DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

Control Strategies and Validation in the Hybrid Optimization and Performance Platform (HOPP)

Genevieve M. Starke, Jonathan Martin, Kaitlin Brunik, and Parangat Bhaskar

National Renewable Energy Laboratory

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308

This report is available at no cost from NREL at www.nrel.gov/publications.

Technical Report
NREL/TP-5000-88069
November 2025

Control Strategies and Validation in the Hybrid Optimization and Performance Platform (HOPP)

Genevieve M. Starke, Jonathan Martin, Kaitlin Brunik, and Parangat Bhaskar

National Renewable Energy Laboratory

Suggested Citation
Starke, Genevieve M., Jonathan Martin, Kaitlin Brunik, and Parangat Bhaskar. 2025. *Control Strategies and Validation in the Hybrid Optimization and Performance Platform (HOPP)*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5000-88069. https://www.nrel.gov/docs/fy26osti/88069.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated under Contract No. DE-AC36-08GO28308

This report is available at no cost from
NREL at www.nrel.gov/publications.

Technical Report
NREL/TP-5000-88069
November 2025

15013 Denver West Parkway
Golden, CO 80401
303-275-3000 • www.nrel.gov

NOTICE

This work was authored by NREL for the U.S. Department of Energy (DOE), operated under Contract No. DE-AC36-08GO28308. Funding provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Wind Energy Technologies Office. The views expressed herein do not necessarily represent the views of the DOE or the U.S. Government.

Cover photos (clockwise from left): Josh Bauer, NREL 61725; Visualization from the NREL Insight Center; Getty-181828180; Agata Bogucka, NREL 91683; Dennis Schroeder, NREL 51331; Werner Slocum, NREL 67842.

Executive Summary

The Hybrid Optimization and Performance Platform (HOPP) is a tool designed to simulate various configurations of hybrid power plants and hybrid energy systems, as well as evaluate their financial feasibility. This report provides an overview of HOPP, detailing the energy storage dispatch strategies it offers, along with field testing and validation studies conducted for the tool. We present two scenarios that demonstrate different applications of HOPP. The first scenario examines the efficacy of hybrid power plants across various locations in the United States. The second scenario features an energy system that produces both hydrogen and electricity, showcasing a dispatch strategy that selects the most profitable energy vector based on price signals. Additionally, we validate HOPP using operational data from both unit-scale and utility-scale power plants. The validation process demonstrates that HOPP can accurately simulate the power output of both wind and solar photovoltaic plants at both scales, with comparable fidelity to an existing commercial software tool. Finally, HOPP is used in a field test at the National Renewable Energy Laboratory, applying an optimal dispatch strategy to a physical battery in a unit-scale hybrid plant. In this real-world setting, HOPP's optimal dispatch strategy improved the hybrid plant's ability to meet a load signal while minimizing operational costs.

Acknowledgments

We thank the National Renewable Energy Laboratory (NREL) Advanced Research on Integrated Energy Systems (ARIES) staff who helped us set up the field test for this project: Emanuel Mendiola and Przemyslaw Koralewicz. We also thank Caitlin Murphy of NREL and her group and Arvind Tiwari of GE Global Research for providing some of the validation data sets and insight on the project. In addition, we would like to thank Caitlin Murphy, Paul Denholm, Chris Bay, Patrick Moriarty, Brian Smith, and Paul Veers for their thoughtful review and Sheri Anstedt and Amy Brice for editorial support.

List of Acronyms and Abbreviations

AC	Alternating current
AOI	Angles of incidence
ARIES	Advanced Research on Integrated Energy Systems
BESS	Battery Energy Storage System
BOS	Balance of system
CapEx	Capital expenditure
DC	Direct current
DHI	Diffuse horizontal irradiance
DNI	Diffuse normal irradiance
DOE	Department of Energy
ERS	Essential reliability services
FLORIS	FLOw Redirection and Induction in Steady State
GE	General Electric
GHI	Global horizontal irradiance
HES	Hybrid Energy System
HOMER	Hybrid Optimization of Multiple Energy Resources
HOPP	Hybrid Optimization and Performance Platform
HPP	Hybrid Power Plant
IEC	International Electrotechnical Commission
LCOE	Levelized cost of energy
MAE	Mean absolute error
ME	Mean error
MIDC	Measurment and Instrumentation Data Center
MVDAS	Medium Voltage Data Acquisition System
MW	Megawatt
MWh	Megawatt-hour
NOAA	National Oceanic and Atmospheric Administration
NPV	Net present value
NREL	National Renewable Energy Laboratory
PPA	Power purchase agreement
PV	Photovoltaic
REopt	Renewable Energy Integration & Optimization
SAM	System Advisor Model
SOC	State of charge

Table of Contents

List of Figures

List of Tables

1 HOPP Overview

The National Renewable Energy Laboratory (NREL) developed the Hybrid Optimization and Performance Platform (HOPP) as an open-source tool to model hybrid power plants and hybrid energy systems (NREL 2023b). For this report, we use the definition of a hybrid power plant (HPP) developed by the International Energy Agency (IEA) Wind Technology Collaboration Programme Task 50, which defines a hybrid power plant as "a combination of two or more electricity generation and/or storage technologies, used to provide electrical power services that are coordinated at a single point of connection" (Kotarbinski et al. 2024). A hybrid energy system (HES) is defined here as "one or more HPPs, which provide(s) an energy and/or nonenergy product such as electricity, hydrogen, heat, or fresh water to accommodate specific end-use needs". This tool was designed to answer the following research question: When and where are HPPs profitable financially and feasible operationally, and how can we ensure their optimal design? HOPP serves as a design tool for HPPs, assessing the performance and overall profitability of hybrid system designs.

HOPP consists of a simulation component and a financial modeling component. The simulation component supports multiple technologies that can be combined in user-defined configurations. Currently supported technologies include wind, solar, battery storage, concentrating solar power (both tower and trough configuration), and wave energy, with additional technologies in development. The simulation component also includes a grid model that imposes transmission constraints. An off-grid version of HOPP is also in development. The tool provides a general structure for each technology model, allowing users to either select the existing model or add their own. Inputs for the simulation include the power rating of the technology and specific parameters that are outlined in the respective technology modules. HOPP currently supports NREL's Python package of the System Advisor Model, pySAM, for simulation and financial modeling(SAM) (Danecek et al. 2021). Additionally, for wind resource modeling, HOPP supports the FLOw Redirection and Induction in Steady State (FLORIS) model, an open-source wind plant optimization and modeling tool (NREL 2021). HOPP also features hydrogen modeling capabilities using an open-source stand-alone proton exchange membrane electrolyzer model (NREL 2023a).

The financial component within HOPP uses financial models derived from SAM's SingleOwner model to assess the profitability of the plant design. It calculates the net present value of the plant based on the plant lifetime and the levelized cost of energy or hydrogen, primarily focusing on quantifying the energy value stream. Recent research has also highlighted additional non-energy value streams that contribute to the plant's revenue potential (Ericson et al. 2022; Schleifer et al. 2023) which are being considered for future versions of HOPP. The financial component supports SAM capabilities and a general costing structure for user-specific financial models. The SAM financial model is highly customizable—it takes inputs such as capital costs, operations and maintenance costs, a grid electricity price signal for specific locations, an average inflation rate, discount rates due to shared infrastructure within a hybrid plant, and policy inputs.

By integrating the simulation and financial components, HOPP can design HPPs or HESs that provide hourly output and dispatch information along with a financial value estimates. This capability allows renewable energy developers to optimize projects, potentially resulting in better financing terms, reduced capital and operational expenses, improved plant performance, and increased energy density per unit of land area compared to single-technology variable renewable energy generators.

1.1 Dispatch Strategies

The dispatch of hybrid systems is more complex than that of single-technology variable renewable energy generators as it requires the coordination of all the included technologies along with the grid. In HOPP, this coordination is accomplished by implementing a modular framework where each technology operates as a separate module (W. Hamilton et al. 2022). These modules are interconnected using Pyomo, which is a programming language developed by Sandia National Laboratories (Hart et al. 2017). The technology modules are dispatched according to centralized information collected by the Pyomo structure, as well as technology-specific information and constraints. More information on this structure and the dispatch of the systems, such as the grid system and the photovoltaic system, is available in W. Hamilton et al. (2022) and W. T. Hamilton et al. (2020). In the following sections we focus only on the battery dispatch and hydrogen dispatch models that are relevant to the considered scenarios. Although HOPP can also optimize for maximium gross profit from the hybrid plant, in this study we focus on load-following strategies.

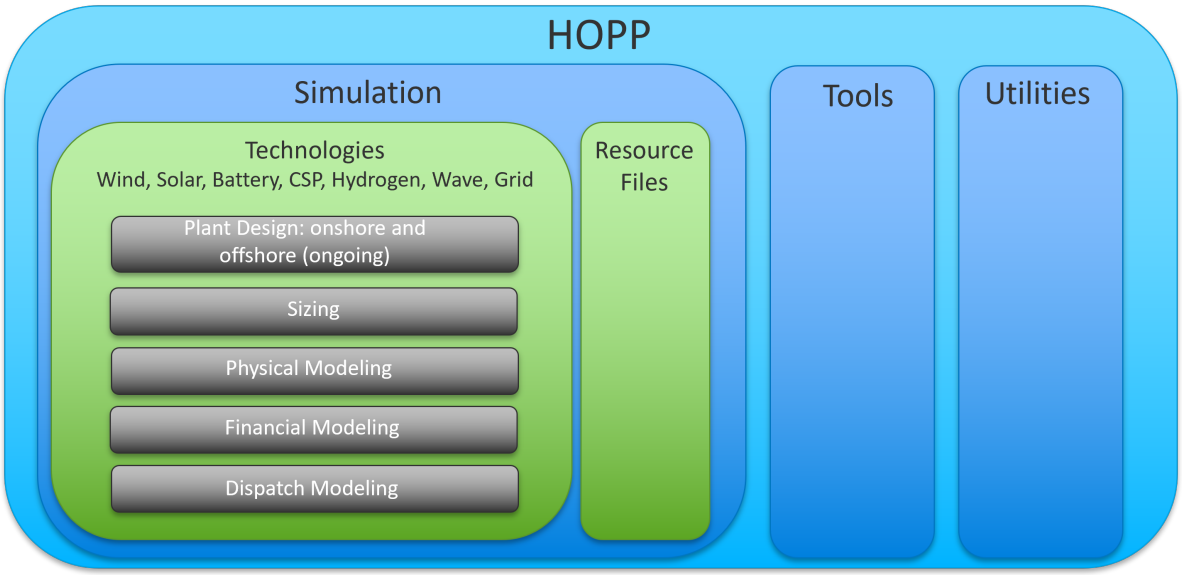

Figure 1. A representation of the architecture of the HOPP code

1.1.1 Minimize Operating Costs While Following a Load Demand

This dispatch strategy aims to track a load profile while minimizing the operating costs of the plant in response to load demand. This strategy uses perfect forecasting and incorporates a look-ahead horizon over which the dispatch is optimized. Currently, HOPP operates on an hourly timescale using steady state models, but this can be adjusted to a sub-hourly timescale if sub-hourly resource data is provided. The battery's operation is governed by the following equation, where τ represents the look-ahead time period, and t represents a time index. These dynamics are detailed in W. Hamilton et al. (2022) and are summarized here.

$$b_t^{soc} = b_{t-1}^{soc} + \Delta\left(\frac{\eta^+ \cdot \dot{w}_t^+ - \dot{w}_t^- / \eta_-}{C^B}\right) \quad \forall t \in \tau : t \geq 2 \tag{1.1}$$

Equation 1.1 represents the update to the battery state of charge, b_t^{soc}, and is dependent on the previous state of charge; the time step, Δ; the charge efficiency, η^+; the discharge efficiency, η^-; the power in and out of the battery at time t, which is $\dot{w}_t^+$ and $\dot{w}_t^-$, respectively; and the battery manufactured specified capacity, C^B. The state of charge is bounded by the equation $\underline{S}^B \leq b_t^{soc} \leq \overline{S}^B, \forall t \in \tau$, where $\underline{S}^B$ and $\overline{S}^B$ are the minimum and maximum operation bounds of the battery state of charge, respectively.

The power flow in and out of the battery is also bound by constraints that have the following form:

$$\underline{P}^B y_t^- \leq \dot{w}_t^- \leq \overline{P}^B y_t^- \tag{1.2}$$

$$\underline{P}^B y_t^+ \leq \dot{w}_t^+ \leq \overline{P}^B y_t^+ \tag{1.3}$$

where $\underline{P}^B$ and $\overline{P}^B$ are the battery minimum and maximum power ratings, respectively, and y_t^- and y_t^+ are indicators of whether the battery is charging or discharging during time period t. These indicators have a value of 1 if the battery is discharging (y_t^-) or charging (y_t^+) during the time period, and 0 otherwise. The indicators are also subject to the constraint $y_t^+ + y_t^- = 1$ to ensure the battery is only performing one activity over each time period. Finally, the cycle count of the battery is calculated using

$$b^c \geq \frac{\Delta}{C^B} \sum_{t \in \tau} \dot{w}_t^- \tag{1.4}$$

Bounds on some of the variables are enforced subject to the following constraints:

$$b^c,\ b_t^{soc},\ \dot{w}_t^+,\ \dot{w}_t^- \geq 0 \qquad \forall t \in \tau \tag{1.5}$$

$$y_t^+,\ y_t^- \in 0, 1 \qquad \forall t \in \tau \tag{1.6}$$

This battery model is used in the systemwide optimization that is performed in HOPP, which combines all of the technology dispatch models to create the system hybrid dispatch model represented by $(\mathscr{H})$. The systemwide optimization takes into account all of the technologies in the simulation, such as inputs from the grid, battery, photovoltaic field, and wind farm (or any present combination of these technologies). The hybrid configuration is included by adding an exponential time-hybrid weighting factor to weight the different parts of the objective function, which is represented by γ. The load is imposed on the system through the time-varying transmission limit in the grid module W_t^g, which is used as the upper bound that the hybrid system can produce. Since the load is used as the upper bound, the hybrid system may underproduce when following this dispatch, which may affect the financial calculations in this case. A sample of an objective function that includes the grid, a photovoltaic (PV) field, a battery, and a wind plant is shown below.

$$\begin{aligned}(\mathscr{H}) \quad \text{minimize} \sum_{t\in\tau} \Big[& \gamma^t \Delta \left(P_s^t (W_t^g - \dot{e}_t^s) + P_t^p \dot{e}_t^p\right) + \varepsilon y_t^g \\ & + \gamma^t \Delta C^{pv} \dot{w}_t^{pv} \\ & + \gamma^t \Delta C^{wind} \dot{w}_t^{wind} \\ & + \gamma^t \Delta \left(-C^{bc} \dot{w}_t^{+} + -C^{bd} \dot{w}_t^{-}\right) \Big] + C^{bt} b^c \end{aligned} \tag{1.7}$$

We will now explain the variables line by line in Equation 1.7. The first line is the grid requirement, where the electricity selling price, P_t^s, is used to penalize the system for missing load, which tries to match the electricity sold to the load profile. The variable P_t^p stands for the electricity purchased from the grid, and $\dot{e}_t^s$ and $\dot{e}_t^p$ stand for the electricity sold to and purchased from the grid at time t, respectively. The final term in the first line εy_t^g is used to determine if the system is net-generating at time t by using y_t^g as a binary that is equal to 1 if the system is net-generating and 0 otherwise. The second line represents the contribution from the photovoltaic field and includes the hybrid factor, the generation cost of the photovoltaic field, C^{pv}, and the power being generated by the photovoltaic field at time t, represented by $\dot{w}_t^{pv}$. The third line represents the wind farm contribution and is very similar to the solar term. The variable C^{wind} represents the generation cost of wind, and $\dot{w}_t^{wind}$ is the wind produced at time t. Finally, the last line is the contribution from the battery model, where C^{bc} and C^{bd} are the operating costs for charging and discharging the battery, and C^{bl} is the term denoting the life cycle cost for the battery. This is the minimization problem solved to find the optimal battery dispatch to follow a load while minimizing operating costs. Default values for the cost parameters are given in W. Hamilton et al. (2022).

1.1.2 Following a Load Demand

This dispatch strategy focuses on meeting load demand without considering cost. It uses a simpler battery model than the one described in the previous section. While the generation technologies follow the same dispatch strategy outlined in the previous section, the battery dispatch uses a simplified heuristic method, making decisions at each time step without accounting for future information. The logic starts by taking the difference between the overall generated power and the load requirement as $\Delta w_t = (W_t^g - \dot{w}_t^{hybrid})$, borrowing W_t^g as the load requirement from the previous section, $\dot{w}_t^{hybrid}$ as the total hybrid power production, and Δw_t as the power difference, all at time t. The change in the power from the battery is then decided according to:

$$\dot{w}_t = \begin{cases} -\min(|\Delta w_t|, \Delta\dot{C}^-) & \text{if } \Delta w_t < 0 \\ \min(|\Delta w_t|, \Delta\dot{C}^+) & \text{if } \Delta w_t > 0 \\ 0 & \text{otherwise} \end{cases} \tag{1.8}$$

where $\dot{C}^-$ and $\dot{C}^+$ are the discharge and charge rates of the battery, respectively. If the generation is equal to the load, then the battery does nothing.

The state of charge is updated according to the following equations:

$$b_t^{soc} = \begin{cases} \min\left[(b_{t-1}^{soc} + \dot{w}_t),\ \underline{S}^B\right] & \text{if } \Delta w_t < 0 \\ \max\left[(b_{t-1}^{soc} + \dot{w}_t),\ \overline{S}^B\right] & \text{if } \Delta w_t > 0 \\ b_{t-1}^{soc} & \text{otherwise} \end{cases} \tag{1.9}$$

where b_{t-1}^{soc} represents the state of charge at the previous time step. This logic also enforces the bounds on the state of charge of the battery: $\underline{S}^B \leq b_t^{soc} \leq \overline{S}^B, \forall t \in \tau$. Because this method is a faster heuristic method, it operates without minimizing the operating costs and only considers meeting the load profile for the system.

1.1.3 Hydrogen Dispatch Strategy

The hydrogen dispatch strategy is implemented under the assumption that the electricity produced by the plant can either be sold to the grid as electricity or used to produce hydrogen, but not both in the same hour. The decision on which action to take is dependent on the current prices of electricity and hydrogen. This scenario represents an initial study that seeks to examine the opportunities of a plant that can produce and export both electricity and hydrogen.

The plant simulation is first given a load demand that is equal to 15% of the hydrogen plant rating in order to maximize the number of hours that the plant is able to produce hydrogen. An electrolyzer is typically not operated when the electricity available is less than 5% of its power rating (Electric Power Research Institute (EPRI) 2022). A value of 15% was chosen for the load demand profile to provide a slight buffer above this cutoff limit.

At every hour, profit of the energy sale, represented by $Pr_{E,t}$, for exporting electricity to the grid is compared to the profit of selling the energy as hydrogen. The profit from selling electricity to the grid is calculated using the expression

$$Pr_{El,t} = LMP_t \; E_{plant,t}, \tag{1.10}$$

where $Pr_{El,t}$ is the profit from selling the electricity, LMP_t is the local marginal price of electricity at time t, and $E_{plant,t}$ is the electricity produced by the plant at time t.

The potential profit from hydrogen that could be produced by the plant at time t is calculated using the expression

$$Pr_{H_2,t} = LP_{H_2,t} \; H_{plant,t}, \tag{1.11}$$

where $Pr_{H_2,t}$ is the profit from selling the hydrogen, $LP_{H_2,t}$ is the local price of hydrogen at time t. $H_{plant,t}$ is the hydrogen produced by the plant at time t, which is calculated using the conversion that an electrolyzer system uses about 57 $kWhr$ of electricity per kilogram of hydrogen produced.

The dispatch is then decided using

$$Pr_{E,t} = \begin{cases} Pr_{El,t} & \text{if } Pr_{El,t} > Pr_{H_2,t} \\ Pr_{H_2,t} & \text{otherwise} \end{cases} \tag{1.12}$$

This relation obviously depends on the local marginal price, meaning that if the local marginal price exceeds some threshold, then it becomes more economical to convert the hydrogen and sell it to the grid. These peaks in the local marginal price tend to occur around hours of peak electricity usage.

1.2 Dispatch Analysis

To examine the effectiveness of the plant design and the dispatch strategies, dispatch analysis is performed after the simulation. The dispatch analysis considers how effective the plant is at meeting or exceeding a specified power setpoint.

The final output is the percentage of time the plant is at or above the power setpoint and the percentage of the total power setpoint that is provided by the plant, also referred to as the compliance percentage. The percentage of time the plant achieves the power setpoint is calculated using:

$$\%_{time\ met} = \frac{hr_{met}}{hr_{total}} * 100 \tag{1.13}$$

where hr_{met} is the total number of hours the plant achieves the power setpoint and hr_{total} is the total number of simulated hours. Additionally, it can be valuable to know how much load was missed. HOPP calculates the percentage of the load that was missed during the simulation by evaluating the following expression

$$\%_{missed\ load} = \frac{\sum_{t=1}^{N_t}(\max[(P_{stpt,t} - P_{plant,t}), 0])}{\sum_{t=1}^{N_t}(P_{stpt,t})} \tag{1.14}$$

where $P_{stpt,t}$ is the power setpoint for the simulation and $P_{plant,t}$ is the generation of the plant at time t.

Additionally, the power setpoint can have a compliance standard that allows a tolerance of error in the power setpoint. This is added into the analysis by adjusting the power setpoint value such that $P_{stpt} = (1 - tol)P_{stpt,0}$. The percentage of the total power requested from the plant that is fulfilled is calculated as:

$$\%_{power,met} = \frac{\sum_{t=1}^{N_t}(\min[P_{plant,t}, P_{stpt,t}])}{\sum_{t=1}^{N_t}(P_{stpt,t})} \tag{1.15}$$

where $P_{plant,t}$ and $P_{stpt,t}$ are the power produced by the plant and the power setpoint at time t, respectively, and N_t is the total number of simulation time steps. In this calculation, the plant output is capped at P_{stpt} in order to keep this value from being skewed from when the plant produces more than the power setpoint.

2 Hybrid Controls Scenarios

Two representative scenarios illustrate the capabilities of HOPP and show how the capabilities can be leveraged in hybrid system design and analysis. The first scenario (Section 2.1) examines the percentage of the year that a HPP can meet or exceed a given power setpoint at five different sites in the United States and investigates what configurations could satisfy this requirement. The second scenario (Section 2.2) examines a HES that decides between selling electricity or producing and selling hydrogen based on pricing information.

2.1 Scenario 1: Meeting a Power Setpoint

In Scenario 1, we investigate the ability of a HPP to meet or exceed a given power setpoint. The metric used to measure performance is plant compliance factor, which is given as a percentage of the year that the plant can supply the power setpoint, as described in Section 1.2. This scenario highlights the capabilities of the tool to meet a power setpoint, but it does not address the fluctuating value of energy. First, we examine the performance of combinations of different technologies. Then, we compare the performance of a wind-solar-storage HPP in five locations in the contiguous United States.

2.1.1 Comparing Technology Combinations

First, we investigate the performance of various combinations of wind, solar, and storage in achieving a power setpoint schedule. The technology capacities are kept constant throughout all scenarios; only the inclusion of various technologies changes. We use 100 MW of wind, provided by turbines based on the 5-MW NREL reference turbine (Jonkman et al. 2009). We also use a 100 MW solar photovoltaic farm and a 4-hour, 50-MW battery. The interconnection size is set at 100 MW. Kansas was used as the reference location to examine the performance of the selected technologies in a strong resource area. The chosen latitude is 38.0, the longitude is -100.8, and the elevation is $2,838$ feet. The wind and solar resources used for this study are based on historical weather data from the year 2013.

Table 1. The different cases and compliance metrics for the firm power technology comparison

	Wind	Solar	Storage	Compliance (%) Optimized Dispatch	Compliance (%) Heuristic Dispatch
Case 1	X	-	X	67	66
Case 2	-	X	X	55	55
Case 3	X	X	-	72	72
Case 4	X	X	X	92	92

The results are shown in Table 1. In this scenario, we examine four different cases. Case 1 includes co-located wind and storage, and Case 2 includes co-located solar and storage. Case 3 includes co-located wind and solar, and Case 4 includes all three technologies together. In Cases 1, 2, and 4 we consider the compliance factors for the cost-optimized dispatch and the heuristic dispatch. Case 3 is the same in both the cost-optimized and heuristic dispatch because the dispatch strategies focus on the battery dispatch, and Case 3 does not include storage technology. Thus, the compliance factor is the result of the combined solar and wind generation profiles.

As shown in the table, the compliance percentages between the optimized and heuristic dispatches are nearly identical, differing by less than one percentage point. The cost-optimized dispatch adheres to the load signal while minimizing operating costs, resulting in longer computational times due to the consideration of more design variables. In contrast, the heuristic dispatch focuses solely on meeting the load requirement. Given the similarity in results and the significantly faster computation time of the heuristic method, we will use the heuristic dispatch for the remaining scenarios.

We now examine the differences in compliance among the various technology combinations. Case 2, comprising solar and storage, has the lowest compliance factor. This outcome may be due to the low capacity factors of solar technology and the specific component sizing used in this study. While solar PV generally has a higher energy capacity density than wind—allowing greater solar capacities to be deployed within the same area as 100 MW of wind—we maintained identical capacity sizing for a direct comparison. However, this approach may not reflect the most realistic scenario. The case with the next lowest compliance factor is Case 1, which consists of wind and storage. While wind technology has a higher capacity factor, it is less consistent than solar, which can diminish the effectiveness of shorter-term storage (1-4 hours) in the energy market. Case 3, which includes wind and solar,

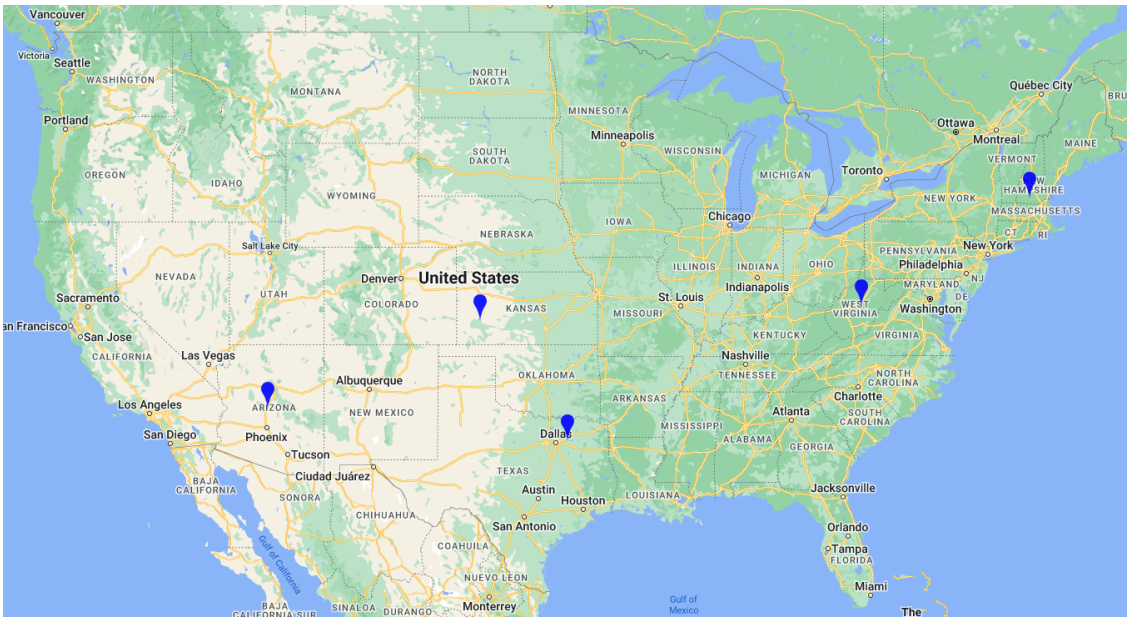

Figure 2. A map showing the locations of the five sites used in the hybrid plant comparison

achieves a 72% compliance factor, even without storage. This suggests that a co-located HPP, consisting solely of generation technologies can, in some instances where complementarity is strong, more consistently meet a power setpoint compared to single technologies like wind or solar paired with an energy storage method. It should be noted that these results are based on fixed component sizing and results may vary with optimized component sizing. Finally, Case 4, which incorporates all three technologies, and significantly outperforms the other cases, reaches a compliance factor of over 90% in both dispatch scenarios. This suggests that a HPP combination of wind, solar, and storage has the potential to reliably meet a power setpoint for the grid.

2.1.2 Comparing Sites

We now examine the viability of hybrid plants meeting a power setpoint in five different locations in the United States. The locations were chosen to represent a range of conditions and geographic areas, with a focus on onshore locations. The latitude, longitude, and elevation are provided for each site in Table 2.

Table 2. Descriptions of the five sites used in Scenario 1

Site	Latitude	Longitude	Elevation (feet)
Kansas	38.0	−100.8	2,838
Arizona	34.27	−112.0	3,570
Texas	32.86	−96.2	555
West Virginia	38.60	−80.7	1,493
New Hampshire	42.90	−72.1	1,358

The sites are also shown in Figure 2 on a map of the contiguous United States. The locations were chosen according to the following criteria:

- **Kansas**: The site was selected for its strong wind resources and relatively high solar potential, making it a good candidate for a HPP location. National analyses have shown this area to have high hourly complementarity, where wind and solar resources peak at different times, enhancing the ability to meet a power setpoint (Harrison-Atlas et al. 2022; Murphy et al. 2023).
- **Arizona**: The site was chosen as an example of an area with strong solar resource and relatively weak wind resource. This case explores how a hybrid plant might perform under less favorable wind conditions and how its design could adapt accordingly.

- **Texas**: The site was selected as an example of a hybrid plant in eastern Texas, where wind resources remain significant but are less abundant than in western Texas. This case evaluates the performance of a hybrid plant under these less favorable conditions while being situated closer to an urban center.
- **West Virginia**: The site was selected to evaluate the performance of a hybrid plant in an area with limited wind and solar deployment. Overall, the Eastern United States has weaker wind and solar resources compared to the central regions. However, research suggests that the complementarity between these resources could be leveraged to meet a power setpoint (Harrison-Atlas et al. 2022; Murphy et al. 2023). The West Virginia site provides an example of hybrid plant performance in a region where the strength and complementarity of solar and wind resources are more variable and site-dependent.
- **New Hampshire**: The site was selected to examine the effectiveness of a hybrid plant in a region with multiple major urban centers and less abundant resources than the central United States. New Hampshire faces notable spatial constraints, making site selection particularly critical.

Next, we outline some of the financial assumptions used for this study. The installed cost of wind, solar, and battery storage are outlined in Table 3. Technology costs are from the NREL Annual Technology Baseline (ATB) 2025 moderate case (National Renewable Energy Laboratory 2022). The solar installed costs and the battery installed costs are based on the 2025 Moderate Utility-Scale PV-Plus-Battery technology case, where the PV and battery share an inverter. The wind installed costs are based on the 2025 Moderate Land-Based Wind technology case and assumes a 10% cost reduction from HPP balance-of-system and shared infrastructure synergies with the other technologies. These costs were used as the basis for the financial assessment, along with default numbers supplied by SAM, from which HOPP draws its financial models (National Renewable Energy Laboratory 2021; *System Advisor Model Version 2020.11.29 (SAM 2020.11.29).* 2020). Additionally, price files based on Cambium 2023 Mid-Case scenario generation and emissions assessment regions are used for each location for year 2025 (Gagnon et al. 2024). For Scenario 1, no policy incentives were considered, such as the investment tax credit or the production tax credit.

Table 3. Installed costs for the technologies used in Scenario 1

Cost Name	Cost Value
Wind Installed Cost	$1,085,400 / MW
Solar Installed Cost[1]	$746,560 / MW
Storage Installed Cost	$534,018 / MW

[1] $/MW-DC

Each of the five sites in Scenario 1 is examined using the same HPP configuration to compare resources across the sites. The initial configuration contains 100 MW of capacity from wind, 100 MW of capacity from solar energy, and a 200-MWh battery with a power rating (discharge rate) of 50 MW. The plant is connected with a 100-MW interconnection to the grid. A constant load of 30 MW was set as the setpoint for the plant. Based on the similarity of the results between the optimized load following dispatch and the heuristic dispatch in Section 2.1.1, we use the heuristic dispatch strategy for the remaining studies in this section. The results from this study are shown in Table 4 which include the net present value (NPV) of the plant, the average price of electricity in that region (pulled from the Cambium dataset), and the compliance percentage for the load.

The primary criterion for evaluating the HPP design at each location was its compliance percentages. Among the locations considered, the Kansas site emerged as the best candidate for the hybrid plant, as it can already meet the firm power requirement for over 90% of the year with the initial design configuration. This result agrees with the resource assessment of the United States presented in Clark et al. (2022). Furthermore, the Kansas site has the second highest (least negative) NPV value, primarily due to its higher electricity generation. However, the New Hampshire site, despite generating less electricity, benefits from higher electricity sales prices. While the NPV values are interesting for comparison, they are not the primary focus of this scenario, which is concerned with the ability of a hybrid plant to meet a power setpoint. The plant designs are not optimized for profitability, and the financial analysis does not account for non-energy value streams, tax credits, or incentives that could be applied based on the region of construction. Additionally, the heuristic battery strategy is not aimed at maximizing profit, as it does not consider costs. Among the five sites, the Arizona and West Virginia locations perform the worst in meeting the power setpoint, each reaching the setpoint less than 80% of the year.

Table 4. LCOE, NPV, and Compliance Factor for the five sites using the initial hybrid plant design

	Kansas	Arizona	Texas	West Virginia	New Hampshire
NPV ($M[1])	−47	−124	−109	−80	−42
Average electricity price ($/MWh)	35.5	43.9	33.7	52.1	47.5
Compliance (%)	92	76	87	73	85

[1] Million

In view of the variations outlined above between the plant locations, we explore hybrid plant configurations that are optimized to each of the five locations selected. The NREL tool REopt was used to determine the revised configurations (Mishra et al. 2022). The limiting factors in the optimizations were the percent of power provided by renewable energy generation, which was set to the range of 90% – 100% and the total energy capacity of the battery, which was capped at 200 MWh—the same as in the initial plant configuration. Grid charging of the battery was not allowed. The location of the site was entered in REopt, and a residential electricity rate within the region was chosen for the optimization. A constant load of 30 MW was given to the system for an analysis period of 25 years. The default discount rate and electricity cost escalation rate were used. The wind and solar capacity were not capped, and the large commercial wind turbine category in REopt was chosen.

Table 5. Customized plant configurations for each site

	Kansas	Arizona	Texas	West Virginia	New Hampshire
Wind Capacity (MW)	125	135	115	175	170
Solar Capacity (MW)	130	176	111	153	203
Storage Capacity (MWh)	200	200	200	200	200
Storage Rated Power (MW)	31	34	30	34	36

The optimization results are shown in Table 5. While the storage rated power is similar across all sites, the duration of the battery increased from a 4-hour battery in the initial plant to approximately a 6-hour battery in the optimized plants. While it is currently uncommon for a battery duration to exceed 4 hours in utility applications, the battery for the design still has the same overall energy capacity as the previous cases. This means that the storage capacity in these systems could also be achieved by using a 4-hour 50 MW battery by operating the battery differently, though there are differences in the costs associated with changing the rating of the battery which should be taken into account. Across the sites the wind and solar capacity vary widely. The three more western sites (Kansas, Arizona and Texas) all have a much lower wind capacity than the more eastern sites in West Virginia and New Hampshire. The eastern sites also have higher solar capacity than the Kansas and Texas sites. Notably, the Kansas, Arizona, and New Hampshire sites all have a higher solar capacity value than wind, accounting for three out of the five sites studied. The sites in Texas and Kansas were the sites with the most even split between the solar and wind capacities.

Table 6. NPV and Compliance Factor for the five sites using the customized hybrid plant design

	Kansas	Arizona	Texas	West Virginia	New Hampshire
NPV ($M)	−21	−157	−107	−90	−20
Compliance (%)	91	81	87	84	93

The NPV and compliance factor were recalculated using the custom plant configurations for each site, with a constant load set at 30 MW with an interconnection of 100 MW. The results from these runs are shown in Table 6. Note that these results are calculated using the variable price signal based on the Cambium Mid-case scenario. Compared to the results in Table 4, the compliance percentages in Arizona, West Virginia, and New Hampshire improved, with the West Virginia site seeing the largest increase. The compliance factor for the Kansas plant decreased by one percentage point, though it remained above the 90% goal, and its NPV improved significantly. The Texas site saw no

change in its compliance factor, but its NPV improved slightly. This configuration was also the closest to the original plant design. With the optimized configurations, every plant now achieves compliance above 80% for the year studied, though only two plants exceed the 90% threshold.

The goal of this study was to examine the variability and possible configurations of hybrid plants in different locations in the United States. We chose sites in Kansas, Arizona, Texas, West Virginia and New Hampshire to represent different conditions and resource availabilities across the contiguous United States. We then examined the performance of the same plant across each of these locations with a plant optimized to each location, with the goal of meeting a power setpoint for more than 90% of the year. While only two of the locations achieved the setpoint goal, all of the sites were able to meet the power setpoint for more than 80% of the year. Three of the sites significantly improved in the compliance factor with the optimzed design, while the other two plants improved in NPV. This study gives an example of hybrid plant designs across multiple locations in the United States and demonstrates the site-based differences across the designs.

2.2 Scenario 2: Deciding Between Hydrogen and Electricity Production

The second scenario models a hybrid plant that can make a decision between the output of hydrogen or electricity to the wider energy market. This scenario considers a HES comprising an offshore wind farm, a battery and an electrolyzer plant. The plant is placed off the coast of Texas and Louisiana, shown in Figure 3(a). The wind farm was placed here to take advantage of existing natural gas pipeline infrastructure in Louisiana to transport the hydrogen produced, where the hydrogen could be blended with natural gas up to a certain concentration, or perhaps use refinished pipes for pure hydrogen transportation..

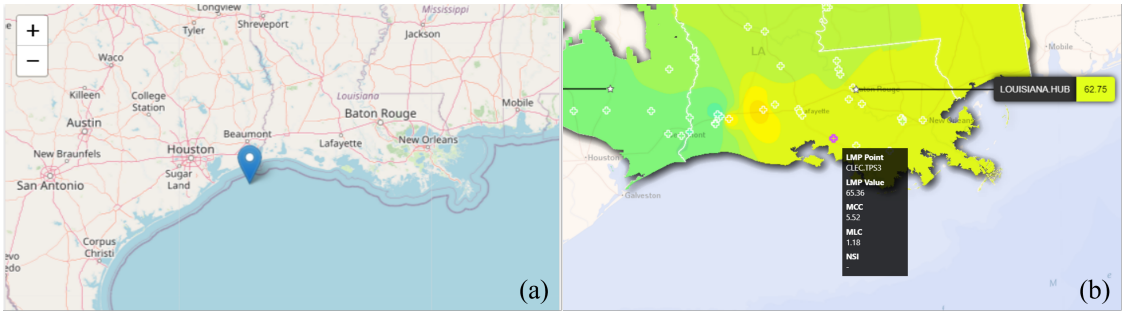

Figure 3. (a) The location of the offshore wind farm in Scenario 3 and (b) the area the local marginal prices (LMP) in $/MWh for electricity were chosen from (*LMP Contour Map* 2022), where the color gradient represents the LMP values, MCC represents the Marginal Congestion Cost of the selected node, which is the cost of congestion compared to the load average of the system prices, and MLC is the Marginal Loss Cost, which is the cost of marginal losses that occur as a result of long-distance transport to a specific node which varies by location.

The HES consists of a 100-MW offshore wind farm, using Technical University of Denmark 10-MW reference turbines (Bortolotti et al. 2019). The configuration also uses a 50-MW, 4-hour battery. The plant includes 100-MW electrolyzer capacity. The price for hydrogen was set at $2.00/kg hydrogen as a mid-range value from information from the Global Hydrogen Review 2021 (International Energy Agency (IEA) 2021) that would make hydrogen produced from renewable energy sources competitive to hydrogen produced using natural gas. The electricity prices used were local marginal prices (LMP) from the site highlighted in Figure 3(b) from 2019. A 200-hour sample of the hourly LMP from this location in the market of the Midcontinent Independent System Operator (MISO) is shown in Figure 4. The hydrogen electrolyzers were modeled using the setup in (Clark et al. 2023), based on a proton exchange membrane (PEM) electrolyzer.

The simulation was run for a year and the profit from the different scenarios was be compared. The simple heuristic dispatch model was used for the battery, and a setpoint signal of 15% of the electrolyzer capacity was given for the plant to achieve. Since PEM electrolyzers are not operated on less than 5% of their power capacity, the value used is the lower limit of electrolyzer operation range plus a 10% buffer.

Figure 5 shows a 200-hour snapshot from the year of simulation. Figure 5(a) shows the total hydrogen that could have been produced (black line) and the hydrogen was sold (red dashed line). Figure 5(b) shows the profile of the electricity sold to the grid instead of hydrogen. The hydrogen and electricity sold profiles are inversely proportional

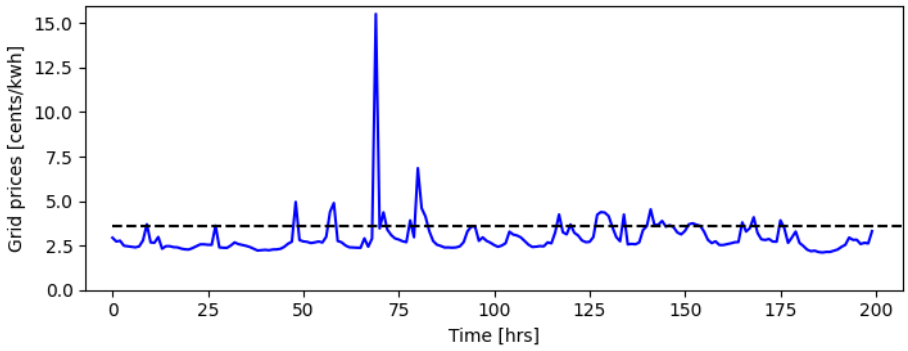

Figure 4. Electricity prices (blue) for 200 hours in 2019 compared to the cut-off point (black dashed) above which it would be more profitable to sell electricity.

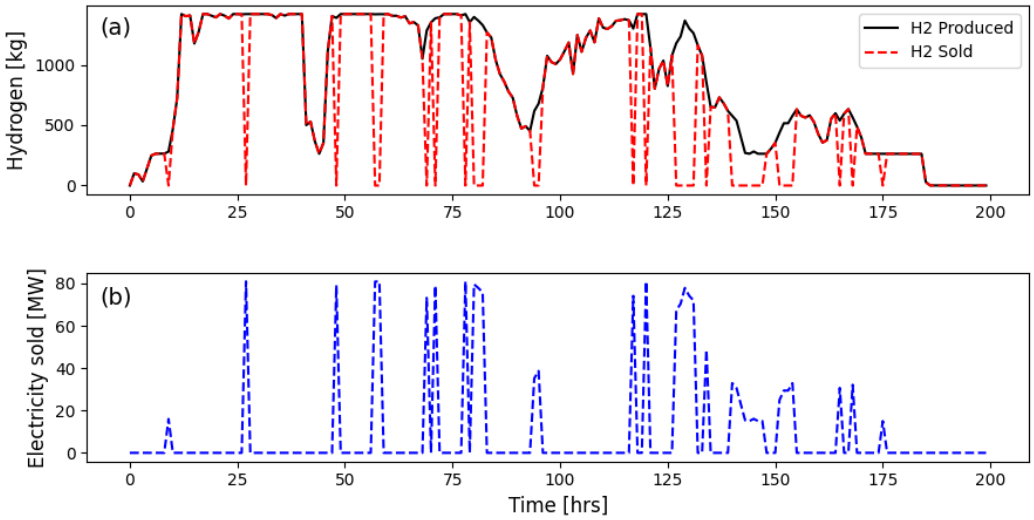

Figure 5. (a) Hydrogen that was sold (red dashed) compared with the total kilograms of hydrogen that could have been produced (black), and (b) the electricity sold in megawatts to the grid.

to each other because it is assumed that the plant is always exporting either hydrogen or electricity. The periods where electricity is sold instead of hydrogen also correspond with higher electricity price spikes above the cutoff point shown in Figure 4, showing that the plant chooses the option that gives the higher profit.

Table 7. Profit Values for the Different Hydrogen and Electricity Export Options

Case	Yearly Revenue [$M]
Only Hydrogen	8.7
Only Electricity	6.3
Hydrogen + Electricity	9.1

The purpose of this scenario was to maximize the profit of the plant by choosing between exporting hydrogen or electricity to the energy market. Table 7 shows the profit values for the year if the plant only exported hydrogen, if the plant only exported electricity, and if the plant exported a combination of hydrogen and electricity, in millions of dollars. Adding electricity export to the hydrogen-only case results in a 5% increase in the profit for the year, and adding hydrogen to the electricity-only case results in a 44% increase in the profit for the year. This is a proof of concept study that uses HOPP to model a situation where a hybrid plant needs to choose between the production

of hydrogen or electricity with the goal of maximizing profit. Future work in this area will involve more in-depth modeling of the hydrogen system and incorporation of forecasting into the decision to sell hydrogen or electricity.

3 HOPP Validation

The results presented in the previous two sections have been from simulation studies. To gain more confidence in their findings, it is important to validate the modeled PV array and wind turbine power output of HOPP against other sources and quantify the error in the modeling. To this end, HOPP was validated against two different types of data:

1. **Unit-scale data** from ARIES (Advanced Research on Integrated Energy Systems), NREL's hybrid energy system research platform
2. **Utility-scale data** from confidential sources.

Unit-scale data validates HOPP's modeling of power output from individual generating units feeding into a larger utility-scale plant, i.e., a single wind turbine or a PV array feeding a single inverter. These are the assets that will be available for field testing of the battery dispatch, and HOPP's modeling of them must be validated in advance of the field test. The following ARIES assets, located at NREL's Flatirons Campus, were modeled in HOPP:

- A 0.43-MW research PV array built by NREL in conjunction with First Solar consisting of First Solar Series 4 105-W modules. This PV array was previously detailed in an NREL report (Gevorgian et al. 2022).
- A 1.5-MW research wind turbine developed by the U.S. Department of Energy (DOE) based on the platform of GE's 1.5-MW SLE commercial model. The turbine was previously detailed in an NREL report (Mendoza et al. 2015).

Utility-scale data validates HOPP's modeling of power output from an entire power plant, including losses due to ground cover ratio in large PV plants and wake losses in large wind plants. These are commercial assets that are owned by utilities, generating and selling power to a local grid. The location and identity of assets must be kept confidential because their data are proprietary. No time-history data from the assets will be published; only the results of the validation will be published. These assets are a ∼50-MW unidentified solar farm and a ∼100-MW unidentified wind farm.

There is inherently a slight gap in performance between an ideal modeled asset and a real-world asset that has been in service for years and experienced some degradation. One goal of the validation task was to build in the ability of HOPP to "tune" its modeled power output to match measured output by adjusting model coefficients that reflect non-ideal performance. To tune HOPP's PV and wind modeling to each asset, the following "tuning coefficients" were adjusted:

- For PV: Inverter efficiency—the percent efficiency at which DC power output is converted into the measured AC power output at a PV inverter. The default value for this parameter is 96% in HOPP.
- For wind: Turbine-specific loss—the average percent of power that each turbine underproduces vs. the supplied power curve. The default value for this parameter is 0% in HOPP.

Once HOPP is tuned and validated against each asset's measured performance, the simulated output of HOPP will also be compared against the output of HOMER, a state-of-the-art hybrids simulation software from UL Solutions. We will compare HOPP's simulated output with simulations produced by HOMER using the same resource data and simulation parameters. The same tuning parameters will be adjusted in HOMER, and their final values compared between HOPP and HOMER. Once these parameters have been adjusted to minimize error in both simulations, the relative accuracy of both will be compared.

To evaluate the error between simulated power output, P_{sim}, and the measured data, P_{meas}, two metrics are used: mean error and mean absolute error. Mean error or average error, represented by ME, is simply the average of the error at each data point i, with N total data points reported for each asset. These data points consisted of average power measured over 1 hour, matching the intervals at which HOPP and HOMER were simulating power profiles. The final formula for ME is thus:

$$ME[\text{kW}] = \frac{\sum_{i=1}^{N}(P_{sim} - P_{meas})}{n} \tag{3.1}$$

To scale *ME* to the relative output of each asset, *ME* is divided by the average measured power output, producing an *ME* in units of percentage rather than kW:

$$ME[\%] = \frac{\sum_{i=1}^{N}(P_{sim} - P_{meas})}{\sum_{i=1}^{n}(P_{meas})} \tag{3.2}$$

ME is the metric used to adjust the tuning coefficients. HOPP simulations were run iteratively, with the tuning coefficient incremented until *ME* (%) had an absolute value of less than 0.1%. However, even an *ME* of zero does not mean HOPP simulated the asset perfectly. For example, if one data point was simulated to produce 5 kW, and a second data point was simulated to produce 3 kW, but the true output was 4 kW at both, the *ME* would be zero. The model would be perfectly tuned, but not ideally modeling the asset. Thus, a metric is needed to measure *absolute* error once the model is tuned. The mean absolute error, *MAE*, applies a small alteration to the *ME* equation by changing the error calculation to an absolute value:

$$MAE[\text{kW}] = \frac{\sum_{i=1}^{N}|P_{sim} - P_{meas}|}{n} \tag{3.3}$$

$$MAE[\%] = \frac{\sum_{i=1}^{N}|P_{sim} - P_{meas}|}{\sum_{i=1}^{n}(P_{meas})} \tag{3.4}$$

MAE is the metric used to evaluate the overall fidelity of the model to the measured data. Ideally, the *MAE* values produced by HOPP should be about the same or lower than those produced by HOMER. However, *ME* is still useful for breaking down the errors into "bins" to see if there is model bias based on any of the input parameters. When *ME* is binned, the error is only averaged over the subset of the data in which a parameter p is within limits $binstart$ and $binend$. This subset of of data will only contain N_{bin} data points, with the equation for *ME* becoming:

$$ME_{binned} = \frac{\sum_{p \geq binstart}^{p < binend}(P_{sim} - P_{meas})}{N_{bin}} \tag{3.5}$$

For example, if the binning parameter is the year in which the data were taken, the equation becomes:

$$ME_{2021} = \frac{\sum_{year \geq 2021}^{year < 2022}(P_{sim} - P_{meas})}{N_{2021}} \tag{3.6}$$

ME_{binned} indicates how consistently HOPP models power output across changing input parameters. Positive ME_{binned} indicates HOPP is overestimating power given a certain input parameter, and negative ME_{binned} indicates HOPP is underestimating power. The input parameters that will be used to bin *ME*, where applicable, are:

1. Year
2. Month
3. Time of day
4. Power output level
5. Ambient temperature
6. Wind direction (wind only).

The following summarize these error evaluation parameters:

- *ME* (Mean Error): Used to evaluate the tuning of the model. Tuning coefficients are adjusted to bring *ME* close to zero.
- ME_{binned}: Used to evaluate consistency of the model across changing input parameters.
- *MAE* (Mean Absolute Error): Used to evaluate the overall fidelity of the model to the real measured power output and compare the fidelity of HOPP and HOMER.

The ability to take in validation data from the available assets, compute the error, and make these breakdowns is being built into HOPP as a new feature as part of this project. This will allow further validation of HOPP against other assets in the future.

3.1 Unit-Scale Data From NREL ARIES

3.1.1 0.43-MW PV Array

The first asset that HOPP was tuned to was the First Solar PV array at ARIES. Although global horizontal irradiance (GHI) data were available for the PV array since 2018, a direct normal irradiance (DNI) sensor was only installed at the site in June 2021. Thus, the output of the array could only be accurately modeled from June 2021 onward. Incremental adjustments to the tuning coefficient of inverter efficiency are shown in Figure 6, with the value of 95.1% minimizing the *ME*.

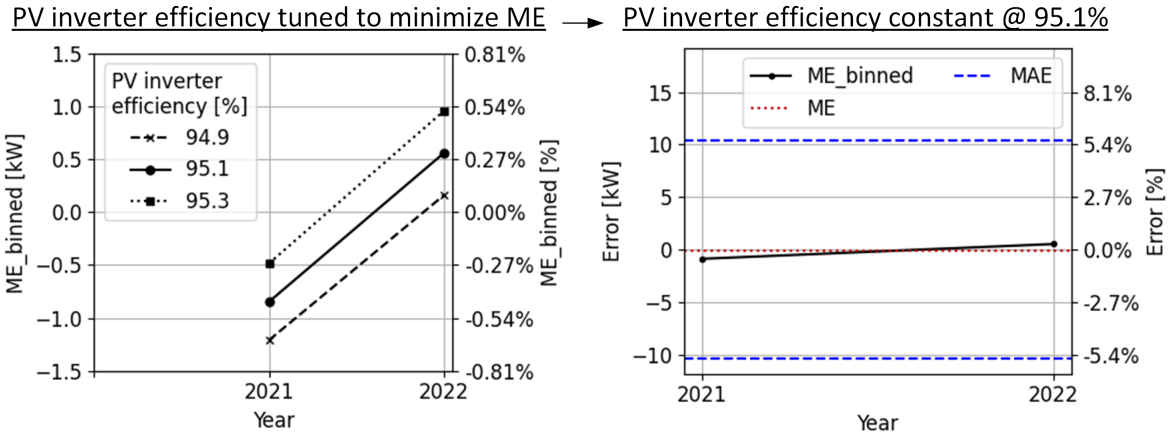

Figure 6. Tuning of HOPP-modeled PV inverter efficiency, with *ME* broken down into yearly bins

Figure 6 also gives an example of how *ME* breakdowns will be shown in following figures. Even though the absolute value of *ME* over the entire data set is less than 0.1%, it will be significantly more when broken down into yearly bins, as seen on the left side of Figure 6. However, the yearly *ME* breakdown is still very small when compared to the *MAE*, as seen on the right side of Figure 6. This indicates that there was no significant bias in the model between years.

However, performing the same comparison of ME_{binned} to *MAE* on the other parameters indicates more bias, as shown in Figure 7. In the monthly breakdowns (top left), HOPP underestimated power during the summer months, but overestimated it during the spring/fall months (no data were available from the winter months of January and Feburary). HOPP also underestimated power during the morning and overestimated power in the afternoon (top right). The measured output level (bottom left) also biased the *ME*, with higher output levels being underestimated and lower output levels being overestimated. All three of these seem to correlate with the temperature bias (bottom right), which shows that power was overestimated at lower temperatures and underestimated at higher temperatures.

To determine whether these errors are due to problems with the HOPP modeling, we compared them with the same results from HOMER. The PV inverter efficiency was retuned for HOMER, with a very similar value of 94.7% minimizing the *ME*. As shown in Figure 8, the shapes of the HOMER *ME* breakdowns are very similar to those from HOPP, indicating the two models are in agreement. Moreover, the *MAE* from HOMER (8.32%) is greater than that from HOPP (5.58%), indicating that HOPP achieved higher fidelity to the measured data, at least with this one asset.

The difference in MAE may be due to HOMER having a more pronounced temperature bias, at least with its default settings, compared to HOPP. HOMER's temperature ME_{binned} ranges from +18 kW to −10 kW, whereas HOPP's only ranges from +8 kW to −4 kW. HOPP, via SAM, features a default tool to calculate PV cell temperature using the Sandia PV Array Performance model (King, Kratochvil, and Boyson 2004), which was used in the validation process. More advanced tools are available in SAM that can be activated through additional integration with HOPP in the future (Gilman et al. 2018). HOMER also features a tool to calculate temperature effects on power, but this

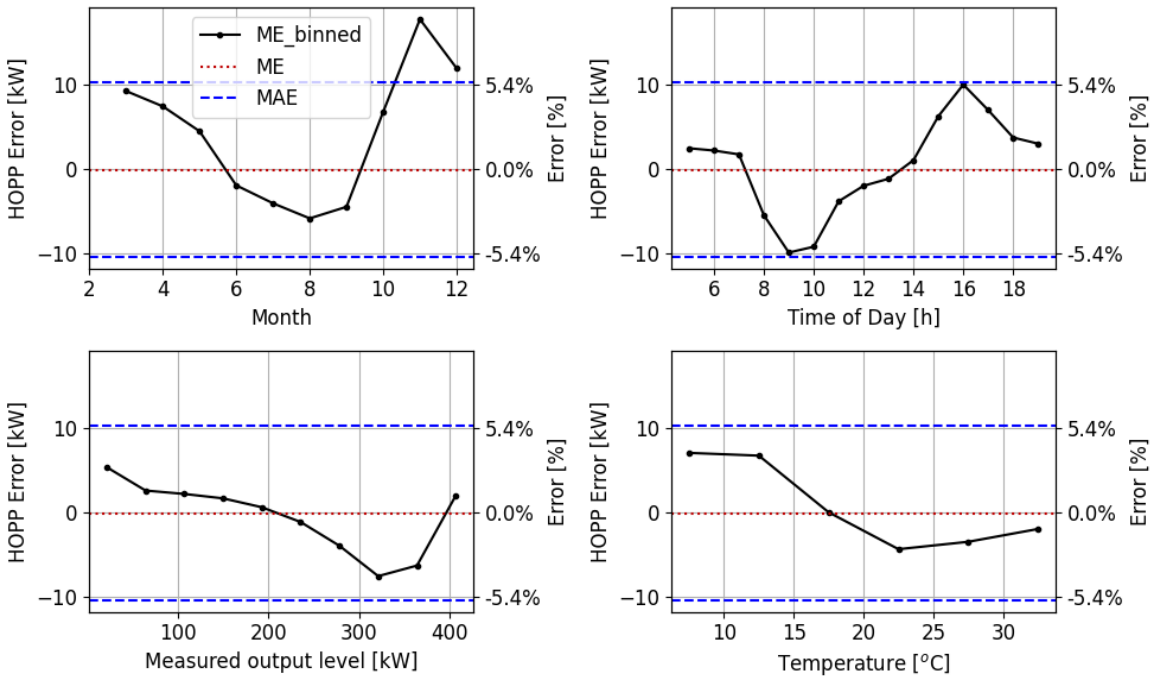

Figure 7. Breakdown of HOPP-simulated PV errors, unit-scale data

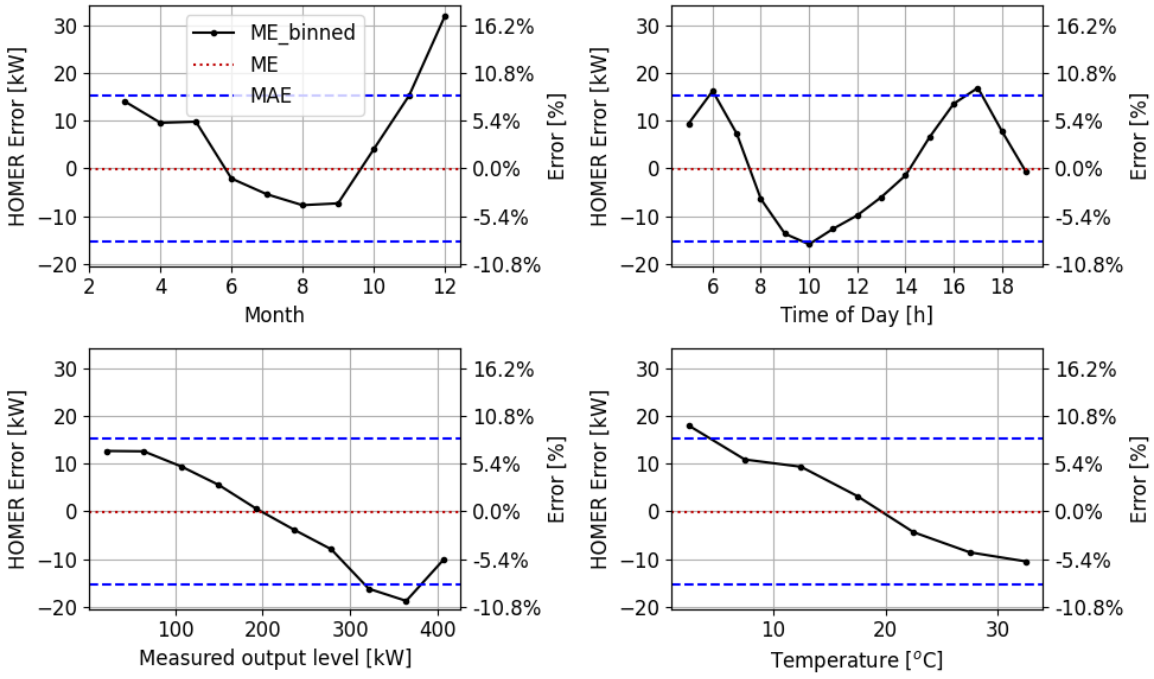

Figure 8. Breakdown of HOMER-simulated PV errors, unit-scale data

consists of a manually entered table of percent power loss per degree Celsius away from a nominal cell temperature. For this validation, we used HOMER's default values of 0.5%/°C and 47°C nominal temperature. Further work could have been done to optimize these parameters and improve the performance of the HOMER modeling, but this was outside the scope of the work. The goal was only to evaluate the baseline performance of both models, and given

HOPP's comparable performance to HOMER, we consider HOPP's PV modeling at the unit scale to be validated by these results.

3.1.2 1.5-MW Wind Turbine

The same tuning procedure was conducted on the turbine-specific loss coefficient for HOPP modeling of the 1.5-MW ARIES wind turbine. HOPP and HOMER tuning both found relatively high loss coefficients of 1.4% and 1.7%, respectively, which was surprising given the performance of the turbine. The *MAE* was found to be relatively high for both models, above 15%. To investigate, the binned *ME* is broken down in Figures 9 and 10.

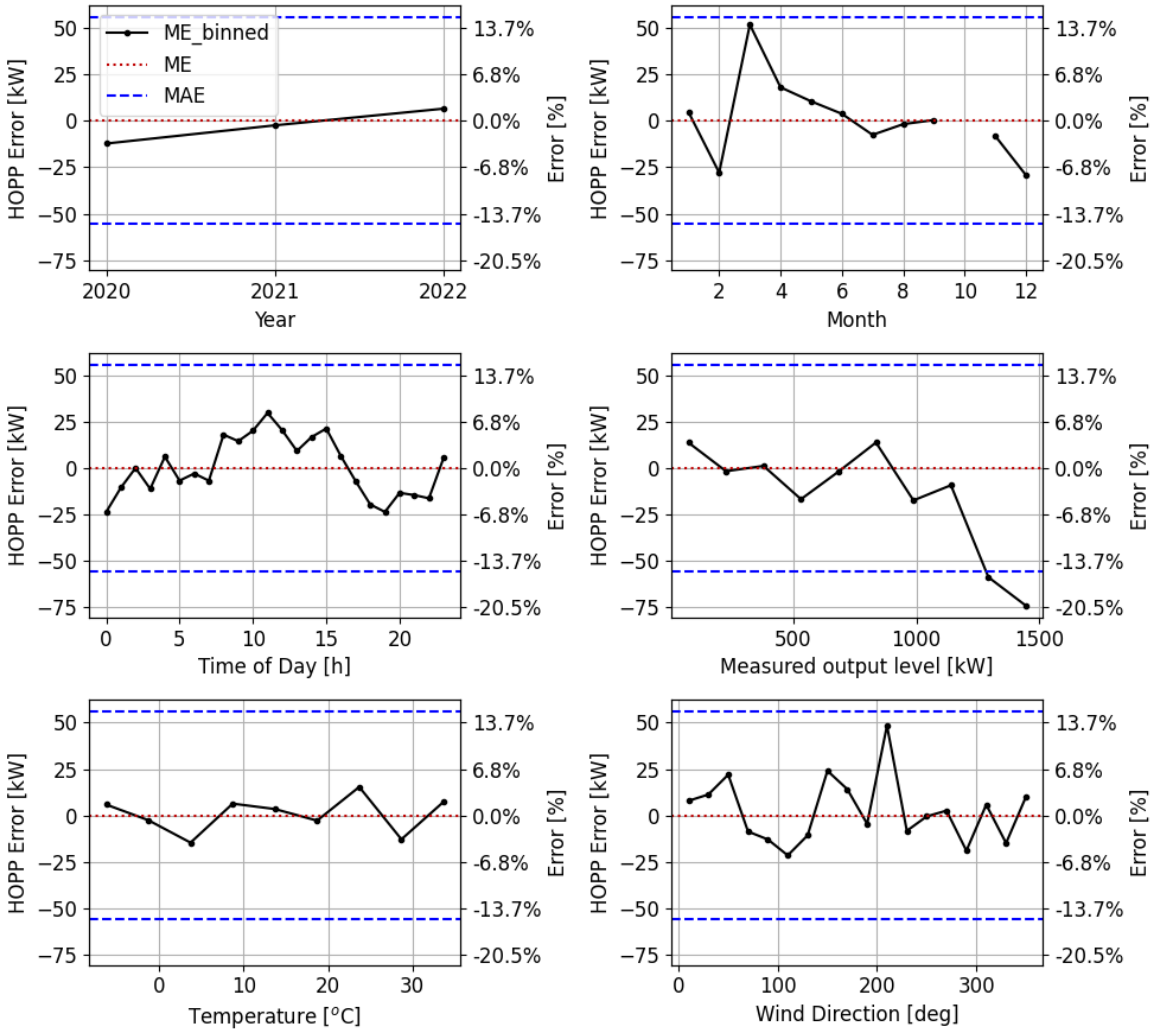

Figure 9. Breakdown of HOPP-simulated wind errors, unit-scale data

Once again, as in Figures 7 and 8, HOPP and HOMER appear to be in good agreement with each other, as the shape of their error breakdowns is extremely similar. The wind direction is a large potential source of modeling error due to the nature of the geography of the Flatirons Campus. For example, the spike in error in the bin from 200 to 220 degrees is because three other turbines lie directly upwind from the 1.5-MW turbine in this direction. The chaotic climate of Colorado also comes into play, with errors being much larger during the gusty winter months than in the relatively calm summer months. Regardless, the simpler explanation for the inability to model the turbine's performance is the complexity of the research being performed with it. Although the data were filtered with a status code meant to indicate whether the data were suitable for International Electrotechnical Commission (IEC) power testing, there were still some research activities that affected the turbine's performance that were only recorded in

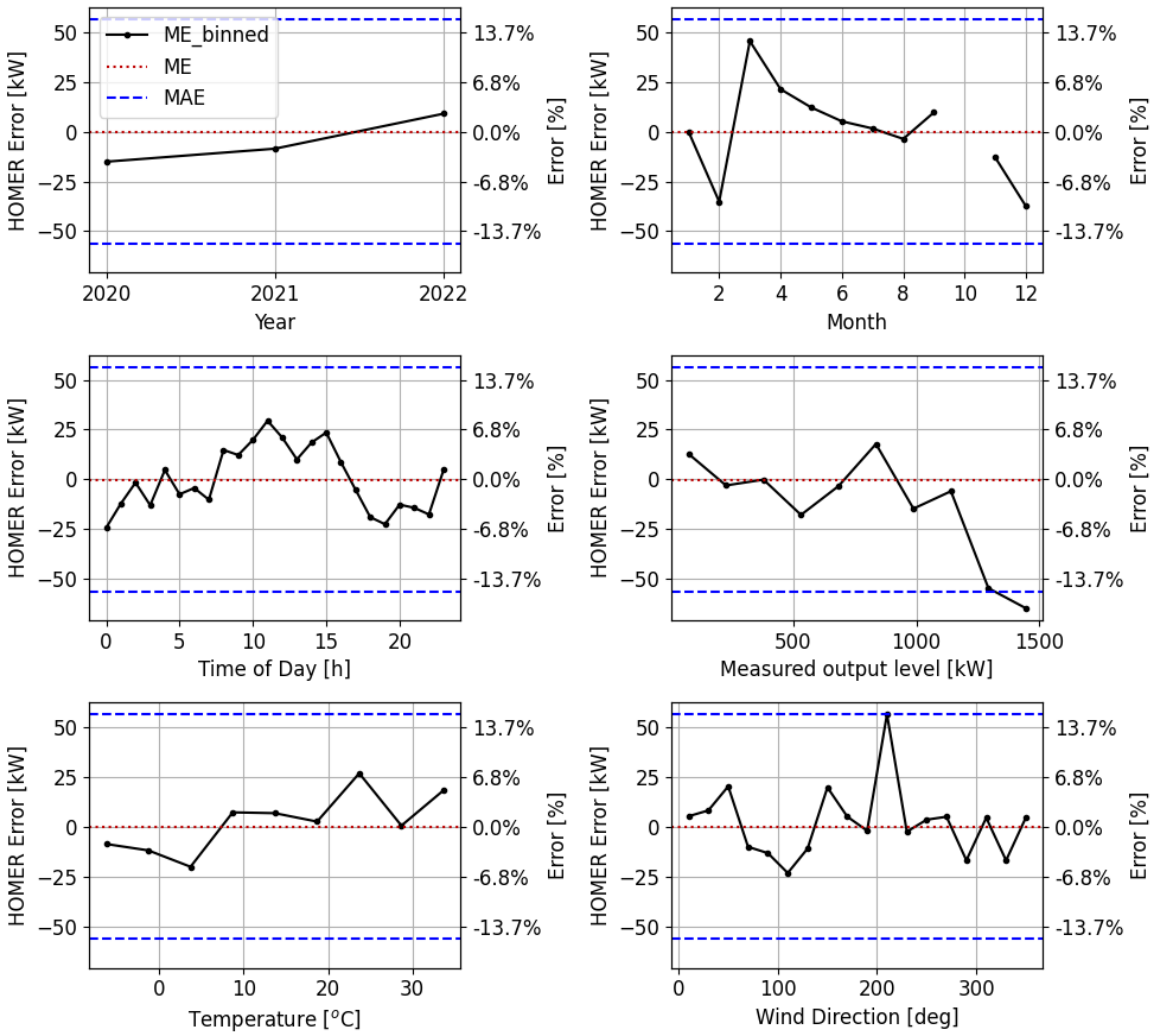

Figure 10. Breakdown of HOMER-simulated wind errors, unit-scale data

manual logs (Roadman 2023).This is reflected by the noticeable overestimation of the turbine's performance during the working hours of 8 a.m. to 4 p.m. (hours 8 through 16 on the x-axis of the middle left graph in Figure 9).

Another way to view the data is by plotting the wind speed data against both the measured and modeled power output, also known as the power curve. Two power curves are shown in Figure 11, with the modeled power output in blue dots and the measured output in orange x's. In the full set of data available since the turbine's commissioning (left), it can be seen how most of the measured data surround the modeled power curve, but with some "dirty" outlier data falling well below the curve. These outliers could not be automatically filtered out, and they caused the tuned turbine loss to be larger than it realistically should have been. However, we were able to identify a smaller, clean subset of data (right) in July and August of 2022. Here, the turbine was operating consistently with no modifications to its configuration to cause it to dip beneath the power curve.

The smaller subset of data, occurring in summer months with calmer winds, had very few points where the hourly power output was above 800 kW. Concerned that insufficient data at high power levels would influence the tuning coefficient, we only used data from 800 kW and below for the final tuning and validation.

With this smaller data set, there were insufficient data to break down the error into bins as with the full data set, so the comparison with HOMER is shown in another form. In Figure 12, the measured wind power output is plotted on the x-axes while the simulated power output is plotted on the y-axes. Each hourly data point is plotted as a blue dot.

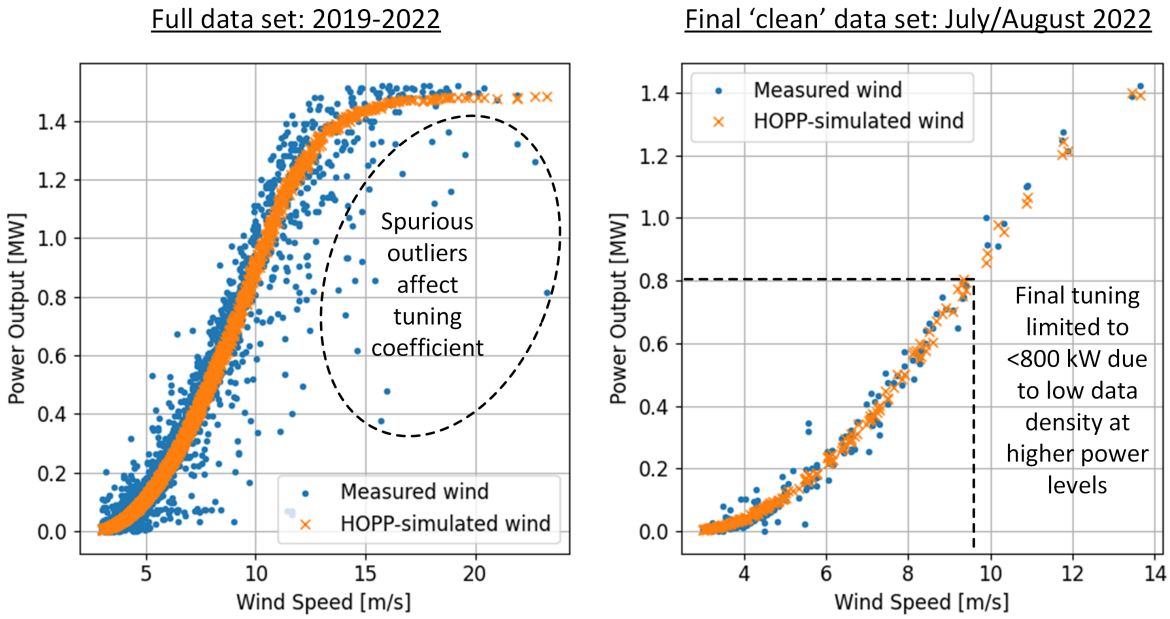

Figure 11. Power curves for NREL 1.5-MW turbine, with (left) full data set and (right) down-selected "clean" data set

A perfectly simulated data point will fall on the solid green line. A point in which the power is overestimated will fall above the solid line, and a point in which power is underestimated will fall below the solid line. The dashed line represents a linear fit to the plotted points, and due to the adjustment of the tuning coefficients, this line lies perfectly on the diagonal for both HOPP (left) and HOMER (right).

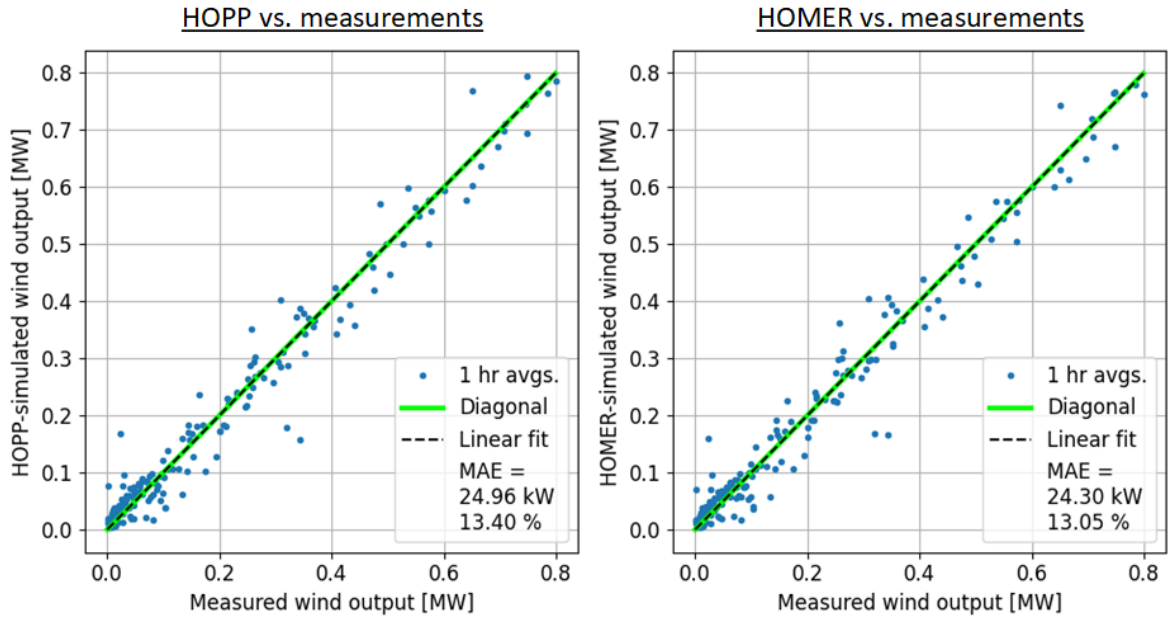

Figure 12. Comparison of measured vs. simulated power output from NREL 1.5 MW turbine, with cleaned data set. Simulated with HOPP (left) and HOMER (right).

Figure 12 shows how HOPP and HOMER both achieve a very similar performance with this reduced data set from the ARIES 1.5-MW turbine. Another way to to show a portion of this data is in a time series. Although the whole data set would be unreadable if shown on the width of one page, a small subset of 5 days is shown in Figure 13. The figure shows how the simulated power output from HOPP and HOMER track well with each other, and even though

they deviate at times from the output of the turbine, the deviations tend to be small. We consider this a sufficient validation of HOPP's modeling capabilities of an individual wind turbine, and the ARIES hybrid wind turbine/PV array system. This system will later be used for field testing of the controls optimization strategy from Section 1.1.1, using wind and solar data from the time period shown in Figure 13.

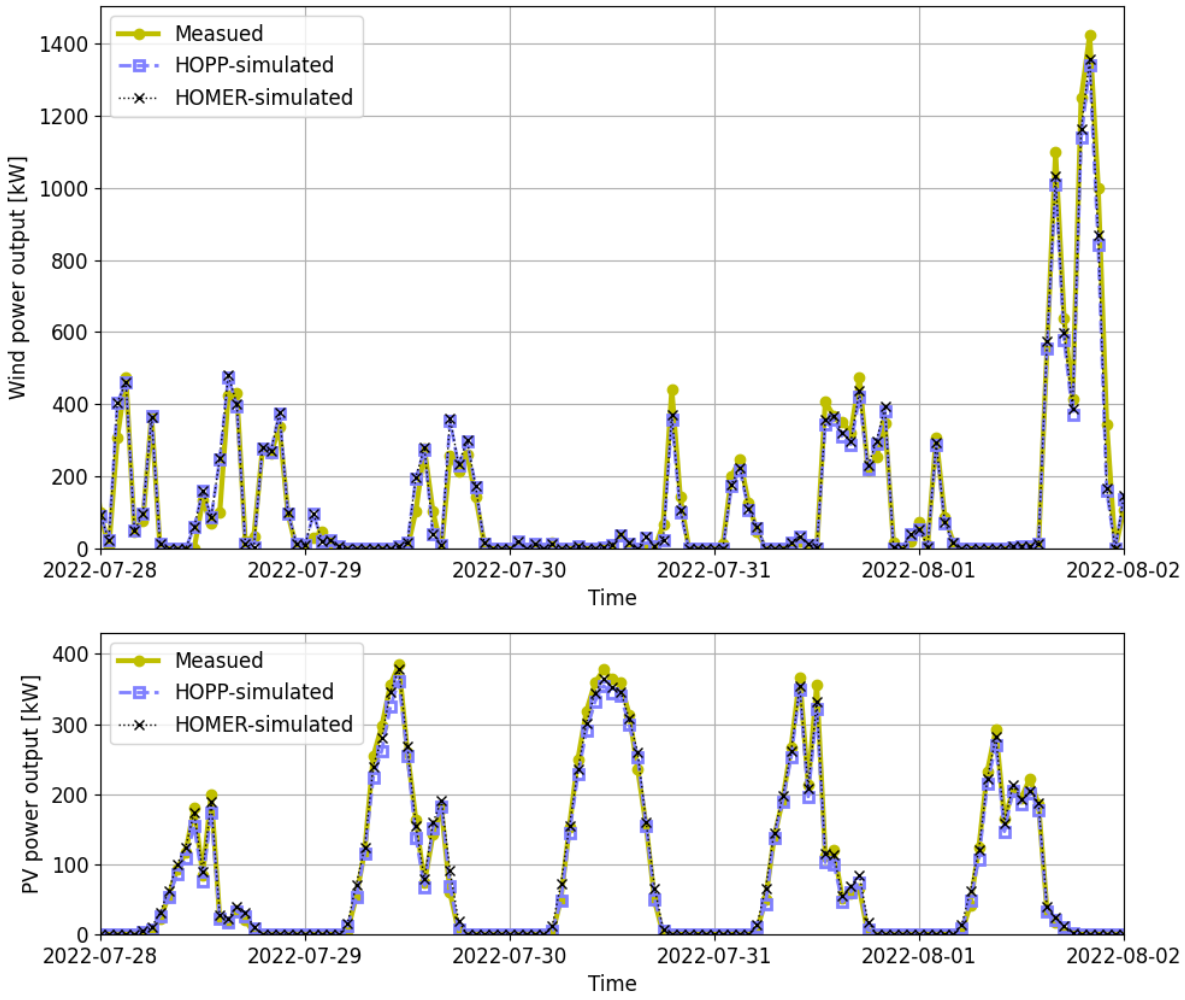

Figure 13. Time-history comparison of measured power output, HOPP-simulated power output, and HOMER-simulated power output for NREL 1.5-MW wind turbine (top) and 0.43-MW PV array (bottom)

3.2 Utility-Scale Data

Once HOPP's power output simulation was validated on a unit scale, we sought data from outside sources to validate the scale of a full utility-scale power plant. Ideally, a co-located wind and solar hybrid plant would have been used, but such data could not be obtained. Instead, separate sources were found for a utility-scale PV plant and wind plant. These data are proprietary—the data sources will not be disclosed, but they are both currently operating power plants within the continental United States providing commercial power.

3.2.1 Solar Farm

The utility-scale solar farm source cannot be specified further than to say it had a capacity of approximately 50 MW. We obtained 4 months of power data from dozens of ∼1-MW inverters, as well as direct irradiance measurements from throughout the plant. We were therefore able to accurately model the plant's output in both HOPP and HOMER, although once again, there appears to be some temperature bias, as shown in Figures 14 (for HOPP) and 15 (for HOMER). However, unlike with the unit-scale NREL PV array, there appears to be no appreciable difference in the temperature bias or in the overall *MAE*, which came out to 4.2% of the mean output for both models. The

temperature bias appears to be the root cause of the monthly bias, with the cooler ambient temperatures and reduced incident angles in the winter months leading to lower cell temperatures. Note: the ambient temperatures in the lower right of Figures 14 and 15 have been redacted because the range of these temperatures and their correlation with the month could be used to identify the location of the PV plant.

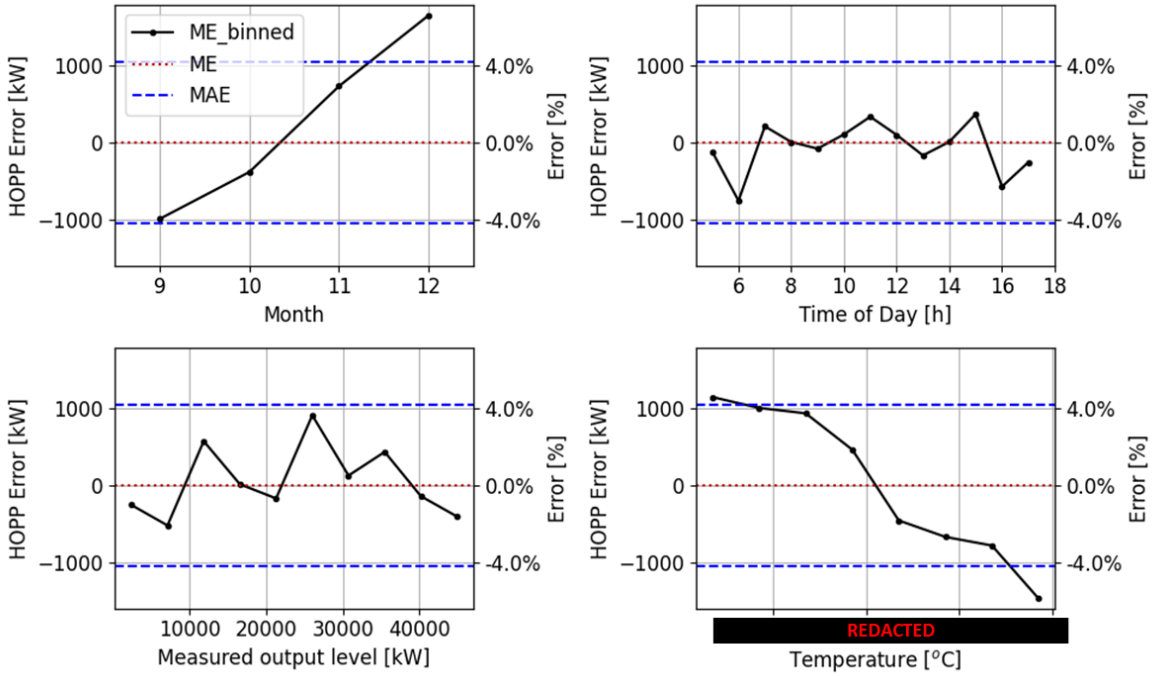

Figure 14. Breakdown of HOPP-simulated PV errors, utility-scale data

The only ME_{binned} breakdown that was noticeably different between HOPP and HOMER was in the dawn/dusk hours of 6–7 a.m. and 4–5 p.m. (indicated as hours 6 and 16 on the x-axis of Figure 14, upper right). There is noticeable underestimation of power output in these hours by HOPP, but not by HOMER. We sought to investigate whether this could be due to conflicting models of the solar angle of incidence (AOI) used by HOPP and HOMER. The AOI calculation done by HOPP is detailed by the SAM user manual (Gilman et al. 2018), but it is unknown how the private software HOMER calculates its AOI.

As a check on HOPP's AOI calculation, the HOPP-calculated AOI values for this data set were compared to those calculated by the NOAA Solar Calculator (NOAA ESRL Global Monitoring Laboratory 2023). The time-of-day ME_{binned} in the AOI between HOPP and the NOAA calculator is shown in Figure 16. The pattern tracks very close to the pattern in HOPP's time-of-day power output error breakdowns in Figure 14. Consequently, we suspect that the AOI modeling in HOPP could be improved to eliminate the time-of-day bias.

However, the power produced in these dawn and dusk hours is a very minimal share of the total power produced by PV. To get an overview of the quality of the PV modeling of HOPP and HOMER side by side, the same measured-vs.-simulated power output plots that were first shown in Figure 12 are shown again here in Figure 17, this time for the utility-scale PV plant. These plots show a much tighter fit to the diagonal, consistent throughout the range of power output. There does not appear to be a noticeable difference in the fidelity of the HOPP and HOMER power modeling on the grand scale, and we consider the HOPP model validated for utility-scale PV modeling.

3.2.2 Wind Farm

Utility-scale wind farm data with on-site wind resource measurements that could be used for this validation effort were difficult to come by. However, our contacts were able to provide a confidential source of 10 days' wind speed and power output data from a commercial wind farm. The combined capacity of these turbines was approximately 100 MW, but the exact capacity, number, and model of turbines will not be given to keep the farm anonymous.

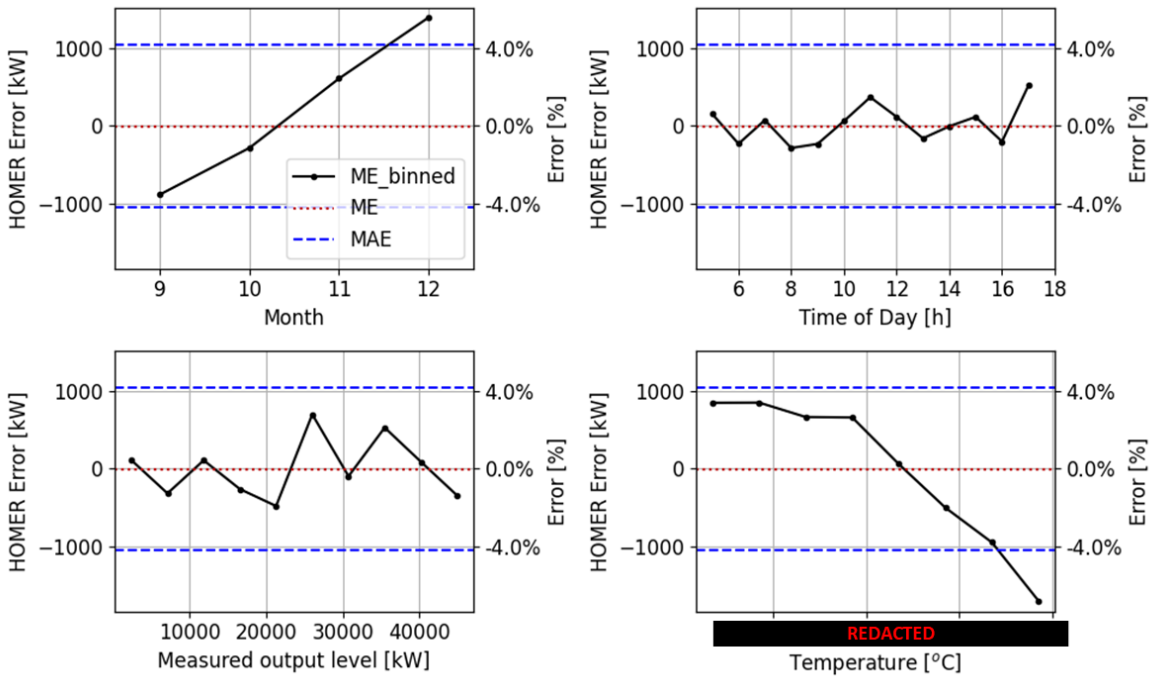

Figure 15. Breakdown of HOMER-simulated PV errors, utility-scale data

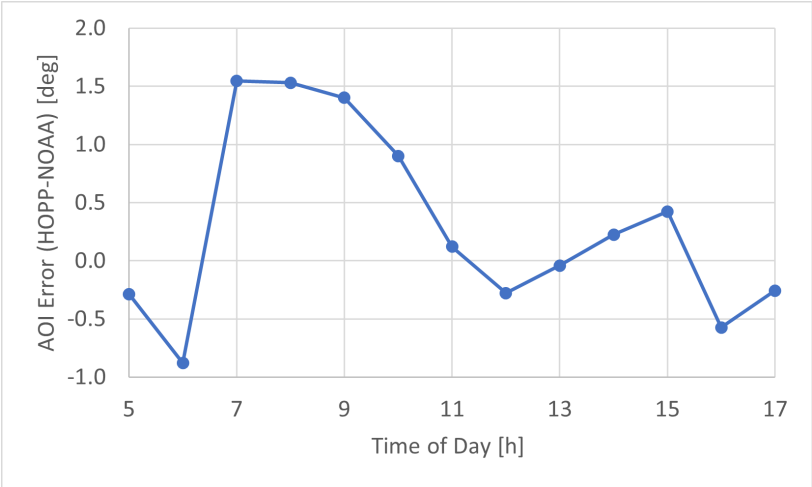

Figure 16. Breakdown of angle of incidence error when comparing HOPP to NOAA Solar Calculator (NOAA ESRL Global Monitoring Laboratory 2023)

Likewise, although specific numbers are given for the *ME* and *MAE* in terms of percentage, those numbers in terms of kilowatts have been redacted.

Unlike the unit-scale individual turbine modeling, the utility-scale wind farm modeling needed to model wake losses as well as a turbine-specific loss coefficient. The locations of the individual turbines were not known, so an assumed turbine configuration was used based on the overall area of the plant and a turbine location optimization routine built into HOPP from previous work (Stanley and King 2022). HOPP, via SAM, features a default tool to calculate wake loss using the Simple Wake Model (Freeman et al. 2014), an adaptation of a 1996 University of Wisconsin model (Specialists 2013), which was used in the validation process. HOMER did not feature a similar tool for calculating wake loss, only a single field for users to input their own calculated wake loss coefficient. Thus, for the validation, we used the HOPP-calculated wake loss coefficient of 19.76% as an input to HOMER. On top of the wake loss, turbine-specific loss also had to be tuned to each model to minimize *ME*. The turbine-specific loss was found to

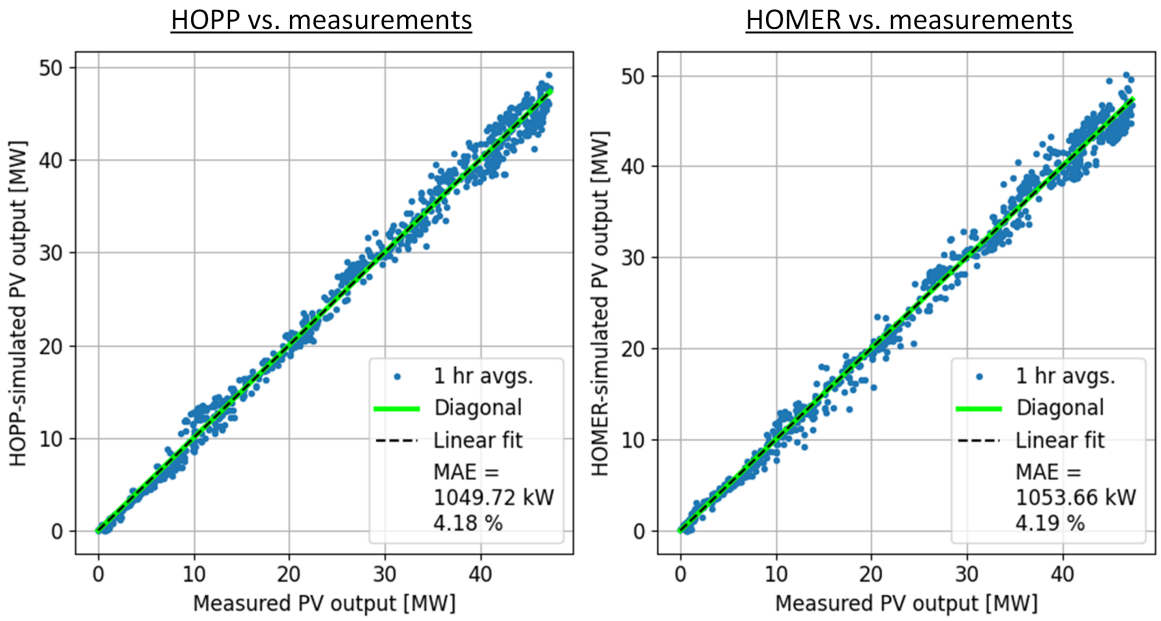

Figure 17. Comparison of measured vs. simulated power output from utility-scale PV plant. Simulated with HOPP (left) and HOMER (right).

be 0.2% for HOPP and 0.6% for HOMER. With the tuned model, the lowest *MAE* values (in terms of percentage) for the entire study were found—3.2% for HOPP and 2.6% for HOMER. However, these numbers should be taken cautiously given the limited amount of data (10 days' worth) over which the validation occurred.

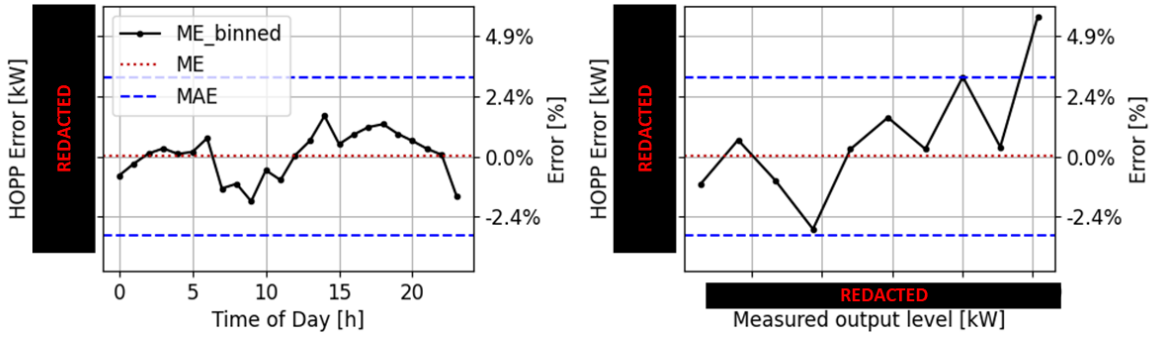

Figure 18. Breakdown of HOPP-simulated wind errors, utility-scale data

Although the breakdowns of ME_{binned} that can be conducted on this data set are limited compared to the unit-scale data, there are still two valuable comparisons: time of day and power output level. The time-of-day-based ME_{binned} does not appear to differ noticeably between HOPP and HOMER, which helps validate the modeling capabilities of HOPP. However, the output-level-based ME_{binned} slopes upward with power output in HOPP, shown in Figure 18 (right), whereas the same slopes downward in HOMER in Figure 19. This means that while HOPP is overestimating the power output of the plant at its peak output, HOMER is underestimating power output at the peak, which may be due to the incomplete modeling of wake losses, mentioned earlier. HOPP does not have complete information on the true position of the turbines to make its calculation of wake loss; thus, there may be higher wake losses at higher power levels than HOPP is calculating. HOMER, on the other hand, only takes in a single wake loss coefficient, and it is not clear how complex of a wake loss calculation it is performing.

However, when plotting the measured-vs.-simulated data side-by-side in Figure 20, it becomes clear that the divergence between HOPP (left) and HOMER (right) only occurs at a handful of points. The divergence of ME_{binned} at higher power levels in the previous figures is not a robust trend, and more than 10 days' worth of data is needed to confirm it. But what is also clear from Figure 20 is that this asset can be modeled with high fidelity by HOPP, more

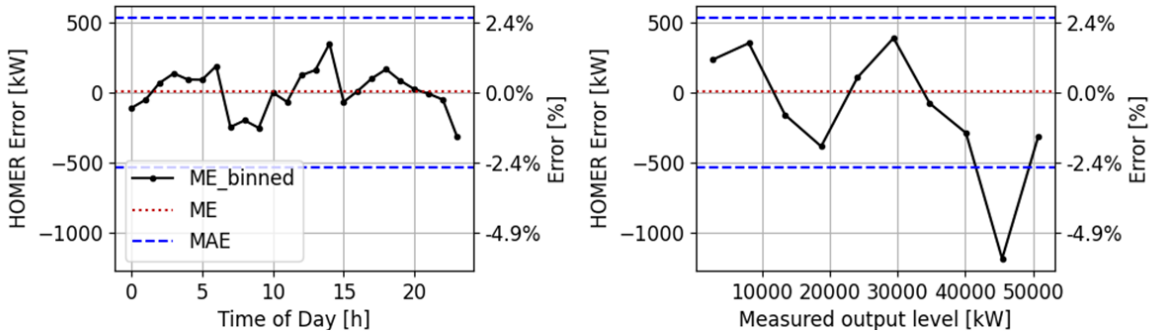

Figure 19. Breakdown of HOMER-simulated wind errors, utility-scale data

so than the utility-scale PV plant, even without detailed turbine location data. We consider this a strong validation of HOPP's abilities to model the power output of utility-scale wind power plants.

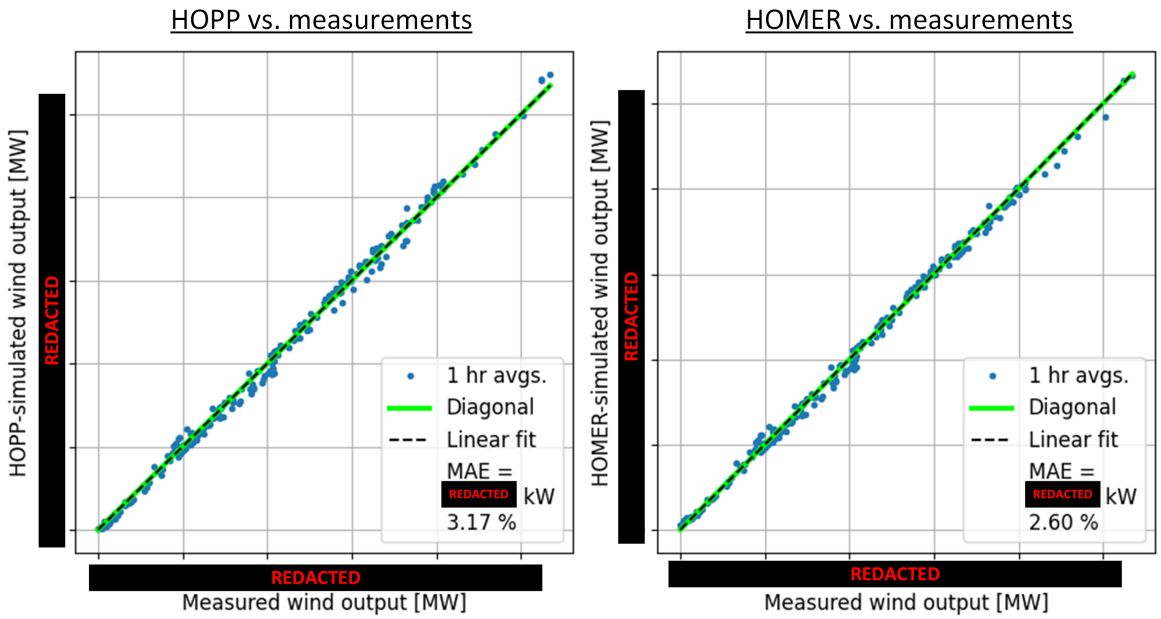

Figure 20. Comparison of measured vs. simulated power output from utility-scale wind plant. Simulated with HOPP (left) and HOMER (right).

3.3 Validation Summary

The final values for the tuning coefficients and the resulting *MAE* values for all the discussed assets are summarized in Table 8 for the PV assets and Table 9 for the wind assets.

In Table 8, when the NREL data were reduced to the subset, it appears that HOPP performed just as well on the unit-scale data as the utility-scale data, with both having an *MAE* of 4.2%. However, this may be an artifact of limiting the data to one season (summer of 2022) and thus limiting the range of temperatures that this data set covers. Note that the tuned inverter efficiency for this subset is much higher (98.9%) than for the full data set (95.1%), which may be compensating for the incomplete modeling of temperature effects. The full data set, which contains the full range of temperatures, has a higher *MAE*. HOPP was able to improve upon the *MAE* of HOMER for both versions of the NREL unit-scale data set, but not for the utility-scale data set, where its *MAE* was exactly the same as HOMER's. More data from more locations would be needed to see if HOPP can robustly outperform HOMER, but there is enough here to say that HOPP can perform on the same level as HOMER when it comes to modeling power output.

For the wind data in Table 9, the reduced NREL data set did lower the *MAE* in terms of kilowatts by more than half, but since this subset had a much lower mean output, the *MAE* in terms of percentage did not change drastically.

Table 8. Summarized Validation Results for All PV Assets

	PV	HOPP			HOMER		
	Capacity	Inverter efficiency	MAE		Inverter efficiency	MAE	
	kW	%	kW	%	%	kW	%
NREL First Solar - All Data	430	95.1	10.4	5.6	94.7	15.5	8.6
NREL First Solar - Subset	430	98.9	7.5	4.2	99.5	12.1	6.7
Unidentified Solar Farm	∼50,000	96.0	1,049.70	4.2	96.0	1,053.70	4.2

There is a much bigger gap between the percentage *MAE* of the individual turbine and that of the utility-scale wind farm. Although more data from more assets are needed to determine whether this is a robust trend, the initial validation study seems to indicate that both HOPP and HOMER can model the output of full wind farms with greater fidelity than they can model the output of an individual turbine. Although HOPP was not able to improve upon HOMER's *MAE*, there was not a noticeable gap between the two in either the unit-scale or utility-scale data.

Table 9. Summarized Validation Results for All Wind Assets

	Wind	HOPP			HOMER		
	Capacity	Turbine loss	MAE		Turbine loss	MAE	
	kW	%	kW	%	%	kW	%
NREL GE 1.5 SLE - All Data	1,500	1.4	55.7	15.2	1.7	56.4	15.4
NREL GE 1.5 SLE - Subset	1,500	0.6	25.0	13.4	0.2	24.3	13.1
Unidentified Wind Farm	∼100,000	0.2	**Redacted**	3.2	0.6	**Redacted**	2.6

Although HOPP's performance can be further evaluated with other data, the current work validates HOPP's abilities to model both PV and wind power output with high fidelity. Crucially, the data subset identified and illustrated in Figure 13 can be used for field-testing the battery optimization strategy from Section 1.1.1 in a real-world, unit-scale hybrid plant at NREL.

4 ARIES Field Testing

Having validated HOPP's ability to model power output from real-world renewable generation, we now seek to validate HOPP's ability to optimize the dispatch of a real-world hybrid plant. Unfortunately, HOPP does not have access to control a real-world hybrid plant with utility-scale wind turbines, PV arrays, and battery assets, but it does have access to the modest unit-scale hybrid "plant" at ARIES. This plant consists of the single 1.5-MW turbine and the 0.43-MW PV array from Section 3 as well as a 1-MW, 1-MWh lithium-ion battery energy storage system (BESS). This hybrid system served the entire load of NREL's Flatirons Campus when the campus was unexpectedly cut off from the grid by a transformer failure in the summer of 2020 (Koralewicz et al. 2022). Through NREL's ARIES platform, HOPP researchers can connect to the BESS controls and deploy HOPP's dispatch strategies in real time. HOPP thus optimized the output of the wind/solar/battery hybrid to meet the load of NREL's Flatirons Campus.

4.1 Objectives and Test Configuration

The scenarios previously discussed in Section 2 can optimize the dispatch of a hybrid plant for an entire year simultaneously in a "batch" simulation. These scenarios have perfect information on the wind and solar resources for the entire year and do not have to account for modeling error in the power output. Conversely, the real-world field testing, taking place in real time, needs to update results iteratively with real results that might contradict simulated results.

We therefore integrated the dispatch strategy presented in Section 1.1.1 with a real-time controller for the ARIES hybrid plant. The controller has real-time measurements of solar irradiation, wind speed, power generation, and battery SOC sampled at 100 ms from the ARIES Medium Voltage Data Acquisition System (MVDAS). The MVDAS measurements are sent back to HOPP to replace predicted values from previous simulations. The simulations are thus run iteratively, with the optimized dispatch changing as real conditions deviate from predicted conditions. Rather than being yearlong measurements of *actual* wind speed and solar irradiation, the wind and solar resources given to HOPP for this field test are instead *forecasts* of the same.

As shown in Equation 1.7, the optimization routine will try to meet the load with the wind, solar, and battery assets as often as possible via the penalty imposed by grid electricity purchases, P_t^p. The only control HOPP has over the load-matching is the charge/discharge schedule of the battery. HOPP first predicts the upcoming output of the wind and solar assets, then optimizes the upcoming battery schedule, as shown in Figure 21. The optimized battery schedule is combined with the predicted PV and wind generation into the total plant output, which should meet the load a maximum percentage of the time. The effectiveness of this optimization can be measured by Equation 1.15, with a tolerance of $\pm 5\%$ of the load.

As shown in Figure 21, the final output from a single iteration of HOPP in the field test is total plant power rather than battery schedule. This is because if HOPP's optimized battery schedule were to be sent directly to the BESS, the total plant power generated would not match the load because of error in the predicted generation of wind and PV. Instead, for the sake of the field test, a "real-time balancer" script was written to serve as an intermediary between HOPP and the ARIES plant controller. This script corrects the BESS schedule for errors in the predicted wind and PV output and ensures the combined BESS + wind + PV output will truly match the load when HOPP intends it to. Since this script only has to perform simple arithmetic and input/output operations, it can operate much faster than HOPP—in cycles of 100 ms, the same as the ARIES MVDAS and plant controller. The balancer script thus communicates with ARIES in cycles of 200 ms via Modbus, with the script first sending a battery power set point to ARIES at $t = 0$ ms. It then reads a response from ARIES (including confirmed acceptance of the set point) 100 ms later ($t = 100$ ms). This response also contains measurements of power and battery SOC from the MVDAS, which are used to calculate the next battery set point. This next set point is sent 100 ms after the response was received ($t = 200$ ms), restarting the loop that began at $t = 0$ ms.

In the response from ARIES, the measurements of battery power and battery SOC are always *live* measurements because the nature of the test inherently involves controlling the battery. However, since there is no direct control over the wind or PV assets (some curtailment controls could theoretically be implemented in the future), there is an option for these measurements to be live *or historical*. To implement this historical wind/PV option, the real-time balancer calculates the time gap between the real time that the test was started and the desired start time in the historical data (in this case, 8 a.m. on July 28, 2022). The balancer would subtract this gap from the clock time and

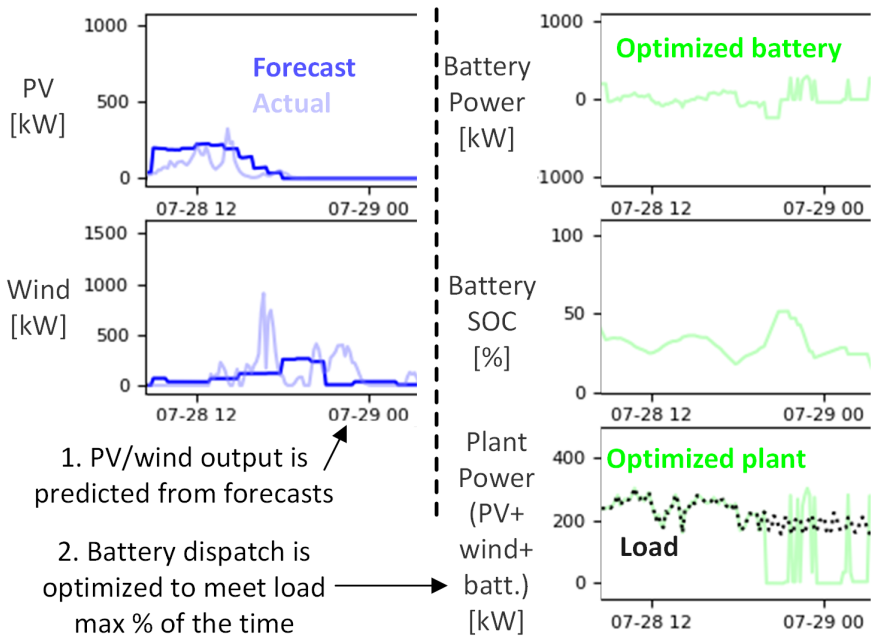

Figure 21. Iterative forecasting/optimization process in field test

send it as a time stamp to the ARIES controller, which would then look up the historical wind and PV power output stored in MVDAS. In this historical configuration, the live power output of the battery is combined with historical measurements of wind and PV output. However, the operation of the plant as a whole is indistinguishable from the fully live configuration, since HOPP is only given forecasts of wind/solar resources to make its wind/PV power predictions and dispatch optimizations. These predicted power outputs are replaced with measurements at the same rate they would be as if they were being measured live, so the optimization happens just as it would in a fully live test.

There were several reasons why the field test was conducted with the historical configuration rather than the live configuration. First, it allowed the field test to be conducted with a segment of historical PV and wind generation data that had already been vetted for data quality—specifically, the segment shown in Figure 13. It also prevented the valuable ARIES research assets in the turbine and PV from being tied up in this project when historical data would serve just as well if not better. Those assets were instead used for other research projects that could only be done with live testing. Additionally, it allowed the entire field test to be simulated offline, with a simulated ARIES controller that could mimic the expected responses to troubleshoot errors. The simulations provided predictions of what would happen using certain time periods of historical data. Finally, given the unproven nature of the HOPP/balancer/ARIES controller system, permission has not yet been obtained for this system to operate continuously, i.e., 24 hours a day. As is, the system can only operate while being attended to during work hours, and testing with live wind/PV data could only be done in discontinous 8-hour daytime/weekday segments. The historical data, on the other hand, can be split across work days with no discontinuities, and the test can even be partially "rewound" if something goes wrong without having to start over from the beginning (as happened multiple times).

In addition to the power and SOC communication occurring with the ARIES controller, HOPP also receives two more pieces of information: weather forecasts and historical load. The weather forecasts, discussed further in the following section, are synchronized to the historical time of the wind/PV power data. Thus, HOPP cannot "see ahead" in the historical tests any further than it could in a live test. For the load data, however, there is currently no forecasting tool that can be used to produce forecasts of the Flatirons Campus load—only measured historical data. Thus, in the case of load, HOPP is given the full historical load data rather than a forecast and can indeed see ahead to what the actual measured load would be at a certain point in the future. We do not believe this has much impact on the results because the load follows a very predictable daily profile. Unlike the wind/PV/BESS power

data, HOPP reads updated wind forecasts and static load data directly from ARIES data storage, rather than via the plant controller, because the data do not need to be updated at a high frequency.

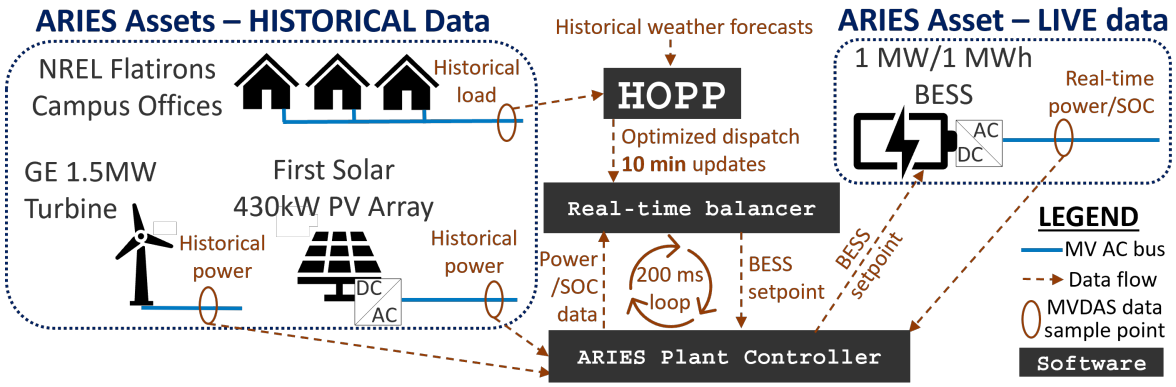

Figure 22. Asset and software configuration used in field test to validate dispatch optimization

The final developed field test configuration, using the historical wind and PV power data, is shown in Figure 22. As part of developing this configuration, HOPP was modified to simulate power output in 10-minute increments and to only optimize the upcoming 48 hours of dispatch. The setup was debugged using a simulated response from the battery to simulate the total plant output using different segments of historical data. This was necessary to prepare a prediction of the BESS's charge/discharge schedule in advance and allow the ARIES operators to schedule the test appropriately and avoid peak demand charges from the grid. The turbine, PV array, and BESS were all connected to a grid-connected medium-voltage AC bus, and if the BESS was to charge more than the combined turbine and PV array generation at any time, electricity would have to be purchased from the local utility. Thus, we sought to schedule a test that would not incur unnecessary electricity costs during peak hours. After simulating several possibilities, the time period of 8 a.m. on July 28, 2022, through 8 a.m. on July 30, 2022, was selected as the source of the wind and PV data ands the Flatirons Campus load data for the field test. The test limited the time that the BESS would spend charging at its maximum rate of 1 MW.

The ability to simulate the entire field test allows comparisons to be made between the real field test results and the simulated field test results. A baseline simulation was run in which the dispatch optimization routine was not used. Instead, the baseline simulation always tried to use the BESS to meet the load, with no anticipation of the future state of the BESS's charging capabilities or the upcoming renewable generation. The baseline simulation indicates how much improvement has been obtained in hybrid plant performance with the dispatch optimization compared to no optimization. Ideally, a real-world baseline test would also have been conducted, but time and budget constraints prevented it.

4.2 Power Forecasting

Another comparison that can be made between different simulated results is between different levels of forecast accuracy. Three forms of forecasting were used to test the range of forecast accuracy: two extremes of baseline and ideal performance, and the forecast that was actually used in the field test:

1. **Baseline forecast performance: the "null" forecast**. This forecast simply used historical averages of weather data for a certain day of the year and time of day. In this case, HOPP tries to optimize performance, but does not have access to any meteorological predictions, only historical averages.

2. **Ideal forecast performance: the "perfect" forecast**. This "forecast" gives HOPP access to all measured data at all points in history, effectively allowing it to see into the future. This simulates what the theoretical limit of performance of the optimization would be.

3. **Field test forecast performance: a "rolling" forecast**. This forecast gives HOPP access to the same forecast information that would be available at the same point in history it is trying to start its optimization from. That is, if it is optimizing a schedule starting at 12 p.m. on July 28, 2022, it has access to:

    - The measured data for 12 p.m. on July 28, 2022 (a "persistence" forecast)

- A 1-hour-ahead forecast for 1 p.m. on July 28, 2022
- A 2-hour-ahead forecast for 2 p.m. on July 28, 2022
- Et cetera, up to...
- A day-ahead forecast for 12 p.m. on July 29, 2022
- The null forecast from there on.

This forecast receives rolling updates every hour, simulating a live-updating forecast.

The measured data that were taken for the perfect forecast and persistence forecasts came from two sources:

- The solar GHI and DNI, temperature, and pressure were measured on the ground at the NREL Flatirons M2 site, with instruments detailed at the NREL Measurement and Instrumentation Data Center (MIDC) website (Jager 1996).
- The wind speed, direction, temperature, and pressure were measured near turbine hub height at a met tower approximately 150 m upstream of the prevailing winds from the turbine (Mendoza et al. 2015).

These data were averaged over several years to create the null forecast.

For the rolling forecasts, a solar forecast had already been created by First Solar based on National Oceanic and Atmospheric Administration (NOAA) data to predict GHI, DNI, diffuse horizontal irradiance (DHI), and temperature (Gevorgian et al. 2022). These forecasts are proprietary to First Solar and will not be shown here. For predicting wind, hourly wind speed and gust forecasts for the upcoming week for the nearby Boulder, Colorado, NOAA weather station were logged into MVDAS. A linear regression was performed between the measured Boulder weather data and the measured wind speed at the turbine met tower to develop a wind speed forecast for the 1.5-MW turbine. For the temperature and pressure, the same temperature and pressure from the solar forecast were used. No forecast of wind direction was created, so the direction was assumed to be constant at the prevailing winds of 292°.

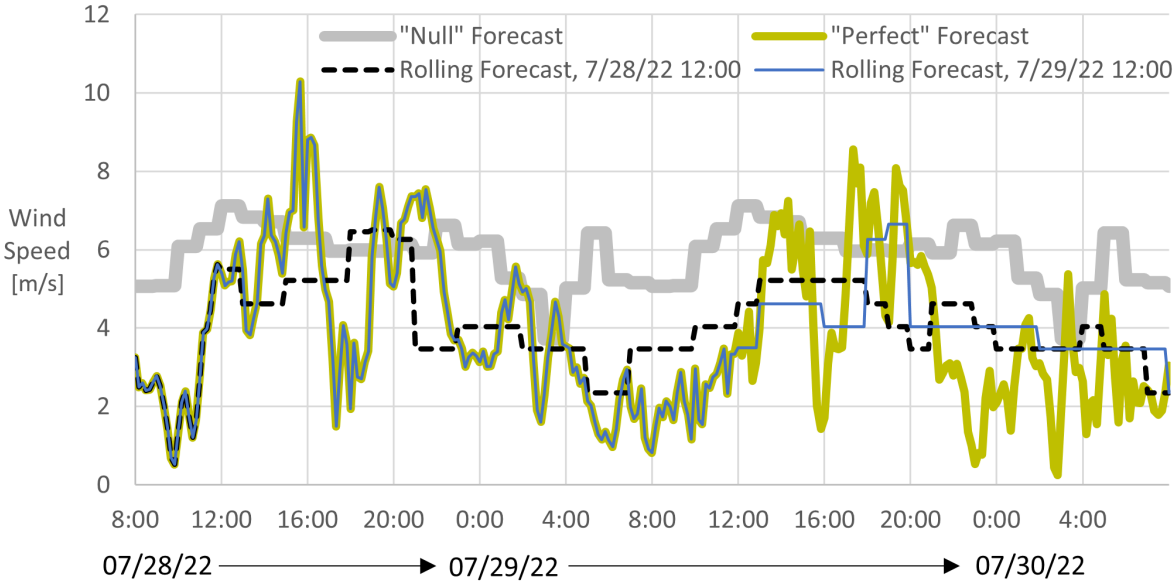

Figure 23. Wind forecasts during simulated field test period

As shown in Figure 23, the null forecast does not deviate much from the average wind speed, while the perfect forecast reflects the highly variable real wind speeds at the Flatirons Campus. The rolling forecast at noon on July 28 correctly predicts the overall wind trend in the long term: higher winds in the afternoon/evening of July 28, lower winds overnight and into the morning, then higher winds in the afternoon of July 29. However, it does not predict a sharp spike in winds around 16:00 (4 p.m.) on July 28, which will produce a sharp increase in wind turbine power

compared to what is expected. The updated forecast on noon of July 29 correctly predicts spikes in wind between 13:00–16:00 and 18:00–20:00, but again underestimates their magnitude. This is important to keep in mind when viewing the results of the field test.

Prior to the field test and simulations, the accuracy of the forecasts is evaluated using the same *MAE* metric used to validate HOPP's modeling of power output in Section 3. The forecast wind and irradiation are used as inputs to HOPP's power modeling for the entire July/August 2022 time period used in the model validation, and the same methods from Equations 3.3 and 3.4 are applied. The forecasts are evaluated both for their accuracy in predicting power output over both 1-hour intervals and 24-hour intervals. The *MAE* for the "perfect forecast" over the 1-hour intervals is the same as the modeling error of HOPP, given in Tables 8 and 9 previously. The summary of all these forecast power *MAE* values is given in Table 10. The *MAE* is given not just as a function of the PV array and wind turbine alone, but also as a function of the PV/wind *hybrid* plant; i.e., this shows the error in the predicted power output of the hybrid PV/wind plant as a whole. It does not show the error the predicted power output of the hybrid PV/wind/*battery* plant as a whole because full simulated field tests are needed in which the battery is dispatched.

Table 10. Summarized *MAE* Values for All Forecast Power Outputs

	First Solar 0.43 MW PV				GE 1.5 MW wind turbine				PV/wind hybrid plant			
	1 hour averages		24 hour averages		1 hour averages		24 hour averages		1 hour averages		24 hour averages	
Forecast	kW	%	kW	%	kW	%	kW	%	kW	%	kW	%
"Perfect"	8	4.2	44	1.0	25	13.4	248	5.3	26	7.0	252	2.8
Hour-ahead	43	22.8	550	12.2	158	80.9	1170	25.3	164	43.7	1293	14.4
Day-ahead	49	26.9	674	14.8	165	86.0	1196	26.0	172	45.9	1373	15.3
"Null"	80	44.4	1598	36.0	169	91.6	1430	30.9	187	49.9	2144	23.8

We start at the left-hand side of Table 10 with the PV forecast of 1-hour average power output. All the *MAE* values will be discussed in terms of percentages of average power output for the particular asset and time period being discussed. Both the hour-ahead and day-ahead irradiance forecasts can give HOPP a much better prediction than the null forecast, with the hour-ahead forecast removing about half of the *MAE* present in the null forecast. The hour-ahead forecast is only marginally better than the day-ahead forecast, and this holds true throughout the table. For the 1-hour average PV power, both or these forecasts are about halfway between the perfect and null extremes, indicating that there is much predictive power in these forecasts, but also more that could be gained. That is not to say that there is anything that could be improved about these PV forecasts; they may already be at the limit of the predictive capabilities possible at this site. Unsurprisingly, when looking at 24-hour average power, the *MAE* for all the forecasts goes down. The gap between hour-ahead forecast and the perfect forecast gets smaller while the gap between the day-ahead forecast and the null forecast gets larger. This indicates that in individual hours, these forecasts have errors in opposite directions that will cancel themselves out over the course of a day, giving them better predictive capability for a full day's power output.

Unfortunately, the errors in the wind turbine power output are much larger, especially with the 1-hour averages. This is unsurprising, given the relatively crude nature of the wind forecast developed for this project (a linear regression from historical data). More predictive wind forecasts would require state-of-the-art numerically based techniques with radar and satellite imagery and possibly machine learning that are beyond the scope of this project. However, there was still at least some predictive capability in the wind forecasts, as indicated by the drops in *MAE* relative to the null forecast, although they are still much closer to the null forecast than the perfect forecast.

When looking at the *MAE* for the combined PV/wind hybrid, the low predictive capability of the wind forecast causes the 1-hour average *MAE* for the combined PV/wind hybrid to be quite high when using the hour-ahead and day-ahead forecasts—much closer to the null forecast than the perfect forecast. However, these numbers are much better when looking at the 24-hour averages. Again, the random errors in the 1-hour averages in opposite directions have a tendency to cancel each other out, leading the 24-hour *MAE* to be much lower. The strong predictive capability of the PV forecast cancels out the weakness of the wind forecast, creating a usable hybrid PV/wind power output forecast for the purposes of this field test.

4.3 Results

The rolling forecasts were successfully implemented in the integrated of HOPP/balancer/ARIES controller system, and the field test was successfully carried out over the course of two-and-a-half weeks in July 2023. There was much troubleshooting of the system once live testing began with the real BESS instead of the simulated BESS; issues were seen that did not occur in simulation. For example, HOPP initially did not account for the fact that the BESS's SOC could overshoot its limits, and the calculated SOC would reach as high as 91.5% despite having a limit of 90.0% in both the ARIES controller and HOPP. HOPP was still internally simulating how its commands would affect the SOC over the upcoming 10-minute time period to ensure that limits were not reached. If HOPP calculated that its power command would bring the SOC outside of the limits, it would cancel out this command and instead send a power command of zero. However, with the overshoot to 91.5%, even if HOPP was trying to discharge the BESS down to 90.5% within the next time period, the command would be cancelled out since 90.5% was still higher than the 90.0% upper limit. This resulted in the battery being stuck at a full charge and never discharging, despite HOPP's optimizer calling for a discharge. Because the test was being run with historical data, once we realized this problem was occurring and resolved it, we could simply "rewind" the test to a point before the error began, rather than having to start a new test (which would have been required using live data).

An example of data from HOPP's real-time balancer handling the communications between HOPP and the ARIES plant controller is shown in Figure 24. This slice of data is a very small segment of 2.5 minutes of the full 48-hour test, but it gives an example of typical operation in between HOPP and ARIES. Four pairs of data can be seen in this figure, starting at the top with load and and plant output. The balancer is attempting to keep plant output as close as possible to load, and it succeeds, with plant output never varying more than 5 kW from load in either direction. The reason there are variances is because the next pair of data down the chart: the wind forecast and wind output. Wind output fluctuates very quickly at points, by more than 10 kW/s, meaning it changes by more than 1 kW in the time it takes data to be exchanged between ARIES and HOPP (100 ms)—and this section of data was just taken for demonstration purposes; there are points where wind output is fluctuating even more rapidly. By contrast, the PV output is very steady and consistent with the forecast (the next pair of data down the figure). To match the plant output to the fluctuating wind, HOPP adjusts the commanded set point, and the BESS power output from ARIES adjusts. However, the adjustment to the power output is delayed by approximately 100 ms–140 ms between when HOPP gets the updated wind and PV measurements and when it sends back its commanded BESS set points, and another 40 ms between when ARIES receives the BESS command and when the BESS starts ramping to the new set point. The delay is what causes the variance between the plant output and the load, even when the dispatch optimizer is asking for the plant to meet the load exactly. The variance is greatest when the wind and PV output change most rapidly. Later, we will conduct a thorough breakdown on HOPP's ability to meet the load within certain tolerances throughout the field test.

We now zoom out to a larger section of the field test in Figure 25. This figure plots all data in 10-minute averages and shows the first 8 hours of the field test. These 8 hours were broken up over two business days, Friday, July 7, 2023, and Monday, July 10, 2023, with all the historical data taken from July 28, 2022. The first segment of this historical data ran from 08:00 to 10:30 on July 28, 2022, but when 10:30 was reached, it was the end of the work week in the present time, and the ARIES operators suspended the system. At this point, the BESS had an SOC of 28.5%. Over the weekend, the BESS did not maintain charge perfectly, but was charged back to exactly 28.5% using grid electricity just before resuming the test on July 10, 2023, with the historical data starting from 10:30 on July 28, 2022. In this way, a "seamless" field test was created despite the test being split over multiple days.

The thick gray line in Figure 25 shows the load demand, and the dashed black line shows the total plant output. Where these two lines overlap, the load is being met by the hybrid plant, but where they do not, the total plant is almost always zero, and the load is being met by the grid (not shown). In this segment of data, the combined wind and PV output (blue line) is mostly below the load, so to meet the load with the hybrid, the optimizer needs to discharge the BESS. The BESS power output (thin yellow line) is positive when the BESS is discharging, but negative when the BESS is charging. When the optimizer decides not to meet the load with the hybrid plant, the hybrid output drops to zero, and the load is met with the grid instead. This allows the BESS to use the wind and PV output to charge itself and prepare to meet the load during upcoming time periods. The SOC is shown in a separate plot at the top of the figure, with highlighted sections where the hybrid plant is meeting the load.

Figure 26 presents the same data as Figure 25, but for the entire 48-hour field test. The data in Figure 25 are much easier to read due to the shorter time period, and serve as a good reference to discern what is happening in the full

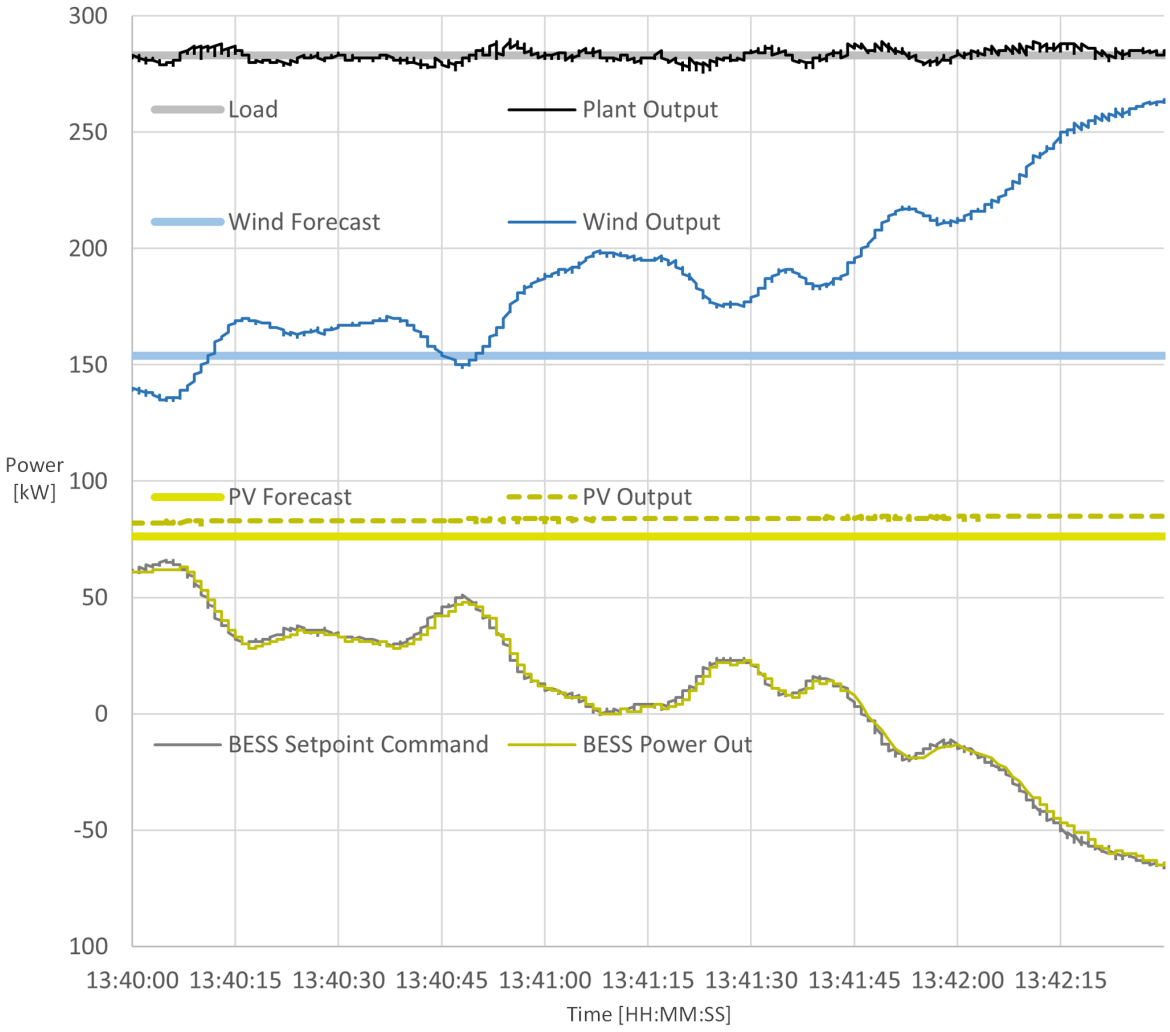

Figure 24. Example of small-scale time history results from the field test: 2.5 minutes showing the real-time balancing of fluctuating wind with the BESS so that the plant meets the load

field test on Figure 26. One important thing to note is the large spike of wind + PV power that occurs around 16:00 on July 28, 2022. Recalling Figure 23, this was due to a spike in wind that was not predicted by the forecast and that was not optimized for. At first, HOPP does not predict the sudden rise in wind + PV output and keeps the plant output at zero, which causes the BESS to charge rapidly. HOPP only starts meeting the load with the hybrid plant when the wind spike is almost over, and by this time the BESS has been charged to its limit. Even though some of the wind and PV output is going toward meeting the load, they are generating beyond the load, and this excess is sold to the grid. If the optimizer had been given an accurate forecast of this spike, it could have started meeting the load sooner instead of overcharging the battery.

At the top of Figure 26 is a strip chart showing where the hybrid plant met the load. This uses the same metric presented in Equation 1.15 and is similar to the compliance factors evaluated in the simulations of Section 2.1, although those scenarios targeted firm power instead of following a load demand. Over the full field test, the $\%_{power,met}$ of the ARIES wind/PV/BESS hybrid was 57.8%. This number on its own does not mean much, considering that this hybrid plant was a combination of research assets that were used for different purposes and not sized to serve this particular load (or sited in a location with predictable and reliable wind/solar resources). What is of more importance is how this number has been *improved* by the dispatch optimization process. For this reason, the field test needs to be compared to a baseline of hybrid performance without the dispatch optimization, as described at the end of Section 4.1. The baseline simulation, which did not use any dispatch optimization, only produced a $\%_{power,met}$ of 51.9%, meaning the dispatch optimization improved $\%_{power,met}$ by 5.9 percentage points. Ideally, the baseline would have also been

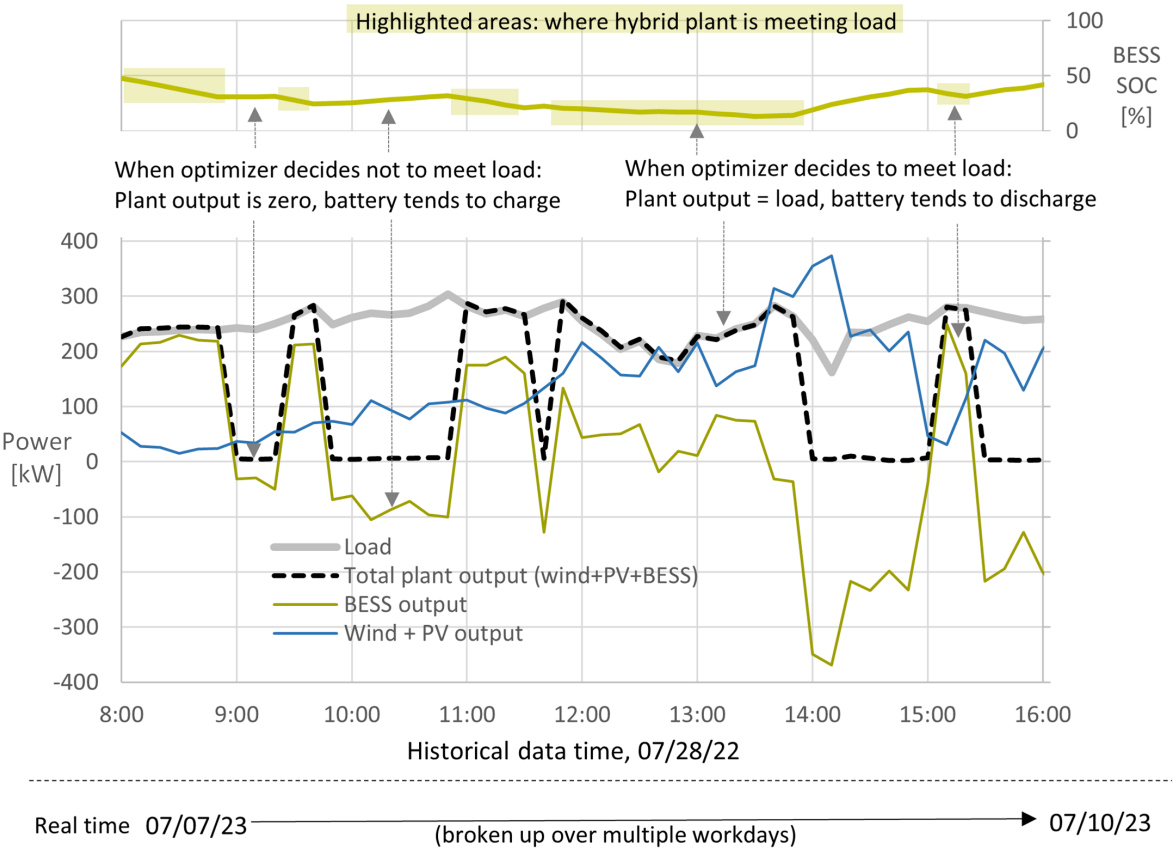

Figure 25. First 8 hours of the field test, highlighting on the SOC profile (top) where the optimizer was meeting the load using the BESS

conducted with a real field test, but budget and time constraints prevented this. For additional validation, the same field test was run in simulation and confirmed there were minimal differences between the simulation and actual field test. This simulation produced a $\%_{power,met}$ of 58.5%, only 0.7 percentage points different from the actual field test, and as Figure 27 shows, their SOC profiles are nearly identical.

Also of interest in Figure 27 is some reasoning as to why the dispatch optimization improved the $\%_{power,met}$. Clearly, in the baseline simulation, when the hybrid plant tries to meet the load whenever it can, the BESS is kept at a low SOC. The BESS can only charge when wind + PV generation is higher than the load, and the plant does not strategize to save excess generation in the BESS to better meet the load later. These profiles also show how the field test may have been advantaged in this comparison due to the relatively low level of wind and PV generation compared to the load. If the test had been conducted on a time period with more wind and PV generation, the improvements produced by optimization may not have been as substantial. The field test effectively only consists of one data point, and more tests and simulations must be conducted in the future to determine if these improvements are robust.

Another comparison that can be made is between the different forecasts, as shown in Figure 28. The null forecast, even with optimization, produces a profile very similar to the baseline simulation, and only produced a $\%_{power,met}$ of 52.9%. The comparison between forecasts confirms that, despite the modest improvements in *MAE* from Table 10, the forecasts were still a big help to the optimizer. However, the forecasts could have been more beneficial—with access to the perfect forecast, the optimizer managed to achieve a $\%_{power,met}$ of 61.6%, 3.1 percentage points higher than with the regular rolling forecasts. The reason can be seen in their SOC profiles, which are very similar except in a few key places, and especially before and during the two afternoon/evening wind spikes. With the perfect forecast, the optimizer knew how to use these sudden influxes of power optimally and never had excess generation to send to the grid with the BESS at full charge.

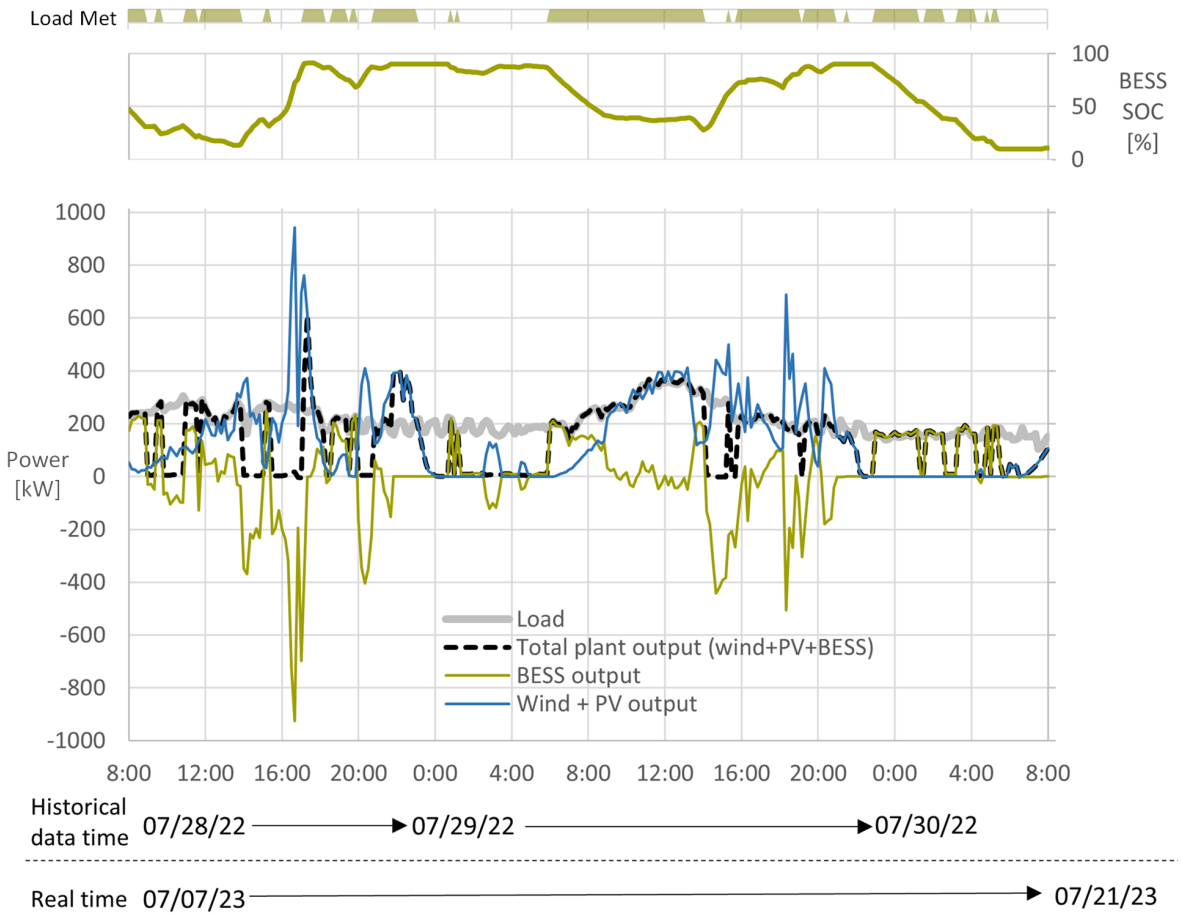

Figure 26. Full field test. On top is a strip chart showing at which points the load is being met by the hybrid plant. The BESS SOC profile below the Load Met strip chart is the easiest to compare with the simulations side-by-side, and will be used in side-by-side comparison in subsequent figures. Refer to the wider Figure 25 to better understand the details of the lower power chart.

The previously mentioned values of $\%_{power,met}$ are collected in Table 11, alongside calculations of operating cost, which did not vary significantly between simulations and the field test. Calculating a full LCOE is not relevant for this system, which consists of research assets; the main objective of the optimizer in this situation was to maximize $\%_{power,met}$.

Table 11. Summarized Operational Costs and $\%_{power,met}$ for Field Test and All Simulations

#	Test Type	Forecast type	Dispatch optimization type	Operating cost [cents/kWh]	% power met (5% tolerance)
1	Baseline Simulation	None	None	3.71	51.9%
2	**Field Test ACTUAL**	**Rolling**	Minimize operating costs while following a load demand (Section 1.1.1)	**3.67**	**57.8%**
3	Field Test Simulation	Rolling		3.68	58.5%
4	Field Test Simulation	Null		3.71	52.9%
5	Field Test Simulation	Perfect		3.70	61.6%

One last analysis performed here is a breakdown of the exact percentage of the load met by the hybrid plant when HOPP was commanding it to meet the load, shown in Figure 29. The percentage of the load met is grouped into 1 percentage point-wide bins, i.e., the number of data points where between 97% and 98% of the load was met divided by the total number of data points where HOPP attempted to meet the load at all, is shown in the"97%–98%" bin. Interestingly, this distribution is not centered at 100%; the most frequent bin is 101%–102%. Unfortunately, this is due to a bug in the system that was not caught until after the field test was completed: With the MVDAS data, when the wind and PV generators were not generating, these power measurements would read slightly negative

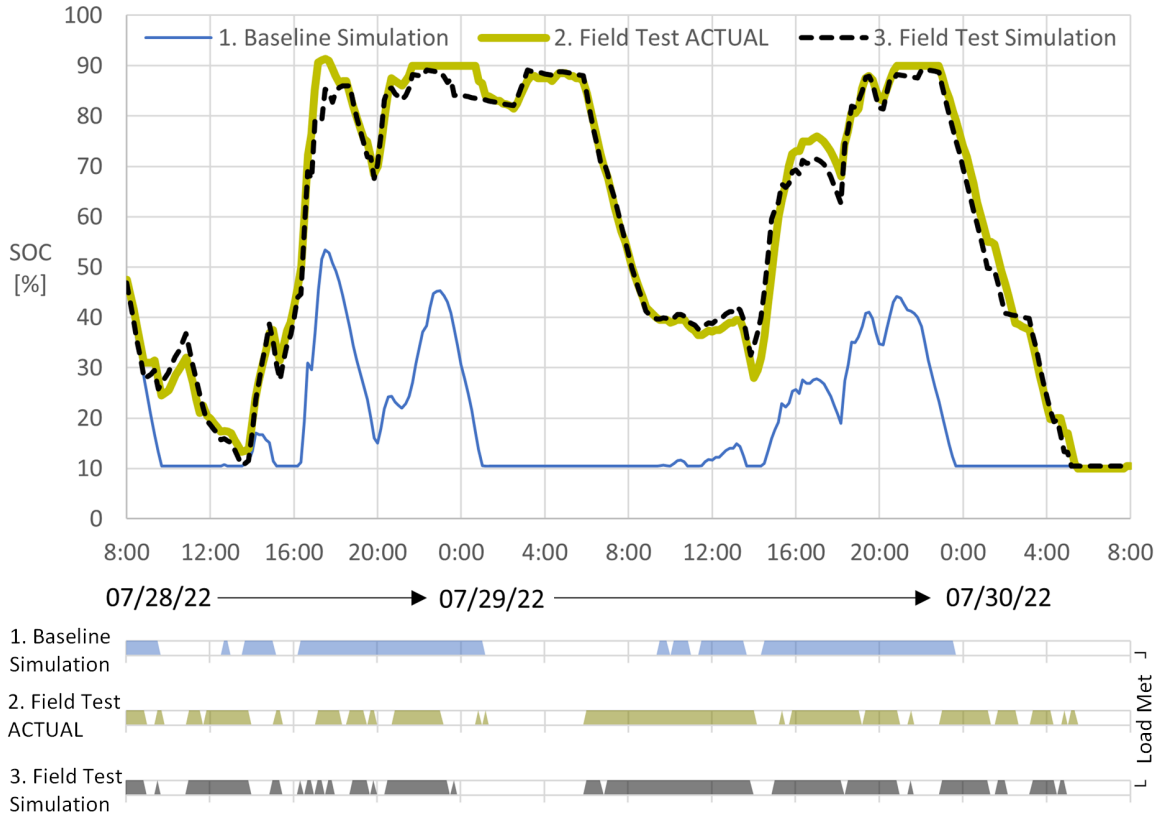

Figure 27. SOC profile comparison of the baseline simulation, actual field test, and simulated field test

numbers. In the main part of HOPP (where optimization was occurring), any wind/PV measurements that were slightly negative were automatically changed to zero to prevent errors from occurring. However, in the real-time balancer, these numbers were not changed to zero, leading the balancer to command slightly more power from the BESS than necessary when the wind or PV output was zero.

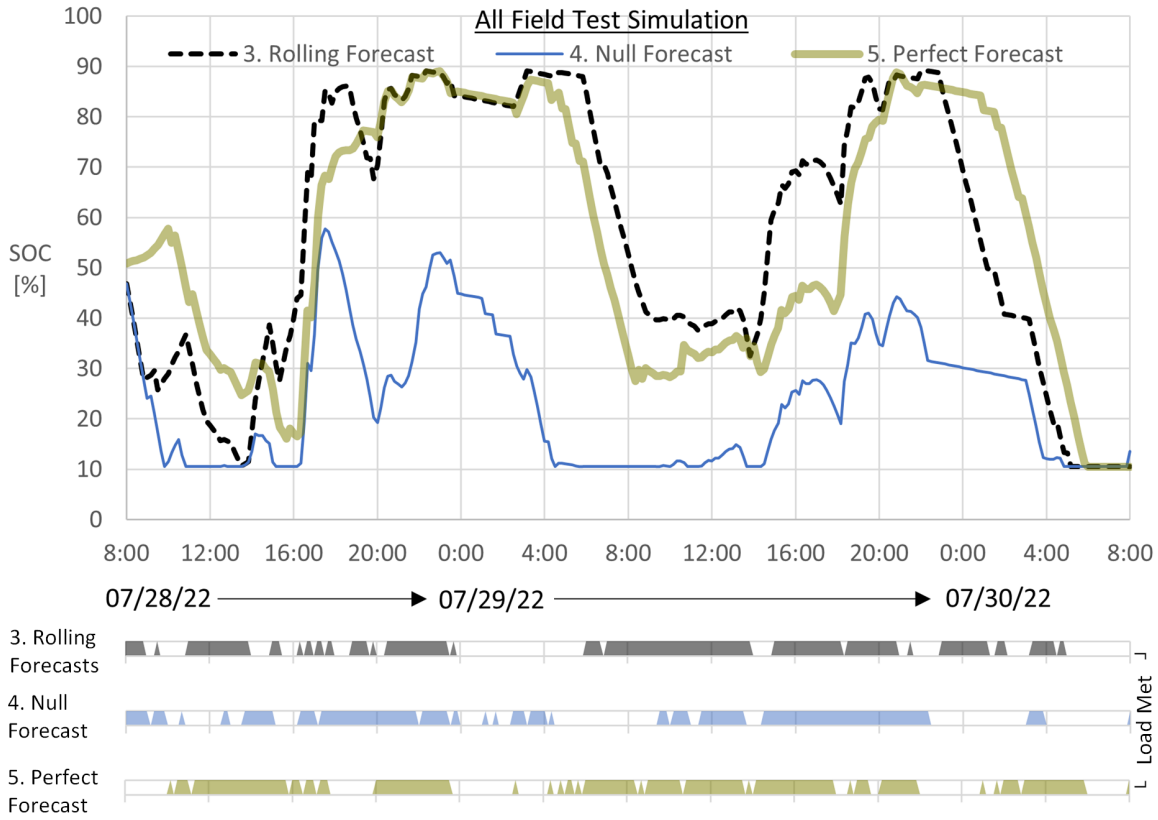

Figure 28. SOC profile comparison of the simulated field test using rolling forecasts, the null forecast, and the perfect forecast

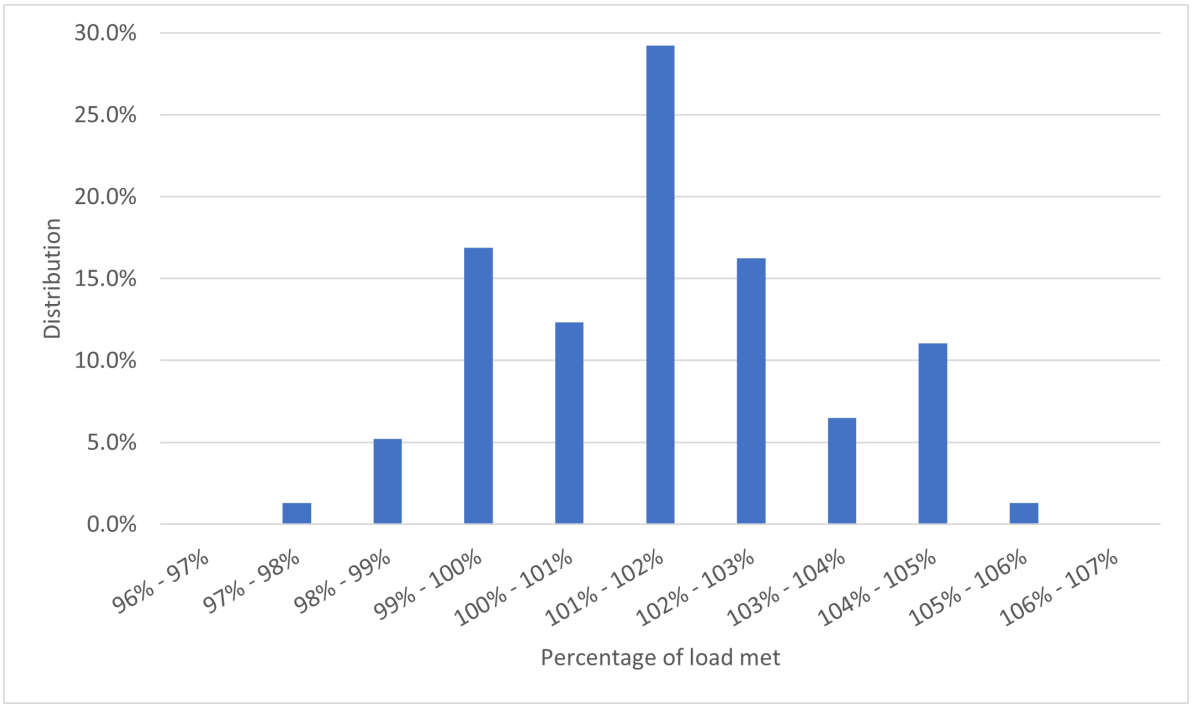

Figure 29. Distribution of the exact percentage of load met by HOPP/balancer/ARIES controller system during the field test

5 Conclusions

This report gives an overview of the Hybrid Optimization and Performance Platform software and demonstrates its capabilities. Capabilities include the presentation of three dispatch case studies, the validation of the tool on both small and commercial scales, and finally the implementation of a field test using an optimized dispatch strategy from the tool.

The HOPP tool is capable of simulating various technologies together in a hybrid power plant. The tool includes wind, solar, battery, concentrated solar power, and hydrogen models, with more models under development. The HOPP tool is also capable of combining the technical simulation of the technologies with financial modules to estimate the profitability of the plant. To demonstrate these capabilities, we presented three dispatch scenarios. The first examined the feasibility of hybrid plants, in general, before looking at HPPs in five locations in the United States. These locations were chosen because of varying resource profiles. The study found that, at the current technology prices, HPPs were most feasible in the Kansas, Texas, and New Hampshire locations. The plants in the Arizona and West Virginia locations may need additional siting considerations (i.e., siting turbines on a ridge). The second scenario examined the hours that HPPs could have available to contribute to energy reliability services. It was found that the HPP with wind, solar, and storage had the best performance, though it did not have the maximum number of hours available for reliability services because some of the power went to smoothing out the variable plant output instead of to ERS. Finally, the third scenario examined a situation where a plant could decide between producing electricity or producing hydrogen based on the prices of these commodities. This preliminary study found that the yearly revenue from the plant could be increased if both electricity and hydrogen were produced, rather than just one.

Next, the HOPP platform was validated against unit-scale data from the NREL ARIES facility and confidential utility-scale data. The solar data exhibited low modeling error from both sources, though small temperature trends can be seen in the error indicating temperature effects that are not currently being captured by the modeling. Relative to solar, the wind modeling errors were greater in the ARIES data, due to difficulties in obtaining a large "clean" data set without interference from research activity, but less than solar in the clean utility-scale data set. Both wind data sets showed less dependence on temperature or time of day than solar, though more dependence on the measured output level of the farm. In any case, the modeling error of HOPP was never significantly greater than that of a commercial modeling software, HOMER (and sometimes less), giving confidence that these models are sufficiently accurate for commercial use.

Finally, HOPP was implemented in a field test using a physical battery in the NREL ARIES facility. An optimized dispatch was fed to the physical 1-MW battery, which used historical data from the ARIES wind and solar assets. The field test followed a load signal based on the NREL Flatirons Campus usage. The field test showed good agreement between the simulated dispatch and the actual dispatch. Between the two dispatch profiles, the period of time where the load was met were extremely similar, and the states of charge never deviated from each other by more than 10 percentage points. This increased confidence that the simulated dispatch optimization scenarios are delivering realistic results. The optimized dispatch also improved the amount of time the load could be met over baseline dispatch, by 5.9 percentage points in the real field test and 6.6 percentage points in the simulated field test. After troubleshooting bugs with the interface between HOPP and ARIES early in the field testing process, HOPP was able to smoothly interact with the ARIES hybrid plant controller, proving that the optimized dispatch strategies can be applied to physical hardware.

HOPP is a versatile tool that can be used to simulate various kinds of HPPs, and it shows good agreement in both the validation and the field tests presented in this report. Ongoing development of the tool will add more functionality, such as onshore and offshore plant component design capabilities, additional supported technologies, and power-to-x capabilities.

References

Bortolotti, P., H. C. Tarres, K. L. Dykes, K. Merz, L. Sethuraman, D. Verelst, and F. Zahle. 2019. "IEA Wind TCP Task 37: Systems Engineering in Wind Energy - WP2.1 Reference Wind Turbines" (June). https://doi.org/10.2172/1529216. https://www.osti.gov/biblio/1529216.

Clark, C. E., A. Barker, K. Brunik, M. Kotarbinski, E. Grant, O. Roberts, J. King, A. P. Stanley, P. Bhaskar, and C. Bay. 2023. "Opportunities for green hydrogen production with land-based wind in the United States." *Energy Conversion and Management* 296:117595. ISSN: 0196-8904. https://doi.org/https://doi.org/10.1016/j.enconman.2023.117595. https://www.sciencedirect.com/science/article/pii/S019689042300941X.

Clark, C. E., A. Barker, J. King, and J. Reilly. 2022. "Wind and Solar Hybrid Power Plants for Energy Resilience" (January). https://doi.org/10.2172/1842446. https://www.osti.gov/biblio/1842446.

Danecek, P., J. K. Bonfield, J. Liddle, J. Marshall, V. Ohan, M. O. Pollard, A. Whitwham, et al. 2021. "Twelve years of SAMtools and BCFtools." 10.1093/gigascience/giab008, *GigaScience* 10, no. 2 (February): giab008. ISSN: 2047-217X. https://doi.org/10.1093/gigascience/giab008. eprint: https://academic.oup.com/gigascience/article-pdf/10/2/giab008/60687743/giab008.pdf. https://doi.org/10.1093/gigascience/giab008.

Electric Power Research Institute (EPRI). 2022. *Water Electrolyzer Stack Degradation.* Technical report. Electric Power Research Institute (EPRI). https://restservice.epri.com/publicdownload/000000003002025148/0/Product.

Ericson, S., S. Koebrich, S. Awara, A. Schleifer, J. Heeter, K. Cory, C. Murphy, and P. Denholm. 2022. *Influence of Hybridization on the Capacity Value of PV and Battery Resources* [in American English]. Other. https://doi.org/10.2172/1854329.

Freeman, J., J. Jorgenson, P. Gilman, and T. Ferguson. 2014. *Reference Manual for the System Advisor Model's Wind Power Performance Model.* Technical report NREL/TP-6A20-67399. National Renewable Energy Laboratory, August. https://doi.org/10.2172/1150800. https://www.osti.gov/biblio/1150800.

Gagnon, P., P. A. S. Perez, K. Obika, M. Schwarz, J. Morris, J. Gu, and J. Eisenman. 2024. "Cambium 2023 Scenario Descriptions and Documentation" [in en]. *Renewable Energy.*

Gevorgian, V., P. Koralewicz, S. Shah, E. Mendiola, R. Wallen, and H. Villegas Pico. 2022. *Photovoltaic Plant and Battery Energy Storage System Integration at NREL's Flatirons Campus.* Technical report NREL/TP-5D00-81104. National Renewable Energy Laboratory, February. https://doi.org/10.2172/1846617. https://www.osti.gov/biblio/1846617.

Gilman, P., N. A. DiOrio, J. M. Freeman, S. Janzou, A. Dobos, and D. Ryberg. 2018. *SAM Photovoltaic Model Technical Reference 2016 Update.* Technical report NREL/TP-6A20-67399. National Renewable Energy Laboratory, March. https://doi.org/10.2172/1429291. https://www.osti.gov/biblio/1429291.

Hamilton, W., J. Martinek, J. Cox, and A. Newman. 2022. "Integrating Concentrating Solar Power Technologies into the Hybrid Optimization and Performance Platform (HOPP)" (August). https://doi.org/10.2172/1884793. https://www.osti.gov/biblio/1884793.

Hamilton, W. T., M. A. Husted, A. M. Newman, R. J. Braun, and M. J. Wagner. 2020. "Dispatch optimization of concentrating solar power with utility-scale photovoltaics." *Optimization and Engineering* 21, no. 1 (March): 335–369. ISSN: 1573-2924. https://doi.org/10.1007/s11081-019-09449-y. https://doi.org/10.1007/s11081-019-09449-y.

Harrison-Atlas, D., C. Murphy, A. Schleifer, and N. Grue. 2022. "Temporal complementarity and value of wind-PV hybrid systems across the United States." *Renewable Energy* 201:111–123. ISSN: 0960-1481. https://doi.org/https://doi.org/10.1016/j.renene.2022.10.060. https://www.sciencedirect.com/science/article/pii/S0960148122015531.

Hart, W. E., C. D. Laird, J.-P. Watson, D. L. Woodruff, G. A. Hackebeil, B. L. Nicholson, and J. D. Siirola. 2017. *Pyomo - Optimization Modeling in Python.* Vol. 67, Springer Optimization and Its Applications. Springer Cham. ISBN: 978-3-319-86482-2. https://doi.org/10.1007/978-3-319-58821-6.

International Energy Agency (IEA). 2021. *Global Hydrogen Review 2021.* Technical report. International Energy Agency (IEA). https://www.iea.org/reports/global-hydrogen-review-2021.

Jager, A., D.; Andreas. 1996. *NREL National Wind Technology Center (NWTC): M2 Tower.* Technical report NREL/DA-5500-56489. National Renewable Energy Laboratory. https://doi.org/10.5439/1052222. https://midcdmz.nrel.gov/apps/sitehome.pl?site=NWTC.

Jonkman, J., S. Butterfield, W. Musial, and G. Scott. 2009. "Definition of a 5-MW Reference Wind Turbine for Offshore System Development" (February). https://doi.org/10.2172/947422. https://www.osti.gov/biblio/947422.

King, D. L., J. A. Kratochvil, and W. E. Boyson. 2004. *Photovoltaic Array Performance Model.* Technical report SAND2004-3535. Albuquerque, NM: Sandia National Laboratories. https://doi.org/doi.org/10.2172/919131.

Koralewicz, P., E. Mendiola, R. Wallen, V. Gevorgian, and D. Laird. 2022. *Unleashing the Frequency: Multi-Megawatt Demonstration of 100% Renewable Power Systems with Decentralized Communication-Less Control Scheme.* Technical report NREL/TP-5000-80742. National Renewable Energy Laboratory, September. https://doi.org/10.2172/1891206. https://www.osti.gov/biblio/1891206.

Kotarbinski, M., B. McDowell, J. Katz, G. Starke, and N. Riccobono. 2024. *A 2023 Perspective: What Is the Value of Hybridization?* [In American English]. Other. https://doi.org/10.2172/2331419.

LMP Contour Map. 2022. Accessed October 5, 2022. https://api.misoenergy.org/MISORTWD/lmpcontourmap.html.

Mendoza, I., J. Hur, S. Thao, and A. Curtis. 2015. *Power Performance Test Report for the U.S. Department of Energy 1.5-Megawatt Wind Turbine.* Technical report NREL/TP-5000-63684. National Renewable Energy Laboratory, August. https://doi.org/10.2172/1215120. https://www.osti.gov/biblio/1215120.

Mishra, S., J. Pohl, N. Laws, D. Cutler, T. Kwasnik, W. Becker, A. Zolan, K. Anderson, D. Olis, and E. Elgqvist. 2022. "Computational framework for behind-the-meter DER techno-economic modeling and optimization: REopt Lite." *Energy Systems* 13, no. 2 (May): 509–537. ISSN: 1868-3975. https://doi.org/10.1007/s12667-021-00446-8. https://doi.org/10.1007/s12667-021-00446-8.

Murphy, C., D. Harrison-Atlas, N. Grue, T. Mosier, J. Gallego-Calderon, and S. Elliott. 2023. *Complementarity of Renewable Energy-Based Hybrid Systems* [in en]. Technical report NREL/TP–6A20-81901, 1972008, MainId:82674. April. Accessed August 30, 2024. https://doi.org/10.2172/1972008. https://www.osti.gov/servlets/purl/1972008/.

National Renewable Energy Laboratory. 2021. *System Advisor Model Version 2021.12.02 (2021.12.02). SAM source code.* Accessed August 18, 2022. https://github.com/NREL/SAM/tree/develop/api/api_autogen/library/defaults.

———. 2022. *2022 Annual Technology Baseline.* Golden, CO. https://atb.nrel.gov/.

NOAA ESRL Global Monitoring Laboratory. 2023. *NOAA Solar Calculator.* website, October. https://gml.noaa.gov/grad/solcalc/index.html.

NREL. 2021. *FLORIS. Version 2.4.* https://github.com/NREL/floris.

———. 2023a. *electrolyzer. Version 0.1.* https://github.com/NREL/electrolyzer.

———. 2023b. *Hybrid Optimization and Performance Platform. Version 2.2.* https://github.com/NREL/HOPP/releases/tag/v2.2.0.

Roadman, J. 2023. personal communication, 31 1.

Schleifer, A., D. Harrison-Atlas, W. Cole, and C. Murphy. 2023. "Hybrid Renewable Energy Systems: The Value of Storage as a Function of PV-Wind Variability: Article No. 1036183" [in American English]. *Frontiers in Energy Research* 11. ISSN: 2296-598X. https://doi.org/10.3389/fenrg.2023.1036183.

Specialists, T. E. S. 2013. *Transient System Simulation Tool (TRNSYS).* website, July. http://www.trnsys.com/.

Stanley, A. P., and J. King. 2022. “Optimizing the physical design and layout of a resilient wind, solar, and storage hybrid power plant.” *Applied Energy* 317:119139. ISSN: 0306-2619. https://doi.org/https://doi.org/10.1016/j.apenergy.2022.119139. https://www.sciencedirect.com/science/article/pii/S0306261922005153.

System Advisor Model Version 2020.11.29 (SAM 2020.11.29). 2020. https://sam.nrel.gov.