective optimization: Theory and applications. *Evolutionary Computation 17*, 135–166.

de Freitas, A. R. R., P. J. Fleming, and F. G. Guimares (2015). Aggregation trees for visualization and dimension reduction in many-objective optimization. *Information Sciences 298*, 288–314.

Emmerich, M.T.M.; Deutz, A. (2018). A tutorial on multiobjective optimization: fundamentals and evolutionary methods. *Nat Comput 17*.

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

Russell, J. M. and A. Allman (2022). Sustainable decision making for chemical process systems via dimensionality reduction of many objective problems. *AIChE Journal*, Under Review.

Saxena, D. K., J. A. Duro, A. Tiwari, K. Deb, and Q. Zhang (2013). Objective reduction in many-objective optimization: Linear and nonlinear algorithms. *IEEE Transactions on Evolutionary Computing 17*, 77–99.

van Eck N.J., T. V. W. L. (2019). From louvain to leiden: guaranteeing well-connected communities. *Scientific Reports 9*(5233).