An Integrated Platform for Wind Power Plant Operations: From Atmosphere to Electrons to the Grid (A2e2g)

Michael Sinner, Dave Corbus, Elina Spyrou, Jennifer King, Sanjana Vijayshankar, Andrew Kumler, Chris Bay, and Vahan Gevorgian

National Renewable Energy Laboratory

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5000-84826
March 2023

An Integrated Platform for Wind Power Plant Operations: From Atmosphere to Electrons to the Grid (A2e2g)

Michael Sinner, Dave Corbus, Elina Spyrou, Jennifer King, Sanjana Vijayshankar, Andrew Kumler, Chris Bay, and Vahan Gevorgian

National Renewable Energy Laboratory

Suggested Citation
Sinner, Michael, Dave Corbus, Elina Spyrou, Jennifer King, Sanjana Vijayshankar, Andrew Kumler, Chris Bay, Vahan Gevorgian. 2023. *An Integrated Platform for Wind Power Plant Operations: From Atmosphere to Electrons to the Grid (A2e2g)*. Golden, CO: National Renewable Energy Laboratory. NREL/TP-5000-84826. https://www.nrel.gov/docs/fy23osti/84826.pdf.

NREL is a national laboratory of the U.S. Department of Energy
Office of Energy Efficiency & Renewable Energy
Operated by the Alliance for Sustainable Energy, LLC

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

Contract No. DE-AC36-08GO28308

Technical Report
NREL/TP-5000-84826
March 2023

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

NOTICE

This work was authored by the National Renewable Energy Laboratory, operated by Alliance for Sustainable Energy, LLC, for the U.S. Department of Energy (DOE) 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.

This report is available at no cost from the National Renewable Energy Laboratory (NREL) at www.nrel.gov/publications.

U.S. Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via www.OSTI.gov.

Cover Photos by Dennis Schroeder: (clockwise, left to right) NREL 51934, NREL 45897, NREL 42160, NREL 45891, NREL 48097, NREL 46526.

NREL prints on paper that contains recycled content.

Acknowledgments

The authors would like to thank Jian Fu and Cynthia Bothwell of the U.S. Department of Energy's Wind Energy Technologies Office for their support and technical input to this project. We would also like to thank Yingchen Zhang, Bethany Frew, and Paul Denholm for providing feedback to earlier versions of this work. In addition, we are grateful to the Electric Reliability Council of Texas market information team for providing data sets in a prompt manner.

List of Acronyms

A2e2g	Atmosphere to Electrons to the Grid
AGC	automatic generation control
APC	active power control
CAISO	California Independent System Operator
DDPG	deep deterministic policy gradient
ERCOT	Electric Reliability Council of Texas
FLORIS	FLOw Redirection and Induction in Steady State
GPR	Gaussian process regression
HRRR	High-Resolution Rapid Refresh
ISO	independent system operator
ISO-NE	ISO New England
MISO	Midcontinent Independent System Operator
MW	megawatt
MWh	megawatt-hour
NREL	National Renewable Energy Laboratory
NWP	numerical weather prediction
PMC	power-maximizing control
RTO	regional transmission organization
SARSA	state-action-reward-state-action
SCADA	supervisory control and data acquisition
SPP	Southwest Power Pool
UTC	Coordinated Universal Time
WPPC	wind plant power control

Executive Summary

The research objective for the Atmosphere to Electrons to the Grid (A2e2g) project was to design a platform that merges forecasting tools with aerodynamic and economic models. The value proposition is that expanding wind power plant operation to include grid services allows those plants to operate in markets for grid services as well as energy markets, increasing revenue streams for wind plant operators while contributing to reliable grid operation.

The A2e2g project comprises a platform that integrates forecasting tools to account for weather uncertainty, aerodynamic wind plant models to account for wake dynamics and wind plant operation, and economic models to advise on operation for a wind power plant that offers grid services in addition to energy, as shown in Figure ES-1.

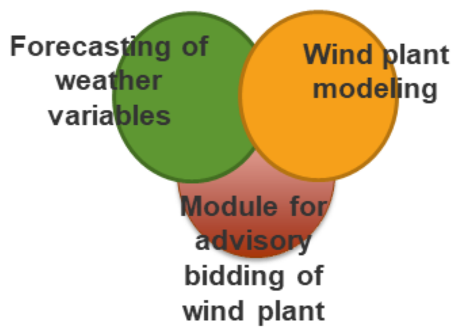

Figure ES-1. The A2e2g integrated platform

The A2e2g platform is a holistic tool with modules that can be run to advise on market participation and control and operate a wind power plant in real time. Figure ES-2 shows the different components of the platform. The A2e2g platform assumes two stages: the first is day ahead, with a focus on managing uncertainty, and the second is real time, with a focus on managing variability. Detailed descriptions of each component can be found in Sections 4–6. The different components have models written and developed in the Python programming language. The code is assembled into a Python package and can be easily downloaded and installed from the A2e2g repository.[1] The code is open source and free to the public.

[1] https://github.com/NREL/a2e2g

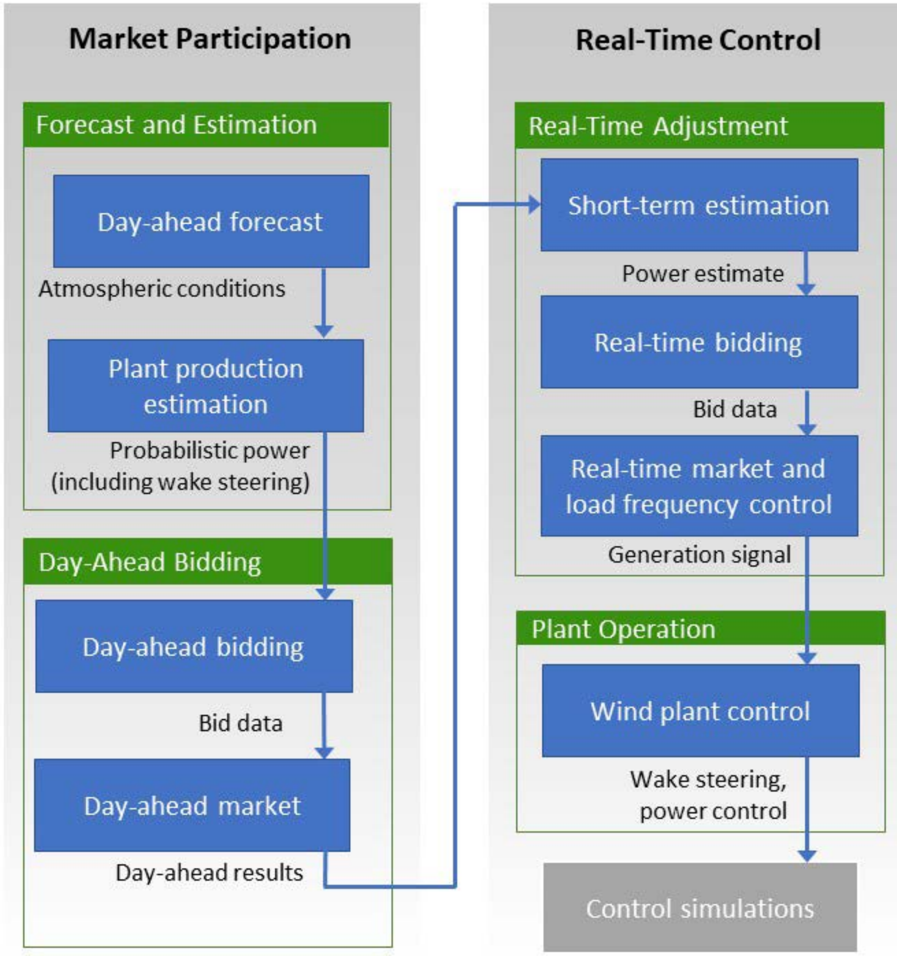

Figure ES-2. A2e2g platform showing market participation (first stage) and real-time control (second stage)

The A2e2g platform can be used as a high-level controller for a wind plant for market participation and real-time wind plant control. The A2e2g controller operates at the 4-second and longer time frame and does not include communication to inverter-based wind plant controls at the second and subsecond time frame. The 4-second time frame corresponds to the supervisory control and data acquisition (SCADA) operations of wind plants, and the A2e2g platform can be used to respond to grid service communications from power system operators for regulation and other slower grid services.

Operational responses of wind plants to grid frequency events in the second and subsecond time frames, such as primary frequency response, are not included in the A2e2g controller as those events would be immediately sensed by the wind plant inverter-based control system. Figure ES-3 gives an overview of the different grid services. Energy and capacity services are included in one category, with the remaining services in an "essential reliability service" category, a term for grid services that has recently been used to include additional grid services beyond those covered by "ancillary services" (the terms used for grid services are not always consistent across different power system operators). Essential reliability services are further subdivided into operating reserves and other services. The A2e2g platform could be used for providing energy and capacity and regulating and ramping reserves, as shown in Figure ES-3.

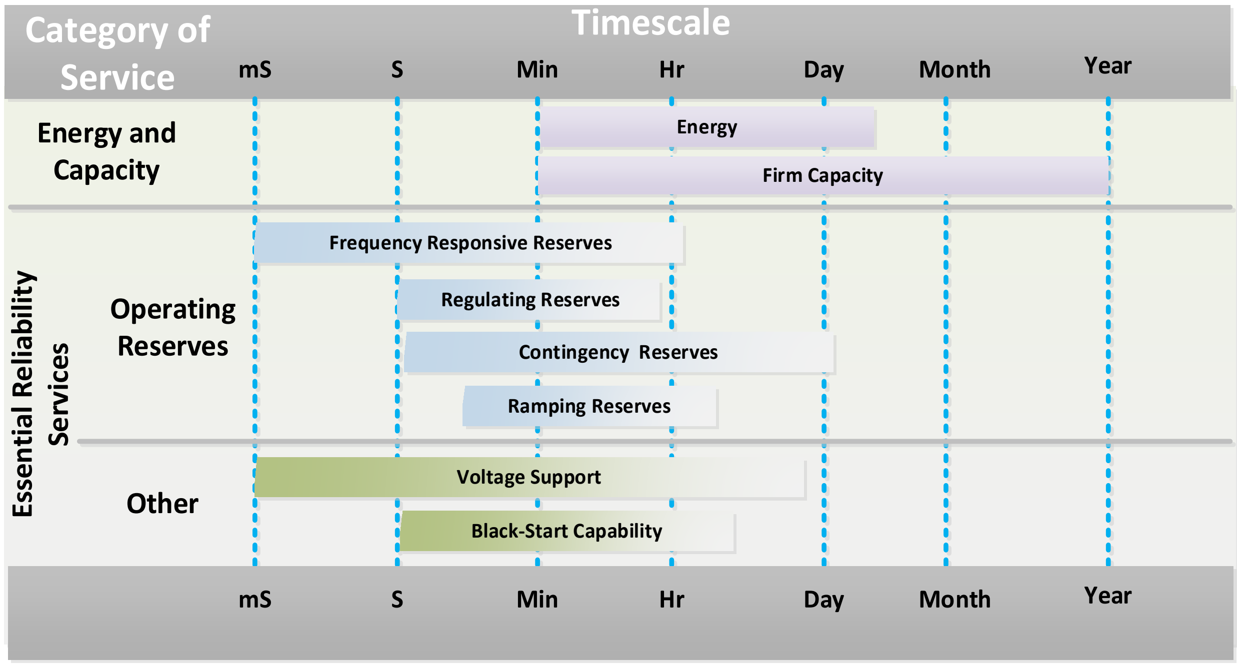

Figure ES-3. Overview of grid services

The project demonstrated the control functionality and communication of the A2e2g platform with the GE 1.5-megawatt wind turbine operating at the National Renewable Energy Laboratory's Flatirons Campus at the 4-second SCADA time frame, but it was not implemented on a wind plant.

Market participation. The A2e2g platform can be used to model different market participation scenarios for wind plants providing grid services. In this project, we focused on regulation as a grid service, as regulating reserves are generally more valuable and have current markets, require fast response, and are deployed in short time frames. However, the A2e2g platform demonstrated in this report can be expanded to include other grid services, including potential future grid services for which markets do not currently exist. We developed the A2e2g platform for both wind only and wind plus energy storage.

The development of the A2e2g analytical formulation for market participation required a detailed understanding of how wind plants participate in energy and ancillary services markets. In this project, we consider uncertainty in the wind resource, because in the majority of organized markets in the United States, a high volume of reserves is traded day ahead. In developing the A2e2g platform for market participation for regulation, we considered how variable energy resources consider independent system operator/regional transmission organization performance targets with respect to the delivery of regulation capacity to decide the regulation capacity to offer. We also evaluated adjustments to resource revenues for imperfect delivery of regulation capacity and how that could affect the opportunity costs the variable resource faces when it offers regulation. We provide case studies in Section 5.3 that illustrate the advantage of offering regulation with the A2e2g platform, the economic attractiveness of the regulation product, and how co-locating wind and energy storage could allow the wind-plus-storage facility to meet delivery targets not only for regulation but also for energy. Note that these case studies depend on the power system market rules, historical prices, wind resource for the time period under study, and related parameters, and would vary for other case studies.

Real-time control. In addition to the market participation component described earlier, the A2e2g platform also includes a wind-plant-level active power controller. A closed-loop power tracking controller is needed to enable wind power plants to respond to active generation control signals from the grid operator in real time (i.e., 4-second SCADA time frame). The controller takes the current power set point being sent by the grid operator and produces individual power reference signals for the wind turbines within the plant. The controller can account for important aerodynamic interactions between the turbines, which occur in a seconds-to-minutes timescale as turbine wakes propagate through the plant (note the similarity to the shortest timescale for A2e2g market participation). To increase the flexibility of a wind plant providing grid services, we include in A2e2g a battery providing short-term energy shifting, which helps to relieve curtailment and provide a faster active power response.

The A2e2g platform represents an important step toward developing "near-firm" power for a wind plant. A near-firm wind plant will be able to deliver on its day-ahead bids for power provision with high reliability. Advanced control of wind power plants with energy storage will provide firm and flexible day-ahead dispatchable power and pave the way for other renewable energy sources to do the same. Figure ES-4 summarizes some of the features and impacts of the A2e2g platform, along with the key takeaways from this report and the wind plant stakeholders who could benefit.

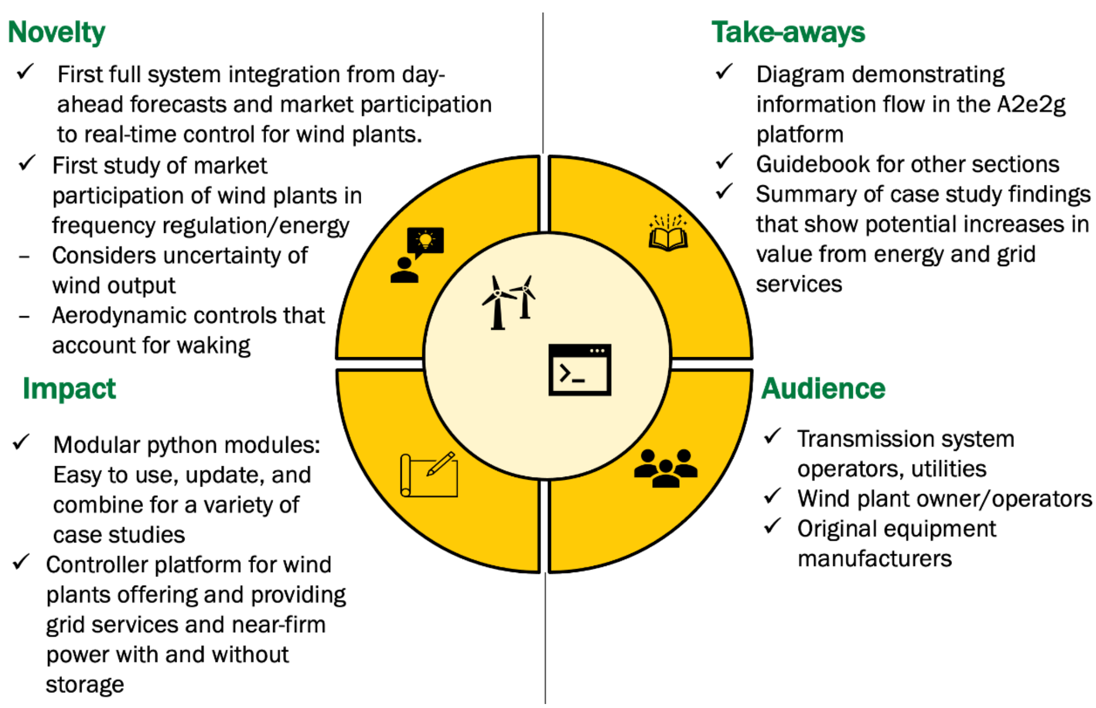

Figure ES-4. The A2e2g platform features and impacts

Table of Contents

List of Figures

List of Tables

1 Project Overview

1.1 Background

For a wind power plant, the resource inputs are driven by the planetary boundary layer flow via the atmosphere and the downscaling of that flow into the complex flow within the wind plant. The plant efficiency is a result of complex interactions between plant inflow, mechanical and power system controls for wind turbines and the plant, and power flow based on grid dynamics and power system operation and value streams. The energy capture and performance of a wind plant are driven by multiscale physical processes over different timescales including electromechanical and power system dynamics. Figure 1 shows a breakdown of the timescales for power system operation and how they relate to mesoscale modeling, wind plant operation, and wind turbine controls.

Wind plant cost and performance are considered the primary drivers for sustainable wind energy development and deployment, with a focus on optimizing wind plant technology to achieve future high wind penetration levels by lowering the levelized cost of energy and mitigating technical barriers to deployment. Although individual wind turbine technology has been well-optimized by original equipment manufacturers to meet commercial market demands, wind power plants have not received the same rigorous performance analyses. It is now necessary to examine the cost and performance of the entire wind plant for opportunities to increase the value per megawatt. To capture additional revenue, the wind power plant should meet power system operators' performance criteria for supply of additional services.

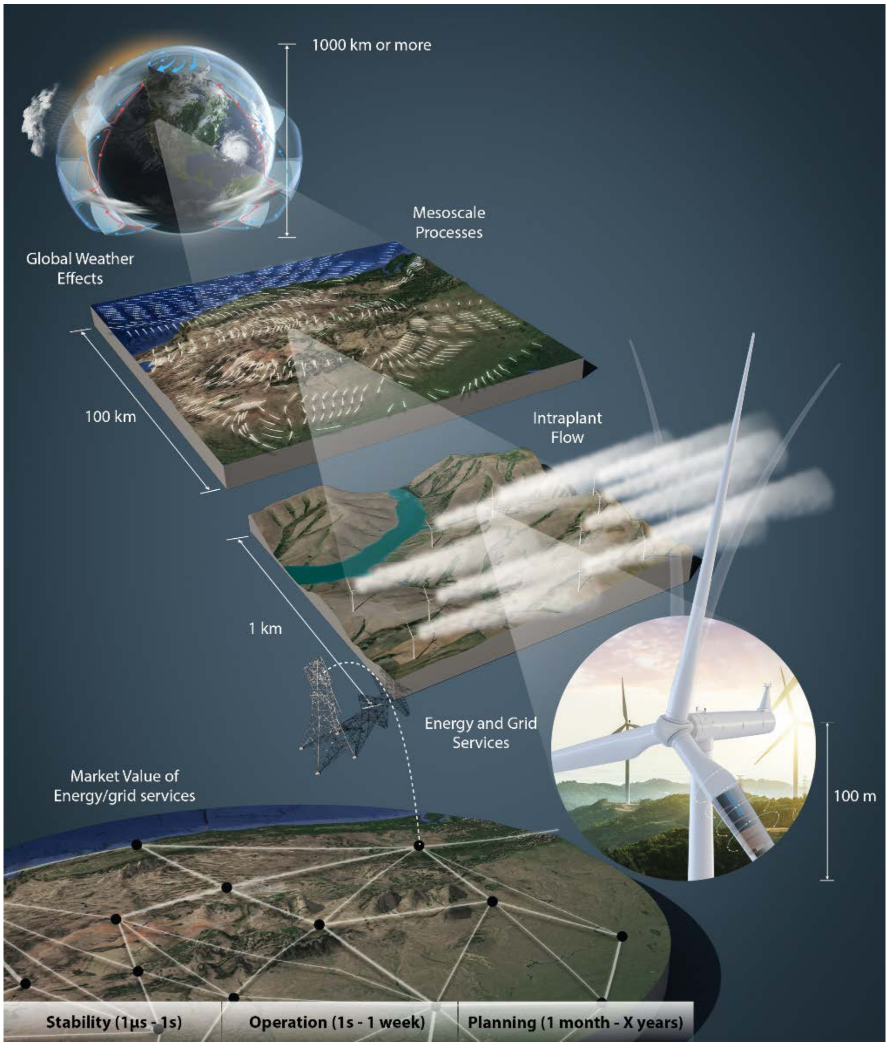

Figure 1. Spatial and temporal scales for a wind plant system. ***Image by Besiki Kazaishvili, National Renewable Energy Laboratory (NREL)***

1.2 Grid Services

As the percentage of wind energy grows, there is an increasing need for wind plants to provide "essential reliability services," also known as ancillary services, to maintain the reliability and stability of the power grid. Historically, these services were provided by larger synchronous generators such as fossil-fuel or hydroelectric generators.

In addition to supplying energy, wind power plants currently supply some grid services as required in several operating areas as well as by many market operators. In some cases, when the provision of certain grid services is a requirement, the wind plant may or may not get explicitly get paid for this service. For example, in the Public Service of Colorado operating region, wind plants must supply regulation down reserves through active power control (APC), which means operating the wind turbines to achieve a set power output. In the Electric Reliability Council of Texas (ERCOT), wind plants may provide ancillary services, but only a handful of wind generators have opted to qualify for this provision and their participation in regulation markets is currently minimal.

As the use of renewable energy increases, the need for grid services will likely increase, and as a result, so will the value of providing grid services. The research highlighted in this report is focused on the future, with the goal of developing tools that can help wind plant operators make informed decisions when providing these additional services.

As future electricity systems integrate higher amounts of variable energy resources, there will be fewer traditional resources that can supply services, thereby driving up the price of service provision. Hence, it would not be surprising if the wind plant of the future operates more like a dispatchable conventional generator, which actively manages the uncertainty and variability of its electric output. For example, the California Independent System Operator (CAISO) is currently designing an imbalance reserve product that will help operators manage uncertainty between the day-ahead and 15-minute market, and in their latest deliberations they added variable energy resources as eligible suppliers.

The Atmosphere to Electrons to the Grid (A2e2g) platform can be used as a high-level controller for a wind plant for market participation and real-time wind plant control. A2e2g operates at the 4-second and longer time frame. The 4-second time frame corresponds to the supervisory control and data acquisition (SCADA) operations of wind plants, and A2e2g can be used to respond to grid service communications from power system operators for regulation and other slower grid services.

Operational responses of wind plants to grid frequency events in the second and subsecond time frames, such as primary frequency response, are not included in A2e2g, as those events would be immediately sensed by the wind plant inverter-based control system. Figure 2 gives an overview of the different grid services. Energy and capacity services are included in one category, with the remaining services in an essential reliability service category (note the terms for grid services are not always consistent across different power system operators). Essential reliability services are further subdivided into operating reserves and other services. The A2e2g platform could be used for energy and capacity and regulating and ramping reserves, as shown in Figure 2.

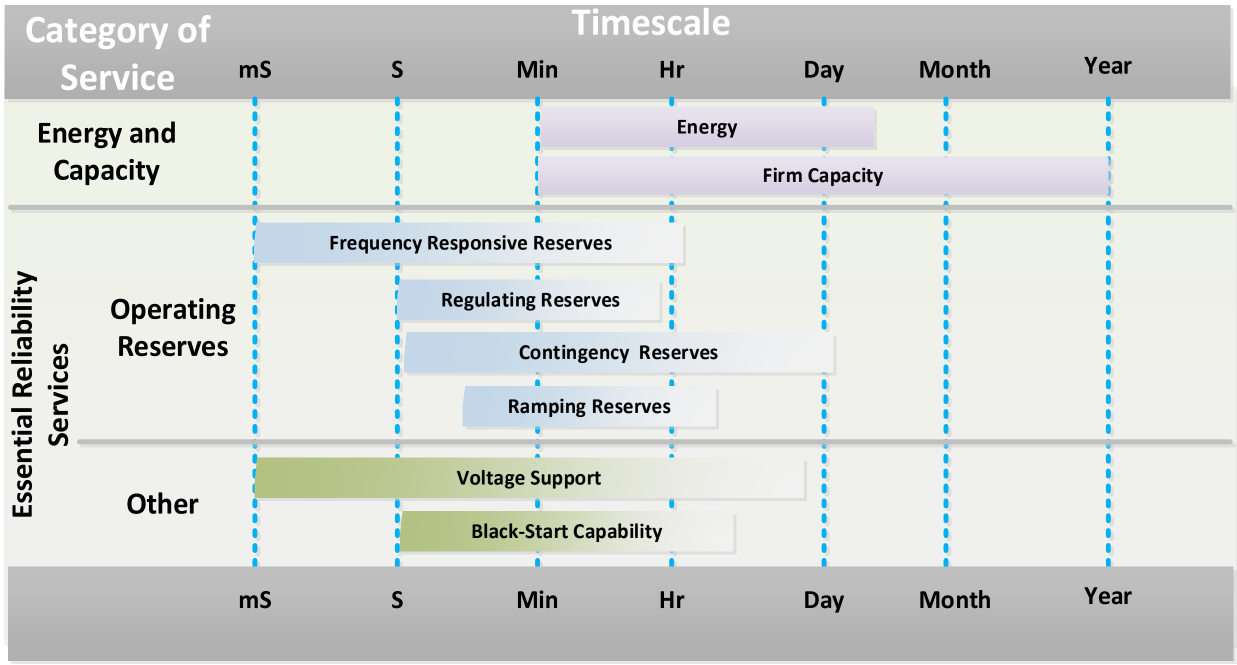

Figure 2. Essential reliability services

1.3 Summary and Project Timeline

The research objective for the A2e2g project was to design a platform that merges forecasting tools with aerodynamic and economic models to maximize a wind power plant's value streams for energy and grid services. The value proposition is that expanding wind plant operation to include grid services allows them to operate in both grid service and energy markets, potentially increasing revenue while contributing to reliable grid operation.

During the first year of the project, we developed a modeling platform for A2e2 to support design optimization and controls development. In that same year, we implemented the A2e2g platform in a LabVIEW controller at the National Renewable Energy's (NREL's) Flatirons Campus for the GE 1.5-megawatt (MW) wind turbine to demonstrate the control of an individual wind turbine and demonstrate the communications with a FLOw Redirection and Induction in Steady State (FLORIS) simulation model.[2] The A2e2g platform was also simulated for a 110-MW wind plant using ERCOT market data to simulate the full wind plant control for the different timescales and the interaction with energy and grid service markets.

Work in the second year included developing probabilistic, dynamic forecasting approaches using the FLORIS model to account for plant aerodynamics in power forecasts, including capturing the value of wake steering for real-time control. The A2e2g advisory module for market participation was developed to yield offers for regulating reserves that consider probabilistic forecasts for wind power, expected imbalance costs, and compliance performance targets (see Section 3) for grid services for a given market footprint (e.g., ERCOT). The forecasting of weather variables included new probabilistic forecasts using state-of-the-art

[2] https://www.nrel.gov/wind/floris.html

methods (see Section 4) developed in the second year. This work was integrated into a novel wind plant scheduling and control platform (referred to as the "scheduler") that provides for the maximum value of energy and full grid services under historical prices for any aerodynamic conditions, as shown in Figure 3.

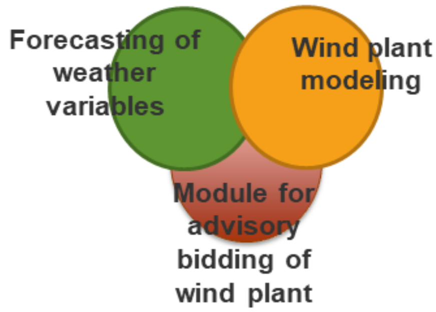

Figure 3. A2e2g integrated platform

Work in the third year included making further improvements to the scheduler, such as:

- Adding energy storage
- Refining the Python code
- Integrating aerodynamic and grid controls
- Developing use case examples including a full end-to-end wind power plant simulation for the ERCOT market area
- Developing and implementing scheduler equations for energy storage
- Comparing the FLORIS day-ahead and real-time models
- Integrating battery state information into the aerodynamic controller.

1.4 Key Contributions

The A2e2g platform represents an important step toward developing "near-firm power" for a wind plant, as shown in Figure 4. A near-firm wind plant will be able to deliver on its day-ahead bids for power provision with high reliability. Advanced control of wind plants with energy storage can provide firm and flexible day-ahead dispatchable power and pave the way for other renewable energy sources to do the same. The major components of A2e2g include the wind plant aerodynamic controller, which manages variability and operates in real time at the wind plant level, and the market participation advisor, which manages uncertainty in the minutes-to-day-ahead time frame and allows the wind power plant operator to consider system signals, such as market prices for energy and grid services.

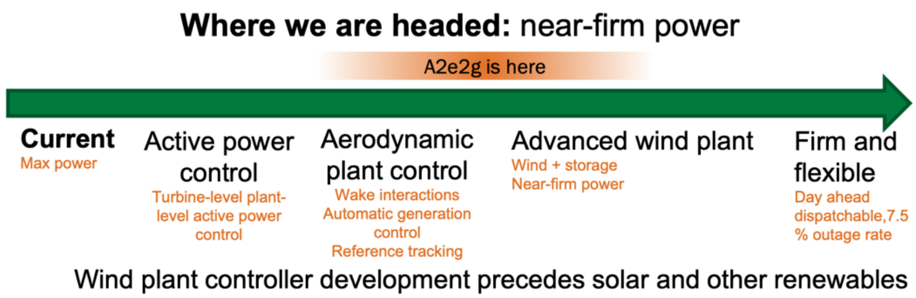

Figure 4. General timeline of advanced wind power plant control

Figure 5 summarizes some of the features and impacts of the A2e2g platform, along with the key take-aways from this report and the intended wind plant stakeholders that could benefit.

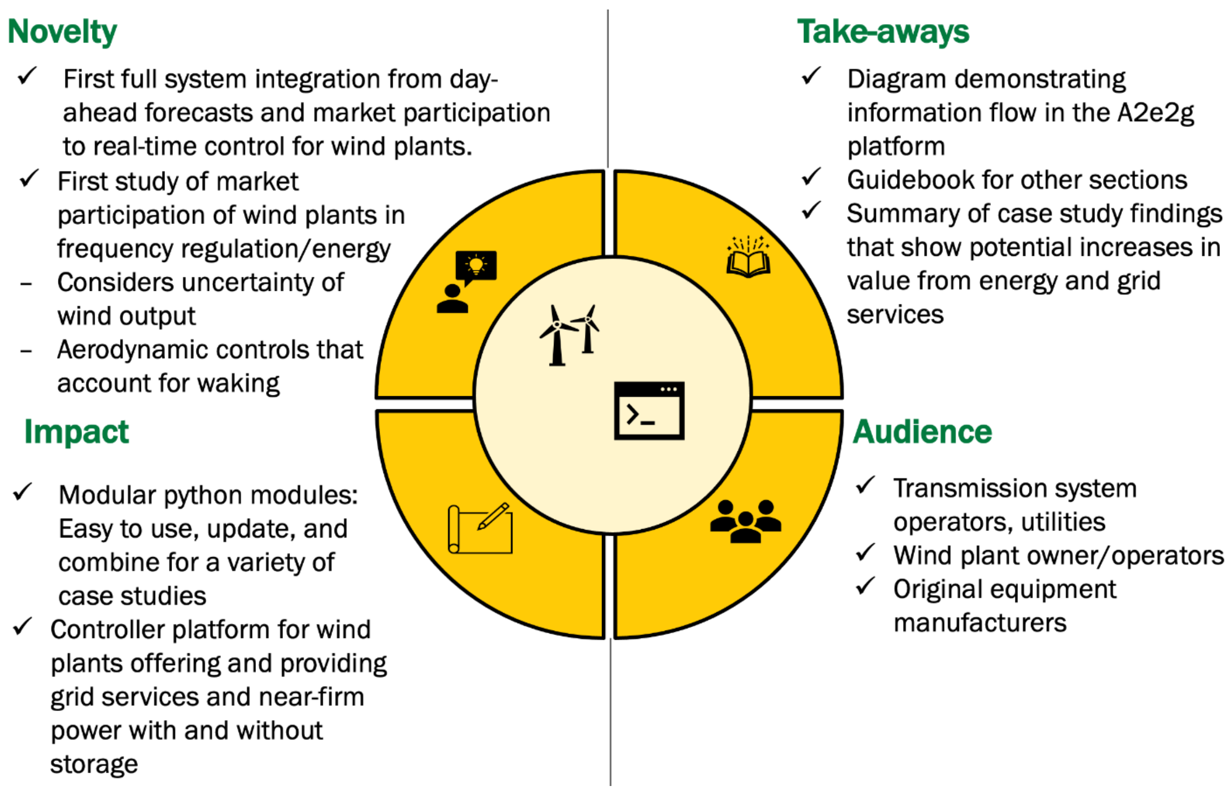

Figure 5. The A2e2g platform features and benefits

1.5 Report Structure

The rest of this report is structured as follows. In Section 2, we give a high-level description of the A2e2g platform, including its modules and code implementation. This section should be most useful to those with a systems viewpoint and those desiring a more general background of the project. Section 3 provides background on regulating reserves disqualification and settlement practices for major U.S. system operators. In Sections 4–6, we provide detailed descriptions of the A2e2g modules as well as results where appropriate. Section 4 describes atmospheric and power forecasting; Section 5 explains how the wind plant participates in energy and ancillary services markets; and Section 6 covers the operation of the wind plant to respond to power set points received from the system operator. Section 7 concludes this report.

2 Architecture of the A2e2g Platform

2.1 A2e2g Overview and Diagram

The A2e2g project developed a platform that integrates forecasting tools to account for weather uncertainty with aerodynamic models to account for wake dynamics and wind plant operation, and economic models to advise on operation for a wind power plant that offers grid services in addition to energy.

The A2e2g platform is a holistic Python tool with modules that can be run to advise on market participation and control and operate a wind power plant in real time. Figure 6 shows the platform components. The purpose of each component will be described briefly in the next subsection, with detailed descriptions that follow.

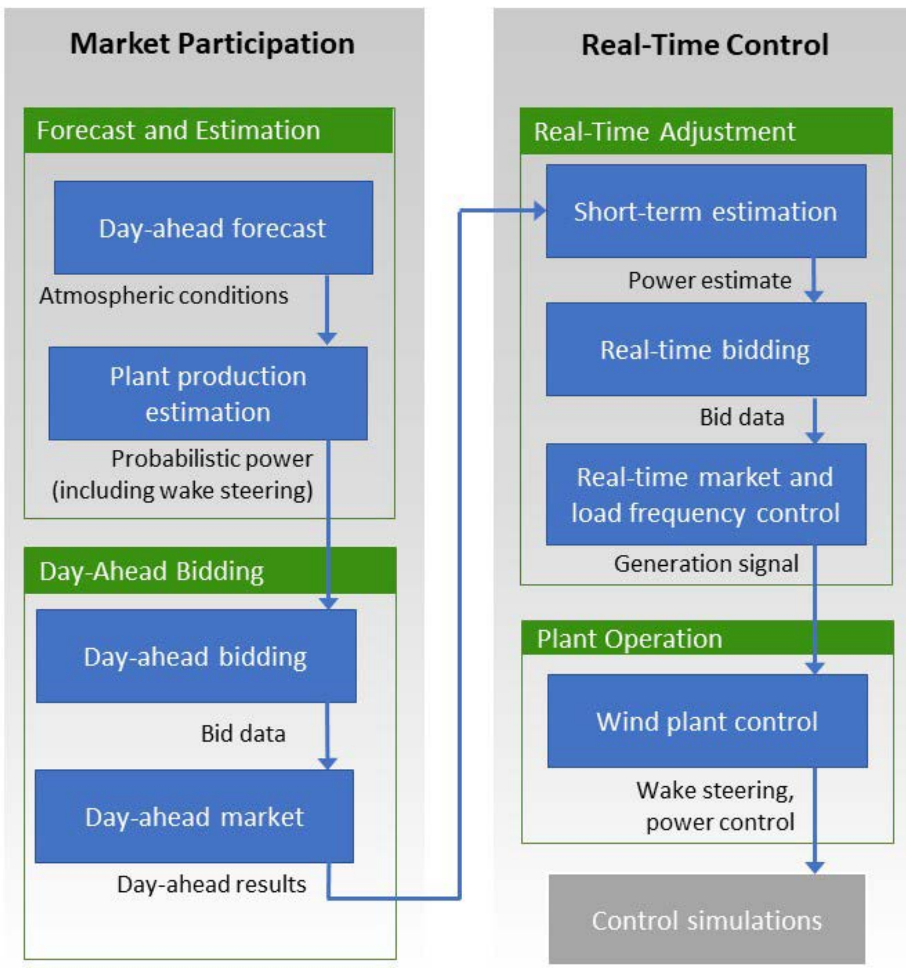

Figure 6. A2e2g platform showing market participation (first stage) and real-time control (second stage)

The current structure of the platform assumes two stages: the first is in day ahead and the second is in real time. We assume that managing uncertainty (i.e., errors of the wind power forecast due to imperfect forecast of weather variables such as wind speed and wind direction) is key in the first stage and managing wind speed and wind direction variability is key in the second stage.

For that reason, the first stage modules comprise probabilistic forecasting of wind power and offering of energy and regulating reserves. Note that the first stage (Figure 6) includes a module to simulate the day-ahead market. In this application, the day-ahead market module reflects the day-ahead market problem the system operator would solve in practice.[3] For this report, we conducted price-taker analyses that assume historical prices would be observed even if the wind power plant used the advisory offering module to participate in the market. In other words, the price-taker analyses assume that the change in the offering approach of a single participant (here, a wind power plant) would not affect the historical prices observed. For future work or follow-up studies that investigate the impact of an updated offering by multiple power plants at the same time, detailed production cost models could be used as day-ahead modules. The second stage modules reflect the participation of the power plant in real-time operations and controls at the system level and a module that controls the wind power plant so that it can follow the "basepoint" signal (i.e., the active power output the plant is instructed to follow by the load frequency control system).

2.2 Platform Components

2.2.1 Day-Ahead Atmospheric Forecasting

The first step for a wind power plant operator to participate in day-ahead energy and service markets is to obtain a forecast of the atmospheric conditions at the plant. The A2e2g platform provides a location-specific forecast based on readily available numerical weather prediction (NWP) model data. The forecast includes wind speed, wind direction, and turbulence intensity, which can each affect the wind plant's power output.

The atmospheric forecasts generated are probabilistic, including both an expectation and standard deviation (assuming a Gaussian distribution for each forecast component). As a result, we can provide a probabilistic treatment of the plant power forecast and finally plant operator bids into day-ahead markets.

2.2.2 Probabilistic Forecast for a Plant's Potential Upper Operating Limit

Next, probabilistic forecasts of the plant's power are generated. This step takes the probabilistic wind speed and direction forecasts as input and, using a steady-state flow model of the wind plants, produces a probability distribution for the power produced by the wind plant at 5-minute intervals. The use of a flow model in the power forecasting step, as well as wind direction input information, ensures that wake interactions between wind turbines are captured when generating the wind plant power estimate. The mean and standard deviation of the distributions (along with an assumption for a Gaussian distribution) are then returned to be used in the day-ahead offering module.

[3] A sophisticated participant could choose to simulate multiple day-ahead market outcomes using production cost models with different assumptions. Using the outputs of the production cost model, the participant could submit an offer that considers overall system conditions.

2.2.3 Day-Ahead Offering

This module estimates the maximum capacity the wind power plant could offer for the ancillary service in question: regulating reserves. Although the actual module for this platform component is straightforward (a single formula), we took multiple steps to derive this formula, starting with the power system operator's performance requirements for delivering regulating reserves. The performance requirements are described in the tariff of each power system operator and are used to decide if resources should disqualify from providing regulating reserves. It is worth noting that this module is specific to the power system operator and aspects such as the length of the performance assessment period along with the plant's ramp rate affect the capacity that can be offered. In practice, a wind power plant operator could extend this module to consider any price forecasts that would help them optimize the objective function of its choice (e.g., the wind power plant's expected profit).

2.2.4 Day-Ahead Market

This module is a price-taker dispatch model that uses historical prices as inputs to maximize the profit of the wind power plant generated from supplying energy and regulating reserves. The price-taker model assumes that the change in the offering of a single wind power plant would not affect the prices observed. However, if the offering of multiple wind power plants changes, future studies could consider using a production cost model in place of this module. For example, NREL's Scalable Integrated Infrastructure Planning Model (NREL undated) is a tool that could be used for more detailed simulations of day-ahead markets.

2.2.5 Short-Term Atmospheric and Power Forecasting

Forecasts are regenerated for balancing in the real-time markets. A deterministic short-term atmospheric forecast is generated 5 or 10 minutes ahead based on persistence of the current wind conditions. The wind plant model is evaluated based on the short-term atmospheric forecast to produce a plant power forecast accounting for waking interactions. This power forecast is provided to the real-time offering module.

2.2.6 Real-Time Offering

For most of the case studies we conducted, the module is simplified so that the wind power plant is not coupled with storage. Hence, the module compares the day-ahead energy award to the real-time available power and offers to buy or sell energy in case of over- or underforecasting, respectively.

2.2.7 Real-Time Market and Signal Generation

Similar to the day-ahead market module, this module could also be replaced by an open-source tool that could simulate real-time market operations and automatic generation control (AGC). Here, we use a price-taker dispatch module to yield the real-time awards and historical AGC signals to derive power reference signals with a 4-second resolution that the wind power plant controller must follow.

2.2.8 Wind Power Plant Aerodynamic Control

The wind plant's ability to generate active power at the level requested by the system operator depends on its real-time control system. Therefore, we provide a wind-plant-level aerodynamic controller for the real-time operation of the plant. The controller provides individual power references to each wind turbine in the plant and balances the references to ensure that the total power being generated tracks the requested active power level. This balancing is needed to account for wake interactions between turbines, which alter the wind resource available to each turbine in the farm and may mean that some turbines may not be able to generate their "share" of the power request. The individual wind turbine controllers track their references using a turbine-level active power controller, which is assumed to use blade pitch control for demonstrative purposes.

The control module also includes operating logic for managing an energy-shifting battery. The operating logic is designed to help the wind power plant in regulating its active power output while minimizing the need for dynamic curtailment of the plant.

2.2.9 Aerodynamic Simulations

We also include a module to mimic the aerodynamic interactions in a real wind power plant using a medium-fidelity aerodynamic model. This model allows us to simulate a wind plant's response to power set points using the real-time aerodynamic controller. The simulation replaces a real wind plant to validate the controller's handling of unknown disturbances and complex flows. It is suitable for short-term simulations (on the order of 1 hour to 1 day of operation) and can be driven using actual site-measured wind conditions and actual system operator active power signals, or synthetically generated conditions and signals. This module allows control system engineers to evaluate and compare competing control system designs prior to deployment on an actual wind power plant. If deploying the A2e2g platform on an operational wind power plant, this module would be replaced by the real wind plant.

The aerodynamic simulator is a code that approximates wake propagation dynamics as well as wind turbine and flow responses to changes in turbine set points while being relatively lightweight (operable on a personal computer) and user-friendly for those familiar with FLORIS. Users can extract the error between the system operator's power signal and the actual dynamic power produced by the plant, which could be used for further analysis of control system performance.

2.3 A2e2g Code

All the models discussed in this report were written and developed in the Python programming language. The code is assembled into a Python package and can be easily downloaded and installed from the A2e2g repository (https://github.com/NREL/a2e2g). The code is open source and free to use. To enable modularity and debugging of the code, each of the modules are coded individually so that they can be swapped in or out of the overall simulation, allowing for use of the models or actual data for each operation. This approach also allows for the effective testing of the code through clear divisions, as well as the future use of different models within each section.

The main operational code is divided into five separate submodules: forecast, power_forecast, market, control, and simulation. The forecast, power_forecast, and market submodules contain the models that produce the weather forecasts, power forecasts, and market signals, respectively. The control submodule contains the various plant-level, real-time aerodynamic control strategies for following power set points from the system operator. Finally, the control_simulation submodule provides tools for testing the aerodynamic controllers interacting with a simulated wind plant. This submodule stands in for a real wind power plant in simulation and would not be used for actual deployment on a real plant. Significant work was also done so that the code can be used in high-performance computing environments, which can be useful when looking at several long simulations across many different configurations.

3 Background on Delivery Targets for Regulating Reserves

In addition to energy and capacity, there are multiple essential reliability services that wind power plants could provide (Denholm et al. 2019). As a result, we decided to focus on regulating reserves because they are generally more valuable, require fast response, and are deployed in short time frames (Ela et al. 2019; Kahrl et al. 2021). We also considered the possibility of imperfect delivery of regulating reserves due to forecast errors at the time of offering. Our literature review revealed that reports that consider supply of ancillary services in U.S. markets generally ignore uncertainty (Kahrl et al. 2021; Loutan et al. 2020; Loutan et al. 2017), whereas academic papers usually consider uncertainty but in a profit-maximizing paradigm accounting only for hypothetical penalties that plants would incur if they did not deliver (Soares et al. 2017; Liang et al. 2011). To address this literature gap, we reviewed targets for delivery of regulating reserves and adjustments to settlements as a result of imperfect delivery and in Section 5 we discuss the implications of our findings for wind power plants participating in regulation markets. We provide a detailed discussion of current (as of April 2021) practices in Syprou et al. (2022) and here we summarize key take-aways from the review of publicly available documents (Southwest Power Pool 2020; New York Independent System Operator [NYISO] Operations Engineering 2020; California Independent System Operator Corporation 2020; NYISO Customer Settlements 2020; NYISO 2021; Midcontinent Independent System Operator [MISO] 2019; California Independent System Operator Corporation 2019; PJM 2016; PJM 2020a; Southwest Power Pool 2015; ERCOT 2020; ISO New England 2015; Henson 2015; MISO 2018; Reedy 2018; PJM 2020b). Our summary is based on a review of publicly available documents, and it has not yet been reviewed by power system operators.

3.1 Performance Targets

Before offering regulating reserves to the market, a resource must qualify as a supplier by undergoing a test. Once the resource is qualified, its performance with respect to delivery of regulating reserves is monitored and decisions on its disqualification are made on a regular basis (e.g., every month), on a case-by-case basis or after completing unannounced tests. In the case of periodic calculation of performance metrics, the wind power plants would aim to offer regulating reserves at a level that would keep them from being disqualified.

Our literature review suggests that system operators who periodically monitor delivery of regulating reserves use two types of metrics. The first type of metric (Type I) assesses how frequently over the assessment period the accuracy of delivery of regulating reserves was worse than a specified tolerance. If the frequency of “unacceptable delivery” is higher than operator-determined target frequency, F_{TARGET}, the plant is disqualified. The second type of metric (Type II) estimates the average realization rate (i.e., the capacity of regulating reserves delivered vs. requested or contracted). If the average realization rate or precision during an assessment period is less than the operator-determined target score, s_T, the resource is disqualified. Note that the performance score could be an aggregate score that tracks, in addition to precision, time delays. For managing uncertainty, the precision score is more relevant. We hereafter refer to metrics that estimate average precision as Type II. Table 1 summarizes the performance metrics and rules for resource disqualification when the current practice involves periodic calculations of performance metrics.

Table 1. Performance Metrics and Rules for the Disqualification of Regulation Providers Used by Five U.S. Independent System Operators/Regional Transmission Organizations That Periodically Estimate Performance Metrics

	Metric Name or Short Description	Assessment Period	Disqualification Rule
ERCOT (Type I)	Number of 5-min intervals with generation resource energy deployment performance within tolerance	Calendar month	Metric is below 85% of intervals with regulation awards
New York Independent System Operator (Type II)	Regulation performance index	Calendar month	Metric is below 0.85
Southwest Power Pool (Type II)	Compliance rating	Calendar month or 5 best tests	Metric is below 75%
CAISO (Type II)	Historical regulation performance accuracy	Calendar month	Metric is below 25%
PJM Interconnection (Type II)	Performance score (delay, correlation, precision)	100 hours	Metric is below 40%

3.2 Settlement Schemes

The supplier of regulating reserves is commonly paid a price for regulation, which is an output of an optimization problem solved by the power system operator. For co-optimization of energy and reserves, this price considers opportunity costs that suppliers might face when they are asked to supply regulation instead of energy. The regulation price is multiplied by the capacity that the participant is asked to reserve for regulating reserves to yield the payment for the participant. In the case of CAISO, in addition to capacity payments, there could be payments associated with mileage. When regulation is being deployed, adjustments to energy payments are common. For example, regarding regulation down deployment, the generator produces less energy than the market rewarded them for, and the energy revenues are adjusted accordingly. The adjustments to the energy payments could be considered by wind power plants when they offer regulating reserves. More importantly, in this project, we aim to understand the impact of imperfect delivery of regulating reserves on settlements.

Based on our review, it is clear that in all seven U.S. independent system operators (ISOs)/regional transmission organizations (RTOs) payments are adjusted when the delivery of regulating reserves is imperfect. In summary, there are at least four ways in which the settlements are adjusted. Under the first scheme, the payment for regulating reserves is calculated upon delivery and is proportional to the delivery performance or availability. Under the second scheme, the payments for regulating reserves are fully rescinded when the delivery is below an acceptable threshold. The third scheme combines the first two because it fully rescinds payments when performance is below a threshold and pays proportional to the performance score when performance is above a threshold. Lastly, under the fourth scheme, the credit for providing regulating reserves is not adjusted. However, the suppliers are subject to over-/undergeneration penalties with stricter tolerances compared to energy-only resources. In Table 2, we summarize our observations and further consider the implications of the different schemes for offering of regulating reserves by wind power plants in Section 5.2.

Table 2. Credits and Penalties for Regulation Settlement Schemes Across All U.S. ISOs/RTOs

	Award Credit	Penalty
ERCOT (Scheme 4)	At market prices	Under-/over-generation penalty with stricter tolerance than for energy resources
New York Independent System Operator (Scheme 1)	Regulation performance index—adjusted	Undergeneration penalty with different penalty than for energy resources
Southwest Power Pool (Scheme 2)	At market prices	Rescission of award credits if outside tolerance for uninstructed deviations
CAISO (Scheme 1)	Adjusted for unavailable capacity (CAISO 2009), (CAISO 2012)	[The penalty is implicit because the credit is adjusted based on availability]
PJM Interconnection (Scheme 3)	Performance score adjusted	No credit if score is below 25%
ISO-New England (Scheme 3)	Performance adjusted	No credit if automatic response rate or regulation capacity deviates more than 20% or 15%, respectively
Midcontinent Independent System Operator (Scheme 2)	At market prices	Rescission of awards and prorated share of regulation market cost if unit fails to follow the set point or provide acceptable mileage for four consecutive intervals of 1 hour

4 Forecasting

Forecasts are needed for the power plant operator to bid into day-ahead and real-time markets. The forecast module in the A2e2g platform provides atmospheric and power forecasts to meet this requirement.

4.1 Day-Ahead Atmospheric Forecasting

4.1.1 Meteorological Data

Before a forecast can be made and deemed "accurate," observations are needed to validate the forecast. For meteorological variables at hub height, tall-tower measurements are often the best observations available to validate a forecast. Unfortunately, these measurements are hard to come by, and users are often left to verify the observations based on scarce metadata. Despite this, multiple efforts were made to find a data set worthy of inclusion in this study, and we ultimately decided on the high-frequency measurements of the 200-meter (m) meteorological tower at the National Wind Institute at Texas Tech University (Kelley and Ennis 2016), as shown in Figure 7. Observations were obtained at two heights (75 and 116 m) at a frequency of 50 hertz. We performed quality control on these data after acquisition based on Kelley and Ennis (2016) and resampled to 1-minute resolution for the period of August 1, 2019, to July 27, 2020. The primary variables of interest are wind speed, wind direction, and turbulence intensity.

Data for the day-ahead forecast comes from the High-Resolution Rapid Refresh (HRRR) numerical weather prediction model (Benjamin et al. 2016). At the time of this study, the HRRR model had a spatial resolution of 3 kilometers and a temporal resolution of 1 hour, with the forecast horizon extending to 36 hours. The HRRR is run hourly, with 36-hour forecasts made four times a day (00, 06, 12, and 18 Coordinated Universal Time [UTC]), and 18-hour forecasts at all other times. Because of the requirement of a day-ahead forecast by a particular time each day for this project, the forecast run of 12 UTC is obtained. Once the HRRR is downloaded, the forecasts are cubically interpolated to 5-minute resolution to better align with observations and power forecast requirements.

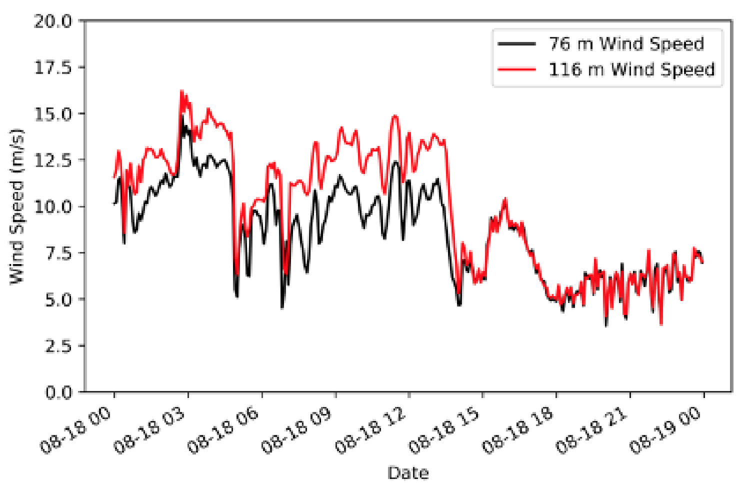

Figure 7. Example wind speed data from the 200-m meteorological tower at the National Wind Institute

4.1.2 Forecasting Method

For forecasting on the day-ahead timescale (and longer), NWP is the go-to method. In the last 5−10 years, NWP has been combined with a variety of postprocessing methods, such as machine learning and other statistical methods to further improve upon the NWP forecast (Foley et al. 2012; Giebel and Kariniotakis 2017). Here, we use the well-tested method of analog ensemble forecasting (Monache et al. 2013; Kumler and Lundquist 2021; Monache et al. 2011). The analog ensemble technique leverages past observations and historical forecasts to create a library of validation data. Once a new forecast is made, the technique analyzes the historical forecasts and their corresponding observations. The 10 best historical analogs (five historical forecasts and their corresponding observations, for a total of 10) are selected similar to the metric introduced in Monache et al. (Monache et al. 2013), and the average of the 10 analogs is taken for a 6-hour period. This process is continued until the entire day-ahead forecast is made. A probabilistic forecast is generated by taking the standard deviation of the 10 analogs for each time stamp.

Due to the forecast-horizon limitation of 36 hours, and the need to download the 12 UTC model run, the entire day-ahead forecast cannot be obtained for the Texas site. Despite this, we can leverage the strengths of the analog ensemble technique to fill the gap in the forecast by analyzing the analogs for the entire day ahead. Based on these results, the remaining 6 hours are filled, and proven to be accurate despite the absolute absence of an NWP forecast as shown in Figure 8.

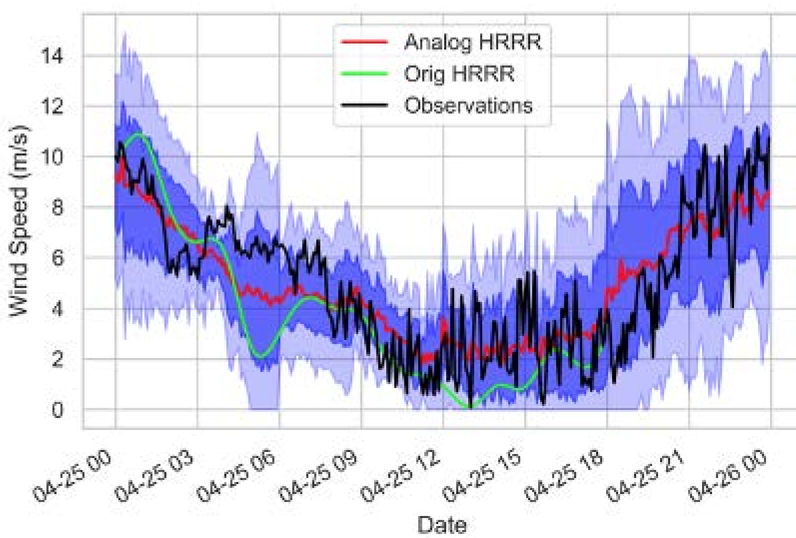

Figure 8. Example day-ahead forecast for the 80-m wind speed using the analog ensemble technique (the dark purple region represents one standard deviation while the lighter purple represents two standard deviations)

4.2 Probabilistic Plant Capacity Forecast

With the wind speed and direction forecasts generated as described in Section 4.1, we proceed to generating the day-ahead forecast for the power output of the wind plant. To do so, we used FLORIS (NREL 2020). FLORIS is a wind power plant flow modeling tool that is computationally efficient and can be used for making real-time controls decisions. FLORIS models time-averaged turbine-to-turbine waking interactions and their effect on the power output of the plant via velocity deficits at the wind turbine locations. Additionally, FLORIS models the effects of turbine yawing (the rotational motion of the wind turbine around a vertical axis, used to align the turbine rotor with the oncoming wind) both on turbine power and wake propagation. This is useful if wake steering (Fleming et al. 2017) is implemented at the plant.

To generate wind plant power forecasts from the probabilistic atmospheric forecast, we use a pregenerated "wind plant power curve." The FLORIS model of the wind plant is evaluated for a full sweep of operational wind speeds and wind directions, generating a lookup table for the power output of the farm in any given atmospheric condition. An example of such a wind plant power curve is provided in Figure 9. The figure demonstrates the plant power (on the y-axis) as a function of both wind speed (plotted on the main x-axis) and wind direction (which varies from 0 to 360 degrees for each 0.4-meter-per-second step in wind speed), highlighted in the inset. Wind speed is the main factor in determining plant power output, but wind direction also plays an important role due to wind-turbine-to-turbine waking, as demonstrated in Figure 9.

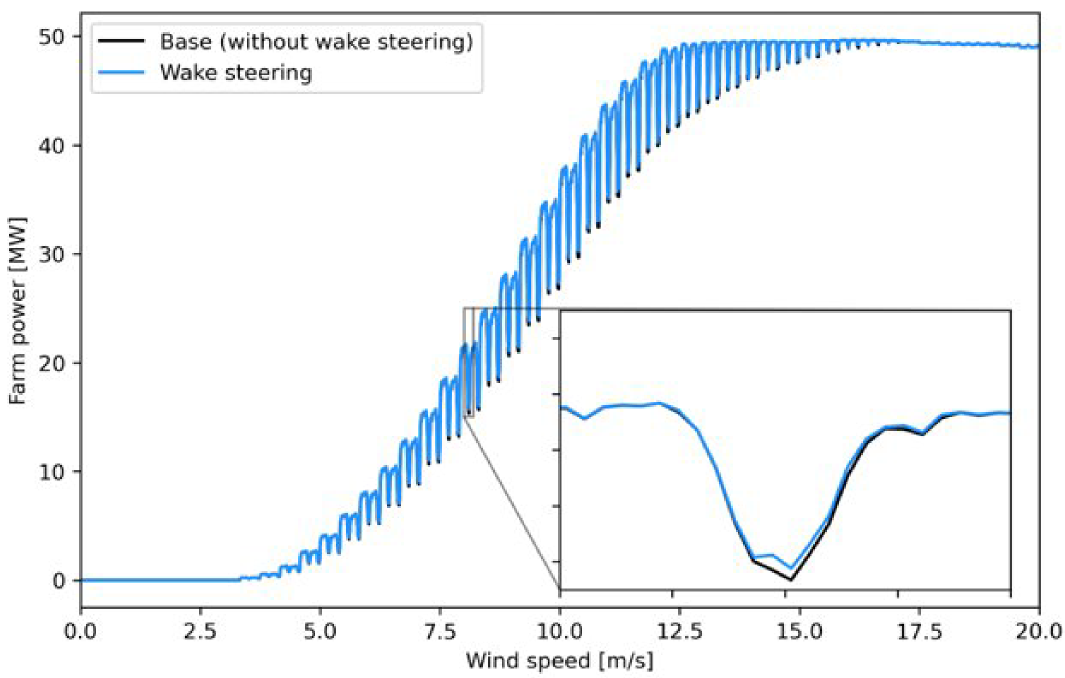

Figure 9. Power curve for a 50-MW wind power plant accounting for both wind speed and wind direction with wake steering active (blue line) and inactive (black line)

This wind plant power curve is used to generate a probabilistic plant power forecast as follows. We assume that the wind speed and wind direction can be represented as independent Gaussian random variables, and for each 5-minute time step in the day-ahead atmospheric forecast, probabilities are assigned to each of the pairs of lookup table inputs (wind speed and direction) according to how likely each condition is based on the forecasted wind condition. The independent Gaussian assumption is a simplification, but if a forecast with a higher degree of accuracy (such as wind speed/wind direction correlation information) is available, the assigned probabilities can be replaced with those from an arbitrary distribution.

The mean wind plant power, μ_P, and the power standard deviation, σ_P, are then computed according to the following canonical expressions, where $\mathcal{W}$ is the set of wind conditions used to generate the plant power curve (so that $w \in \mathcal{W}$ is a wind speed, wind direction pair); $p(w)$ is the probability of that pair occurring given the atmospheric forecast distribution (after normalizing the total probability to 1), and $P(w)$ is the FLORIS-modeled power output of the farm for that wind condition.

$$\mu_P = E[P(w)] = \sum_{w \in \mathcal{W}} P(w) p(w) \tag{1}$$

$$\sigma_P^2 = E[(P(w) - \mu_P)^2] = \sum_{w \in \mathcal{W}} (P(w) - \mu_P)^2 p(w) \tag{2}$$

This approach is repeated at every time step to obtain a mean power and standard deviation in power at 5-minute intervals for the following day to be used in day-ahead bidding.

Note that the approach we present here is based on a power curve that is computed offline. This approach allows for fast online implementation by simply multiplying the pregenerated powers, $P(w)$, by their probability of occurring. We provide the pregenerated power curve both with and without wake steering activated (see Figure 9) to account for both operational scenarios. However, these pregenerated curves cannot account for unforeseen changes in the plant, such as a wind turbine being shut off for maintenance. An alternative would be to apply the mean and standard deviation calculations mentioned earlier to powers computed online. In this paradigm, the FLORIS model is updated based on the status of all turbines; wind condition pairs, w, are sampled from the distribution provided by the atmospheric forecast; the power of the plant for each sample, $P(w)$, is found by evaluating the updated FLORIS model; and the power samples are collected to find an empirical mean and standard deviation. Although we ensured that the A2e2g platform can handle this approach, we have not fully implemented it because the computational time needed to evaluate FLORIS online at each sample time is prohibitive. We consider this an area for future investigation.

An alternative to the model-based power forecasting described here is to use a surrogate model to provide the mean power and standard deviation. This surrogate model method is included in the A2e2g platform for interested users but is not our main approach; details are provided in Appendix A.1.

4.3 Short-Term Atmospheric and Power Forecasting

Real-time markets require a short-term forecast of power for the wind plant. In this context, “short term” refers to a 10- or 5-minute-ahead forecast of the power production. Within A2e2g, we provide a persistence-based, deterministic short-term forecast using atmospheric conditions. For the purpose of this description, we will assume a real-time market update period of 5 minutes. At any given real-time market, the x and y components of wind speed are averaged over the last 5 minutes and these averages are used as a prediction of the wind speed and wind direction 5 minutes in the future. The deterministic wind speed and wind direction predictions are passed through the FLORIS plant model to produce a power estimate 5 minutes ahead of real time. The advantage of using the atmospheric conditions and FLORIS model, as opposed to simply applying persistence to the current wind power plant output itself, is that the FLORIS model inherently provides an expected power condition, whereas the plant output may contain dynamic effects that will have changed by the time the 5-minute-ahead forecast comes to bear.

5 Market Participation

Several articles studied the offering of regulating reserves by wind power plants, anticipating benefits for plant and power system operators (Nock et al. 2014; Soares et al. 2017; Kahrl et al. 2021). For plant operators, benefits are anticipated because of an increase in the value streams the plants have access to. For system operators, benefits are anticipated because of an increase in the number of suppliers for regulating reserves, a product essential to grid reliability. In particular, in simulations of forward-looking systems with high integration of variable renewable energy, researchers observe that prices for ancillary services are higher than today when variable energy resources are not supplying ancillary services (Seel et al. 2018). Therefore, including variable energy resources as suppliers could potentially lower prices for ancillary services in systems with high integration of variable energy resources (Kahrl et al. 2021).

However, to the best of our understanding, there is a research gap, as a limited number of papers consider the real-time availability of reserves offered by wind power plants in day-ahead markets (Hosseini et al. 2020) and no article has studied how existing U.S. ISOs/RTOs performance targets and settlement schemes affect the participation of variable energy resources with uncertain output in regulation markets. Therefore, in this work, we consider the delivery targets outlined in Section 3.

For reference, the nomenclature used in this section is as follows:

π_P	Penalty for undergeneration of energy
π_{PPA}	Out-of-market payments (e.g., from power purchase agreements that credit electricity production at a price, or a production tax credit)
$E(s)$	Expected value for performance score s
F_g	Probabilistic forecast (cumulative distribution function) of wind power plant electric output
f_g	Probabilistic forecast (probability mass function) of wind power plant electric output
F_{TARGET}	Operator-determined minimum frequency of acceptable precision (% of time intervals during which the plant output deviates within acceptable tolerances from the instructed by the AGC output)
p_{EN}	Energy price
p_{REG}	Regulation price
q_O	Upper operating limit
q_{EN}	Energy award
q_{REG}	Regulation award
s_T	Operator-determined target precision score
T	Operator-determined target precision score

5.1 Quantity

Under assumptions discussed in Spyrou et al. (2022), we show that a wind power plant that aims to achieve ISO/RTO performance targets of Type I/II (see Section 3.1) could submit a specific forecasted percentile as its upper operating limit or its total regulation capacity in both an upward and downward direction. The exact percentile is determined by the performance targets ISOs/RTOs use for disqualifying regulation suppliers. Eq. (3) and (4) show the exact formula for a conservative bound of the upper operating limit for performance targets of Type I and II, respectively. Performance targets of Type I require a minimum frequency of intervals, F_{TARGET}, with acceptable performance during an assessment period. Similarly, performance targets of Type II require a minimum average performance score, s_T, during an assessment period. For all targets reported in Table 1 of Section 3.1, Eq. (3) and (4) yield a percentile lower than the median.

$$\boldsymbol{F_g^{-1}(1 - F_{TARGET})} \tag{3}$$

$$\boldsymbol{F_g^{-1}(1 - s_T)} \tag{4}$$

One of the assumptions underpinning Eq. (3) and (4) is that the assessment period over which the performance metrics are calculated is large enough for the histogram of observed wind power potential capacity to converge to the probabilistic forecast. However, in practice, the assessment period is finite and includes n time intervals. If the wind power plant operator aims to meet the performance target with probability, $PCON$, then they can solve the inequality (5) to obtain the upper operating limit.

Calculation considering length of assessment period	
$\sum_{i=0:x} \binom{n}{i} * F_g(q_O)^i * \left(1 - F_g(q_O)\right)^{n-i} \geq PCON$	(5)
Where: x =	
$floor((1 - F_{TARGET}) * n)$ For Type I target	$floor((1 - s_T) * n)$ For Type II target

Applying Eq. (5) for different choices of n and s_T, we obtain several insights. For example, we obtain the forecasted percentile that can be submitted as q_O for three values of n (8, 96, 720) and $PCON = 0.97$. In Figure 10, we show the percentiles for $s_T = 0.85$ and 0.75, respectively. By applying Eq. (3) and (4) for $s_T = 0.85$ and 0.75, we obtain the two lines shown in Figure 10. In summary, we observe that the longer the assessment period (higher values of n), the higher the percentile, $F_g(q_O)$, thereby the higher the upper operating limit, q_O, the variable resource can offer. For large n, the percentile from Eq. (5) converges to the percentile from Eq. (4). In addition, we observe that the lower the threshold the performance score must exceed ($s_T = 0.75$ vs. 0.85), the higher the upper operating limit the variable resource can offer for the same length of the assessment period. In other words, the longer the performance assessment period, the easier it is for the wind power plant to offer higher capacity for regulating reserves.

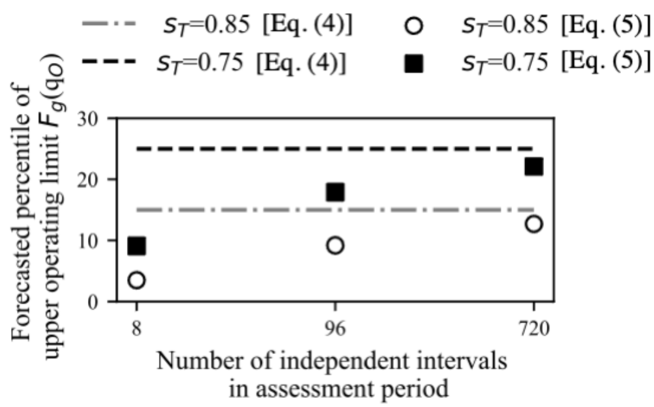

Figure 10. Percentile of the upper operating limit for different performance targets and number of intervals in the assessment period

5.2 Price

We assumed that the wind power plant incurs no additional wear and tear costs to provide regulating reserves. Hence, the wind power plant is facing: lost opportunity costs for any energy revenues it foregoes by offering regulating reserves, and penalties for imperfect delivery of regulation. Lost opportunity costs related to energy revenues associated with a power purchase agreement appear in all schemes (see Section 3.2). However, penalties vary by scheme. For example, in Scheme 1, resources are penalized by rescinding a portion of their regulation award payments. Similarly, in Scheme 2, resources are penalized by rescinding their regulation award payments when the delivery of regulation is below a threshold. In Scheme 4, resources pay over-/undergeneration penalties based on the deviation of their actual output and the requested output (i.e., the basepoint signal of the AGC system).

Table 3. Credits and Penalties Under Various Settlement Schemes

Scheme	Resources With Regulation Awards
1	Receive credits proportional to a score that depends on availability or performance
2	Receive credits only if performance exceeds a threshold
3	Receive credits proportional to a score when the performance score exceeds a threshold
4	Incurs over-/undergeneration penalties with stricter tolerances than energy-only resources

Both types of costs could be reflected as terms in the formula for the offer price of regulating reserves (second column of Table 4). There is a term for foregone energy revenues in all schemes. However, the term for penalties appears only in Scheme 4 because in Schemes 1–3, the penalties to suppliers are not expected to exceed the revenues for regulating reserve awards.

Table 4. Lower Bounds for Offer Prices of Energy and Regulation Under the Four Settlement Schemes. (The first and second columns provide the offer prices for energy and regulation without considering that the variable energy resource might be scheduled to provide energy and regulation awards equal to the upper operating limit. The third column provides the offer price for regulation considering that the total awards from the resource might equal its upper operating limit (Spyrou et al. 2022).

Scheme	c_{EN}	c_{REG}	c'_{REG}
1	$\pi_P * F_g(q_{EN}) - \pi_{PPA} * (1 - F_g(q_{REG} + q_{EN}))$	$\frac{\pi_{PPA} * F_g(q_{REG} + q_{EN})}{E^*(s)}$[9] $- \frac{\pi_{PPA} * F_g(q_{REG})}{E^*(s)}$	$\frac{p_{EN} - c_{EN}}{E^*(s)} + c_{REG}$
2	$\pi_P * F_g(q_{EN}) - \pi_{PPA} * (1 - F_g(q_{REG} + q_{EN}))$	$\frac{\pi_{PPA} * F_g(q_{REG} + q_{EN})}{1 - F_g(s^{-1}(T))} - \frac{\pi_{PPA} * F_g(q_{REG})}{1 - F_g(s^{-1}(T))}$	$\frac{p_{EN} - c_{EN}}{1 - F_g(s^{-1}(T))} + c_{REG}$
3	$\pi_P * F_g(q_{EN}) - \pi_{PPA} * (1 - F_g(q_{REG} + q_{EN}))$	$\frac{\pi_{PPA} * F_g(q_{REG} + q_{EN})}{E^*(s \mid s \geq T)} - \frac{\pi_{PPA} * F_g(q_{REG})}{E^*(s \mid s \geq T)}$	$\frac{p_{EN} - c_{EN}}{E^*(s \mid s \geq T)} + c_{REG}$
4	$\pi_P * F_g(q_{EN} + q_{REG}) - \pi_{PPA} * (1 - F_g(q_{REG} + q_{EN}))$	$\pi_P * F_g(q_{EN} + q_{REG}) + \pi_{PPA} * (F_g(q_{REG} + q_{EN}) - F_g(q_{REG}))$	$p_{EN} - c_{EN} + c_{REG} + \frac{\pi_P}{q_{REG}} * (\varepsilon * q_O * F_g((1 - \varepsilon) * q_O) + \int_{(1-\varepsilon)*q_O}^{q_O} (q_O - g) * f_g(g)dg)$

In the third column of Table 4 we show the regulation price that the wind power plant would be willing to accept considering that curtailment leads to foregoing revenues in the energy market. Co-optimization of energy and regulating reserves in the market would yield a regulation price that is equal to $p_{EN} - c_{EN} + c_{REG}$. This price would be equal to the one yielded by the third column of Table 4 if the delivery of regulation is perfect (i.e., $E(s) = 1 \; or \; q \geq q_O$). However, with imperfect delivery of regulation, the opportunity cost that the wind power plant faces is higher than the one endogenously considered in the co-optimization because when delivery is imperfect, the wind power plant returns at least part of the regulation award payment (Schemes 1–3) and the wind power plant has to pay under- or overgeneration penalties (Scheme 4).

Wind power plants that prefer to offer regulation considering ahead of time the impact of imperfect delivery on settlements have at least two options. First, they could forecast energy prices and embed their opportunity costs in the regulation offer price. Second, they could offer regulation capacity they can deliver almost perfectly, thereby ensuring that $p_{EN} - c_{EN} + c_{REG}$ would adequately reflect their opportunity costs. Additional analytical formulations for this section can be found in Spyrou et al. (2022).

5.3 Case Studies

We now provide numerical examples to illustrate how the equations derived in Sections 5.1 and 0 could inform the supply of regulating reserves by wind power plants to discuss the economic attractiveness of providing regulation and new capabilities enabled by co-locating wind with storage. In Section 5.3.1, we compare the delivery of regulation when a plant offers regulation based on Section 5.1 vs. an approach similar to offering a fixed percentage of nameplate capacity (Rebello et al. 2019). In Section 5.3.2, we estimate how often it would be profitable for a wind power plant to provide regulating reserves in ERCOT over different years. Last, in Section 5.3.3, we present a general case where a wind power plant is concerned about the delivery of capacity awarded in the day-ahead market; this is why a battery is co-located with the wind power plant.

5.3.1 Case Study on Delivery of Regulation Capacity

In this case study, we analyze the performance of an offering approach following equations from Section 5.1 and compare it with another offering approach that consistently offers a fixed percentage (e.g., 10%) of the nameplate capacity for regulating reserves (Rebello et al. 2019).

The hypothetical wind power plant is an aggregate representation of ERCOT's Northern zone. The 20th and 50th percentile of the forecast is publicly available (ERCOT Market Information System undated) and we assume that the forecast follows a normal distribution to estimate other percentiles. We call the approach following equations of Section 5.1 "performance-driven." This approach offers the lower of the 9th percentile and 10% of the installed capacity. Applying Eq. (5), we are 96% and 97% confident ($PCON$), respectively, that offering the 9th percentile helps us meet either of two performance targets: F_{TARGET} 85% over a month[4] or $s_T = 75\%$ over 100 hours[5] (see Figure 10). We call the reference offering approaches "5%" and "10%" because they offer the lower of the 50th percentile and 5% or 10% of the installed capacity, respectively.

According to results reported in Table 5, for F_{TARGET} 85% over a month, the performance-driven approach never disqualifies, whereas the 10% approach would disqualify in 8 months and the 5% approach only in 1 month. Similarly, when the performance target requires an s_T of 75% over 100 hours, only the 10% approach could potentially lead to disqualification of the wind power plant as a regulation provider (see Table 6).

Table 5. Monthly Frequency of Underperformance (Metric Type I) in 2019 for Three Offering approaches: "performance-driven", proposed here, 10%, and 5%. (Given that the performance target, F_{TARGET}, is 0.85 in our example, we highlight in red all months when the performance target is less than 0.85 [i.e., the frequency of underperformance is greater than 15% of hours with positive regulation awards]).

Month	Performance Driven	10%	5%	Month	Performance Driven	10%	5%
January	0.04	0.21	0.13	Jul	0.02	0.18	0.09
February	0.04	0.15	0.08	Aug	0.05	0.22	0.14
March	0.03	0.16	0.11	Sep	0.02	0.18	0.11
April	0.04	0.12	0.06	Oct	0.01	0.14	0.08
May	0.02	0.15	0.09	Nov	0.05	0.22	0.14
June	0.11	0.24	0.16	Dec	0.04	0.17	0.13

[4] Assuming that a month includes at least 60 independent hours.
[5] Assuming that at least eight independent observations are included in 100 hours.

Table 6. Number of Hours in 2019 With Scores Calculated Using the Most Recent 100 Hours With Positive Offers (Metric Type II) Under Three Offering Approaches. (Considering that the performance target, s_T, is 0.75, we highlight in red the hours with rolling scores less than 0.75).

	Performance-Driven	10%	5%
(0,0.745]	0	151	0
(0.745,0.8]	0	523	251
(0.8,0.9]	200	2,734	1,442
(0.9,1]	8,437	5,234	6,949

Besides the performance with respect to F_{TARGET} or s_T, it is worth noting the average capacity each offering approach yields. Overall, the 10% approach appears to be offering the highest capacity and the 5% approach the lowest capacity (see Table 7).

Table 7. Average Capacity Under Three Regulation Offering Approaches

Approach Name	Average Regulation Capacity Offer in MW
Performance-driven	109
10%	138
5%	72

In conclusion, the performance-driven approach clearly dominates the 5% approach because it offers more regulation capacity and disqualifies in fewer months. However, when the performance-driven approach is compared to the 10% approach, there is a trade-off between achieving the performance targets and offering higher amounts of regulation capacity.

5.3.2 Case Study on Economic Supply of Regulation Capacity

In this example, we study the same hypothetical wind power plant as in Section 5.3.1. We further assume that the wind power plant is a price-taker (i.e., the supply of reserves by the hypothetical wind power plant does not impact the prices observed in the ERCOT market for regulation). ERCOT procures regulation separately in the upward (increase generation/reduce load) and downward (reduce generation/increase load) directions. Hence, the historical prices are used as input in a stand-alone optimization model for the hypothetical wind power plant. The optimization model maximizes the bid surplus of the wind power plant (second column of Table 8), which is calculated for Scheme 4 (see Table 8) with $\varepsilon = 0,\ \pi_P = \max\left(\frac{\$20}{MWh}, p_{EN}\right)$ and $depl = 0.25$. We chose $\pi_P = \max\left(\frac{\$20}{MWh}, p_{EN}\right)$ to align with ERCOT's over- and undergeneration penalties (ERCOT Market Education 2020). To estimate the impact of regulation capacity on energy, we follow a common assumption, explained in detail in Kahrl et al. (2021), wherein a resource provides energy that is on average equal to 25% of its regulation capacity award. The bid surplus reflects the net profit the wind power plant would yield considering only payments from the organized wholesale markets.

Table 8. Formulas for Offer Price, Marginal Bid Surplus, and Settlements Considered in the Numerical Example

	Offered Price	Marginal Bid Surplus	Settlements
Energy	$c_{EN} = \pi_P * F_g(q_{EN} + q_{REG}) - \pi_{PPA} * (1 - F_g(q_{REG} + q_{EN}))$	$p_{EN} - c_{EN}$	$(p_{EN} + \pi_{PPA}) \cdot \max(0, q_{O,ACTUAL} - q_{REGUP})$
Regulation down	$c_{REGDN} = (p_{EN} + \pi_{PPA}) \cdot depl \cdot \int_{q_{REGDN}}^{\infty} f_g(g)dg$	$p_{EN} - c_{EN} + p_{REGDN} - c_{REGDN}$	$p_{REGDN} \cdot q_{REGDN} + \pi_P \cdot \max(0, q_{REGDN} - q_{O,ACTUAL}) - (p_{EN} + \pi_{PPA}) \cdot depl \cdot \min(q_{O,ACTUAL}, q_{REGDN})$
Regulation up	$c_{REGUP} = \pi_P * F_g(q_{EN} + q_{REG}) - (p_{EN} + \pi_{PPA}) \cdot depl \cdot \int_{q_{REGUP}}^{\infty} f_g(g)dg$	$p_{REGUP} - c_{REGUP}$	$p_{REGUP} \cdot q_{REGUP} + \pi_P \cdot \max(0, q_{REGUP} - q_{O,ACTUAL}) + (p_{EN} + \pi_{PPA}) \cdot depl \cdot \min(q_{O,ACTUAL}, q_{REGUP})$

Finding 1: For a subset of hours, supply of regulation capacity is more economically attractive than supply of energy. The number of hours over which supply of regulation capacity by wind power plants is economic and the incremental value a hypothetical wind power plant captures varies from year to year.

As Table 9 shows, the wind power plant would find the supply of "regulation up" capacity economic for a few hundred to a thousand hours in a year and the supply of "regulation down" capacity economic for 20%–50% of the year. The reason for this result is that when the wind power plant provides regulation down capacity, it does not need to precurtail energy compared to its wind resource capacity potential and it can supply energy as usual. On the contrary, when the wind power plant provides regulation up capacity, it cannot supply energy up to its full potential because the capacity is reserved for regulation up deployment. For example, if a 10-MW wind power plant provides capacity for regulation up or down 10 MW, then it could provide 0 or 10 megawatt-hours (MWh) of energy, respectively, in the market. On average, the plant would expect to increase or reduce its energy by 2.5 MWh (25%*10 MWh) when regulation up or down is respectively deployed. Therefore, the plant would provide, on average, 2.5 MWh or 7.5 MWh (10-2.5) of energy, respectively; and the plant would forego 25% or 75% of the revenue it could yield by offering the regulation down or up capacity as energy, respectively. The incremental value that the wind power plant captures by offering regulating reserves in addition to energy also varies throughout the years, increasing the revenues by ~2%-9% compared to a case that the wind power plant only provides energy.

Table 9. Monetary Metrics and Number of Hours With Regulation Awards Over Multiple Years Using HB_NORTH prices. (We report results for two cases: when the plant provides both energy and regulation [left of "/"] and when the plant provides only energy [right of "/"]).

	Total (Energy + Regulation) Settlements ($ million)	Regulation Settlements ($ million)	Number of Hours With Regulation - Up Awards	Number of Hours With Regulation - Down Awards	Hours Analyzed	Share of Hours Where Available Capacity in Real Time Lower Than Offered in Day Ahead	Actual Output Over Analysis Period (gigawatt-hours [GWh])
September–December 2017	22.54/21.64	1.41/0	109/0	915/0	2,928	13%	1,119
2018	102.68/100.45	3.89/0	458/0	1,775/0	8,758	13%	4,217
2019	108.49/99.41	13.06/0	901/0	3,731/0	8,742 (17 hours quantile crossing)	11%	4,444
2020	85.50/80.14	8.58/0	893/0	5,344/0	8,743 (40 hours quantile crossing)	6%	4,420
2021	386.88/291.51	131.95/0	1,174/0	4,687/0	8,607 (152 hours quantile crossing)	2%	5,619
2021 excl. February	158.691/150.49	12.42/0	973/0	4,471/0	8,070	2%	5,280
January–August 2022	282.29/276.10	11.17/0	355/0	1,867/0	5,674 (156 quantile crossing)	1.5%	5,613

Finding 2: The incremental value that a wind power plant could capture by providing regulating reserves varies not only per year, but per location and plant as well.

As results for three hypothetical wind power plants from the Advanced Research Projects Agency – Energy (ARPA-E) Performance-based Energy Resource Feedback, Optimization, and Risk Management (PERFORM) data set (ARPA-E PERFORM Forecast Data undated) in Table 10 show, the incremental value a wind power plant could capture varies even in the same year. It appears that the incremental value is higher for ERCOT's West region. However, it is worth noting that ERCOT procures reserves at the system level and energy at a nodal level. Prices at ERCOT's West hub are on average lower than the systemwide energy prices, suggesting that generators in ERCOT's western region cannot deliver energy everywhere in the system because of transmission congestion. Also, reserves are procured without accounting for transmission constraints and as a result the reserve prices do not reflect any potential congestion issues. Hence, while the upside from providing regulation might seem higher in ERCOT's West region, it is possible that regulation provided by a wind power plant in this region might not be deliverable everywhere in the system at least for certain hours of the year. Therefore, the highest incremental value in ERCOT's West region could be explained by the asymmetry in the formation of energy and reserve prices and it does not necessarily convey that the system finds more valuable provision of regulation in ERCOT's West region.

Table 10. Monetary Metrics and Number of Hours with Regulation Awards for Multiple Hypothetical Wind Power Plants in 2018. (We report results for two cases: when the plant provides both energy and regulation [left of "/"] and when the plant provides only energy [right of "/"]).

Hypothetical Wind Power Plant	Pricing Node	Total (Energy Plus Regulation) Settlements ($ million)	Regulation Settlements ($ million)	Number of Hours With Regulation-Up Awards	Number of Hours With Regulation-Down Awards	Actual Output Over Analysis Period (GWh)
Cedro Hill	HB_NORTH	12.40/11.77	0.99	563	1,752	488
Trent Mesa	HB_SOUTH	13.65/13.12	0.87	593	1,707	571
McAdoo	HB_WEST	11.28/10.51	1.17	967	1,737	548

Finding 3: Any out-of-market energy revenues the wind power plant receives affect the attractiveness to provide regulation.

We studied the economic attractiveness of the regulation market ignoring any out-of-market energy payments. In practice, though, out-of-market energy payments wind power plants receive, such as production tax credits or energy payments determined by power purchase agreements, can be significant. That is why we further analyzed the supply of regulating reserves by wind power plants for 1 year with different levels of out-of-market energy prices. The results clearly show that out-of-market energy payments greatly influence the supply of regulating reserves by wind power plants. For example, the number of hours with regulation-up (Table 11, column 5) and regulation-down awards (Table 11, column 6) drastically decreases from 4,632 (3,731 + 901) hours to 2,087 (1,713 + 374) hours when the out-of-market energy price (π_{PPA}) increases from $0/MWh to $20/MWh.

Table 11. Monetary Metrics and Number of Hours With Regulation Awards for Different Values of π_{PPA} in 2019

π_{PPA} ($/MWh)	Total (Energy Plus Regulation) Settlements ($ million)	Power Purchase Agreement Revenue ($ million)	Regulation Settlements ($ million)	Number of Hours With Regulation-Up Awards	Number of Hours With Regulation-Down Awards
0	108.49	0	13.06	901	3731
20	107.68	87.19	10.25	374	1713
30	107.09	131.33	9.2	304	1117
40	106.66	175.11	8.5	253	842
50	106.23	219.2	7.87	213	633

5.3.3 Case Study on Delivery of Capacity by Wind Plus Storage

In the case studies, we focused on the delivery of regulation capacity, and observed how a percentile lower than the median must be offered to meet performance requirements related to frequency of underdelivery, F_{TARGET}, or average precision, s_T. In our analysis, we did not discuss the energy wind power plants could provide and we have assumed that wind power plants could supply—in the real-time market—energy in excess of the low percentiles or supply energy through day-ahead virtual bids that reflect their expected balancing costs (Dent et al. 2011).

In this section, we analyze the delivery of energy and discuss how adding storage could affect the firmness and flexibility of a wind power plant. First, we analyze the impact of storage on firmness by identifying battery sizes that could limit the percentage of hours a plant underdelivers (i.e., provides less energy in real time than promised in day ahead). Second, we conducted a sensitivity analysis with a larger battery system in which we varied the energy capacity of the storage system considered in day-ahead scheduling. For this work, we use data for a hypothetical wind power plant included in the ARPA-E PERFORM data set (ARPA-E PERFORM Forecast Data undated). The plant's maximum and average actual output is 135 MW and 56 MW, respectively.

Reliability or Firmness

In this example, we assumed that each hour the wind plant schedules according to a specific percentile of the day-ahead schedule and aims to underdeliver less than 10% of the time. In practice, the plant might have an economic incentive to change its schedule in real time compared to day ahead but for local reliability issues, it might be beneficial for the system operator to know that the plant can be firm on their day-ahead promises with a target probability (here 90%).

The results in Table 12 suggest that even a small amount of storage (2 MW/2 hours) enables the wind power plant to supply ~10% more capacity in day ahead. In practice, such a small battery might not be needed given the tolerance for deviations. While a small battery is needed to schedule capacity at a lower than the median level, a larger battery is necessary for the plant to schedule in day-ahead capacity equal to the median.

Table 12. For Three Cases of Battery Sizes, Results Show the Percentile That Can be Scheduled in Day Ahead To Achieve 10% Frequency of Underdelivery. (All results are for a hypothetical wind power plant in ERCOT's northern region in 2018).

	No Battery	2 MW/2 hours	45 MW/2 hours
Average day-ahead scheduled capacity (MW)	44 ("second percentile")	48.5 ("ninth percentile")	56 (median)
Frequency of underdelivery (% of hours)	10%	10%	10%

Shifting Energy

For economic reasons, firming up the day-ahead energy schedule is valuable when the system is experiencing high balancing costs. Similarly, the flexibility that storage can add in terms of shifting the on-site wind production to hours that the electricity generation would be more valuable for the system or to hours that the transmission system is not congested is likely to be of interest to plant operators as well. Using 2018 historical prices, we first optimized the day-ahead energy schedule of a wind-plus-battery system based on day-ahead energy prices. Then, using the day-ahead energy schedules as inputs for the real-time problem, we optimized the actual schedule of the wind-plus-battery system assuming that the uncertainty with respect to wind has resolved and the second problem's objective is to minimize the average shortfall (shortage of energy in real time compared to the day-ahead energy schedule). We summarized results from a sensitivity analysis we conducted for a 65-MW, 4-hour battery system. While all scenarios use the entire battery capacity in real time, they differ in the energy capacity considered in day ahead.

Table 13. Results on the Total Day-Ahead Energy Revenue and Average Shortfall Over 8,753 Hours for the Five Cases of Battery Capacity Considered in Day-Ahead Scheduling[6]

Storage System Considered for Day-Ahead Optimization	Day-Ahead Energy Revenue Using Prices for HB_NORTH (million $)	Average Shortfall in MW
65 MW/4 hours	17.5	9.6
65 MW/3 hours	16.97	7.82
65 MW/2 hours	15.9,886	5.83
65 MW/1 hours	14.5	3.52
No battery	12.7	1.50

The results in Table 13 show that the higher the amount of storage considered in the day-ahead optimization the higher the day-ahead revenue of the plant. For example, the plant increases its day-ahead energy revenue from $12.7 million to $17.5 million when the plant adds a 65-MW, 4-hour battery to its day-ahead energy scheduling. In other words, the day-ahead optimization uses the battery to shift energy to more valuable hours—as determined by day-ahead prices. However, the more battery capacity is considered for flexibility in day-ahead scheduling, the less battery capacity is available for managing the wind forecast errors in real time. Thereby, the more flexible the plant appears in day ahead, the less firm it appears in real time as the higher average energy shortfall (shortage of energy in real time compared to the day-ahead energy schedule), as indicated in the top rows of Table 13. In practical applications, wind-plus-storage systems would choose between day-ahead flexibility and real-time firmness by using, in addition to wind power forecasts, other inputs such as price forecasts that help them quantify the real-time balancing costs and the arbitrage opportunities in the day-ahead and real-time markets.

[6] The simulated hours are 8,753 instead of 8,760 because we studied the chronological year 2018 at local time and the forecast data are provided for year 2018 UTC time.

6 Aerodynamic Control

A closed-loop power tracking controller is needed to enable wind power plants to respond to active generation control signals from the grid operator in real time. Therefore, we include a plant-level APC in the A2e2g platform. Aerodynamic interactions between wind turbines, such as turbine-to-turbine waking, are an important consideration for plant-level power control. These interactions happen on the timescales of seconds to minutes as the wind propagates through the plant, coinciding with the timescales of important grid services discussed in the Section 5, such as frequency regulation. These timescales are the focus of the A2e2g platform, and the plant-level power controller is designed for second-level operation and longer.

The controller takes the current power set point being sent by the grid operator as an input and produces individual power reference signals for the wind turbines within the farm. It is assumed that individual turbines have turbine-level APCs enabled, so we include a simple closed-loop turbine model in the platform. Turbine-level APC uses a combination of generator torque and blade pitch actuation to control the power output of the turbine to a reference level. As a point of comparison, we also include a farm-level power-maximizing controller (PMC), which mimics current industry operation of wind farms of having each wind turbine produce as much power as possible using predominantly generator torque control.

We also include a battery operating logic that can work with either the APC or PMC to aid in power reference tracking. The inclusion of battery operating logic reflects recent growth in U.S. wind power plants with collocated battery storage, especially for the provision of ancillary services rather than long-term energy arbitrage (Wiser et al. 2021). We also describe a medium-fidelity simulation environment that can be used to test the controllers presented, as well as newly developed controllers.

Finally, a reinforcement learning-based alternative to the main wind plant control paradigm is provided in Appendix A.2.

6.1 Wind-Plant-Level Aerodynamic Power Control

To provide active power services, a wind plant must dynamically alter its power output to track a reference provided by the system operator. Within the A2e2g platform, we provide a closed-loop wind plant active power controller to achieve this behavior. The controller is based on the work of van Wingerden et al. (2017) and uses a plant-level controller that distributes power references to the individual wind turbines (Figure 11). For A2e2g, we assume that the individual turbines have their own closed-loop active power controllers enabled (described further in Section 6.4).

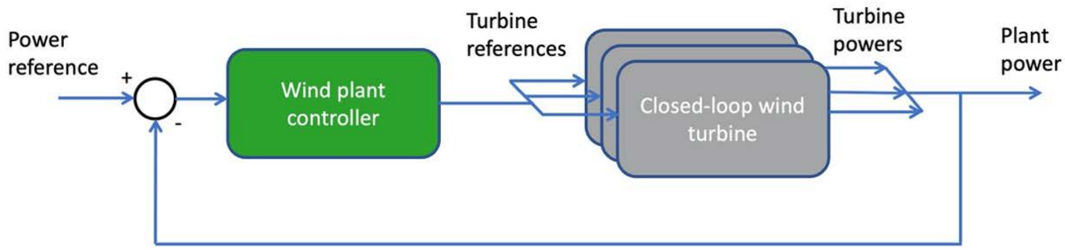

Figure 11. Wind plant power controller structure

As depicted in Figure 11, the wind plant power controller (WPPC) operates on the error between the power reference, $P_{\text{reference}}$, set by the system operator depending on the type of ancillary service being provided, and the current power output of the wind plant, P_{plant} (which is the sum of the individual turbine powers). That is:

$$P_{\text{error}} = P_{\text{reference}} - P_{\text{plant}} \tag{6}$$

The WPPC provides a power reference to each turbine in the wind plant for individual active power control. The individual turbine power references produced by the WPPC are increments on the current power being produced by each turbine; that is, the WPPC generates a single power increment to be applied to all turbines, based on the plant-level power error, and then adds that increment to the current power being produced by each turbine to generate the individual reference at each turbine. By doing so, all wind turbines are asked to increase or decrease by the same absolute amount, relative to their current power production, to meet the plant-level reference.

The need for such a plant-level closed-loop controller is to account for differences in the wind resource available to each turbine. If each turbine is identical and experiences the same free-stream wind input, then tracking a plant-level power reference can be achieved simply by dividing the reference evenly between the turbines in the plant and providing each with the same power reference. This approach is shown in Figure 12, and we include it in the A2e2g platform as a point of reference. In practice, wind turbines are not perfectly identical and the wind resource at each turbine is not the same, particularly due to wake interactions between the turbines (Fleming et al. 2016). Turbine-to-turbine waking means that downstream turbines experience lower wind speeds and cannot achieve the same range of power that upstream (free stream) turbines can. See Figure 13 for a graphic representation of turbine-to-turbine waking.

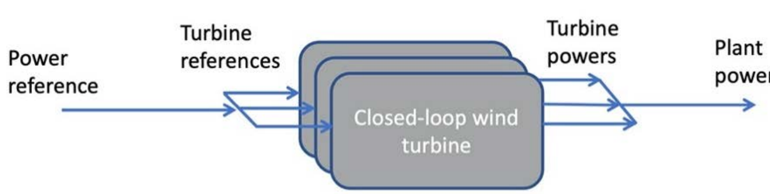

Figure 12. Individual, uncoordinated active power control

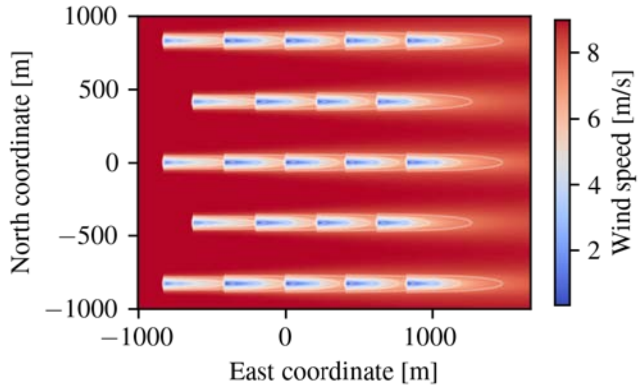

Figure 13. Wind-turbine-to-turbine waking in a wind power plant with a gridded layout under aligned, westerly wind conditions. Figure produced using FLORIS. Westernmost turbines experience free-stream, high-speed winds, whereas downstream turbines operate in the wakes of upstream turbines.

The closed-loop nature of the WPPC means that when a wind turbine can no longer produce more power due to limited wind resource (referred to as "saturation" (van Wingerden et al. 2017), the power deficit is essentially transferred to other turbines by the WPPC. Saturation occurs frequently because turbines are often operating in the wakes of other turbines, meaning that they experience lower wind speeds. Provided that at least one turbine remains unsaturated, the plant will be able to achieve the reference power set point in steady state (once transient effects have died out). If all wind turbines are saturated, the plant is producing the maximum power available given the wind resource, and it will no longer be able to follow a power set point. However, as soon as the resource is increased (or the set point lowered to an attainable value), the WPPC will again begin tracking the reference. For full details of the WPPC as well as a mathematical analysis of its closed-loop performance, refer to Sinner et al. (Sinner et al. 2022).

We also included a wake steering add-on to the WPPC in A2e2g. Wake steering has been shown to increase the power output from wind plants in situations where there are significant wake losses (Fleming et al. 2017). For A2e2g, we find that activating wake steering when all turbines are saturated, and the power set point has not been reached can slightly increase the power production of the wind power plant.

6.2 Power-Maximizing Control

The standard approach for controlling operational wind turbines is to generate as much power as possible within the structural loading limits of the turbine. When wind speeds are below rated, this means extracting as much power as possible from the wind; for higher wind speeds, the rated power is extracted to avoid overspeeds and shutdowns. Such operation does not consider a power reference set point and is generally unsuitable for active power control. Nevertheless, as this is still the industry-standard operation, we include this PMC in the A2e2g platform as a base case; see Figure 14. While PMC alone is not a viable candidate, it can be coupled with energy storage to provide active power services. This is discussed in Section 6.3.

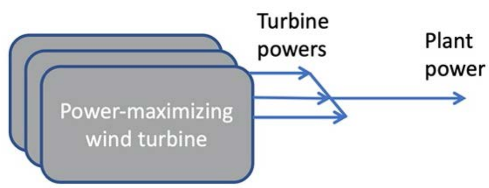

Figure 14. Standard, power-maximizing wind power plant operation

6.3 Battery Operation for Controller Support

As mentioned earlier, PMC has no way of tracking a power reference set point provided by the system operator and is thus not a viable choice for providing active power services from wind power plants. To increase the flexibility of a wind plant using PMC and provide a range of grid services, we include collocated energy storage in A2e2g in the form of a battery. The battery can be used to shift energy and respond to reference set points within its capacity. The combined PMC with battery support is shown in Figure 15.

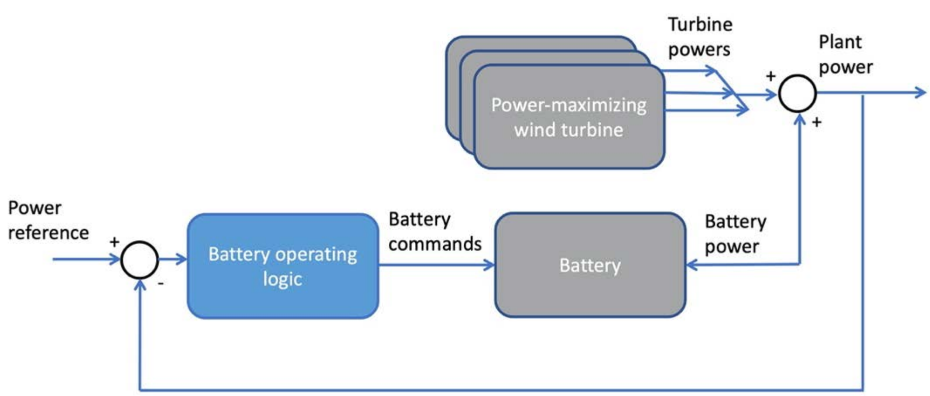

Figure 15. Power-maximizing wind turbine operation with battery support for power reference tracking

The battery charges when the power being produced by the wind turbines exceeds the power reference set point and discharges when the turbine power is lower than the power reference. Charging can happen provided that the battery has capacity, both in terms of its power charging/discharging limit and its total energy storage capacity. In the A2e2g platform, the battery can only be charged with the power being generated by the wind turbines (and not from the grid). We understand the approach of having a battery supporting a power-maximizing wind plant to be that used in the field tests of Watson et al. (2018) and Rebello et al. (2021). The battery's ability to provide support to the plant depends strongly on its power and energy capacities, which is investigated in Sinner et al. (2022).

Wind-plant-level active power control, as described in Section 6.1, generally requires dynamic curtailment of the wind turbines to meet a given power set point. Such curtailment comes at the opportunity cost of the energy that could have been produced. A battery providing short-term energy shifting can both help relieve curtailment and provide a faster active power response. As in the PMC case, the battery is operated to charge when the power produced by the wind turbines is higher than the set point and discharge when the turbines fall short of the power reference. However, using the control system shown in Figure 11, the turbines are rarely producing power more than the reference because they are dynamically curtailed to meet the reference; as a result, the battery has no opportunity to charge.

To allow the wind turbines to produce excess power under WPPC and store that energy in the battery, we augment the power error equation for the WPPC to:

$$P_{\text{error}} = P_{\text{reference}} + P_{\text{available,battery}} - P_{\text{plant}} \tag{7}$$

Therefore, if there is extra storage space in the battery, the reference sent to the WPPC is increased, allowing the turbines to operate at or near their power-maximizing mode when there is sufficient storage available. The available battery power, $P_{\text{available,battery}}$, is limited by both the power capacity of the battery and its energy capacity. If the battery has sufficient storage available and has a high power capacity, the turbines are operated in power-maximizing mode because the reference, $P_{\text{reference}} + P_{\text{available,battery}}$, is higher than the total power available in the wind; on the other hand, when the battery is fully charged, $P_{\text{available,battery}}$ is set to zero and the WPPC resumes its normal curtailing operation. The components of the coordinated WPPC and battery are shown in Figure 16.

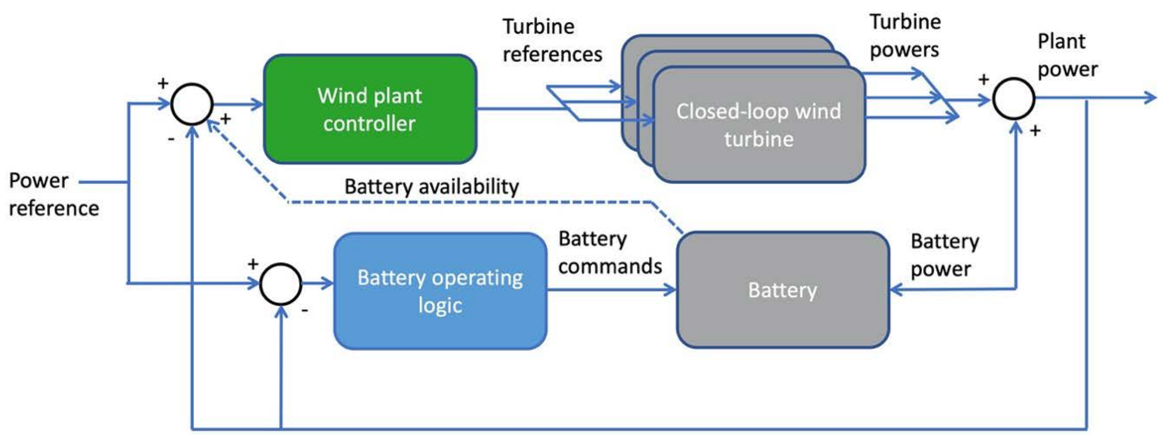

Figure 16. Wind plant power controller with battery support

6.4 Aerodynamic Simulations for Controller Validation

To simulate the behavior of wind power plant controllers in response to power reference set points, we provided a dynamic wind plant simulation code along with the A2e2g platform. The code is based on FLORIS (NREL 2020), which models the wake interactions between wind turbines and assesses the available power. FLORIS also models the effect of yaw misalignments in terms of the effect on both the individual turbine powers and the wake shapes.

Three important changes are made to the standard FLORIS (Version 2.4) package to provide a suitable simulation environment for a dynamic, closed-loop wind power plant. First, dynamic wake propagations are introduced to model the time-marching effect of changes both in the free-stream wind speed and wind direction and in wind turbine set points. Second, FLORIS models turbines using a fixed power and thrust curve, which does not account for curtailment. A modification is made to allow the axial induction factors of the turbines to be set by an external code, which allows us to model the effect of a steady-state curtailment of individual turbines. Finally, dynamic models of the effect of axial induction changes are introduced to ensure a realistic and smooth response to individual wind turbine power reference changes. These modifications to the standard FLORIS package are described in more detail in the following subsections.

6.4.1 Dynamic FLORIS Model

We use the FLORIS package described in Section 4.2 to simulate the wind plant. However, as mentioned previously, FLORIS is a steady-state simulator for time-averaged behavior and is not directly suited to modeling the dynamic response of a wind power plant. In particular, the propagation time of wakes is not modeled in FLORIS—if FLORIS was used directly to model dynamic behavior, wakes would appear to propagate instantaneously throughout the farm. To address this, we use a dynamic modification to FLORIS. We describe the modification here briefly, and refer to Vijayshankar et al. (2021) for more details.

FLORIS calculates the plant power by determining the wind speed (accounting for wakes from upstream wind turbines) at each turbine location within the plant. As a first-order approximation of dynamic wake propagation, the dynamic modification to FLORIS uses Taylor's frozen wake hypothesis to assume that the wake moves at the speed of the free-stream velocity. The wind speed impacting a turbine at a specific time can be found by storing the effects of wakes associated with different wind speeds and directions and applying them at the suitable time that is calculated based on the frozen-wake hypothesis. The buffer accounts for not only the free-stream wind speed, but also the free-stream wind direction and individual turbine yaw angles and axial induction factors. A demonstration of the modified dynamic FLORIS model responding to a wind speed change is shown in Figure 17.

Finally, it should be noted that in A2e2g we only approximate dynamic wind speed effects based on free-stream wind speed and wind direction. Second-order effects due to speed reductions and deflections within turbine wakes are not modeled in the wake propagation—only the free-stream properties are used to calculate the propagation time to keep simulations relatively efficient and easy to run. For further details about the dynamic modification to FLORIS, please refer to Vijayshankar et al. (Vijayshankar et al. 2021).

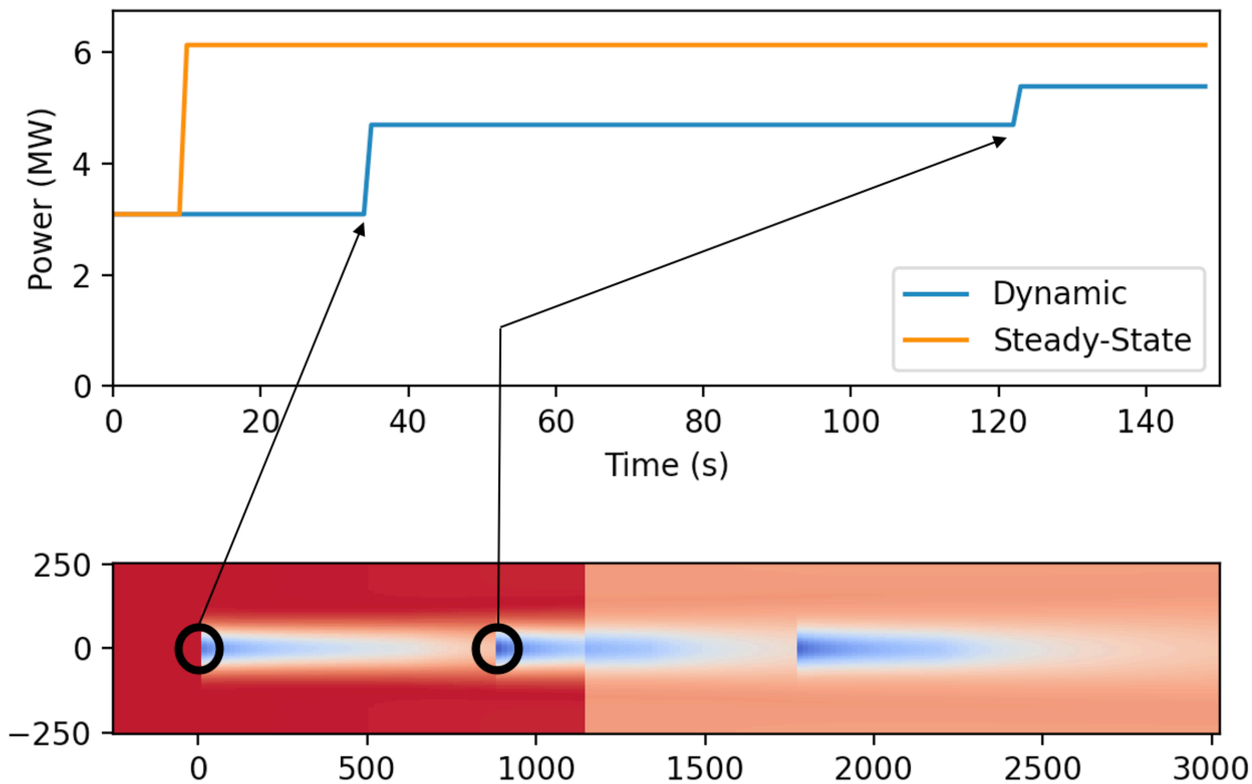

Figure 17. Demonstration of dynamic FLORIS (blue curve in upper plot), in which an increase in the wind speed takes time to impact downstream wind turbines as compared to the standard FLORIS (orange curve in upper plot), which assumes wind speed changes happen instantaneously

6.4.2 Actuator Disk Model Dynamics

FLORIS is designed assuming power-maximizing control, wherein the turbine axial induction factors are automatically chosen to maximize power production during below-rated operation (i.e., set to 0.33) (and reduced to maintain rated power in above-rated operation). To model the reduced axial induction factors necessary for dynamic power output control, we make some changes to the dynamic FLORIS model that allow us to override the automatic power-maximizing axial induction factor and set the turbine axial inductions based on commands from the wind plant controller described in Section 6.1. However, the axial induction factor represents the response of the wind turbine in steady state, rather than its instantaneous response. To model the dynamic response, we include two important factors: the response of the turbine actuators to a change in the power reference, and the settling of the air flow to a new operating axial induction factor.

We model the turbine response by assuming a simple closed-loop model for the wind turbine active power controller. As the power produced is the product of the generator speed and generator electrical torque, altering the turbine power can be achieved either by altering the rotational speed or by changing the generator torque (or a combination of the two). For our model, we assume that the generator torque is held constant, and the blade pitch angle is varied to alter the aerodynamic torque and therefore the rotational speed of the turbine. This allows the turbine controller to alter the power output via the blade pitch actuators. The closed-loop pitch controller is assumed to follow the design of Hansen et al. (2005) for generator speed regulation, resulting in a second-order closed-loop model of the wind turbine controller. The frequency response of this model is depicted by the red line in Figure 18. Because the torque is constant, this frequency response also models the power response of the turbine to a power reference set point, from which a dynamic axial induction factor can be obtained.

In addition to the turbine closed-loop blade pitch response, it also takes time for the wake to respond to a change in the wind turbine aerodynamic properties. We use the first-order model proposed by Knudsen and Bak (2013) to update the axial induction factor used in wake calculations. This model is depicted as the blue line in Figure 18. For simplicity, we use a time constant of 12 seconds across all wind speeds for the first-order flow model.

The power output from each turbine is modeled as the response to only the turbine closed-loop model, whereas the axial induction factor applied to dynamic FLORIS is the combined turbine and flow response (represented by the black line in Figure 18).

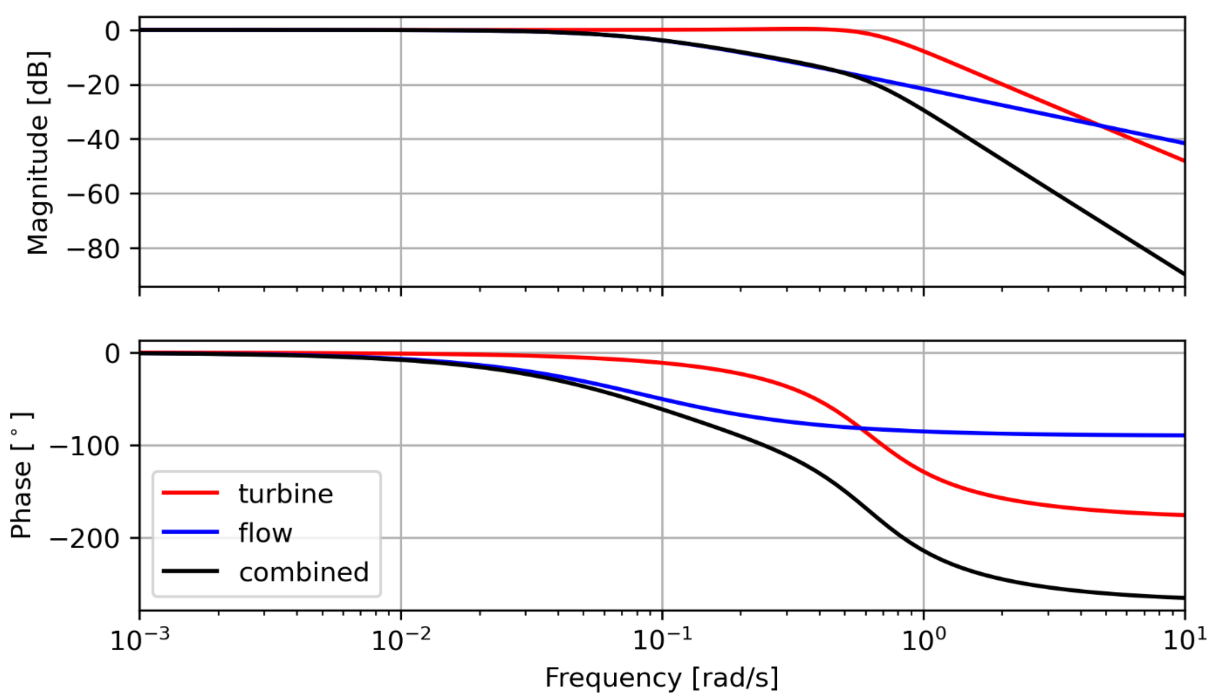

Figure 18. Frequency response of the separate wind turbine and flow models, as well as their combined behavior, used to model dynamic effects in the axial induction factors.

Note: dB = decibels and rad/s = radians per second

6.4.3 Battery Simulator

The battery model we include in the A2e2g platform is a simple state-of-charge model. Energy stored in the battery is the integral of the power supplied to or drawn from the battery, accounting for self-discharge losses (which are small) and charging and discharging inefficiencies. The charging and discharging efficiencies are set at 92%, resulting in a round-trip efficiency of approximately 85%. The battery is modeled as having a fixed charging and discharging capacity throughout its entire range of operation and a fixed total energy capacity. Degradation of the battery is not modeled, as the current simulation platform is not designed for long-term simulations. A2e2g users may specify a reduced-capacity battery as a proxy for some level of degradation. The battery-operating logic only allows the battery to be charged from the wind plant and not from the grid. Further details can be found in Sinner et al. (2022).

6.5 Example Cases

To demonstrate the closed-loop behavior of the wind power plant responding to a signal from the grid operator, we considered two cases: AGC and a possible future behavior of wind plants to commit to their day-ahead bids and produce power at a steady level over each hour, reducing variability in the power supply and thereby reducing the requirement for system operators to procure frequency regulating capacity. We used a fictious 50-MW wind plant with the layout shown in Figure 19, and data from a meteorological mast at Texas Tech University to provide the atmospheric conditions.

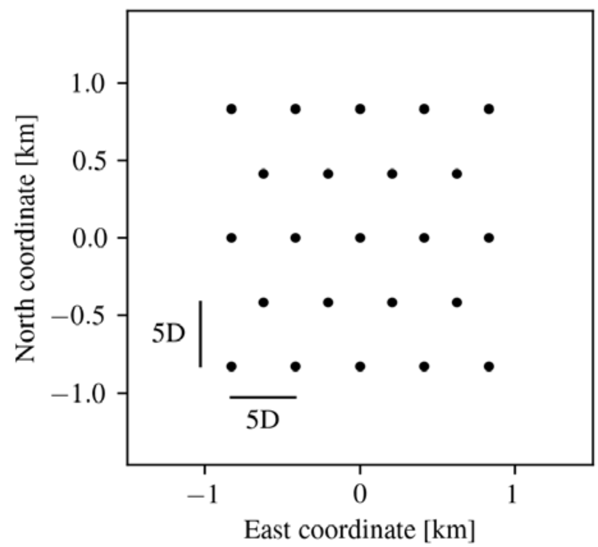

Figure 19. Fictitious power plant layout used for simulations, with turbine-to-turbine spacing of 5 rotor diameters (5D)

6.5.1 Active Generation Control

We first demonstrated the wind plant power controller providing AGC. To do so, we simulated the plant for a 1-hour period (with 15 minutes of simulation warm-up) using data from September 28, 2019. In this interval, the market advisor module (see Section 5) elects to provide AGC services. The wind conditions, power forecasts, AGC signal, and wind power plant response are shown in Figure 20. Our demonstration included four configurations: PMC as described in Section 6.2, WPPC as described in Section 6.1, and both PMC and WPPC including a 16-MWh, 4-MW (i.e., 0.25 C rate) battery with operation described in Section 6.3.

The first plot of Figure 20 shows the day-ahead forecasted mean wind speed along with a one-standard-deviation uncertainty band (shaded region) and the actual wind speed used for simulation. Similar data were obtained for wind direction (not shown here for simplicity). The second plot demonstrates the translation of atmospheric forecasts to plant power capacity forecasts, again showing the mean and one-standard-deviation uncertainty. Also included is the updated 5-minute-ahead forecast (in black). The third plot shows the AGC signal provided to the wind plant after service market participation results in a bid of approximately 4 MW of regulating capacity in black, along with the responses of the four control configurations following the AGC reference signal. The final plot shows the error between the AGC set point and the plant power to clarify differences between the controllers' performance.

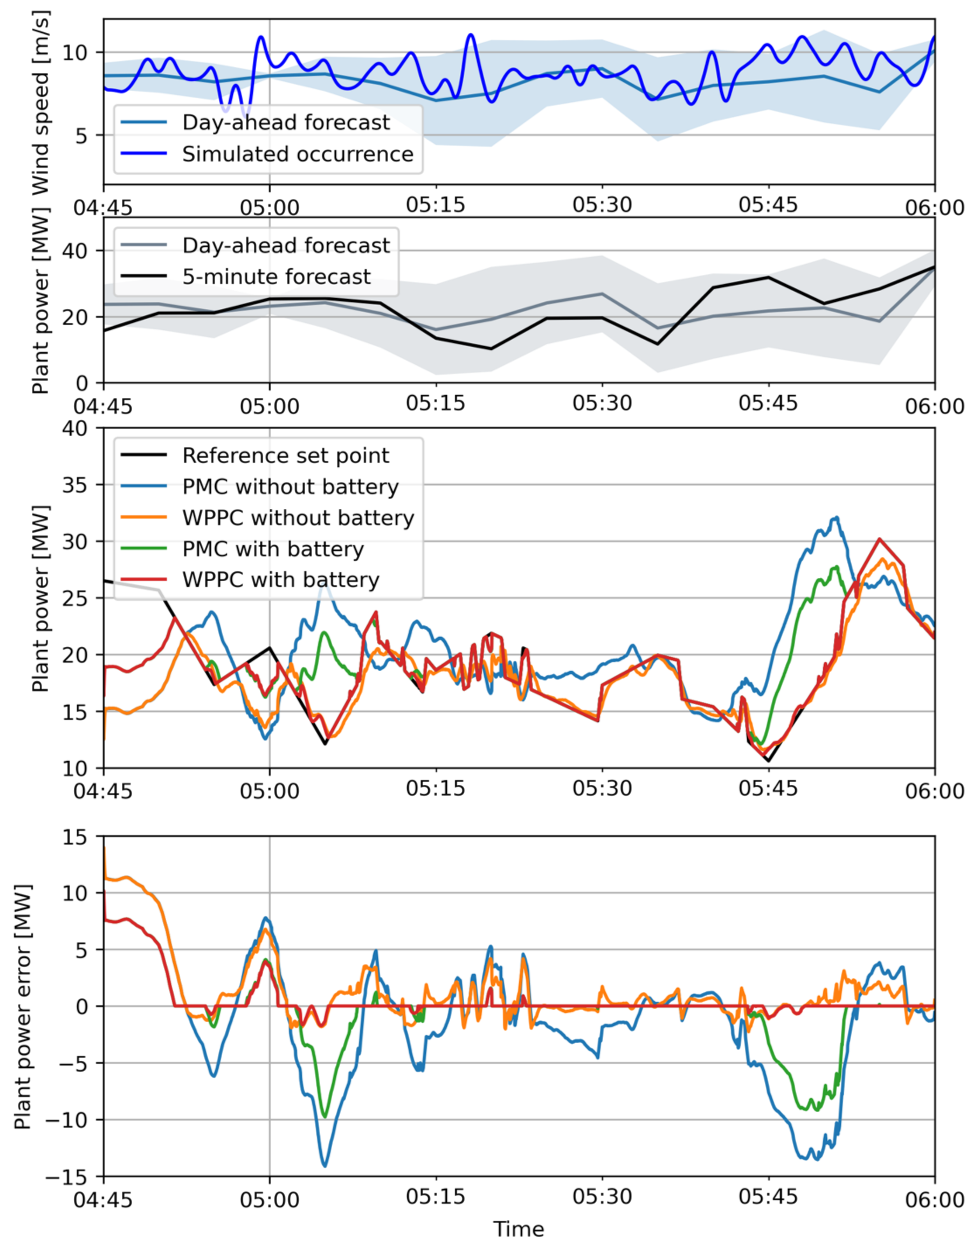

Figure 20. Simulated controller response to an AGC signal provided by the system operator

PMC alone (depicted with a blue line in the last two plots of Figure 20) has no mechanism for adjusting the active power produced by the plant to match the system operator-provided reference signal and performs poorly. On the other hand, the closed-loop WPPC (orange line) responds to changes in the AGC signal by altering the power generated by individual wind turbines while accounting for turbine-to-turbine interactions. This control method is effective for following the lower-frequency (longer duration) content in the AGC signal; however, when the AGC signal changes rapidly, the WPPC is limited by the response time of the individual turbine system. Note that this response time can likely be improved by implementing a more complex individual turbine control logic that incorporates torque control, as opposed to the blade pitch controller used in the current A2e2g platform. Moreover, the AGC signal simulated here is likely to be an extreme case in terms of rapid changes in the reference signal (4 MW changes over 4 seconds, which may be up to 10% of the nameplate capacity of the plant and 50% of the instantaneous capacity). Adding a battery helps both PMC (green line) and WPPC (dark red line). WPPC with a battery added is very effective in following the AGC signal except when there is a significant lack of wind resource, which occurs during the simulation spin-up period from 04:45−04:50. PMC with a battery, while significantly better than PMC alone, still has performance limitations when there is an excess resource available (e.g., 05:05−05:10).

6.5.2 Smooth Hourly Power Provision

As a second use case, we considered a wind power plant's ability to provide steady power generation for an hour based on the hourly mean day-ahead forecast. This use case serves two purposes. First, this type of operation, wherein a wind plant commits ahead of time to produce a fixed amount of power, reduces variability of electricity supply and therefore could reduce the need for system operators to procure other, traditionally fossil-fuel-based, generation for regulation. Second, the provision of power at a steady level followed by a change to a new level demonstrates both the WPPC's step response and its ability to reject disturbances in the form of varying wind conditions. The response of the same four control configurations described in Section 6.5.1 under this "smooth hourly power" operating mode is shown in Figure 21. We included in the simulation a change in the hourly power set point at 04:00 to demonstrate the controllers' step responses.

When there is no battery present, the wind power plant is only able to meet the power reference set point when there is sufficient resource (e.g., 03:20−03:25, 03:55−04:15). During these periods, the WPPC can meet the power reference with errors of approximately 1−2 MW. Adding the small battery reduces the errors to close to zero except when there is enough of a deficit between the reference and the resource available that the battery meets its discharging power limit of 4 MW (e.g., 03:45−03:55). Using a larger power capacity battery (i.e., a battery with a higher C rate) would help to maintain the reference power during these periods. The WPPC alone can respond quickly to the step change in reference around 04:00 and quickly settles to the new set point (within approximately 2 minutes). For a full examination of the WPPC's ability to provide smooth, hourly power, and a study on the performance sensitivity to the battery size, refer to Sinner et al. (2022).

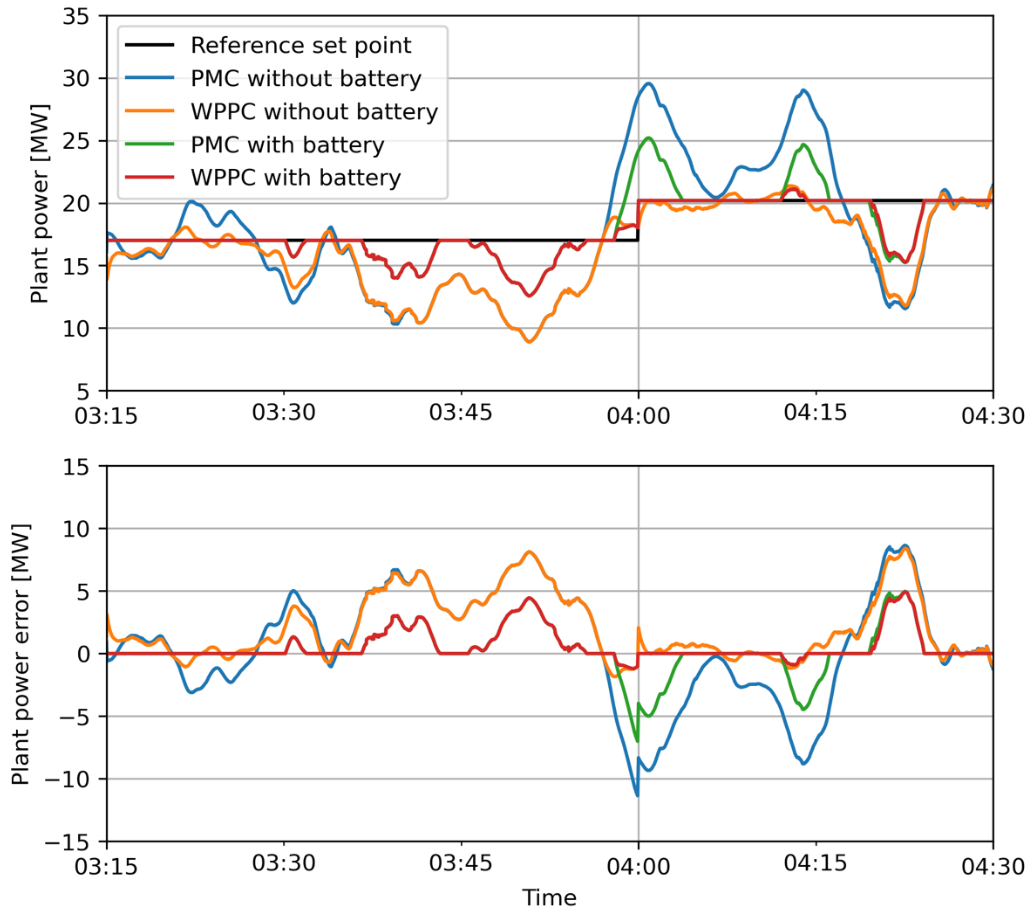

Figure 21. Controlled wind power plant responses to a steady hourly production requirement

6.5.3 Control Adequacy

As the time series results of the example cases demonstrate, the WPPC described in Section 6.1 is suitable for tracking an active power reference signal on the timescale of seconds and longer, which correspond to the timescale of the aerodynamic interactions between wind turbines, provided that the wind resource is sufficient. Moreover, even a relatively small battery can provide support for the wind turbines during periods of insufficient resource and reduce dynamic power-tracking curtailments when there is excess resource. Nonetheless, we anticipate that more advanced control approaches could further improve performance and handle secondary control objectives such as structural load alleviation and battery cycling considerations. The A2e2g platform enables users to specify other controller modules as desired.

7 Conclusions and Next Steps

The A2e2g platform provides the first full system wind power plant integration from day-ahead forecasts and market participation to real-time control for wind plants. It is the first study of market participation of wind plants in frequency regulation/energy considering uncertainty of wind output and aerodynamic controls that account for waking. The findings of a limited amount of case studies for the ERCOT market show potential increases in the value from both energy and grid services. The case studies demonstrate the value of the A2e2g platform as a tool for evaluating future wind plant operational modes, existing and future market structures for energy and grid services, and the range of valuations for wind plants operating in these market structures with and without storage.

The A2e2g platform can be further developed for controlling wind plants offering and providing additional grid services beyond regulation, and further evaluating the market participation of wind plants for existing and future grid services. The case studies presented in this report offer a first step into using the A2e2g platform for a specific market region and historical year for energy and regulation.

The market participation framework of A2e2g has developed analytical methods for wind plants providing regulation that can be used by operators. The use cases in our project indicate how out-of-market energy payments affect the economic attractiveness of regulation markets and suggest that historical regulation prices are sometimes high enough to compensate for foregone energy revenue due to operation in curtailed mode for regulation supply. Our analytical insights are useful for market designers and analysts who might be interested in understanding how different aspects of market design—such as performance targets, performance-related revenue adjustments, and length of assessment period—affect the levels of participation by variable resources in regulation markets. A2e2g offers new analytical conservative bounds for the regulation offer quantity and price of a variable energy resource that consider the impact of the imperfect delivery of regulation capacity on resource profits and the resources' ability to meet the system operator's targets for the physical delivery of regulation capacity.

The A2e2g project demonstrated the control functionality and communication of the A2e2g platform with the GE 1.5-MW wind turbine operating at NREL's Flatirons Campus at the 4-second SCADA time frame, but it was not implemented on a wind power plant. Future work could include demonstrating the A2e2g controller on a commercial wind plant with communications with the power system operator. The A2e2g platform could also be integrated into NREL's grid integration capabilities to add a high-level controller for demonstrating wind plant operation in markets using wind plant simulation capabilities and coupled with the emulation capabilities with hardware in the loop using the GE turbine and energy storage. It could serve as a baseline high-level controller for wind plant grid integration work by the U.S. Department of Energy's Wind Energy Technologies Office.

The A2e2g platform is scalable and can be expanded for use as a regional platform for wind-based operation that can seamlessly integrate geographic diversity for land-based and offshore wind energy across a region, time horizons ranging from subsecond to day ahead (by linking with other tools developed by the Wind Energy Technologies Office for fast-acting essential reliability services), and optimal use of energy storage with advanced control solutions and energy storage sized for various energy and grid service scenarios and different regional markets and price structures. Future efforts would include further development of the computational requirements to facilitate applications in large-scale systems accounting for operational uncertainty and regional diversity and fusing classical horizon-based controls with state-of-the-art deep learning. This platform could be demonstrated for a regional systems-level analysis using the hardware and emulation capabilities of the Advanced Research on Integrated Energy Systems.

The A2e2g controller represents an important step toward developing near-firm power for a wind plant. Advanced control of wind plants with energy storage will be able to provide firm and flexible day-ahead dispatchable power and pave the way for other renewable energy resources to do the same.

References

ARPA-E PERFORM Forecast data. (Undated). Retrieved February 10, 2022, from https://registry.opendata.aws/arpa-e-perform/

Benjamin, S. G., Weygandt, S. S., Brown, J. M., Hu, M., Alexander, C. R., Smirnova, T. G., Olson, J. B., James, E. P., Dowell, D. C., Grell, G. A., Lin, H., Peckham, S. E., Smith, T. L., Moninger, W. R., Kenyon, J. S., & Manikin, G. S. (2016). A North American hourly assimilation and model forecast cycle: The rapid refresh. *Monthly Weather Review*, *144*(4), 1669–1694. https://doi.org/10.1175/MWR-D-15-0242.1

California Independent System Operator (CAISO). (2009). *BPM - CG CC 6524 Non Compliance Regulation Up Settlement*. From https://bpmcm.caiso.com/Pages/SnBBPMDetails.aspx?BPM=Settlements+and+Billing

CAISO. (2012). *BPM - CG CC 6624 Non Compliance Regulation Down Settlement*. From https://bpmcm.caiso.com/Pages/SnBBPMDetails.aspx?BPM=Settlements+and+Billing

California Independent System Operator Corporation. (2019). *Fifth Replacement Electronic Tariff Section 8—Ancillary Services as of Aug 12, 2019*. URL: https://www.caiso.com/Documents/Section8-AncillaryServices-asof-Aug12-2019.pdf

California Independent System Operator Corporation. (2020). *Fifth Replacement Tariff Section 11—California ISO Settlements and Billing as of Nov 18, 2020* (Vol. 5, Issue 1). From http://www.caiso.com/Pages/DocumentsByGroup.aspx?GroupID=8CE57609-7A51-498B-B248-5CD312512ABF

Denholm, P. L., Sun, Y., and Mai, T. T. (2019). *An Introduction to Grid Services: Concepts, Technical Requirements, and Provision from Wind* (NREL/TP-6A20-72578). URL: https://www.nrel.gov/docs/fy19osti/72578.pdf

Dent, C. J., Bialek, J. W., and Hobbs, B. F. (2011). Opportunity Cost Bidding by Wind Generators in Forward Markets: Analytical Results. *IEEE Transactions on Power Systems*, *26*(3), 1600–1608. https://doi.org/10.1109/TPWRS.2010.2100412

Ela, E., Hytowitz, R. B., and Helman, U. (2019). *Ancillary Services in the United States Technical Requirements, Market Designs and Price Trends* (No. 3002015670). Electric Power Research Institute. From: https://www.epri.com/research/products/000000003002015670

Electric Reliability Council of Texas (ERCOT). (2020). *ERCOT Nodal Protocols* (March 2020). URL: https://www.ercot.com/mktrules/nprotocols/current

ERCOT Market Education. (2020). *Settlements 301 Module 6: Real-Time Operations*. URL: https://www.ercot.com/files/docs/2020/05/12/2019_09_Set301_M6_-_RT.pdf

ERCOT Market Information System. (Undated). *Wind Power Production—Hourly Averaged Actual and Forecasted Values by Geographical Region*. From https://www.ercot.com/mp/data-products/data-product-details?id=NP4-742-CD

Fleming, P., Aho, J., Gebraad, P., Pao, L., & Zhang, Y. (2016). Computational fluid dynamics simulation study of active power control in wind plants. *American Control Conference*, *2016-July*, 1413–1420. https://doi.org/10.1109/ACC.2016.7525115

Fleming, P., Annoni, J., Shah, J. J., Wang, L., Ananthan, S., Zhang, Z., Hutchings, K., Wang, P., Chen, W., and Chen, L. (2017). Field test of wake steering at an offshore wind farm. *Wind Energy Science*, *2*(1), 229–239. https://doi.org/10.5194/wes-2-229-2017

Foley, A. M., Leahy, P. G., Marvuglia, A., and McKeogh, E. J. (2012). Current methods and advances in forecasting of wind power generation. *Renewable Energy*, *37*(1), 1–8. https://doi.org/10.1016/j.renene.2011.05.033

Giebel, G., and Kariniotakis, G. (2017). Wind power forecasting—A review of the state of the art. *Renewable Energy Forecasting: From Models to Applications*, 59–109. https://doi.org/10.1016/B978-0-08-100504-0.00003-2

Hansen, M. H., Hansen, A., Larsen, T. J., Øye, S., Sørensen, P., and Fuglsang, P. (2005). *Control design for a pitch-regulated, variable speed wind turbine* (No. Ris0-R-15OO). URL: https://backend.orbit.dtu.dk/ws/portalfiles/portal/7710881/ris_r_1500.pdf

Henson, B. (2015). *Regulation Market. Web-based Training.* (Jan 2015). ISO-NE. URL: https://www.iso-ne.com/static-assets/documents/2015/02/2015_regulation_market_ppt_slides.pdf

Hosseini, S. A., Toubeau, J. F., De Grève, Z., and Vallée, F. (2020). An advanced day-ahead bidding strategy for wind power producers considering confidence level on the real-time reserve provision. *Applied Energy*, *280*(July), 115973. https://doi.org/10.1016/j.apenergy.2020.115973

ISO New England. (2015). *Regulation Market Settlements.* URL: https://www.iso-ne.com/static-assets/documents/2015/02/regulation_market_settlements_02_19_2015.pdf

Kahrl, F., Kim, J. H., Mills, A., Wiser, R., Montanes, C. C., and Gorman, W. (2021). *Variable Renewable Energy Participation in U.S. Ancillary Services Markets Economic Evaluation and Key Issues* (Issue October). https://doi.org/10.2139/ssrn.4082609

Kelley, C. L., and Ennis, B. L. (2016). SWiFT site atmospheric characterization. In *Sandia Report* (No. SAND2016-0216). URL: https://www.osti.gov/biblio/1237403

Knudsen, T., and Bak, T. (2013). Simple model for describing and estimating wind turbine dynamic inflow. *Proceedings of the American Control Conference*, *1*, 640–646. https://doi.org/10.1109/acc.2013.6579909

Kumler, A., and Lundquist, J. K. (2021). *Probabilistic Day-Ahead Forecasting using an Analog Ensemble Approach for Wind Farm Grid Services A2e2g: Atmosphere to Electrons to Grid* (NREL/PR-5D00-80142). URL: https://www.nrel.gov/docs/fy21osti/80142.pdf

Liang, J., Grijalva, S., and Harley, R. G. (2011). Increased Wind Revenue and System Security by Trading Wind Power in Energy and Regulation Reserve Markets. *IEEE Transactions on Sustainable Energy*, *2*(3), 340–347. https://doi.org/10.1109/TSTE.2011.2111468

Loutan, C., Gevorgian, V., Chowdhury, S., Bosanac, M., Kester, E., Hummel, D., Leonard, R., Rutemiller, M., Fregoe, J., Pittman, D., and Kosuth, C. (2020). *Avangrid Renewables Tule Wind Farm Demonstration of Capability to Provide Essential Grid Services*. CAISO. From: https://www.esig.energy/download/avangrid-renewables-tule-wind-farm-demonstration-of-capability-to-provide-essential-grid-services-march-2020/

Loutan, C., Klauer, P., Chowdhury, S., Hall, S., Morjaria, M., Chadliev, V., Milam, N., Milan, C., and Gevorgian, V. (2017). *Demonstration of Essential Reliability Services by a 300-MW Solar Photovoltaic Power Plant* (NREL/TP-5D00-67799; Issue March). National Renewable Energy Laboratory. URL: https://www.nrel.gov/docs/fy17osti/67799.pdf

Midcontinent Independent System Operator (MISO). (2018). *Business Practice Manual No. 005 Market Settlements* (BPM-005-r18). From https://www.misoenergy.org/legal/business-practice-manuals/

MISO. (2019). *Business Practices Manual Energy and Operating Reserve Markets* (Manual No. 002 *Energy and Operating Reserve Markets*; Issue 002). From https://www.misoenergy.org/legal/business-practice-manuals/

Monache, L. D., Anthony Eckel, F., Rife, D. L., Nagarajan, B., and Searight, K. (2013). Probabilistic Weather Prediction with an Analog Ensemble. *Monthly Weather Review*, *141*(10), 3498–3516. https://doi.org/10.1175/MWR-D-12-00281.1

Monache, L. D., Nipen, T., Liu, Y., Roux, G., and Stull, R. (2011). Kalman Filter and Analog Schemes to Postprocess Numerical Weather Predictions. *Monthly Weather Review*, *139*(11), 3554–3570. https://doi.org/10.1175/2011MWR3653.1

Nock, D., Krishnan, V., and McCalley, J. D. (2014). Dispatching intermittent wind resources for ancillary services via wind control and its impact on power system economics. *Renewable Energy*, *71*, 396–400. https://doi.org/10.1016/j.renene.2014.05.058

National Renewable Energy Laboratory (NREL). (Undated). "Scalable Integrated Infrastructure Planning Model." https://www.nrel.gov/analysis/siip.html

NREL. (2020). "FLORIS version 2.4.0." https://github.com/NREL/floris

New York Independent System Operator (NYISO). Customer Settlements. (2020). *Manual 14 Accounting and Billing Manual Prepared By: NYISO Customer Settlements* (Version 5.1).

NYISO. (2021). *Manual 14 Accounting and Billing Manual*. URL: https://www.nyiso.com/documents/20142/2923231/acctbillmnl.pdf/

NYISO Operations Engineering. (2020). *Manual 2 Ancillary Services Manual* (Version 6.0; Issue May). NYISO. URL: https://www.nyiso.com/documents/20142/2923301/ancserv.pdf

PJM. (2016). *PJM Manual 12: Balancing Operations Revision 40* (March 26, 2020). From: https://www.pjm.com/library/manuals

PJM. (2020a). *PJM Manual 11: Energy and Ancillary Services Market Operations* (Revision 111 Effective Date: November 19, 2020). From: https://www.pjm.com/library/manuals

PJM. (2020b). *PJM Manual 28: Operating Agreement Accounting* (Revision: 84). From: https://www.pjm.com/library/manuals

Rebello, E., Watson, D., and Rodgers, M. (2019). Ancillary services from wind turbines: AGC from a single Type 4 turbine. *Wind Energy Science Discussions*, 1–17. https://doi.org/10.5194/wes-2019-26

Rebello, E., Watson, D., and Rodgers, M. (2021). Practical aspects and results—Providing secondary frequency regulation from a single Type 4 wind turbine and a battery storage system. *IEEE Green Technologies Conference*, *2021-April*, 105–111. https://doi.org/10.1109/GreenTech48523.2021.00027

Reedy, S. (2018). *Simulation of Real-Time Co-Optimization of Energy and Ancillary Services for Operating Year 2017*. POTOMAC ECONOMICS. URL: https://www.ercot.com/files/docs/2020/05/19/IMM_Simulation_of_Real-Time_Co-optimization_for_2017.pdf

Seel, J., Mills, A., Wiser, R., Deb, S., Asokkumar, A., Hassanzadeh, M., and Aarabali, A. (2018). *Impacts of High Variable Renewable Energy Futures on Wholesale Electricity Prices, and on Electric-Sector Decision Making*. In *Lawrence Berkeley National Laboratory* (Issue May). From: https://emp.lbl.gov/publications/impacts-high-variable-renewable

Sinner, M., Spyrou, E., Bay, C. J., King, J., and Corbus, D. (2022). *Coordinated Wind Power Plant and Battery Control for Active Power Services*. Submitted to Wind Energy. (Under review. Preprint available upon request.)

Soares, T., Pinson, P., Jensen, T. V., and Morais, H. (2017). Optimal offering and allocation policies for wind power in energy and reserve markets. *IEEE Transactions on Sustainable Energy*, *7*(3), 1036–1045. https://doi.org/10.1002/we.2125

Southwest Power Pool. (2015). Minutes of Southwest Power Pool Operating Reliability Working Group meeting. URL: https://www.spp.org/documents/27182/orwg%20020515%20minutespdf.pdf

Southwest Power Pool. (2020). *Market Protocols SPP Integrated Marketplace: Vol. Rev 59* (Revision 75). From: https://www.spp.org/spp-documents-filings/

Spyrou, E., King, J., Corbus, D., Zhang, Y., and Gevorgian, V. (2022). Offering of Variable Resources in Regulation Markets With Performance Targets: An Analysis. *IEEE Transactions on Sustainable Energy*, *13*(3), 1620–1630. https://doi.org/10.1109/TSTE.2022.3165150

van Wingerden, J. W., Pao, L., Aho, J., and Fleming, P. (2017). "Active Power Control of Waked Wind Farms." *IFAC World Congress*, *50*(1), 4484–4491. https://doi.org/10.1016/j.ifacol.2017.08.378

Vijayshankar, S., Stanfel, P., King, J., Spyrou, E., and Johnson, K. (2021). "Deep Reinforcement Learning for Automatic Generation Control of Wind Farms." *Proceedings of the American Control Conference*, *2021-May*, 1796–1802. https://doi.org/10.23919/ACC50511.2021.9483277

Watson, D., Hastie, C., Gaudette, B., and Rodgers, M. (2018). "Demonstrating Stacked Services of a Battery in a Wind R&D Park." *IEEE Transactions on Power Systems*, *33*(2), 1411–1419. https://doi.org/10.1109/TPWRS.2017.2718512

Wiser, R., Bolinger, M., Hoen, B., Millstein, D., Rand, J., Barbose, G., Darghouth, N., Gorman, W., Jeong, S., Mills, A., and Paulos, B. (2021). *Land-Based Wind Market Report: 2021 Edition*. URL: https://www.energy.gov/sites/default/files/2021-08/Land-Based%20Wind%20Market%20Report%202021%20Edition_Full%20Report_FINAL.pdf

Appendix A. Alternate Modules

A.1 Machine-Learned Power Forecast

As an alternative to the FLOw Redirection and Induction in Steady State (FLORIS)-based forecasting approach described in Section 4.2, we also include a machine-learning-based method in the Atmosphere to Electrons to the Grid (A2e2g) platform. The surrogate model approach we use is a machine-learning technique called Gaussian process regression (GPR). Compared to other regression models, GPR is a nonparametric model that has the benefit of naturally providing probabilistic estimates that are updated based on the available data, which makes it well-suited to forecasting applications. We use the GPyTorch software package to implement GPR for this work, although several other libraries are also available in Python and other programming languages. An example of the surrogate model power predictions compared to exact FLORIS predictions is provided in Figure A-1.

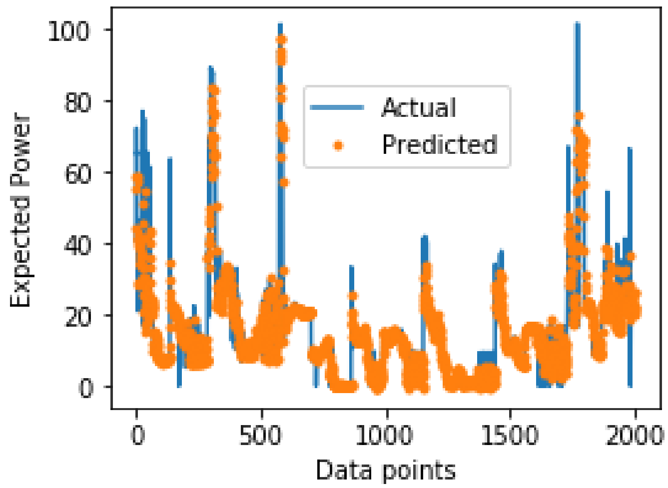

Figure A-1. Expected plant power using FLORIS (blue) and the GPR surrogate model (orange)

A.2 Reinforcement Learned Control

As an experiment, we sought to develop a controller to ensure that the wind plant power output follows the automatic generation control (AGC) signal using yaw control. This controller would take in the current wind condition (wind speed) and produce wind turbine yaw commands to achieve the desired power set point rather than using blade pitch (axial induction). Conceptually, the process of generating yaw commands from wind conditions and the AGC signal could be completed in a model-based fashion, using a model (such as FLORIS) to determine optimal yaw angles online. However, this is impractical for larger farms due to the computations involved in optimizing for every scenario that the controller experiences.

Instead, we use reinforcement learning to determine a map between a situation (wind condition and AGC signal) and action (turbine yaw angles) that maximizes a reward function (minimizes the error between the AGC signal and wind plant power). Once this mapping has been learned, the current situation can be mapped directly to an action without the need for expensive online model evaluations. As such, reinforcement learning methods determine optimal control actions based on interactions with the environment, rather than on physics-based models.

The specific reinforcement learning algorithm that we use is deep deterministic policy gradient (DDPG), which can learn optimal control policies in high-dimensional settings such as the one posed by our wind farm control problem. DDPG uses neural networks to approximate both the optimal value function and current policy. We also considered a somewhat simpler reinforcement learning approach referred to as deep state-action-reward-state-action (SARSA) (which uses a deep neural network to approximate the optimal value function but does not maintain an explicit policy function, instead following a near-greedy search method to explore the action space). To train each controller, we simulated many situation-action trajectories using dynamic FLORIS, with the controller learning optimal policies/actions based on the rewards it experiences in the simulation.

Although this is a computationally expensive process, it can take place offline before the control actions are needed in real time, making it more suitable than optimizing behavior online. This training procedure is shown in Figure A-2 (with the "agent" referring to the learned controller). The action for the wind power plant control problem we consider is the yaw angle of each of the individual wind turbines, whereas the state is the wind speed at each turbine location. The reward that the agent receives is maximum when the AGC signal is perfectly achieved and decreases according to a truncated Gaussian curve as the absolute deviation from the AGC signal increases. Figure A-3 demonstrates how the average rewards evolve during training the DDPG and Deep SARSA.

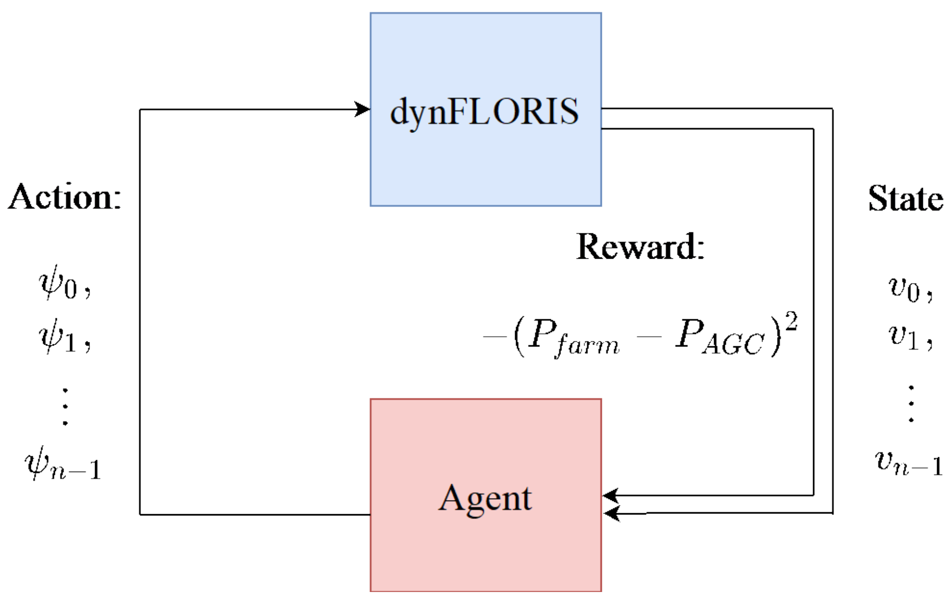

Figure A-2. The learning agent generates an action based on a temporal difference method

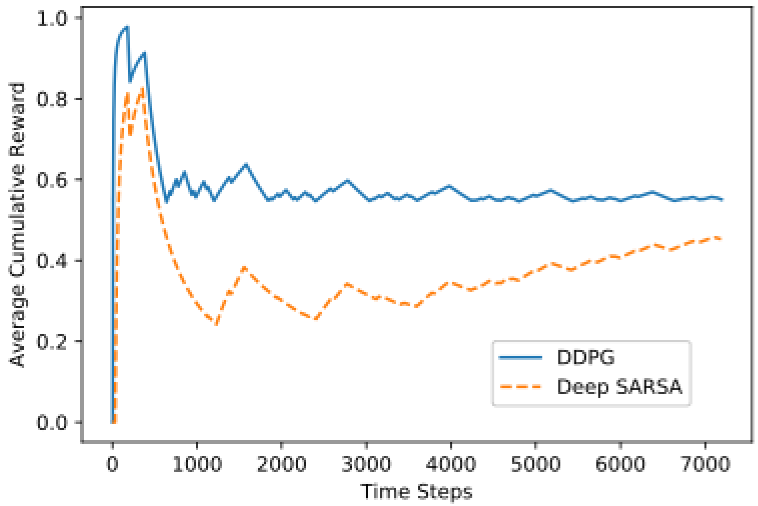

Figure A-3. Average cumulative reward during training

To demonstrate the behavior of the reinforcement learning controller for the purpose of this report, we used a simple wind farm that comprises nine turbines, rated at 1.5 megawatts each, arranged in a 3-by-3 grid. As a baseline controller to use as a reference, we used a proportional-integral controller that adjusts the yaw angle as a function of the error between the AGC signal and power but does not seek out optimal behavior.

Figure A-4 shows the main controller results. On training the controller with uncertain wind speeds and the same reference AGC signal for both DDPG and Deep SARSA, the controller successfully learns to adjust the yaw angles. The power output of the farm during the testing phase is in good agreement with the reference signal. We note that the performance of the baseline controller is better than that of Deep SARSA and is worse than DDPG. On one hand, bearing in mind that we are given only wind speed measurements, the reinforcement learning controller can make better decisions to track the given AGC signal in a model-free setting than the baseline. When there is abundant data available from a wind farm, DDPG has great potential. On the other hand, training the reinforcement learning controller is time-intensive, whereas the baseline controller does not require any training at all. Thus, the choice of the controller is a trade-off between performance and (offline) time spent on training. DDPG may be able to be improved further by continuing to learn online using the data collected from a real wind farm.

Full details about the reinforcement-learned controller can be found in Vijayshankar et al. (2021).

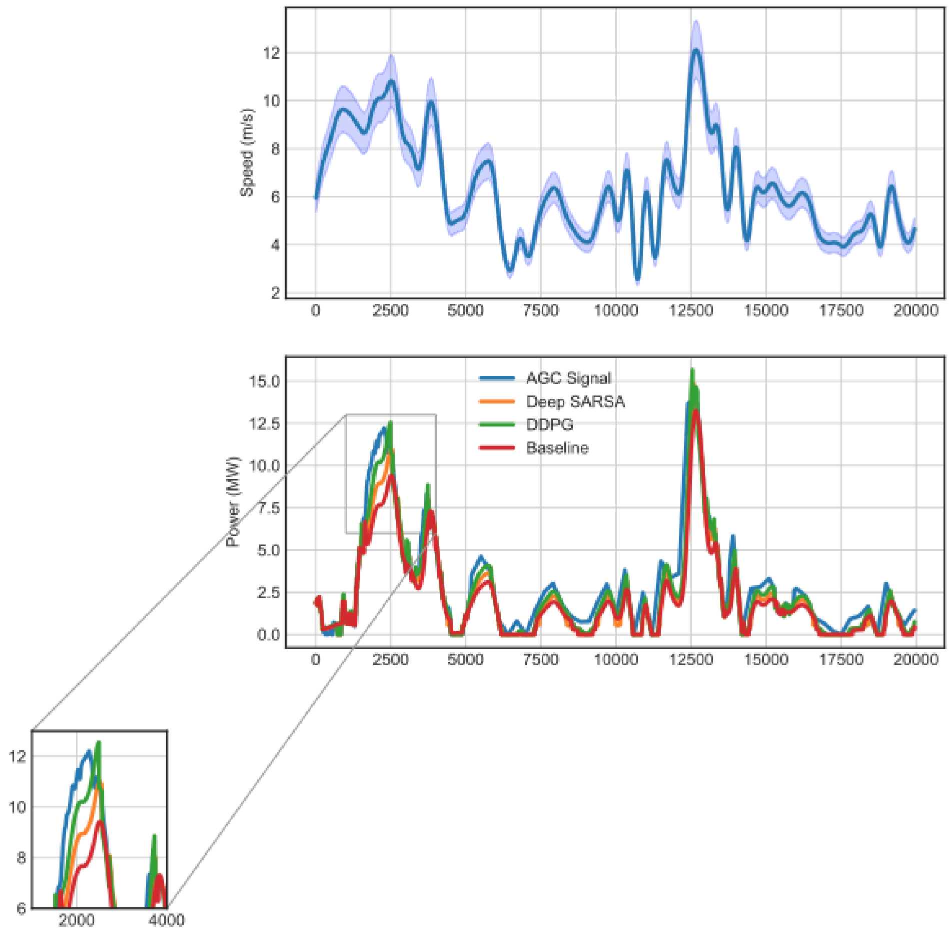

Figure A-4. Performance of DDPG, Deep SARSA, and the baseline control algorithms when trained with uncertain wind forecasts. We allow the wind speed forecast (top plot) to vary by 10% and train the agent in the presence of this uncertainty. Results from the validation step (lower plot) show that both algorithms perform well in tracking the AGC signal.