

Combining Balancing Areas' Variability: Impacts on Wind Integration in the Western Interconnection

Michael Milligan

National Renewable Energy Laboratory
1617 Cole Blvd., Golden, CO 80401
michael.milligan@nrel.gov

Brendan Kirby

National Renewable Energy Laboratory
Consultant
2307 Laurel Lake Rd
Knoxville, TN 37932
kirbybj@ieee.org

Stephen Beuning

Xcel Energy, Market Operations
550 15th Street, 12th Floor
Denver, CO 80202
Stephen.J.Beuning@xcelenergy.com

Abstract

Interest in various wide-area balancing schemes to help integrate variable generation sources such as wind and solar have generated significant discussion. As we have shown in past work, large balancing areas not only help with wind integration, but can also increase the efficiency of operations in systems without wind. Recent work on the Western Wind and Solar Integration Study (WWSIS) has found that cooperation between balancing areas over the WestConnect footprint will increase the efficiency of commitment and dispatch at wind penetrations ranging from 10-20% of annual electricity demand, and will be essential for high penetrations and small balancing areas. In addition the Northwest Wind Integration Action Plan recommended balancing area cooperation as a method to help integrate the large potential wind development. In this paper we will investigate the potential impact of balancing area cooperation on a large-scale in the Western Electricity Coordinating Council (WECC). We will utilize data adapted from the WWSIS for the Western Interconnection. The analysis uses time-synchronized wind and load data to evaluate the potential for ramp requirement reduction that could be achieved with combined operation in the Pacific Northwest. Chronological analysis and ramp duration analysis quantify the benefit in terms of not only the ramp sizes, but the frequency of the potentially avoided ramps that must be managed by the non-wind generation fleet. Multiple approaches that can be used to achieve these benefits will also be suggested in the paper, along

with a description of the proposed Energy Imbalance Service, under study by WECC's Seams Issues Subcommittee. We also suggest other approaches that can help achieve much of the benefit of full consolidation without requiring the physical consolidation of balancing areas.

Introduction

Over the past several years there have been a significant number of new approaches for managing variability in the power system. Some of the interest has been inspired by the rapid expansion of wind energy and the increased level of variability that wind brings to the power system. Because individual loads are generally uncorrelated with each other, and because the same is true for wind energy over short time scales, there is less than perfect coincidence between loads, wind, and load and wind together. The variability per unit in load, wind, and net load (load less wind) generally falls as a larger area is examined. For example, in Milligan and Kirby (2007) we showed the impact of combined balancing area operation with high levels of wind energy penetration in Minnesota.

Combining balancing areas (BA) can be done in many ways. Physical consolidation is the simplest conceptually, but there may be reasons why physical consolidation is not desirable. In these cases, there are other mechanisms that are available that allow two or more BAs to pool their variability. In the United States, the most common approach is to run large energy markets that operate on a fast time scale, typically every 5 minutes. Other methods include the ACE (area control error) Diversity Interchange (ADI) project, run by Northern Tier Transmission Group (see www.nttg.biz) that calculates the net regulation requirement among several balancing areas and sends out revised signals to participating generators on automatic generation control. Similarly, the Joint Initiative is rolling out a dynamic scheduling system (DSS) which will allow subscribing parties in the Western Interconnection to quickly set up dynamic schedules. Although dynamic schedules are of limited value if only applied to individual loads or resources, setting up ACE or other system-wide balancing metric to be shared via dynamic schedule may hold promise for areas that do not consolidate, and yet want an ADI-like product over a longer time period.

Methods such as ADI and DSS are components of what might be called virtual balancing areas (VBA), although we use the term VBA to describe the full spectrum of options that may be available to help share variability without requiring physical consolidation.

In our view, multiple BAs can achieve much, or even all of the benefit of physical consolidation *without* physically consolidating. In this article we describe some variations on approaches to achieve significant benefits from various forms of variability sharing, and analyze potential impacts of variability sharing in the Pacific Northwest with a 16% wind energy penetration, based on the Western Wind and Solar Integration Study database. We discuss operational

benefit, and provide a simple analysis of the impact of variability sharing on resource acquisition as it might be carried out in the planning time horizon.

Alternative Approaches for BA Cooperation

Approaches to share variability can vary significantly, but are ultimately tied to one or more time frames corresponding to the usual power system operational cycle. These are shown in Figure 1.

The unit commitment process involves starting up thermal units, which typically take several hours to achieve operating temperature. Although there can be some variation in the details, unit commitment is typically done a day in advance, based on the load and wind forecast day-ahead. In some areas, the unit commitment is done once per day, in others, the commitment schedule may be updated several times a day.

Load following typically spans from 10's of minutes to several hours. Load following encompasses potentially large swings in demand such as the morning load pickup or evening load drop-off. During times like these, there must be sufficient committed generation that has the ability to ramp quickly enough to meet the changing demand. In areas with fast markets, units bid alternative levels of generation as a function of price. As market conditions change, raising or lowering prices, the economic unit on the dispatch stack increases or decreases output accordingly.

The final time scale is regulation, which usually covers time periods from seconds to minutes. Regulation demand is covered by units on automatic generation control (AGC), which cover all of the variability between the economic dispatch changes. Regulation does not follow a trend to any significant degree, and over suitably long periods of time, regulation will mathematically integrate to zero. Therefore there is no energy component in regulation.




Figure 1. Power systems operation time scales

To integrate large amounts of wind, it is important that there is an institutional mechanism to tap the available flexibility that is inherent in the generation fleet. As discussed in Milligan and Kirby (2010) the structure of the energy market, along with generation scheduling rules, are the most important factors to access this flexibility. There are two primary characteristics that BA cooperation should have to obtain maximum flexibility: (1) if a generating unit can physically change its output, scheduling rules should not prevent this if it is economic; and (2) spreading variability across more units requires less ramping per unit, so having a large pool of generation is advantageous.

Large, Fast, Integrated Energy Markets

These objectives can be accomplished with large markets that operate over a broad electrical footprint at 5-minute intervals. Fast economic dispatch intervals ensure that economic units are allowed to respond to changing system conditions, as needed. They also ensure that a large pool

of generators are not restricted from responding, such as would be the case with hourly markets that are prevalent in much of the Western part of the United States. With economic dispatch occurring only hourly, units that are capable of responding within the hour are restricted to hold their setpoints until the top of the next hour. This implies that the regulation burden, which is much larger over a long time frame like an hour, must be carried by a subset of the generation fleet. Viewed in the context of an optimization problem that minimizes cost subject to balancing loads with resources, the hourly dispatch becomes a binding constraint that is added to the problem, increasing cost and making the optimization problem more difficult.

Robust electricity markets in the U.S. involve a combination of unit commitment and load following, the latter carried out by the 5-minute energy markets. In addition, regulation is typically provided as an ancillary service thru a separate market. In the absence of market and significant transmission constraints, a large market can provide most or all of the benefit of a combined balancing area.

In a market such as implemented in most of the Eastern U.S., any generator can bid itself into the market at prices and quantities that allow it to make a profit. At prices that are below cost, the generator is not obligated to respond; thus the market is purely voluntary. We argue that demand response is a potentially valuable resource also, and even though it is not currently widely used in ancillary service markets or energy markets in the U.S., from the perspective of the power system operator, demand response could be indistinguishable from generation in the bid stack. As different generators have different properties (ramp rates, start-up times, notification times, etc.) so too do different forms of demand response. See Kirby and Milligan (2010).

Energy-only Markets

In the Western Interconnection, areas outside of California and Alberta do not presently have a widespread energy market, although there are some regional hourly markets in operation today. The Seams Issues Subcommittee of the Western Electricity Coordinating Council (WECC) is currently investigating an Energy Imbalance Service (EIS) that would achieve most of the benefits of a large-scale energy market, but without a coordinated unit commitment or regulation market.

The proposed EIS would utilize two tools: the Seams Coordination Tool (SCT), which can determine transmission service curtailment priority for flow components on the grid. This is based on the level of transmission service that is used to deliver the energy. Both tagged and untagged flows are considered. The SCT would need to be modified so that it can pass relevant information to the second tool, the Energy Imbalance Service Tool (EIST).

The EIST uses a security-constrained economic dispatch to provide two functions:

- **Balancing Service:** This service redispatches to balance deviations from schedules in generator output and errors in load schedules.
- **Congestion Redispatch Service:** This will redispatch to relieve overload constraints on the grid

The current approach that is used by WECC BAs for balancing services comes from Tariff Schedules 4 and 9. The proposed EIS replaces part of the BA services, and results in a “virtual consolidation” that results from a wide-area security constrained economic dispatch that covers imbalances. The congestion redispatch service is new to the non-market portions of WECC.¹

These services would be provided by any resource that voluntarily offered responsive capability, and would be cleared by the security-constrained economic dispatch process. The transmission service that would be used by the EIS function would continue to use traditional reserved transmission service, but the EIS would not use pre-reserved transmission. Instead, the EIS flow would receive the lowest service priority, and therefore would not displace any reserved service. EIS flows would be compensated by paying after-the-fact to participating transmission providers. Figure 2 illustrates the operations timeline for the toolkit.

¹ The California Independent System Operator (CAISO) and Alberta Electric System Operator (AESO) each provide this service to their respective BA.

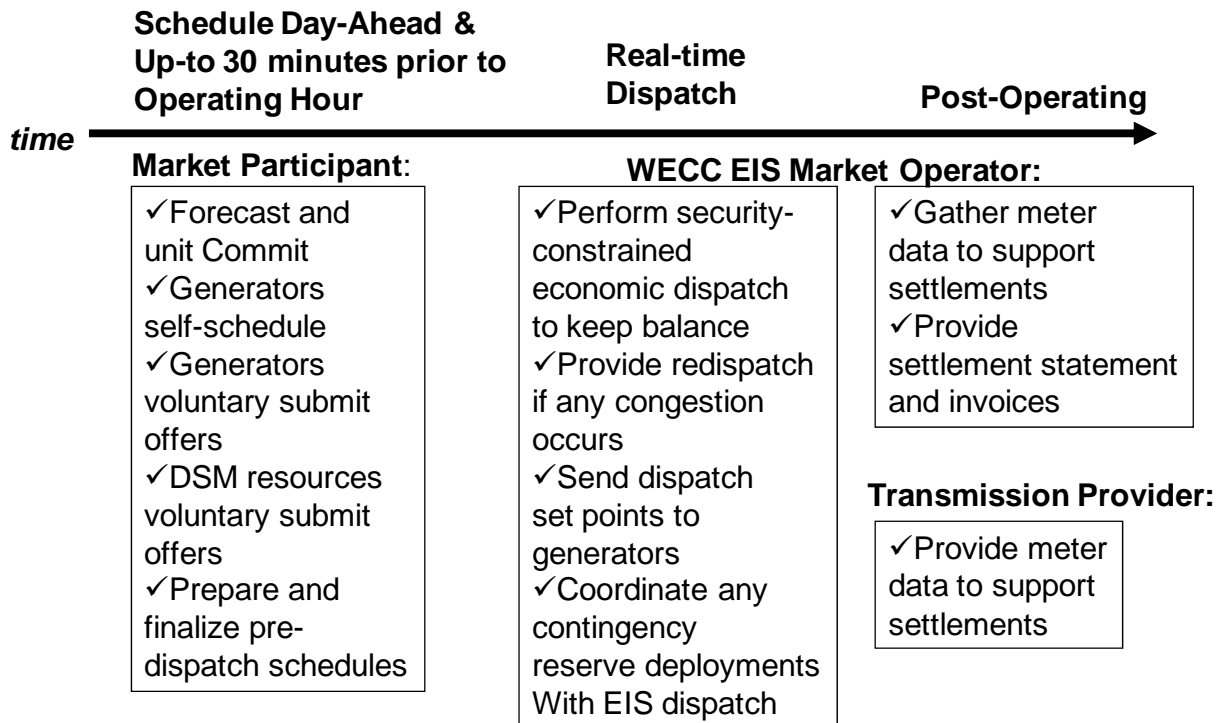


Figure 2. Operation timeline for the WECC EIS toolkit

The EIS would effectively implement one form of a virtual BA across the Western Interconnection (California and Alberta would not be included because they already have energy markets). Imbalances would be netted out, much as they would be in a single BA. The EIS does not result in a coordinated unit commitment, nor does it pool regulation. However, the netting of energy imbalance, which would include impacts of load and wind, are expected to be significant. Figure 3 illustrates the concept, with each of the small bubbles representing a single BA. The arrows between the BAs indicate energy flows that would need to be managed under the current operational approach. However, under the EIS, only the footprint net imbalance would need to be managed, resulting in less variability and less required ramping across the footprint.

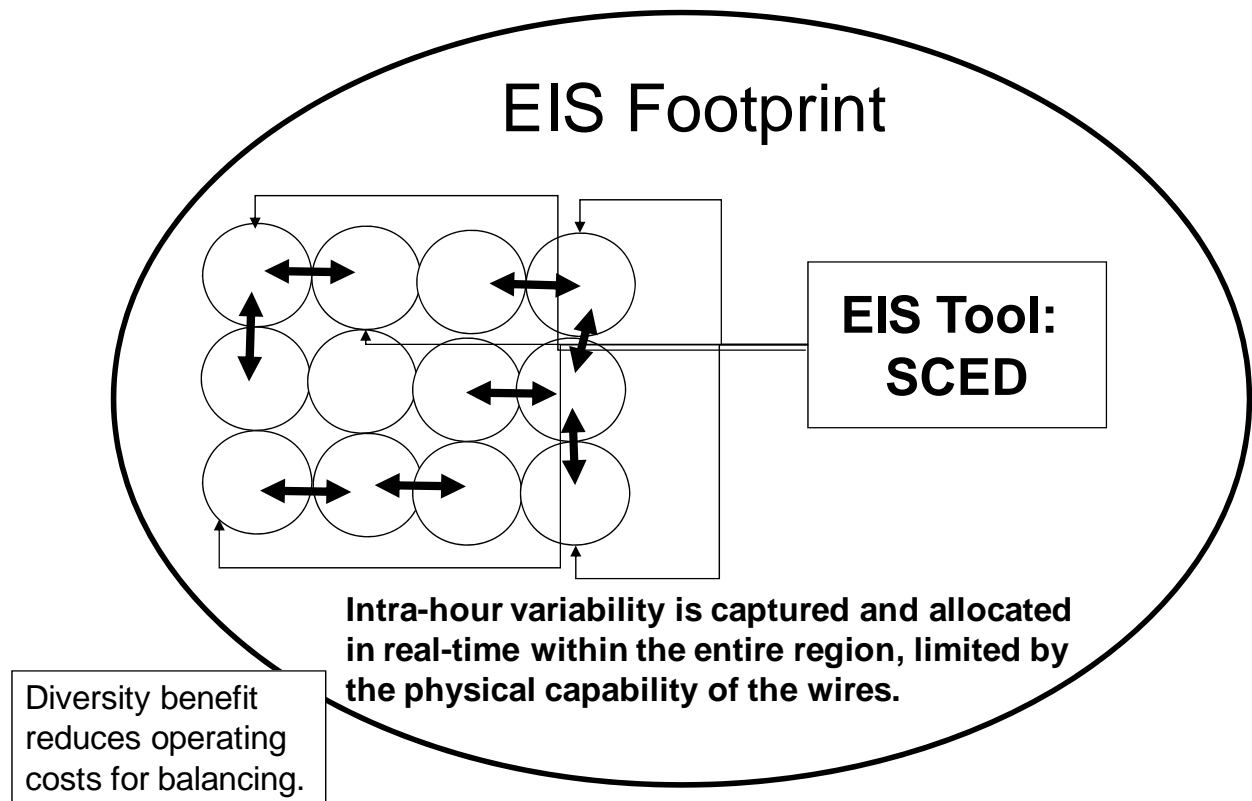


Figure 3. The EIS would effectively pool variability within the operating footprint, similar to a single BA

Virtual Consolidation

The term “virtual consolidation” does not have a widely-accepted definition. As such, we use this term (or its variant “virtual balancing area”) broadly, and apply it to any form of variability sharing across BAs that allow the BAs to retain their autonomy. In our usage, it is then possible that any given form of VBA will achieve only a partial benefit of full consolidation, although other forms of VBA could achieve the full benefit, or nearly so, of full consolidation.

There are several initiatives in various stages of developments in the West that implement some form of variability sharing. These initiatives have come out of the Northern Tier Transmission Group (NTTG) and the Joint Initiative, which is comprised of NTTG, Columbia Grid, and WestConnect. They include:

- ACE Diversity Interchange (ADI). The ADI pilot operates by pooling the ACE signals of participating balancing areas, netting out the variability and sending revised AGC signals to regulating units in the various BAs. This is a regulation-only service, and does not

necessarily reduce the regulation burden on expensive units, and may reduce the regulation burden on relatively inexpensive units.

- **Dynamic Scheduling System (DSS).** This initiative allows subscribers to quickly set up dynamic schedules. This product can be used to electrically move a generator, load, or combination to a different balancing area. In our view, the full benefit of dynamic scheduling would involve pooling ACE or other measure of net imbalance. If implemented over longer time frames that are consistent with load following, a modified form of DSS (or even a persistent dynamic schedule) could potentially achieve a very large benefit.
- **Intra-hour Transaction Accelerator Platform (ITAP).** ITAP allows hourly transactions to be executed sub-hourly. The transaction period is still hourly, but entities have the ability of setting schedules more frequently than in the usual paradigm.

There may be other forms of variability sharing that are not covered here, and some of them may be ad hoc arrangements between entities on an as-needed basis. One such example is described below.

Example: Inter-BA Schedules

Kirby and Milligan (2009) showed that hourly scheduling between BAs with wind deliveries from one area to another results in inefficient dispatch and regulation for both the host and receiving balancing areas. There is also an additional capacity requirement that is imposed on the host BA. This capacity requirement affects both the installed capacity required to meet load (plus the wind exports) and the receiving BA, which must have sufficient capability to meet its own internal loads. The additional capacity requirement that falls on the host is tied to the nature of transactions between BAs, and is not due to a physical capacity need, however. There are institutional mechanisms that can be used to reduce or eliminate this capacity need on the host.

Figure 4 illustrates the example. Wind generation in one BA is being delivered to another BA. Just before 9:30, the wind generation drops over a 15-minute period. The region only allows inter-BA schedules to change at the top of each hour, however, so the physical host BA must continue delivery of the scheduled wind energy from its non-wind fleet until the top of the next hour. In this case, the host BA generation must exceed its load for the duration of the market period; this represents a capacity requirement on the host.²

² The actual amount of “extra” capacity required depends on the maximum credible wind reduction during the hour. In this simple example the wind drops 500 MW from 1200 MW to 700 MW and that is the largest drop ever expected. Actual wind performance is more complex but wind output does not normally drop from full to zero within an hour for a large wind fleet.

Wind Serves External Load

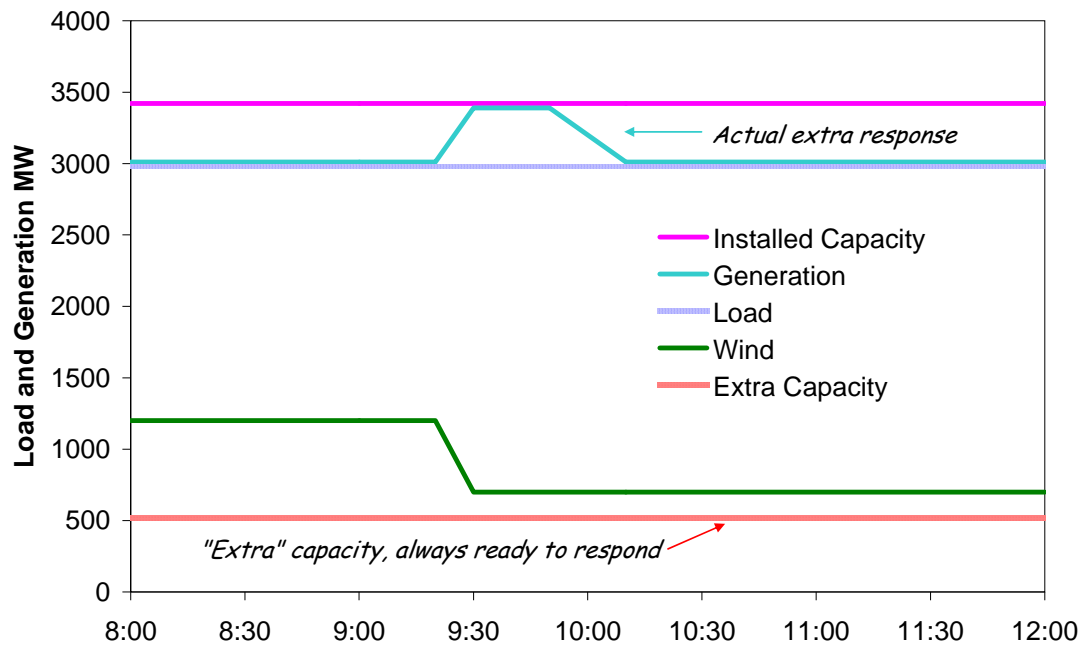


Figure 4. The exporting BA has a capacity obligation to keep the wind schedule.

Unfortunately, from the point of view of the receiving BA, the capacity used by the host to maintain the wind schedule is of no benefit. The receiver cannot use this capacity, and the only real impact is a delay in response speed for the generators in the receiving BA. This is shown in Figure 5. Because the market period clears hourly, the delivered wind power differs from the actual wind power as a result of the host covering the hourly schedule. The receiver does need to respond to the change in wind power, but this response does not happen until the next market period. Therefore, there is a delay in the required response from the receiver, but there is no capacity advantage for the receiver.

External BA Receives Wind




Figure 5. The receiving BA does not benefit from the extra capacity obligation held by the host.

This extra capacity impact can be calculated using an actual example from the Pacific Northwest. Bonneville Power Administration (BPA) delivers approximately 80% of the wind that is physically located within its BA off-system. WECC scheduling practice is to change schedules at the top of the hour, allowing for a 20-minute ramp period. Using public data from BPA's web site, we calculated the impact of the extra capacity requirement on BPA with the hourly scheduling change. We also compared this to the capacity impacts with a 30-minute schedule change, and a 10-minute schedule change (both calculated 10 minutes before real-time). Figure 6 illustrates this for 2009. During the year, additional wind capacity was coming online, and maximum wind output is 1,886 MW. Using the existing hourly schedule, set 2 hours ahead, the maximum annual capacity obligation for BPA is 617 MW; the minimum is -956 MW. Although a negative capacity obligation appears attractive, BPA has indicated in its rate proceedings that minimum-run issues during periods of high wind can be problematic. The average capacity obligation is not a good metric to measure the differences because it is near zero in all cases. However, the sum of the absolute differences for the hourly schedule is 876,013 MW-hours, whereas the same metric for the 10-minute schedule change is 156,100. This represents more than an 80% improvement in overall schedule deviations for the year.




Figure 6. Frequency and amounts of excess capacity under alternative scheduling periods.

To examine a sample week, Figure 7 illustrates the impact on BPA of 10-minute schedules. From this example it is clear that faster scheduling practices between BAs can reduce the capacity obligation on the host BA, improving efficiency and reducing cost.

The improvement in wind delivery can be realized in several ways:

- A dynamic schedule or pseudo-tie that dynamically moves the wind generation to the receiver's BA. Because the receiver must stand ready to accept the wind at the top of the hour under hourly scheduling protocols, accepting the wind on a faster schedule will result in smaller step changes in the schedule that are spread over the hour
- Faster market-clearing periods
- Bi-lateral agreements between the BAs.
- Combined operation of the two BAs.

At present BPA is pursuing a ½-hour schedule adjustment with the California Independent System Operator (CAISO) during hours that BPA runs short on reserves.




Figure 7. BPA's capacity obligation declines significantly if schedules change every 10 minutes.

Aggregation greatly benefits wind integration by reducing variability and reducing forecast errors as shown by BPA and WECC data. Several options are available for capturing this benefit. While BA consolidation has advantages other methods can realize most of the benefits without requiring consolidation. Several methods are being physically implemented and tested.

Analysis of the Northwest Balancing Areas

To further investigate the impact of variability sharing in the Pacific Northwest, we extracted wind and load data from the database used for the Western Wind and Solar Integration Study (WWSIS), managed by the National Renewable Energy Laboratory for the U.S. Department of Energy.³ We used 10-minute wind data, aggregated to hourly and longer time intervals. The data was mapped from transmission zones to balancing areas. The footprint is shown in Figure 8. The wind energy penetration is approximately 16% of annual energy demand.

Because of temporal seams in the underlying numerical weather prediction (NWP) modeling that was used to create the wind data set, we eliminated every third day from our analysis to avoid the

³ The final report for this project will be released May 20, 2010, and will be available at www.nrel.gov.

potential of overstating ramping impacts in the Northwest footprint. More information can be found at <http://wind.nrel.gov/public/WWIS/>.

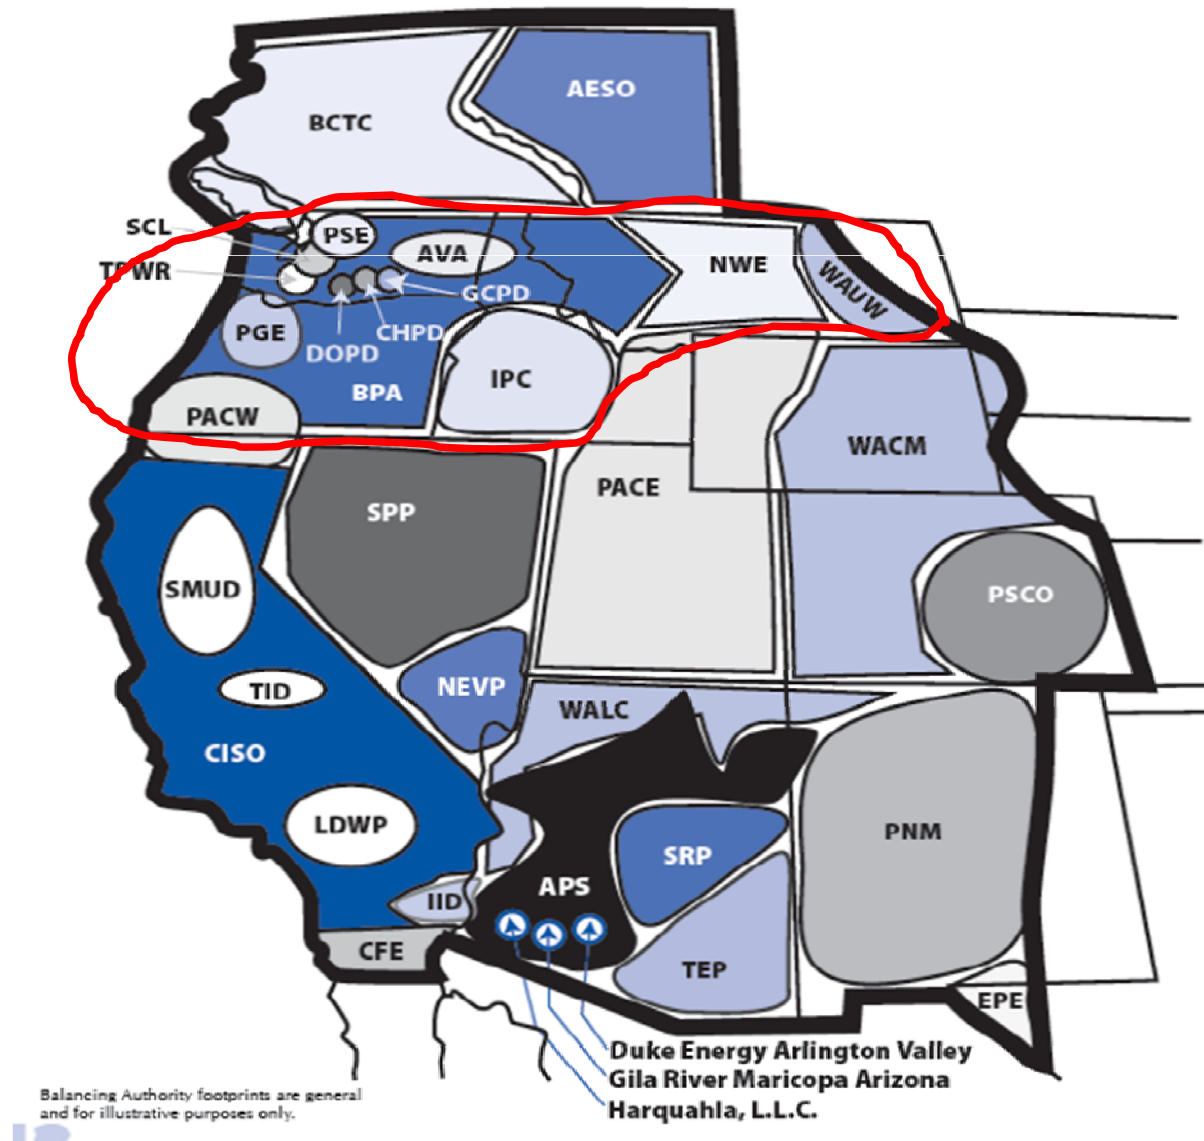


Figure 8. Northwest balancing areas in this analysis

Our analysis does not specify *how* variability pooling is accomplished, whether by actual BA consolidation, energy and ancillary service markets, or some type of VBA mechanism. Our analysis is based on full variability pooling, and is indifferent as to how that is done.

We also ignore transmission constraints. Although this is an obvious simplification in our analysis, this allows us to explore the full benefit of pooling. As transmission constraints are alleviated by building new lines to serve new loads and deliver energy from new resources such

as wind, constraints will become less of an issue. In fact, before the region arrives at a 16% wind energy penetration, we believe that some new transmission will be required.

Ramp analysis

The approach used for the ramping analysis is based on Milligan and Kirby (2007). The analysis proceeds in several steps. First, the net load (load less wind) is calculated for each BA separately, which represents the operational target for the remaining generation fleet assuming that all wind is utilized to meet various renewable energy objectives. Using this hourly data, the ramp requirement is then calculated. The calculations are carried out chronologically, and when one compares the need in two or more BAs at the same time, there is frequently the need for up-ramping capability in one area at the same time there is a need for down-ramping capability in another area. Figure 9 is taken from our analysis and serves as a good example.




Figure 9. Example of full variability pooling in the Pacific Northwest, with and without wind energy

The graph shows a sample week, drawn somewhat arbitrarily from our data set. Individual BA ramps were partitioned into positive and negative ramps, and these were separately combined into estimates of the up-ramps and down-ramps that would be met if the systems were to operate separately (no variability pooling). The wind and load data were then pooled into a single hypothetical balancing area, and the ramp requirements were re-calculated. The difference between the no-pooling case and the full-pooling case is represented in the graph. The upper panel shows the impact of variability pooling in the Northwest in the absence of wind. For this

sample week, the maximum ramp saving is approximately 250 MW, which occurs near hour 100. The symmetric nature of the graph is due to the simultaneous saving of +250 MW and -250 MW in the various BAs at the hour in question.

The lower trace of the graph shows the same information for the case that includes wind energy. Clearly, variability pooling in the wind case is far more valuable than in the no-wind case. For this sample week, more than 1,500 MW/hour up and down can be avoided at hour 75 with full variability pooling. The graph also indicates that 100s of MW/hour ramping can be avoided during other times, and that there are also times that there is no saving.

Figure 10 illustrates the results for the full year in the no-wind case. We can see from the graph that there is a benefit to pooling most hours of the year, ranging upwards to 400 MW/hour, but averaging significantly less than that. The upper panel shows the data chronologically, whereas the lower panel shows a frequency curve using the same data.




Figure 10. Variability pooling is valuable even without any wind.

Figure 11 illustrates the benefit of variability pooling with 16% wind energy penetration. We note that the maximum ramp saving is approximately 2,000 MW/hour in the wind case, compared to 400 MW/hour without wind. Clearly, variability sharing with wind is much more significant than without it, although we emphasize that there is a significant benefit even without wind.




Figure 11. Variability sharing is 5x more valuable when there is significant wind energy in the generation mix.

It is also useful to examine the variability statistics for load, wind, and net load under both the no-pooling and full-pooling cases. Although the statistical distributions of these variables are not necessarily normal, it is common to use the standard deviation as a measure of variability.⁴

Figure 12 shows the standard deviation for load, wind, and net load. For each variable, the standard deviation is calculated for both the pooled case and the individual case. Not surprisingly, all of the combined cases show a lower variability index than the separate cases. For net load, which is an indicator of what the power system must be operated to, the variability index difference between combined and separate BA cases is significant – approximately a 30% reduction in the footprint.

⁴ Non-normality of the underlying distribution implies that the usual significance tests of the standard deviation are compromised. However, the basic utility of the standard deviation as a metric of variability still holds.




Figure 12. Per-unit variability is reduced by pooling for load, wind, and net load.

Variable-ramp analysis and implications for resource acquisition

As more wind energy is developed, planners will need to respond by ensuring that the non-wind generation mix has sufficient flexibility: quick ramping capability and low turn-down levels. For this analysis we focus on the former. There is considerable work in progress to quantify these needs, including the North American Electric Reliability Corporation (NERC) Integrating Variable Generation Task Force (IVGTF) on Flexibility Requirements and Metrics for Variable Generation, and the International Energy Agency (IEA) Grid Integration of Variable Generation Resources (GIVAR) project.

Flexibility is important in the operating time frame, and as we have shown above, there are benefits in pooling variability. We now turn to a discussion of planning for sufficient ramping capability, comparing possible outcomes for separate, uncoordinated planning and coordinated planning that might be carried out by a consolidated BA. As before, we make no assumptions about the type of consolidation, real or virtual, that drives the combined case. In the discussion and analysis that follows, we focus on the extreme cases of no cooperation vs. full cooperation. In reality there might be some coordinated efforts that would result in an outcome between the extremes.

To estimate the impact of separate vs. combined planning for ramping capability, we proceed in a similar fashion as before. First, each BA is analyzed separately to determine the ramping needs it will have, given the load and wind data (for this analysis we used 10-minute wind and load data from the WWSIS). Ramping requirements were calculated for various time steps that ranged from 10 minutes to 12 hours. Each ramp occurs entirely within a single calendar day, and as before, every third day is excluded from the analysis. That is, we calculate the maximum daily ramping requirement in terms of ramp size (MW) and ramp rate (MW/min) for each ramp duration (10 min, 30 min, 1 hr, 2 hr, 4 hr, 8 hr, and 12 hr). Naturally, the ramps are not independent. The maximum 1 hr ramp is likely a part of the maximum 2 hr ramp on the same day. Still, the ramp size and rate metrics for different durations provides insight into the flexibility requirements imposed on the conventional generation and demand response fleet.

Ramps are classified by size, duration, and direction. From that information, maximum, minimum, and average ramps can be calculated, retaining the classification by size, duration, and direction. “Average” in this case means the average daily maximum ramp of the specific duration. To obtain the total ramping needs based on separate planning, the curves are added together. Estimating the combined ramp requirements follows the same basic algorithm as the separate analysis. The only exception is that the load and wind data are combined first, and then the various ramp statistics are calculated. The procedure is carried out for load alone, and for net load. Figure 13 illustrates. While the previous analysis focused on real-time aggregation benefits where an up-ramp in one BA is countered by a simultaneous down-ramp in the opposite direction this analysis focuses on the ramping capacity that each BA needs to have access to. Similar to the analysis of noncoincident peak loads, if one BA needs 100 MW of 2 hr up-ramp capability on one day and another BA needs 200 MW of 2 hr up-ramp capability on another day then they separately need 300 MW of 2 hr up-ramp capability since they have no way of sharing the resource. When combined they only require 200 MW of 2 hr up-ramp capability since they can share the resource.




Figure 13. Ramp duration curves for load alone show the benefit of combined planning for ramping adequacy.

The four traces on the graph appear for both positive and negative ramp requirements. The upper quadrant of the graph shows the aggregated maximum and un-aggregated maximum ramps (top two traces), which represent the separate (un-aggregated) and combined (aggregated) ramping needs. The difference between the curves shows the maximum ramping capability that could be avoided by coordinated planning and subsequent operations. Interestingly the maximum requirements are not that much greater than the average (average of the daily maximums) ramping requirements (13% higher in the un-aggregated 8 hr ramp case), especially in the aggregated case (9% higher in the aggregated 8 hr ramp case). This is because the daily load shape is relatively repeatable. The graph in Figure 14 shows the same information, but as ramp rates in MW/min. It is clear from both graphs that there is a significant reduction in ramping capability needed if future resource needs can be combined across BAs.




Figure 14. Ramp rate savings from cooperative planning for load alone

These benefits increase quite dramatically when 16% wind energy penetration is added to the system as shown in Figure 15 and Figure 16. The 8 hr un-aggregated maximum up-ramp capacity increases by 46% with the addition of this much wind while the un-aggregated average daily maximum up-ramp capacity increases by 30%. Aggregation helps considerably. The aggregated average daily maximum 8 hr up-ramp only increases by 11% though the 8 hr aggregated maximum up-ramp capacity increases by 36%, still a significant savings from the un-aggregated case. If each BA operates independently and acquires its own sufficient ramping capability, there will be much more flexibility in the ground than is actually needed.




Figure 15. There is a significant difference in ramping capability required with a 16% wind energy penetration.




Figure 16. Ramp rate differentials of ramping capability with wind, cooperative planning vs. separate planning.

To help assess the difference in the various cases, we can collect key results from the previous graphs. To create Figure 17 we subtract the aggregated cases from the un-aggregated cases. This provides us with the avoided ramping capability that is needed under the cooperative planning scenario.

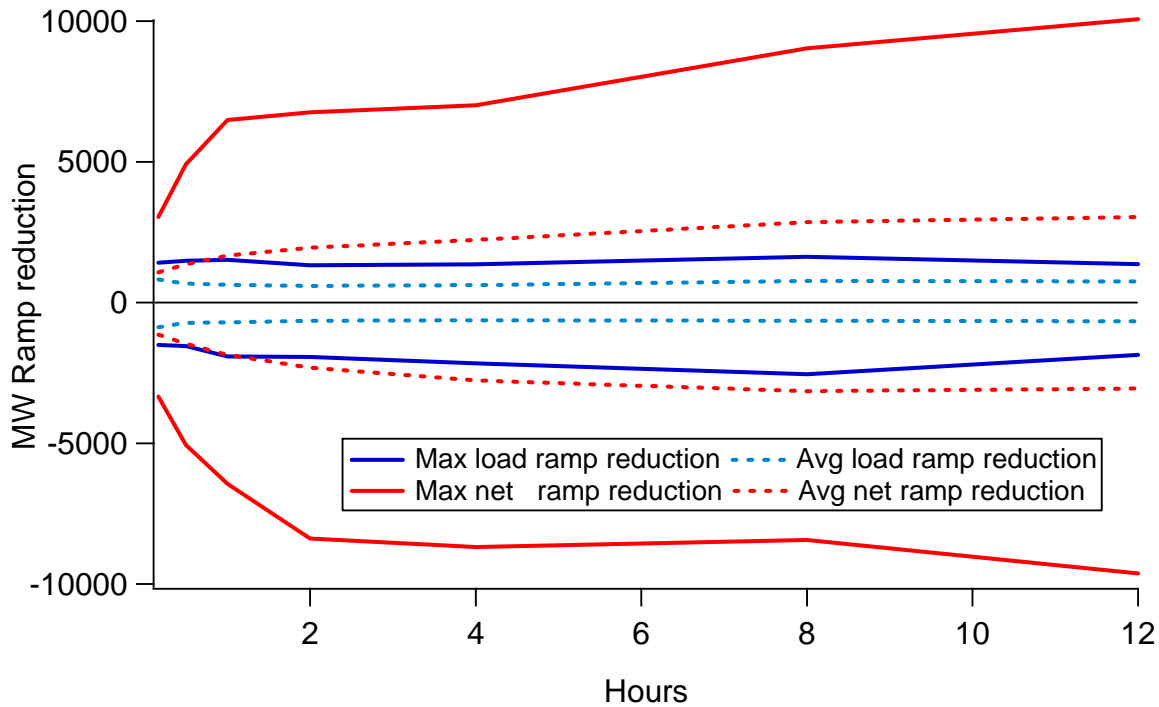


Figure 17. Avoided ramping capability with single-BA planning vs. separate planning

It is clear from the graph that a relatively small number of time periods drive the required ramping capability because of the large difference between the maximum net ramp reduction curves and the average net ramp reduction curves. If each area were to install sufficient ramping capability to meet all of their internal ramping needs, significantly more flexibility will be needed than if the areas operate as a single BA with planning (or markets) providing the required level of ramping. This is further illustrated in Figure 18 which presents duration curves for the ramping requirements. The un-aggregated system is presented in the top two graphs while the aggregated system is in the bottom two. Load requirements are presented in the left two graphs and load-net-wind requirements are in the right two. This somewhat unusual figure is used to facilitate comparing the requirements. The choice of vertical scale is deliberate, also to facilitate comparison. Most significantly, and expectedly, the ramping requirement in all time frames increases with the addition of 16% wind; the right two curves are higher than the left two curves. More to the point of this paper, aggregation reduces the ramping requirements, especially in the 16% wind case (right) but even in the load alone case (left). In fact, the system ramping requirements for the aggregated system with wind are about the same as the requirements for the current un-aggregated system with load alone. One further insight can be drawn from the duration plots. Load alone (left) shows a flatter pattern than load with wind (right) for both the un-aggregated (top) and aggregated (bottom) cases. The maximum ramping requirements with wind occur relatively infrequently. This implies that the ramping reserves used to meet the maximum requirements might be different than those used for the average ramping requirements.

Maximum wind ramps are similar to conventional contingencies in that they are large but infrequent. A reserve similar to supplemental operating reserve but that can sustain response for a longer duration might be appropriate. Its primary characteristic would be that it was inexpensive to hold as a reserve though it could be relatively more expensive to deploy since deployment would be infrequent. Similarly, wind might be curtailed for the most severe and least frequent wind up-ramps if other resources were not available and able to respond at lower cost.

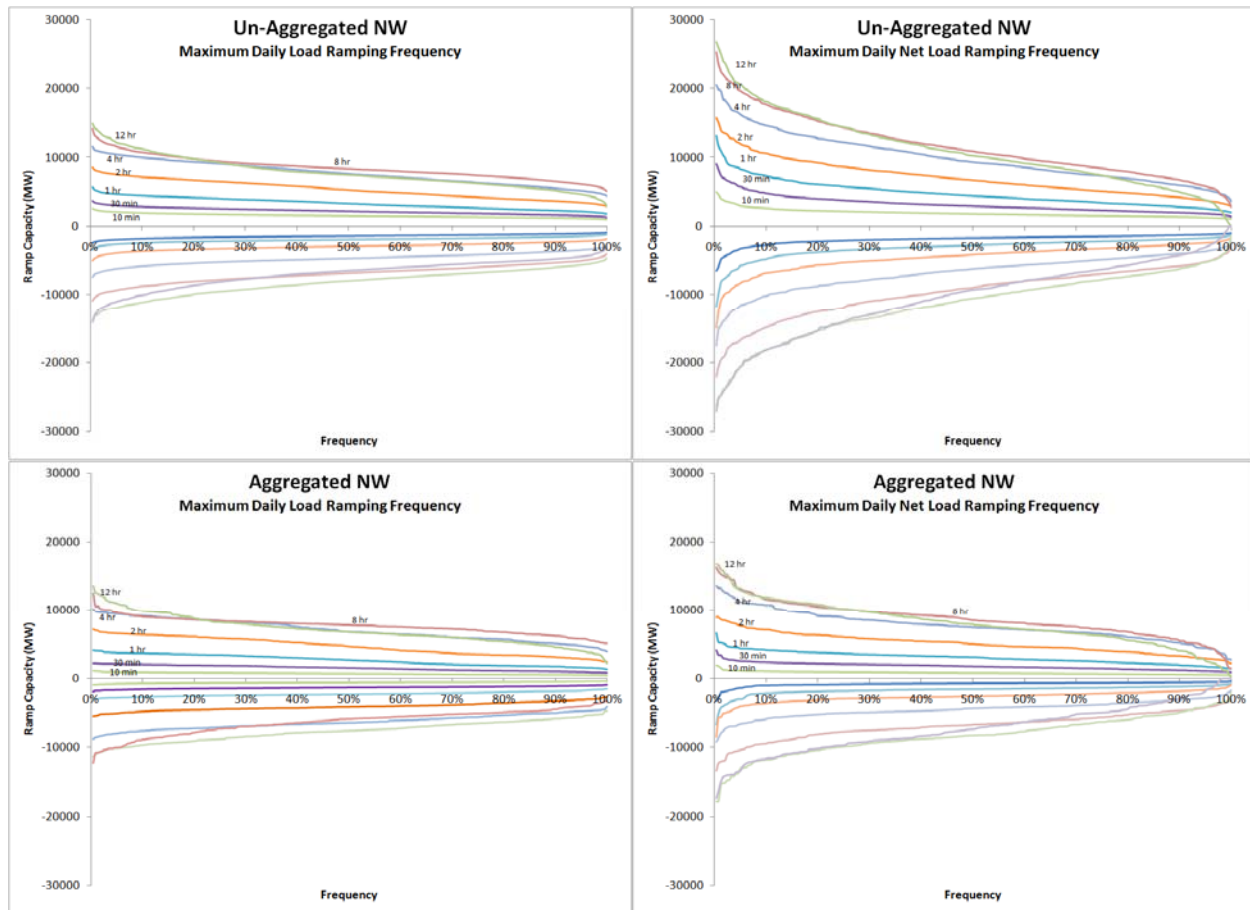


Figure 18 Aggregation always reduces ramping requirements but especially so with 16% wind penetration.

Conclusions

In this paper we illustrate several ways in which BA cooperation can improve the system's ability to integrate wind energy. We focus on a description of an Energy Imbalance Service that has been proposed by WECC's Seams Issues Subcommittee, and illustrate the potential benefits of the EIS or other form of BA cooperation in the Pacific Northwest. We show that the two primary power system characteristics that help with wind integration are balancing area size and fast (5 minutes) economic dispatch and scheduling. Several approaches can be utilized to acquire

these characteristics, including large, fast energy markets (with or without coordinated unit commitment), physical or virtual consolidation, or innovative dynamic scheduling or related measures. Our analysis shows that coordination reduces ramping requirements both operationally and in the planning domain; coordination will result in the need for less flexible generation and lower cost. Further, we find that the approaches discussed in this paper would be efficient even without significant development of wind or other variable generation resources.

References

Milligan, M.; Kirby, B. (2007) “The Impact of Balancing Areas Size, Obligation Sharing, and Ramping Capability on Wind Integration,” presented at WindPower 2007, Houston, TX. Available at <http://www.nrel.gov/docs/fy07osti/41809.pdf>.

Milligan, M.; Kirby, B. (2010). “Market Characteristics for Efficient Integration of Variable Generation in the Western Interconnection.” NREL Technical Report. To appear.

Kirby, B.; Milligan, M. (2010). “Utilizing Load Response for Wind and Solar Integration and Power System Reliability.” Presented at WindPower 2010, Dallas, TX. Available at www.nrel.gov/publications.

Kirby, B., Milligan M. (2009) Capacity Requirements to Support Inter-Balancing Area Wind Delivery. National Renewable Energy Laboratory Technical Report. NREL/TP-550-46274. Available at <http://www.nrel.gov/docs/fy09osti/46274.pdf>.