Analysis of the correlating or competing nature of cost-driven and emissions-driven demand response

Hongxuan Wang[a], Andrew Allman[a,*]

[a]*Department of Chemical Engineering, University of Michigan, Ann Arbor, MI 48109, USA*

Abstract

Multi-objective optimization is an important tool for sustainable decision making. In process operations, the nature of tradeoffs between objectives may change over time, due to, for example, changes in power cost and emissions signals. This work seeks to understand the correlating vs. competing nature of demand response objectives, and their variation in time, through a time-varying ammonia production case study. Results show that correlation between cost and emissions is near perfect when the time-varying cost and emissions factor profiles are in phase, and low when they are anti-phase. Furthermore, analysis of real historical data from CAISO and ISO-NE suggest that for the majority of the time, cost and emissions objectives are strongly correlated. Trends observed are robust over different times of the year and different locations in the same ISO; however, considering ISOs in different parts of the country can provide different results due to variation in generator mix.

Keywords: multi-objective optimization, demand response, network theory, sustainability

1. Introduction

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. While the three pillars of sustainability, which state that a sustainable process should have positive economic, environmental, and social outcomes can provide guidance in of formulating objectives for sustainable decision making, rigorous analysis requires looking at a larger number of objectives corresponding to finer details due to the possibility of tradeoffs within each of the pillars (i.e. greenhouse gas emissions and water usage within the environmental pillar). Recent research by our research team has developed a method for scalable analysis of such "many-objective" optimization problems via the identification and grouping of correlated objectives which point towards similar solutions (Russell and Allman, 2023). While solving a multi-objective problem results in a set of tradeoff solutions known as a Pareto frontier, rather than a single decision, analyzing the correlating vs. competing nature of objectives is valuable as it allows for better understanding of the sustainability tradeoffs existing in decision making for process systems, and enables the determination of informed values towards different sustainability outcomes, development of effective policies which promote certain sustainability outcomes, and advancement of new technologies which can break through existing Pareto curves.

*Corresponding author
Email address: allmanaa@umich.edu (Andrew Allman)

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 sustainability tradeoffs and 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 Pareto frontier for sustainable decision making should not be static, but instead evolve dynamically based on the sensitivity of objective correlations with dynamically varying problem parameters. Of particular interest are the effects of the cost of purchasing power from the grid, which is a known driving force for intermittent demand response operation of flexible chemical processes, as well as the associated greenhouse gas emissions of that grid power, which could also be a driving force for demand response operation if an environmental objective is considered. While much prior work, some of which we review in the following section, has analyzed demand response operation of chemical process systems in intermittent electricity markets, analysis of emissions-driven demand response and its relationship with the cost-optimal solution has not been explored in depth.

In this work, we present an initial effort towards understanding the correlating or competing nature of cost-driven and emissions-driven demand response in chemical process systems. This is done within the context of an ammonia production operation case study, which is a good benchmark problem for considering chemical demand response operation. 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 sensitivities 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. In particular, we also include for consideration two practically important time-invariant objectives that are demonstrably correlated: water usage (an important environmental consideration) and hydrogen safety (an important social consideration) in order to understand when optimizing cost and emissions also optimizes these goals, as well as to better quantify the significance of correlation vs. competition between the time-varying objectives of cost and emissions. We note that this renders the problem under consideration many-objective in nature, requiring the use of an objective dimensionality reduction tool, such as that proposed in Russell and Allman (2023), for effective analysis.

The remainder of this paper is organized as follows: in the next section, we provide a brief review of recent and important works in both chemical demand response and multi-objective optimization. Then, we provide details about our approach for systematically grouping objectives in many-objective problems *a priori* to generating a Pareto frontier. Next, we present the problem formulation for many-objective ammonia production. Results of performing our objective correlation algorithm and analyzing process schedules from

both illustrative examples and real historical power price and emmissions data are discussed in the following section, along with an approach for determining compromise schedules when cost and emissions are in competition. Finally, we conclude by summarizing important observations and proposing avenues for future work.

2. Background

2.1. Demand Response Operation

The past several years have demonstrated a transition in global energy production infrastructure away from baseload fossil fuel sources such as coal and towards intermittent renewable sources such as wind and solar. As grid-scale energy storage technologies are still at a nascent stage, the need to manage this intermittency in the supply side of the energy market provides an impetus for intermittency on the demand side, via flexible and time-varying use of power. This idea is referred to as demand response, an idea that has been the subject of much recent research and shown to be a critically important paradigm in improving grid reliability when intermittency is present. Here, time-varying electricity prices provide a driving force for users of energy to shift their loads away from high price times (where presumably supply struggles to meet demand) and towards low price times (where supply is presumably in excess of demand). For industrial-scale chemical production, this idea represents a paradigm shift from traditional modes of operation, whereby for continuous processes, it has traditionally been thought as ideal to attempt to operate at a maximum capacity steady state for as often as possible. In contrast, demand response operation aims to exploit the electricity cost "driving force" to achieve operating cost savings that can overcome the increased capital investment required to build in the requisite flexibility.

Research over the past several years in optimal process operations has demonstrated the potential economic benefit of demand response operation for several industrially significant chemical processes, including cryogenic air separation (Karwan and Keblis, 2007; Zhang et al., 2015; Tsay et al., 2019; Caspari et al., 2019), the chlor-alkali process (Babu and Ashok, 2008; Brée et al., 2019; Otashu and Baldea, 2019; Weigert et al., 2021), and ammonia production (Allman and Daoutidis, 2018; Allman et al., 2019; Kelley et al., 2022; Hochhaus et al., 2023). More recent advances in the chemical demand response literature have analyzed coupled design and operation of flexibly operating processes (Demirhan et al., 2021; Palys et al., 2023), demand response dynamic needs from a grid perspective (Gao et al., 2022; Tang and Baldea, 2023), and coordination of production networks with multiple actors participating in demand response (Leenders et al., 2021; Allman and Zhang, 2022). The need for consideration of demand response scheduling is expected to further increase in the future due to trends towards electrification of chemical processes, which has been proposed as a mechanism for decarbonizing the industry (IEA, 2020).

When considering sustainable process operations, as in this work, economic driving forces cannot be the only factor taken into account in decision making. In most processes, one would expect greenhouse gas emissions to be an essential objective to minimize within the environmental pillar. It is well known that, much like electricity prices, the associated emissions with power purchased from the grid also varies dynamically (Kopsakangas-Savolainen et al., 2017), which inherently provides an environmental driving force for demand response. Importantly, due to spatiotemoporal variabilities in renewables supply and power demand, as well as the fact that the merit order of generator dispatch need not be correlated with the emissions associated with that generator, there is no guarantee that the economic and environmental signals towards demand response will be related to each other. A few recent demand response adjacent studies have considered the potential effects time varying emissions, including one considering the design of a flexibly operating low carbon utility system (Baumgartner et al., 2019), one considering the effect of price and emissions driven demand response on load shifting and carbon emissions at a grid scale (Fleschutz et al., 2021), and one that utilizes time-varying carbon emissions in the expansion planning of a power grid in a

developing country (Khan, 2018). However, to the best of our knowledge, there currently exist no studies that examine purely emissions-driven demand response operation of an industrial process and analyze its competing vs. correlating nature with traditional cost-driven demand response.

2.2. Many-Objective Optimization

Many-objective optimization problems are a subset of the better known multi-objective optimization problems, and typically refer to problems of four or more objectives. These problems require special attention as they are particularly challenging to deal with in comparison to smaller, two or three objective problems. First, humans inherently struggle to understand manifolds in higher than three dimensions, making the Pareto frontiers of many-objective problems challenging to interpret and use. More practically, common computational approaches for rigorously generating the Pareto frontier encounter tractability challenges for many-objective problems. Typically, a rigorous approach for multi-objective optimization entails using a scalarization method such as weighted sum (Marler and Arora, 2009) or epsilon constraint (Giagkozia and Fleming, 2015) to generate a set of single-objective problems that provide single points along the Pareto frontier, and then solving these single objective problems using the appropriate state-of-the-art global solver, such as CPLEX (Cplex, 2020), Gurobi (Gurobi Optimization, 2018), or BARON (Kılınç and Sahinidis, 2018). However, the number of single-objective problems to solve scales exponentially with number of objectives. This provides such a challenge that heuristic approaches, most commonly 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 reduction *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, 2023). 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 in objective grouping in multiple energy system design and supply chain optimization case studies.

3. Objective Reduction Algorithm

In this section, an algorithm for systematically reducing the dimensionality of (mixed integer) linear programs in a way that preserves tradeoff information *a priori* to solving for the Pareto frontier is presented. The algorithm was originally developed in Russell and Allman (2023), and a graphical representation of the

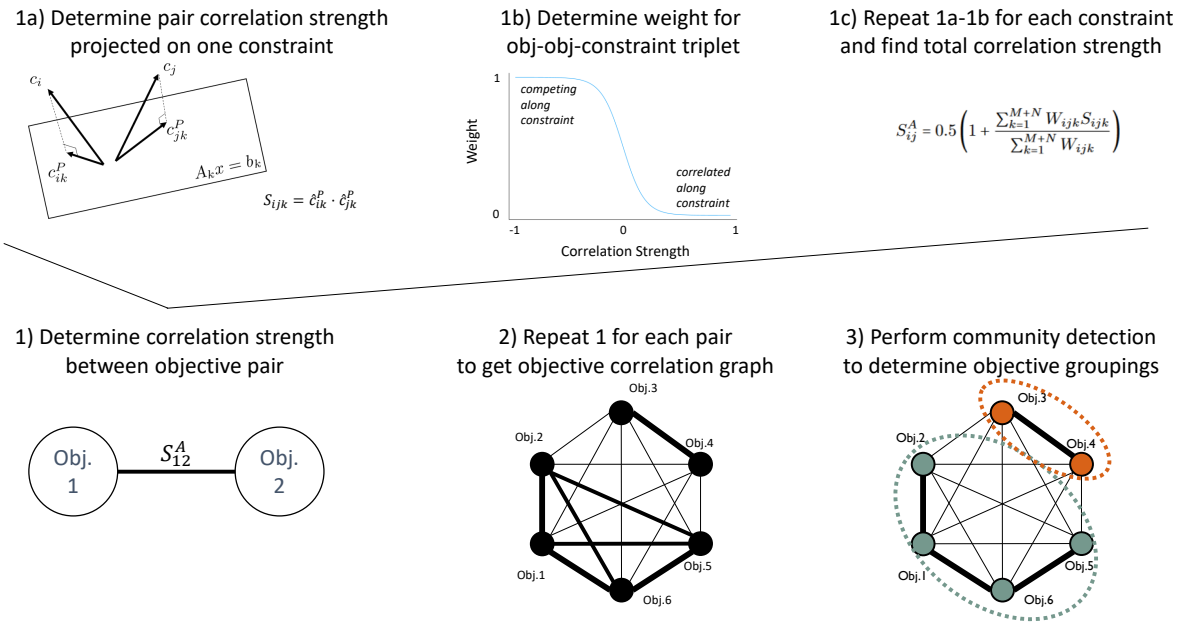

Figure 1: Graphical depiction of objective reduction algorithm described in Section 3.

algorithm is shown in Figure 1. Consider the following general linear N-objective optimization problem:

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

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}:

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

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^A_{ij}, we consider a weighted sum of the constraint interaction strengths. Weights W_{ijk} are zero when both normal vectors c^N_{ik} and c^N_{jk} 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:

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

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 step-like change. In this work, values of $\alpha = 1$ and $\beta = 10$ are used. The total objective interaction strength is determined as follows:

$$S^A_{ij} = 0.5 \left(1 + \frac{\sum_{k=1}^{M+N} W_{ijk} S_{ijk}}{\sum_{k=1}^{M+N} W_{ijk}} \right) \tag{6}$$

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.

For an optimal real-time operation, scheduling, or control problem, the cost vectors c and constraint vectors a are typically parameters that vary in time and will be different for each solve of the problem throughout the moving horizon. Since the objective reduction algorithm is clearly sensitive to these values, it should be expected that the objective correlation graph and will change at each time point in the moving horizon. Thus, in order to reduce objective dimensionality in a way that provides as much information on operating tradeoffs as possible, it is essential to understand the sensitivity of correlation strengths to dynamically varying parameters.

4. Case Study Information

4.1. System Description

The ammonia production system analyzed in this work is adapted from previous works (Allman and Daoutidis, 2018), depicted in Figure 2, 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

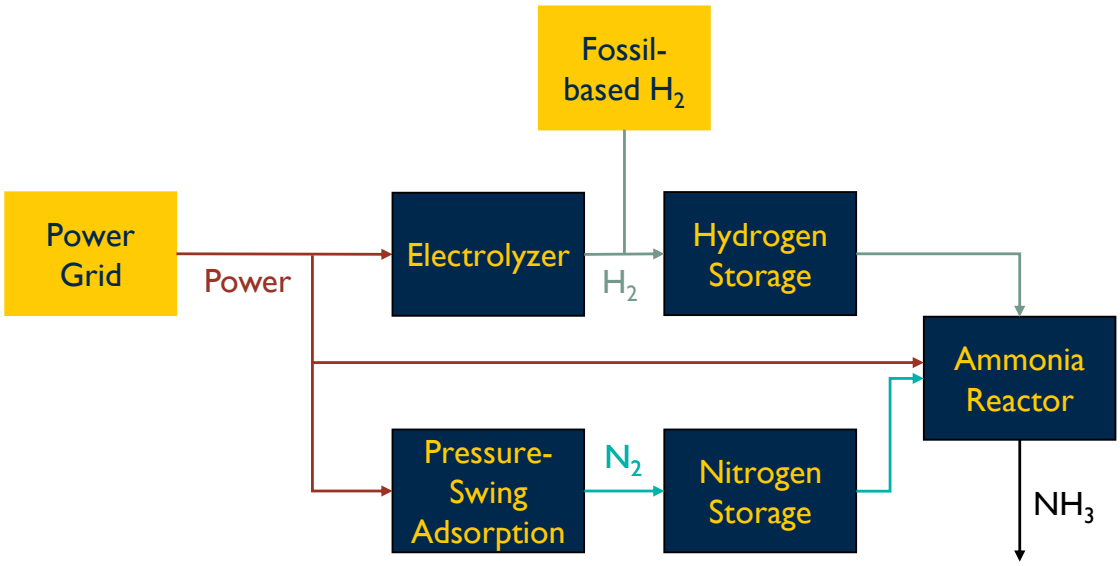

Figure 2: Block diagram of ammonia production process considered in this work.

Haber process. Energy for all process is obtained from purchases from the power grid. Both the costs of these purchases and their associated carbon emissions are time-varying parameters. A four-objective optimization problem is used to determine production rates of nitrogen, hydrogen, and ammonia, as well as purchases of hydrogen from an external steam reformer, on an hourly basis over a 48-hour operation horizon in order to meet a nominal demand of ammonia. The optimization problem is solved repeatedly each hour for the future 48 hour period. In practice, system operating inputs would be determined from the decisions made in the first hour of the 48 hour horizon; however, for this work, we are concerned only with the evolution of correlating vs. competing nature of the four objectives in each individual 48-hour optimization horizon.

4.2. Objective Functions

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:

$$Z = \sum_{t=1}^{48} (\zeta_b(t) p_b(t) + \zeta_{H2,buy} b_{H2}(t) + \zeta_{on} s_e(t) + \zeta_{el} m_{H2}(t) + \zeta_{psa} m_{N2}(t)) \tag{7}$$

Cost contributions come from purchasing power with the grid (p_b), purchasing hydrogen from the steam reforming plant (b_{H2}), starting up the electrolyzer (where s_e is the number of electrolyzers starting up), operating the electrolyzer (where m_{H2} is the amount of hydrogen produced by electrolysis), operating the PSA unit (where m_{N2} is the amount of nitrogen produced by PSA). Note that artificial cost terms, including penalties for deviation from target ammonia production rates and terminal costs for depleting gas from storage, considered in the cost objective in previous works, are not included here. Instead, a constraint which lower bounds the total ammonia produces (described in the next subsection) is included. This enables a more fair comparison of schedules between costs and other objectives, which inherently do not have these artificial terms. 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:

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

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.

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

Water is used by the system in the production of hydrogen, both in electrolysis and 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 (Center for Chemical Process Safety, 1995) :

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

All four objectives considered are of practical significance when considering the economic, environmental, and social performance of the ammonia production system. We do not claim that these objectives are proxies for every possible environmental or social consideration one may wish to optimize with this system, but rather an example of a few objectives that, considered together, enables analysis of various important tradeoffs. The first two objectives (operating cost and carbon emissions) will see time-varying signals that provide a driving force towards demand response operation due to the system obtaining power from the grid. These signals will also change over different solves within a moving horizon scheduling framework. The latter two objectives (water usage and electrolyzer safety) do not see any time-varying signals; as such, many degenerate solutions with equivalent optimal objective values exist. Note that while we are most interested in how the correlating or competing nature of cost and emissions changes over time in this analysis, the inclusion of the latter two objectives enables the use of our systematic objective reduction tool that determines when competition between cost and emissions is statistically significant enough to consider them separately in relation to two known, time-invariant, strongly correlated objectives. In particular, since the parameters for water usage and electrolyzer safety do not change over different solves in the moving horizon framework, their correlation strength acts as a static baseline for comparison for the computational experiments since it will remain unchanged in time. We further expect that the correlation strength between water usage and electrolyzer safety to be high, as both promote not using the electrolyzer whenever possible. In contrast, if one was to only consider the correlation strength between cost and emissions, the cutoff point for when these objectives are considered to be competing versus correlating would be arbitrary.

4.3. Model Constraints

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:

$$p_b(t) = p_{N2}(t) + p_{H2}(t) + p_{NH3}(t) \quad \forall\, t \in \{1, ..., 48\} \tag{11}$$

Here, the amount of power bought from the grid p_b must equal the amount of power 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:

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

For the purposes of determining objective correlation weights, constraints (11) and (12) are used to replace the power bought from the grid with a linear combination of chemical production rates in the cost (7) and

emissions (8) objective functions. This provides an equivalent formulation which helps to prevent obfuscation of objective interactions among constraint surfaces that can occur due to the introduction of an auxiliary variable p_b.

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

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

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:

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

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

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:

$$-d_a(t) R^- \leq m_{NH3}(t) - m_{NH3}(t-1) \leq d_a(t) R^+ \quad \forall\, t \in \{1, ..., 48\} \tag{22}$$

$$\sum_{\tau=t+1-T}^{t} d_a(\tau) \leq D_a \quad \forall\, t \in \{1, ..., 48\} \tag{23}$$

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, in place of the artificial cost penalty for deviating from ammonia target production rate, we require that the total ammonia produced in the 48-hour period exceeds a lower

bound (m^{lo}_{NH3}) determined by the ammonia demand:

$$\sum_{t=1}^{48} m_{NH3}(t) \geq m^{lo}_{NH3} \tag{24}$$

The total four-objective optimization problem is as follows:

$$\begin{aligned} \min \quad & \{Z, H, \Psi, \Xi\} \\ \text{s.t.} \quad & (11) - (24) \end{aligned} \tag{25}$$

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.

4.4. Model Parameters and Data

The majority of the parameters used for the above model are equivalent to those considered in previous work on this system (Allman and Daoutidis, 2018), with exceptions and values of new parameters noted in this section. 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. Finally, the lower bound for total ammonia production in the 48-hour period is set to 45000 kg, equal to a 6.25% decrease from constant operation at the nominal production rate of 1000 kg/h which allows for a slightly larger degree of operating flexibility in the system.

Many different time-varying cost and emissions signals are analyzed in this work. The first 2 instances considered (section 5.1 provide an illustrative example of extreme cases of correlation and competition, using artificial, researcher-generated signals whereby cost and emissions are perfectly in-phase, or perfectly anti-phase. Remaining simulations consider month-long operating periods and analyze correlation strengths and objective groupings for the 48-hour scheduling optimization problems starting at each hour in the month. In all problems, it is assumed that the production rate of the ammonia reactor is initially at its nominal rate of 1000 kg/h, no electrolyzers are turned on, and the storage tanks for hydrogen and nitrogen are empty. Keeping these values constant provides a fair comparison between scheduling results at different time points, although we note that changing these initial conditions has no effect on objective correlations. Costs and emissions factors for grid-generated power within the US are taken from real historical locational marginal price and average grid-wide emissions data from various locations and times in the California Independent System Operator (CAISO) (California ISO, 2023) and Independent System Operator New England (ISO-NE) (ISO New England, 2023). Power prices and emission intensities in Tokyo are obtained from Japan Electric Power eXchange (JEPX) (Japan Electric Power eXchange, 2023) and Electricity Maps (Electricity Maps, 2023). The same data from the UK are also obtained from Nord Pool (Nord Pool, 2023) and Carbon Intensity API (Carbon Intensity API, 2023). Note that although recent works have suggested that marginal, rather than average, emissions provide a better accounting for the emissions effects of demand response at the grid scale (Fleschutz et al., 2021), these data are not readily available from most ISO databases, and we believe that at a process level, which is the consideration of this work, average grid emissions provide a more equitable accounting of emissions incurred between different grid users. However, our algorithmic approach for determining cost-emissions objective correlation is flexible and agnostic as to whether the emissions signals provided are average or marginal emissions, allowing for equivalent analysis to be done

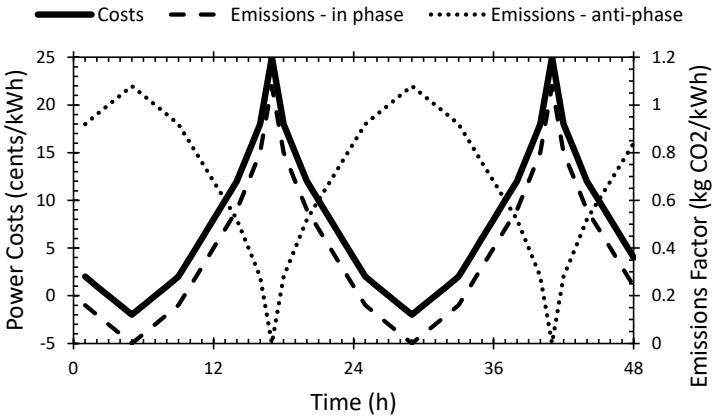

Figure 3: Plots of power costs and emissions factors over time for instances where costs and emissions are in-phase (dashed line) and anti-phase (dotted line).

with calculated marginal emissions signals. An analysis using marginal emissions data is also provided, using data obtained from MISO (Midcontinent ISO, 2023).

5. Results and Discussion

5.1. Illustrative Examples

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 as follows: power costs follow a 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. Plots of power costs and carbon emissions for the illustrative in-phase and anti-phase examples are depicted in Figure 3.

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 4. 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 the average grid emissions increases slightly in the out-of-phase case for this particular profile. Higher emissions from grid energy makes electrolysis look less attractive in comparison to steam reforming, which is well aligned with the water usage and safety objectives which seek to avoid electrolysis. Note also that in the cost-optimal solution in this instance, we somewhat unintuitively still produce hydrogen from electrolysis, taking advantage of the times where power

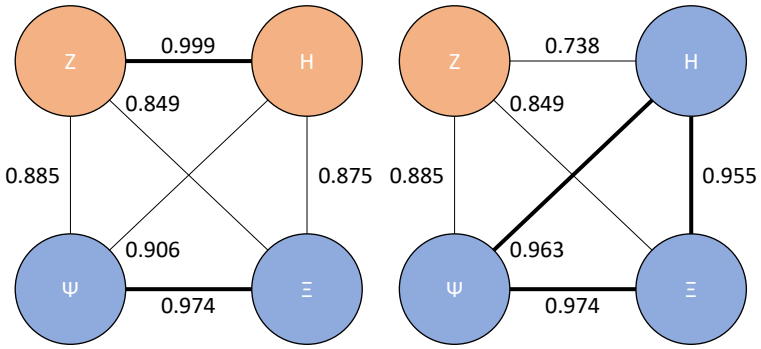

Figure 4: Objective correlation graph with color-coded objective groupings for in-phase (left) and anti-phase (right) instances.

costs are extremely low or negative to produce all the hydrogen needed over the 48 hour period, some of which can be stored for later use.

To make physical sense of the objective groupings, we show the single objective optimal schedules in Figure 5. 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 and consistently purchases the maximum amount of hydrogen from the steam reformer due to higher grid emissions..

5.2. Analysis of Historical Data

Based on the power price and emission data obtained from CAISO during the time period from Aug 1^{st} to Aug 30^{th}, 2022, a general analysis of the correlating versus competing nature of the cost and emission objective is conducted. The specific node used for price data is 0096WD_7_N001. Figure 6 displays the evolution of the correlation strength between the cost and emissions data for the 48 hour operation optimization problem beginning on the corresponding hour. It is evident that, for most of the time, the correlation factor exceeds 0.96. This high correlation strength suggests that the cost and emission objectives exhibit similar operating profiles during August 2022, and little tradeoffs exist between the cost-optimal and emissions optimal solutions. However, the value in the middle of the month is relatively low, and there is a 11 hour period where the grouping results shown from community detection indicate that operating cost and emissions should not be grouped together, with a resulting grouping of [Z],[H,Ψ,Ξ], equivalent to the grouping seen in the anti-phase case of the illustrative example. During all other times, the community suggested objective grouping is [Z,H,Ψ],[Ξ]. Importantly, this suggests that optimizing demand response operation on the basis of cost leads to very little degredation in emissions compared to the emissions-optimal case.

The observed result that cost and emissions are correlated the majority of the time could be due to the fact that solar energy is a large portion of the total energy generated in California, and in day time, solar energy is mainly used as a source of energy with large amount of power generated, making the power price as well as emission relatively low. Interestingly, we see here that water usage and electrolyzer safety, our time-invariant objectives, are grouped separately most of the time despite being highly correlated. This is

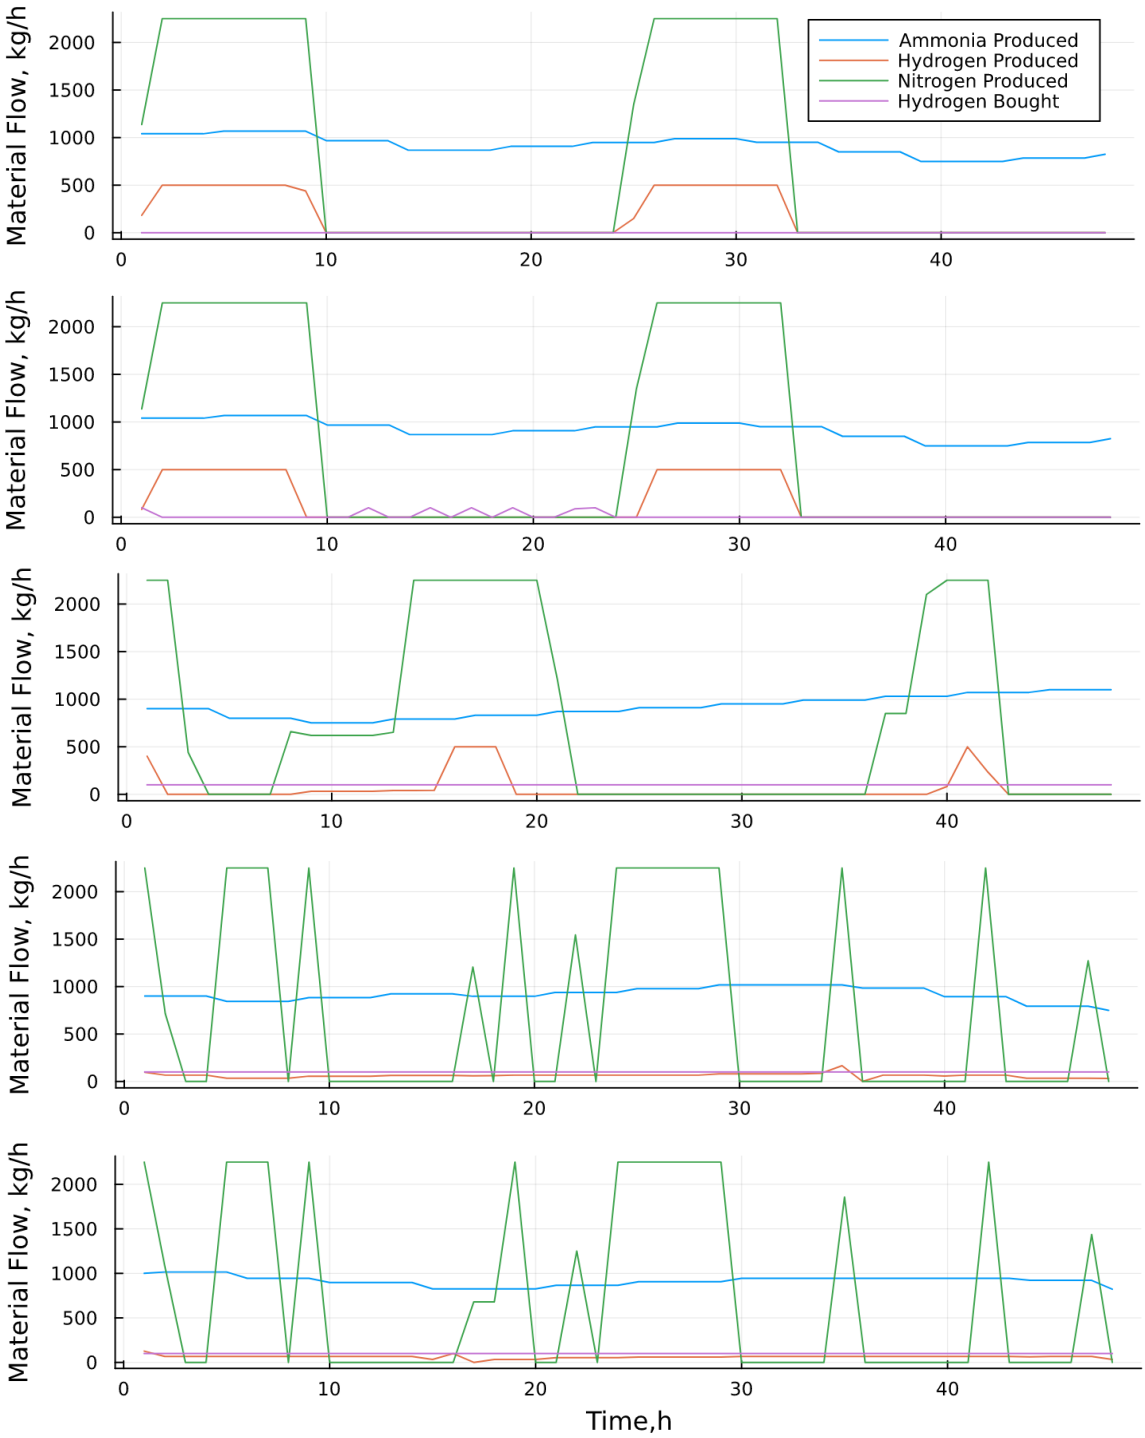

Figure 5: Optimal ammonia system production schedules when optimizing (from top to bottom) operating cost, carbon emissions (in-phase instance), carbon emissions (anti-phase instance), water usage, and electrolyzer safety.

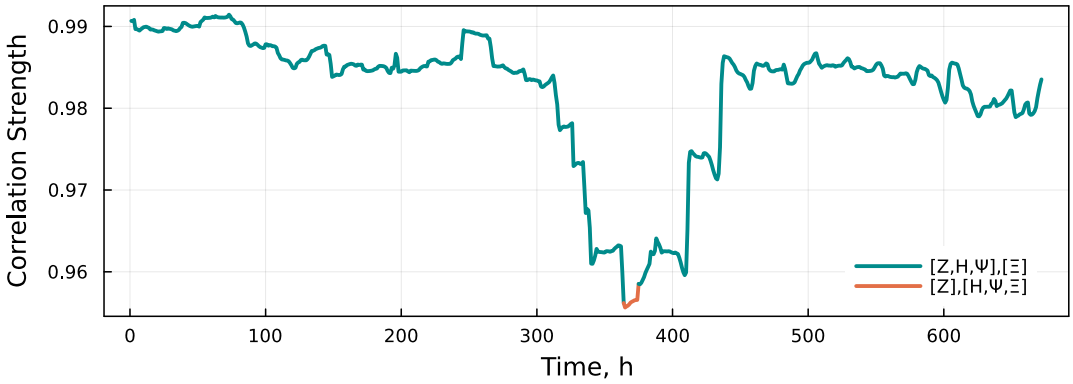

Figure 6: Correlation strength between cost and emission for 48 hour scheduling optimization problems starting between Aug 1^{st} and Aug 30^{th} 2022, with objective groupings determined by community detection denoted by color.

likely an artifact of our tuning of the community detection resolution hyperparameter to ensure that two nonempty groups are returned, and suggests that in the majority of times analyzed, all four objectives are likely to be strongly correlating.

To verify the correlation strength and grouping observations, we next analyze cost and emissions optimal schedules for two representative instances where correlation is observed to be very high, and very low. Figure 7 depicts the schedules for the time period between Aug 3^{rd} 3 pm to Aug 5^{th} 2 pm, which has the largest observed correlation strength of 0.991. The top figure demonstrates the relatively in-phase nature of the cost and emissions signals. The results of using these signals to optimize the operating schedule are shown in the middle (cost-optimal) and bottom (emissions optimal) schedules. It is apparent that the majority of production occurs during periods of low grid prices or emissions, resulting in minimal costs or emissions. With the highest correlation strength, good agreement of the demand response is shown between these two plots, with only slight differences due to the price signal having slightly wider troughs and emissions signal having slightly wider peaks. The observed cost and emissions from each case further support this observation: in the optimal cost case, costs and emission of \$40,974 and 130,247 kg CO_2, respectively, are observed, while in the optimal emissions case, costs and emissions of \$41,500 and 127,506 kg CO_2, respectively, are observed. We see that optimizing for cost results in only a 2.1% increase in emissions from its optimal value, and optimizing for emissions results in only a 1.3% increase in cost from its optimal value, thus supporting the correlation strength observations.

Next, a 48-hour period with low correlation strength is considered. Figure 8 depicts the schedules for the time period between Aug 16^{th} 6 am to 18^{th} 5 am, which has the lowest correlation strength of 0.956 that results in cost and emissions being placed in different objective groupings. Critically, the price and emissions profiles demonstrate the between hours 6 and 10, power price sees a spike while emissions are at a low point, which is the main driver of competition between the two objectives. The results of this are evident in the optimal schedules; while production looks similar most of the time, during hours 6-10 the cost optimal solution scales back most production while the emissions optimal solution continues to make nitrogen, hydrogen, and ammonia at a high level. The results of this competition manifest in the observed objective values, as in the optimal cost case, cost and emissions are \$47,001 and 131,943 kg CO_2, respectively, while in the optimal emissions case, cost and emissions are \$54,115 and 125,202 kg CO_2, respectively. Clearly, tradeoffs between the two objectives are much wider than in the strongly correlated case: optimizing for cost results in a 5.4% increase in emissions from its optimal value, whereas optimizing

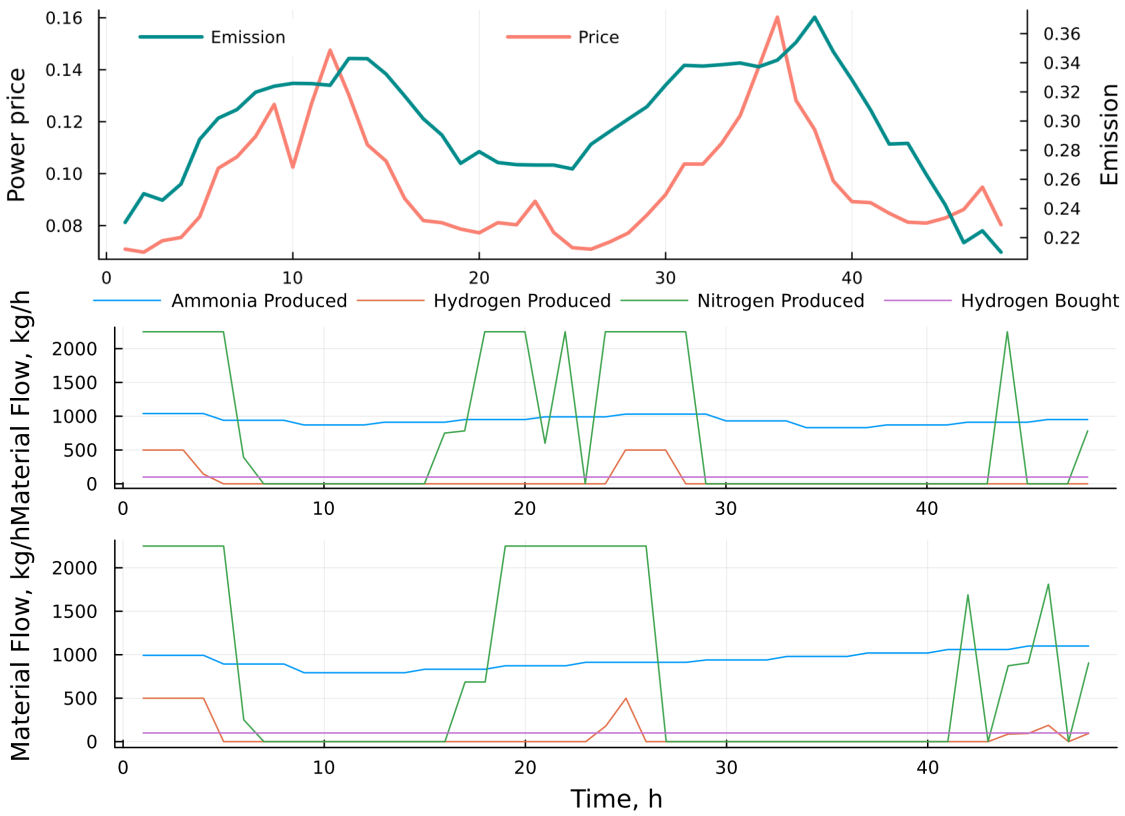

Figure 7: Scheduling results from Aug 3^{rd} 3 pm to 5^{th} 2 pm. (top) Power price (\$/kWh) and emissions factor (kg CO_2/kWh). (middle) Optimal production schedule from optimizing operating cost. (bottom) Optimal production schedule from optimizing carbon emissions.

for emissions results in a 15.1% increase in cost from its optimal value.

All the scenarios demonstrate that the grouping approach using correlation strength can effectively identify the relationships between different objectives. Operating cost and carbon emission objectives are grouped together for more than 90% of the total calculated time period, indicating that the price and emission parameters in California at the specific node analyzed are usually highly correlated with each other. The very high frequency of [Z,H,Ψ],[Ξ] also suggests that, at least during the time period and at the location studied, the ammonia production system under considerations will be able to inherently keep carbon emissions and water usage low at most times simply by minimizing their costs, without any additional economic incentives required to do so.

To validate the utility of our proposed objective grouping strategy, we compare it to the more straightforward approach of using the statistical correlation coefficient R comparing cost and emissions, found using the covariance and standard deviations σ of 48-hour power cost and emissions data as follows:

$$R = \frac{\mathrm{Cov}(\zeta_b(t), \eta_{p_b}(t))}{\sigma(\zeta_b(t))\sigma(\eta_{p_b}(t))} \tag{26}$$

The correlation coefficient R is a measure of linear correlation between the paired cost and emissions signals over the 48-hour optimization period, where a value of one indicates perfect linear correlation (with positive slope) and a value of zero indicates a lack of linear dependence between the signals. A comparative analysis

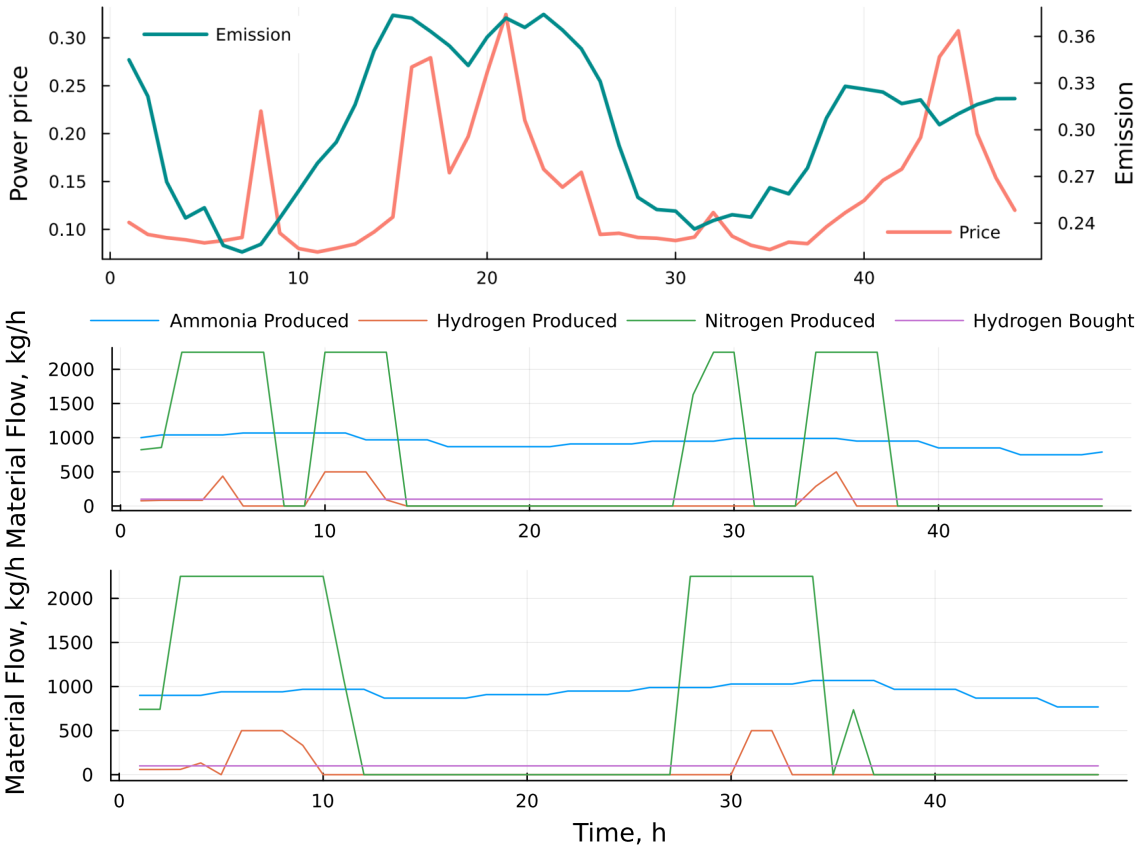

Figure 8: Scheduling results from Aug 16^{th} 6 am to 18^{th} 5 am. (top) Power price ($/kWh) and emissions factor (kg CO_2/kWh). (middle) Optimal production schedule from optimizing operating cost. (bottom) Optimal production schedule from optimizing carbon emissions.

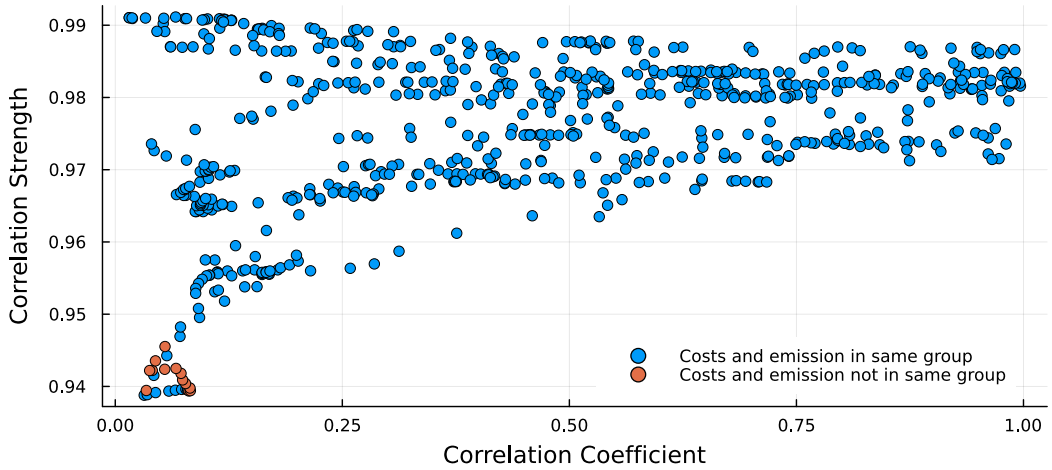

Figure 9: Correlation coefficient vs. correlation strength in January 2022, Los Angeles.

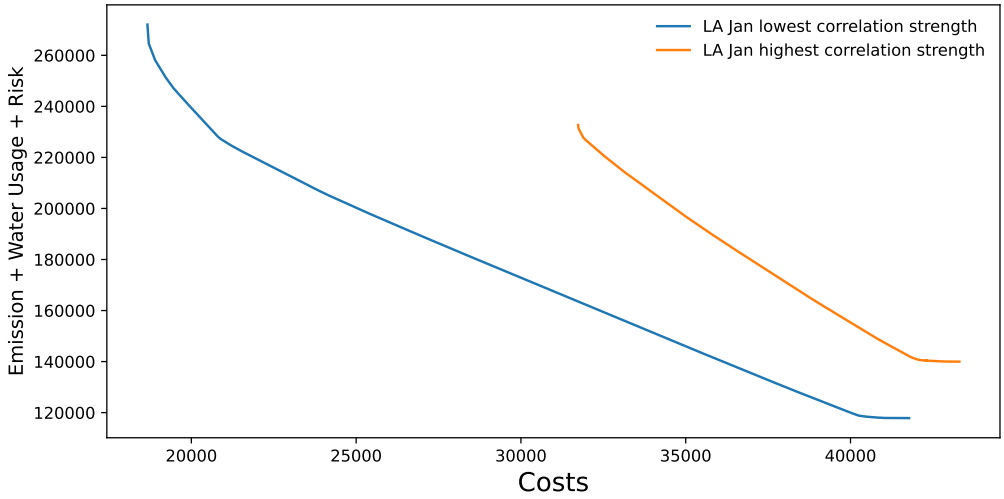

Figure 10: Pareto frontier with the lowest (blue) and highest correlation strengths (orange) in January 2022, Los Angeles.

was conducted by plotting these correlation coefficients with the correlation strength determined through our proposed method for every 48-hour period in January 2022 using prices from the Los Angeles node. This scatter plot is depicted in Figure 9. Notably, points where the two objectives did not align within the same group consistently have low values in both indices. However, it is apparent that correlation coefficients can approach zero even when our metric suggests strong correlation between the objectives, signifying a discrepancy between these indices. To analyze this further, the reduced space Pareto frontier at two points with the highest and lowest correlation strengths, as determined by our proposed metric, was generated. Note that both points have similar correlation coefficients R. As portrayed in Figure 10, it is evident that the our proposed method is effective at identifying when strong competition exists. Specifically, the trade-off between the cost-optimal and emissions-optimal solutions corresponding to the highest correlation strength is considerably less than its counterpart associated with the lowest correlation strength, despite similarity in R values. This observation underscores that the interrelation between cost and emissions optimal solutions is contingent not solely on the correlation between power price and emission intensity but also on the inherent numerical structure of the optimization problem; namely, the projections of the objective vectors onto constraint surfaces.

5.3. Comparing Times of the Year

For further analysis, we sought to understand how correlation and competition between cost and emissions could potentially change during different times of the year, , as the typical profiles for both could conceivably change based on time of year due to, for example, differing availabilities of renewables and consumer energy use patterns. To this end, real locational marginal price data from node MKT_SUB_LNODEONA, located in Los Angeles, is utilized. The following analysis focuses on examining the correlation between cost and emission at this spot during four different months of the year to better understand any seasonality effects: January, April, August, and October 2022. Figure 11 illustrates the correlation strength for each of these months. The curves for January, August, and October consistently exhibit high correlation values, exceeding 0.96 for most of the time. The grouping of [Z,H,Ψ],[Ξ], seen in the previous subsection which suggests that all objectives are strongly correlating, occurs over 90% of the time of these three months and also indicate that operating cost and emission objectives are mostly grouped together in these three months. In contrast,

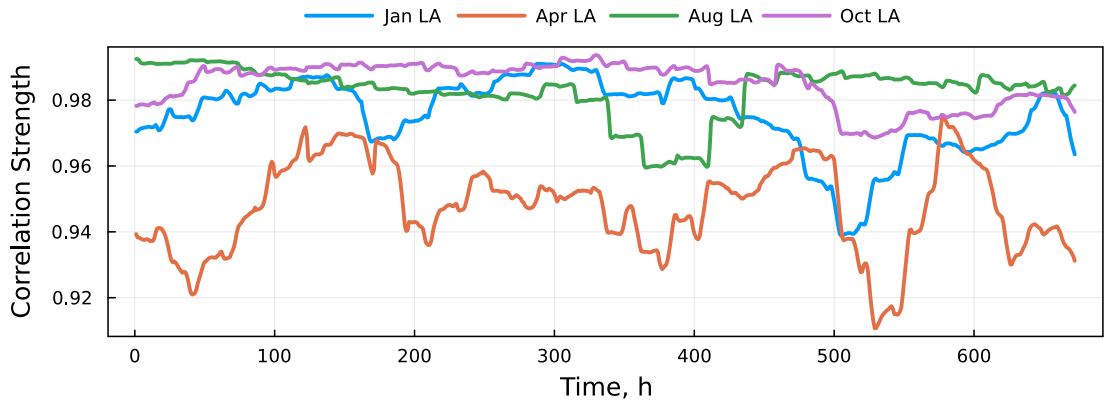

Figure 11: Correlation strength between cost and emission in Los Angeles with four different months in 2022. All are from 1^{st} to 30^{th}

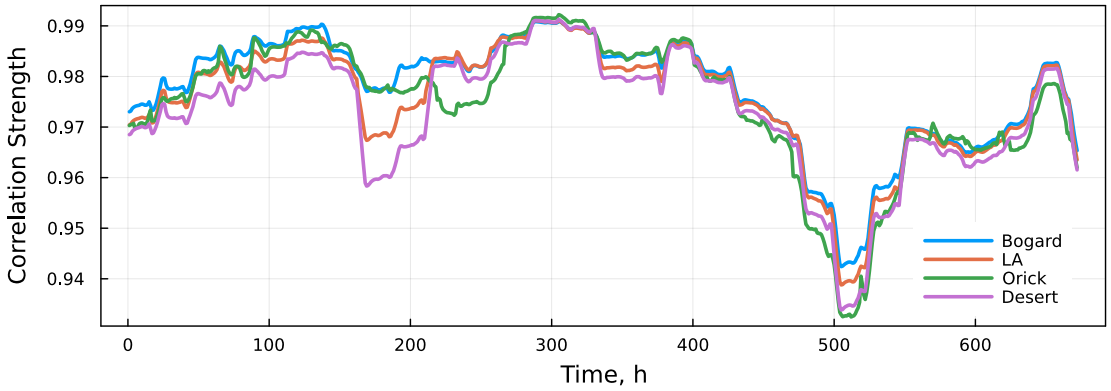

Figure 12: Correlation strength between cost and emission from Jan 1^{st} to Jan 30^{th} using price data from four different nodes in CAISO.

the correlation strength in April is lower, fluctuating around 0.94. Consequently, there are 57 instances where cost and emission are not grouped together, exhibiting a grouping of [H],[Z,Ψ,Ξ]. Moreover, the dominant objective grouping when cost and emissions are grouped together was [Z,H],[Ψ,Ξ], as seen in the in-phase illustrative example, which is different from the other three months. Further analysis of the cost and emissions data reveals that, on average, price and emission values in April were lower than in the other three months analyzed. This helps to explain why water usage is not grouped together with cost and emissions: with lower average signals, the systems tend to use the electrolyzer, which is the largest power user in the system, more for hydrogen production instead of purchasing hydrogen from the steam reformer.

5.4. Comparing Locations

As a final analysis, we sought to understand how correlation and competition between cost and emissions might change between locations within the same ISO, as well as between locations at different ISOs. Correlation strengths in different locations within or without CAISO are displayed in Figure 12. The study considers four specific locations within California: Los Angeles, a world-class megacity in the southern part of the state; Orick, a small coastal city in the northern part of the state; Bogard Buttes, situated at a high

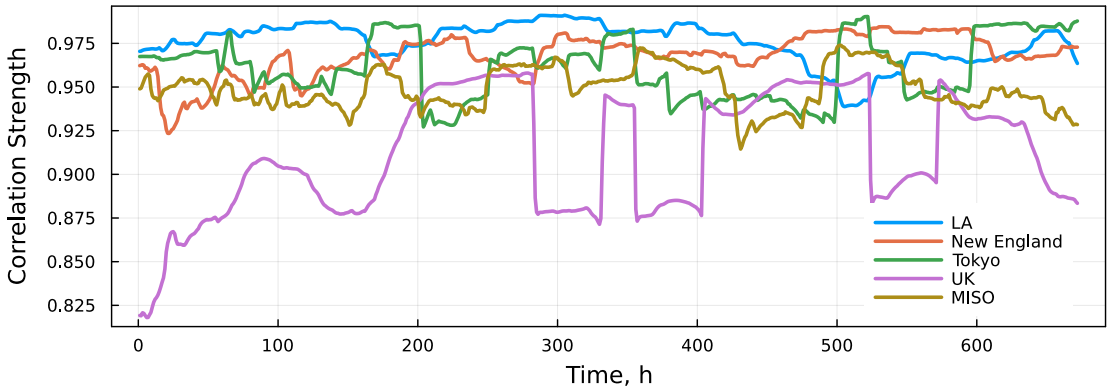

Figure 13: Correlation strength between cost and emission from Jan 1^{st} to Jan 30^{th} around the world.

elevation; and an isolated location in the desert. The diversity in population and climate in the four different locations enables the analysis of a wide range of different local weather patterns and demand profiles that could manifest in differing locational marginal prices. However, we consider that the emissions profile remains constant for all four locations since it is obtained by aggregate CAISO emissions data. Correlation strengths between the cost and emissions objective from Jan 1^{st} to Jan 30^{th} 2022 are calculated. While some slight variations can be observed among the different lines, they all exhibit a similar overall trend. Specifically, their correlation strengths are relatively low around the 22^{nd} of the month, indicating competition between cost and emission that results in cost and emissions being grouped in separate objectives. For the majority of January at all locations, cost and emissions objectives are grouped together. When looking deeper into the locational marginal price data, we note that despite the diversity in locations analyzed, the prices do not differ significantly since all locations are within the same ISO.

However, when performing the same analysis for a location outside of California, the situation becomes different. Correlation strength calculation is conducted using emission data from ISO-NE and power price from node .H.INTERNAL_Hub, located in the middle of Massachusetts. The results, depicted in Figure 13, show a significant difference compared to the results in California. The correlation strength in New England is lower than that in California and the grouping results are also distinct. While cost and emissions are still grouped together for the majority of the time, a substantial number of instances (227 out of 672) indicate competition between cost and emissions. We note that mix of power generation sources in New England differs significantly from California, with generation occurring predominantly from natural gas, which accounts for over 50% of power supplied to the grid. Renewable energy sources have a lower proportion, around 10% to 15%, with solar energy constituting approximately 5% of the entire system. In contrast, in California it is common for renewables generation to account for over two-thirds of the total power generation, with solar accounting for two-thirds of the renewables generation. This discrepancy in resource composition could explain why cost and emissions are not as strongly correlated in New England as they are in California.

The interrelationships among the four defined objectives were further examined in MISO, considering available marginal emissions data, rather than average emissions as used at the other locations. As illustrated in Figure 13, the correlation strength observed in MISO appeared consistently lower compared to that observed in ISO-NE across a significant portion of the data set. Closer analysis revealed that the dominant grouping, occuring in 641 instances, was [Z, Ξ] and [H, Ψ]. Since we know water usage and electrolyzer risk are always strongly correlated and that we are constraining the community detection to give us two groupings, this suggests that most of the time, all objectives are strongly correlated. Among the remaining

31 groupings identified as [Z] and [H, Ψ, Ξ], competition between costs and emissions is discernible.

We further analyze the results from two electricity markets outside the United States, considering data from both the UK and Tokyo, as illustrated in Figure 13. The correlation strength observed in Tokyo closely approximates the MISO curve for a significant duration, punctuated by occasional increments where correlation strength is stronger, During the majority of times, costs and emissions are in competition, with an observed grouping of [Z],[H,Ψ,Ξ] accounting for 515 occurrences out of 672. In the remaining 157 instances, cost and emissions are grouped together with an observed grouping of [Z,H,Ψ],[Ξ]. The difference in grouping outcomes between Tokyo and MISO can be attributed to the correlation strengths with the time-invariant objectives due to the fact that power prices are consistently higher in Tokyo than in the MISO data. Finally, we note that consistently lower correlation strengths are evident in the UK. The two groupings observed the UK do not put cost and emissions in the same group: [Z],[H,Ψ,Ξ] and [H],[Z,Ψ,Ξ] results suggest sustained competition between cost and emissions. Looking deeper at the data, pronounced fluctuations in power prices within the UK significantly influence this diminished correlation strength. In particular, significant spikes in power prices, larger than those observed at any other locations, contribute to the times with the lowest observed correlation strengths.

5.5. Determining Compromise Schedules

The preceding sections provide a comprehensive analysis on the correlation vs. competition of cost-driven and emissions-driven demand response. When both cost and emission objectives align in the same group, the application of optimization methodologies is straightforward as optimizing one objective inherently gives satisfactory performance in the other. However, one key question remains: if costs and emissions are competing, how should one proceed in operation? The answer to this question is not straightforward, as in multi-objective optimization, all points on the Pareto frontier are mathematically equivalent and valid solutions, and choosing which to operate at will depend on the values of the relevant decision making stakeholders. However, here we propose a potential guiding strategy on the basis of identifying "knee points" on the Pareto frontier. These are points along the Pareto frontier where further emissions reduction is relatively more expensive than the emissions reduction achieved so far. Knee points can be quickly identified using the weighted sum approach for generating the Pareto frontier, as the individual solution will persist for multiple values of objective weights. In Figure 14, we illustrate the Pareto frontier depicting the trade-off relationship between cost and emissions in the UK in early January, where competition between cost and emissions was observed to be high. The size of the data points on the graph corresponds to the frequencies in which a Pareto-optimal solution is generated using the weighted sum method. Two of the most frequently occurring data points are strategically chosen for detailed analysis, representing knee points which provide compromise between cost and emissions. The optimal schedules at each of these points are shown in Figure 15, depicting the scheduling results for optimizing only cost, the first selected point, the second selected point, and optimizing only carbon emissions.

Optimizing for cost gives the largest production rates of nitrogen and hydrogen when prices are lowest as depicted in the scheduling results. Conversely, optimizing for carbon emissions prompts a shift in production rates towards conforming to the lowest emission intensity times. The scheduling trend observed for the first selected point, which is in closer proximity to cost, mirrors the behavior observed in cost optimization, albeit with a notable reduction in hydrogen production, providing a reduction in carbon emission reductions, particularly in instances where emission intensity is slightly elevated. Conversely, the second selected point, closer to optimizing for carbon emissions, portrays a scheduling plot akin to the lower graph, but with a few modified decisions which help to further lower operating costs. The scheduling plots derived from the selected points along the Pareto frontier reveal varied compromises between the two objectives, offering valuable insights to strike a balance between cost and emissions in alignment with the values of decision making stakeholders.

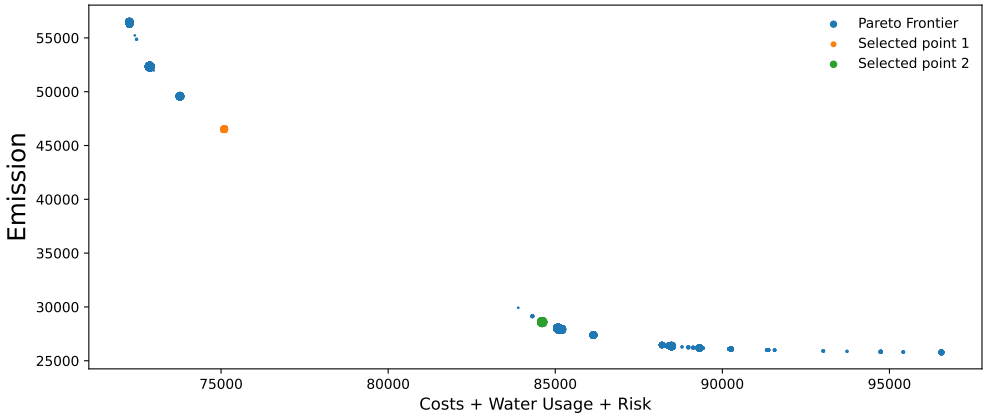

Figure 14: Pareto frontier between cost and emission from Jan 1^{st} to Jan 3^{rd} in the UK. Sizes of points are proportional to their frequencies. Two points are selected for further analysis of optimal schedules, as shown in Figure 15.

6. Conclusion and Future Work

Time-varying demand response for chemical process systems is an important new operating paradigm for enabling the increased prevalence of intermittent renewable energy sources. When considering sustainable decision making, it is possible that operational decisions can be in competition depending on if a cost or emissions driving force is considered, and that the correlating or competing nature of these objectives can change over time. In this work, we applied a recently developed approach for identifying correlation or competition *a priori* to solving many-objective optimization problems to better understand the correlating versus competing nature of cost and emissions driven demand response. 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. Detailed analysis of historical data showed that across a wide variety of both locations and times of the year, emissions and cost objectives are well correlated during a vast majority of the time. Furthermore, when looking at different locations, we determined that correlation trends did not change much between locations within the same ISO; however, different ISO's can have seen vastly different results due to different weather profiles and generator mixes. Finally, we presented an approach for determining a compromise operation strategy in the time periods where cost-driven and emissions-driven demand response are in competition.

In our future work, we will seek to develop tools that incorporate the understanding gleaned in this work to practial application in an online, real-time operation setting. First, we will attempt to build a high-throughput machine learning approach to build classifier models that can predict cost/emissions correlation strengths and objective groupings based on information available *a priori* to online scheduling, such as weather forecasts, price forecasts, and the current state of the grid and process. We will also look to build a tool to identify a single objective grouping that performs robustly well over the full range of uncertain dynamic parameter values, such that we can effectively manage uncertainty that is inherent in predicting parameters relevant to future operation.

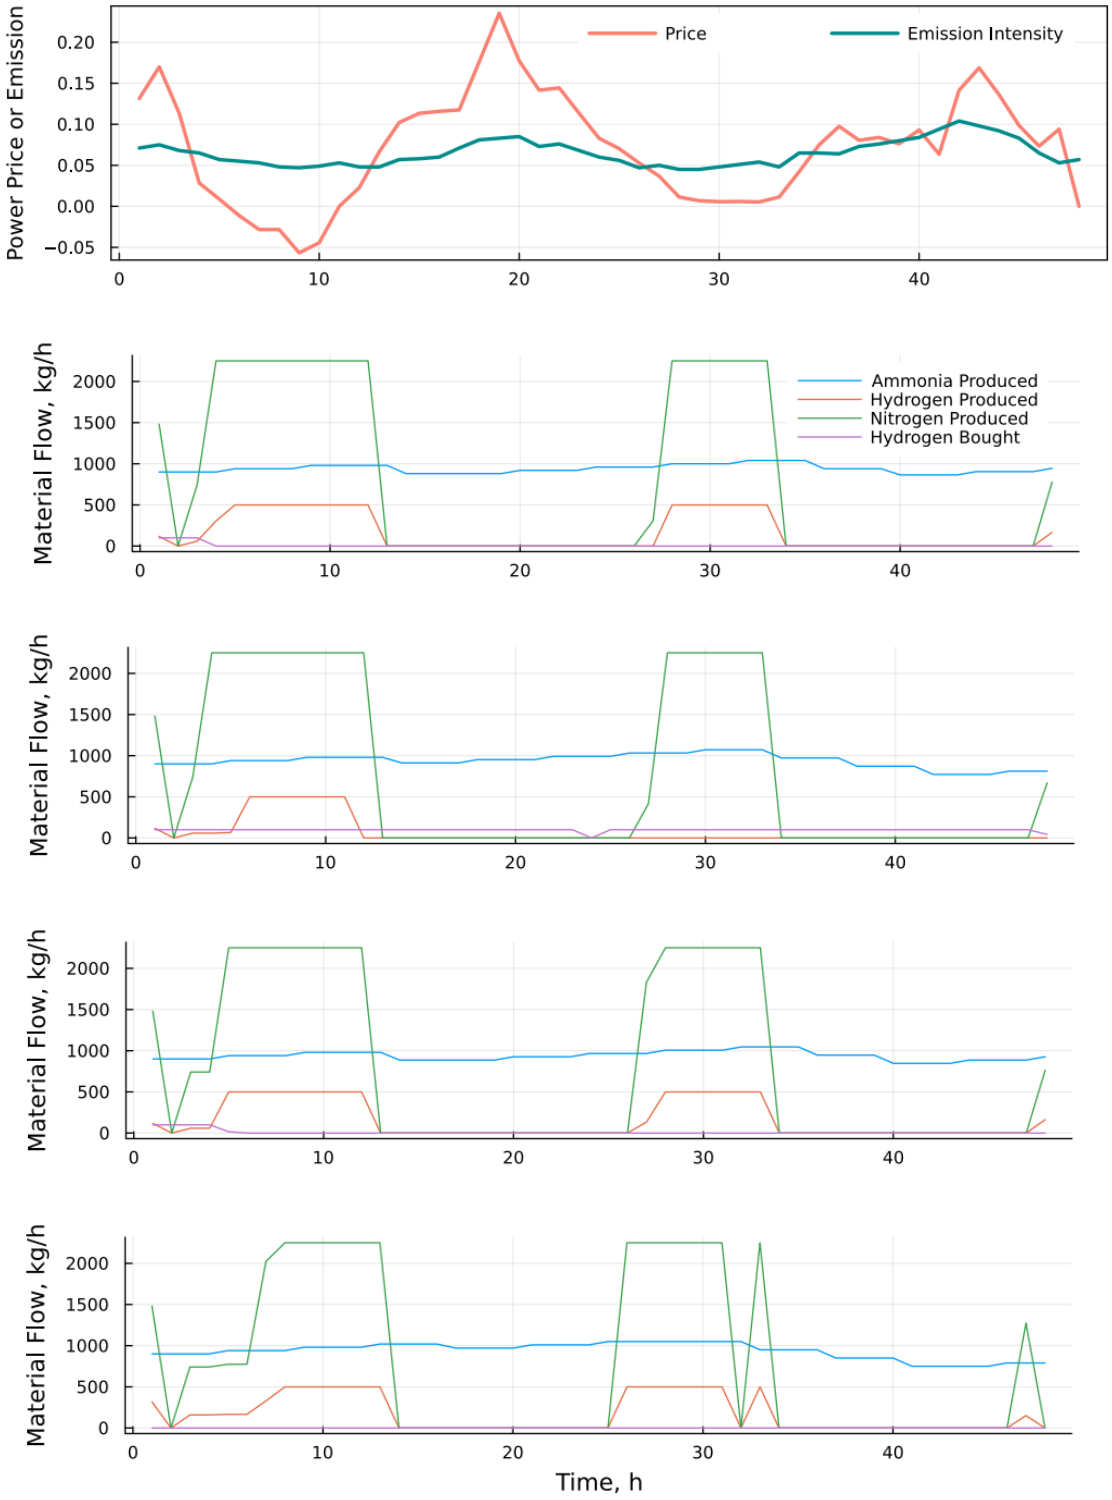

Figure 15: Scheduling results from Jan 1^{st} 12 am to Jan 3^{rd} 12 am in the UK. (top) Power price ($/kWh) and emissions factor (kg CO_2/kWh). (from top to bottom) Optimal production schedules when optimizing operating cost, selected point 1, selected point 2 and carbon emissions.

Acknowledgements

The authors are grateful to the organizers of the FOCAPO/CPC conference (M. Baldea, R. Bindlish, I. Harjunkoski, V. Zavala) for their invitation to participate in this special issue. The authors also acknowledge funding support from the National Science Foundation (NSF) CAREER Award (CBET-2237284).

References

Allman, A., Daoutidis, P., 2018. Optimal scheduling for wind-powered ammonia generation: Effects of key design parameters. Chemical Engineering Research and Design 131, 5–15.

Allman, A., Palys, M.J., Daoutidis, P., 2019. Scheduling-informed optimal design of systems with time-varying operation: A wind-powered ammonia case study. AIChE Journal 65, e16434.

Allman, A., Zhang, Q., 2022. Distributed fairness-guided optimization for coordinated demand response in multi-stakeholder process networks. Computers and Chemical Engineering 161, 107777.

Babu, C.A., Ashok, S., 2008. Peak load management in electrolytic process industries. IEEE Transactions on Power Systems 23, 399–405.

Baumgartner, N., Delorme, R., Hennen, M., Bardow, A., 2019. Design of low-carbon utility systems: Exploiting time-dependent grid emissions for climate-friendly demand-side management. Applied Energy 247, 755–765.

Brée, L.C., Perrey, K., Bulan, A., Mitsos, A., 2019. Demand side management and operational mode switching in chlorine production. AIChE Journal 65, 1–14.

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

California ISO, 2023. Today's outlook - emissions. URL: http://www.caiso.com/todaysoutlook/pages/emissions.html.

Carbon Intensity API, 2023. National data. URL: https://carbonintensity.org.uk/.

Caspari, A., Offermanns, C., Schäfer, P., Mhamdi, A., Mitsos, A., 2019. A flexible air separation process: 2. Optimal operation using economic model predictive control. AIChE Journal 65, 1–14.

Center for Chemical Process Safety, 1995. Guidelines for Safe Process Operations and Maintenance. Wiley-AIChE.

Cplex, I.I., 2020. V20. 1: User's manual for cplex. International Business Machines Corporation .

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

Demirhan, C.D., Tso, W.W., Powell, J.D., Pistikopoulos, E.N., 2021. A multi-scale energy systems engineering approach towards integrated multi-product network optimization. Applied Energy 281, 116020.

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

Electricity Maps, 2023. data-portal - japan. URL: https://www.electricitymaps.com/data-portal/japan.

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

Fleschutz, M., Bohlayer, M., Braun, M., Henze, G., Murphy, M.D., 2021. The effect of price-based demand response on carbon emissions in european electricity markets: The importance of adequate carbon prices. Applied Energy 295, 117040.

Gao, X., Neueven, B., Siirola, J.D., Miller, D.V., Dowling, A.W., 2022. Multiscale simulation of integrated energy system and electricity market interactions. Applied Energy 316, 119317.

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

Gurobi Optimization, L., 2018. Gurobi Optimizer Reference Manual. URL: http://www.gurobi.com.

Hochhaus, T., Bruns, B., Grunewald, M., Riese, J., 2023. Optimal scheduling of a large-scale power-to-ammonia process: Effects of parameter optimization on the indirect demand response potential. Computers and Chemical Engineering 170, 108132.

IEA, 2020. Energy Technology Perspectives 2020. Technical Report. International Energy Agency.

ISO New England, 2023. Real-time maps and charts. URL: https://www.iso-ne.com/isoexpress/.

Japan Electric Power eXchange, 2023. Trading information - trading market data. URL: https://www.jepx.jp/en/electricpower/market-data/spot/.

Karwan, M.H., Keblis, M.F., 2007. Operations planning with real time pricing of a primary input. Computers and Operations Research 34, 848–867.

Kelley, M.T., Do, T.T., Baldea, M., 2022. Evaluating the demand response potential of ammonia plants. AIChE Journal 68, e17552.

Khan, I., 2018. Importance of GHG emissions assessment in the electricity grid expansion towards a low-carbon future: A time-varying carbon intensity approach. Journal of Cleaner Production 186, 1587–1599.

Kılınç, M.R., Sahinidis, N.V., 2018. Exploiting integrality in the global optimization of mixed-integer nwebpagear programming problems with BARON. Optimization Methods and Software 33, 540–562.

Kopsakangas-Savolainen, M., Mattinen, M.K., Manninen, K., Nissinen, A., 2017. Hourly-based greenhouse gas emissions of electricity - cases demonstrating possibilities for households and companies to decrease their emissions. Journal of Cleaner Production 153, 384–396.

Leenders, L., Ganz, K., Bahl, B., Hennen, M., Baumgartner, N., Bardow, A., 2021. Scheduling coordination of multiple production and utility systems in a multi-leader multi-follower stackelberg game. Computers and Chemical Engineering 150, 107321.

Marler, R.T., Arora, J.S., 2009. The weighted sum method for multi-objective optimization: new insights. Structural and Multidisciplinary Optimization 41, 853–862.

Midcontinent ISO, 2023. Market reports - historical real-time fuel on the margin. URL: https://www.misoenergy.org/findapi/MisoWeb_Models_Find_MarketReportItem/_search.

Nord Pool, 2023. Market data - nord pool uk. URL: https://www.nordpoolgroup.com/en/Market-data1/GB/Auction-prices/UK/Hourly/.

Otashu, J.I., Baldea, M., 2019. Demand response-oriented dynamic modeling and operational optimization of membrane-based chlor-alkali plants. Computers and Chemical Engineering 121, 396–408.

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

Palys, M.J., Mitrai, I., Daoutidis, P., 2023. Renewable hydrogen and ammonia for combined heat and power systems in remote locations: Optimal design and scheduling. Optimal Control Applications and Methods 44, 719–738.

Ramsden, T., Steward, D., Zuboy, J., 2009. Analyzing the Levelized Cost of Centralized and Distributed Hydrogen Production Using the H2A Production Model, Version 2. Technical Report. National Renewable Energy Laboratory.

Russell, J.M., Allman, A., 2023. Sustainable decision making for chemical process systems via dimensionality reduction of many objective problems. AIChE Journal 69, e17962.

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

Tang, X., Baldea, M., 2023. A grid view on the dynamics of processes participating in demand response programs. Computers and Chemical Engineering 169, 108070.

Tsay, C., Kumar, A., Flores-Cerrillo, J., Baldea, M., 2019. Optimal demand response scheduling of an industrial air separation unit using data-driven dynamic models. Computers and Chemical Engineering 126, 22–34.

Weigert, J., Hoffmann, C., Esche, E., Fischer, P., Repke, J.U., 2021. Towards demand-side management of the chlor-alkali electrolysis: Dynamic modeling and model validation. Computers and Chemical Engineering 149, 107287.

Zhang, Q., Grossmann, I.E., Heuberger, C.F., Sundaramoorthy, A., Pinto, J.M., 2015. Air Separation with Cryogenic Energy Storage: Optimal Scheduling Considering Electric Energy and Reserve Markets. AIChE Journal 61, 1547–1558.