2.0 The Transactive System

Additional chapter coauthors:
OA Kuchar and C Marinovici − Battelle Memorial Institute

The Pacific Northwest Smart Grid Demonstration (PNWSGD) project featured an innovative transactive system. This chapter discusses the technical performance of that system. Its purpose was to coordinate the dispatch of electric energy with responsive electricity demand in a way that reduced system peaks, reduced costs, and mitigated the challenges from emerging intermittent energy resources like wind. The system partitioned the Pacific Northwest (PNW) power grid into 27 nodes, and these nodes communicated with their nearest neighbors every 5 minutes during the project (1) the delivered cost of electricity and (2) the predicted energy to be exchanged now and during a set of future intervals.

Section 2.1 presents context that the reader might need as system performance is discussed. The project generated many presentations and documents that describe the transactive system. This chapter will not repeat all of the details and concepts from the other presentations and documents.

First, understand that the candidate architectures, advantages, and limitations of transactive systems are under active discussion. The project's system is one example among several candidate system approaches. The Gridwise Architecture Council has become a forum for this technical discussion. So, some of the most general discussion about transactive systems may be found on the Council's Transactive Energy webpage (Gridwise Architecture Council 2015). An important product of its present activities is its Gridwise® Transactive Energy Framework.

For historical context from the Olympic Peninsula Project report that preceded and set the stage for the PNWSGD transactive system, read the GridWise Testbed Projects report by Hammerstrom et al. (2007). Some of the earliest conceptual groundwork specifically for the PNWSGD transactive system may be found in a presentation by Hammerstrom et al. (2009). For publicly available overview presentations about the project's transactive system, refer to Melton and Hammerstrom (2011, 2012, 2014), or Melton (2013).

Perhaps the most detailed discussion about the PNWSGD transactive system design may be found in the Transactive Coordination Signals project report (Battelle Memorial Institute 2013). That report includes much detail about the two classes of transactive signals; the way the project designed and implemented transmission-zone and site nodes; the timing approach used for system signals, including its predictive future intervals; and the functional interfaces between the system and its resources and loads.

A comprehensive list of the technical documents generated by the PNWSGD project is listed in Appendix A. The list includes reports, design specifications, test specifications, and a user guide.

2.1 Context Needed to Discuss Performance of the PNWSGD Transactive System

This subsection presents the context for discussion of the performance of the transactive system that was designed and deployed by the PNWSGD project.

Figure 2.1 is a greatly simplified functional block diagram of a node of the project's transactive system. The large, inclusive block titled "solver/optimizer" represents the algorithmic framework of the calculations that took place at a single system location, a *transactive node*, of the regional transactive system. The project established 14 such transactive nodes (*transmission zones*) to represent large sections of the PNW power grid's transmission and generation, and it defined 13 additional transactive nodes to represent the project's participating utility and university sites. The algorithmic framework at a transactive node was intended to be scalable and self-similar, regardless of the device or group of devices that is being represented by the transactive node.

The main block in Figure 2.1 titled "solver/optimizer" shares a functional responsibility to compute a blended unit cost of energy at this transactive node (marker "3a") and distribute the impact of the blended unit cost through the system. It shares a responsibility to plan for energy balance at the transactive node (marker "6a") and to communicate the impact of that action into the system. Finally, a fundamental responsibility of the block is to accurately balance its energy, including the energy it negotiates to be exchanged with (either imported from or exported to) its transactive neighbors (marker "7").

The transactive node's position within the power system defines its set of *transactive neighbors*. *Transactive neighbors* are the transactive nodes to which the transactive node is electrically connected. A *transactive neighbor* furthermore must be a member of the transactive system, meaning that it has agreed to exchange transactive signals with this transactive node and all of its other transactive neighbors. The blocks at the top of Figure 2.1 represent the transactive node's interface with its transactive neighbors. For the PNWSGD project, these blocks simply implemented an application programming interface. Using extensible markup language, the application programming interface defined the intervals and other contents of the transactive signals and several system management signals. Because the interface was specified the same for all transactive nodes in the project's transactive system, the individual implementations were amenable to conformance testing.

Transactive neighbors necessarily exchanged two paired signals—energy unit cost and energy quantity—with one another. These signals addressed the present and future exchange of energy between the two transactive nodes during a set of future time intervals. The unit-cost-like signal was called the *transactive incentive signal* (TIS, marker "3b") and the energy signal (actually defined as average interval power) was called the *transactive feedback signal* (TFS, marker "6b"). This exchange was bidirectional. Each transactive node was required to both send and receive both signal types to and from each of its transactive neighbors. The project transactive nodes used a common set of 57 sequential future time intervals that ranged in duration from 5 minutes to 1 day. The entire set of intervals predicted cost and quantity for from 3 to 5 days into the future.[1] The signals were exchanged every 5 minutes during the project. More than one signal was sent during some 5-minute intervals if the transactive node determined that its new state differed from the one that was last communicated by an amount that exceeded a defined relaxation criterion.

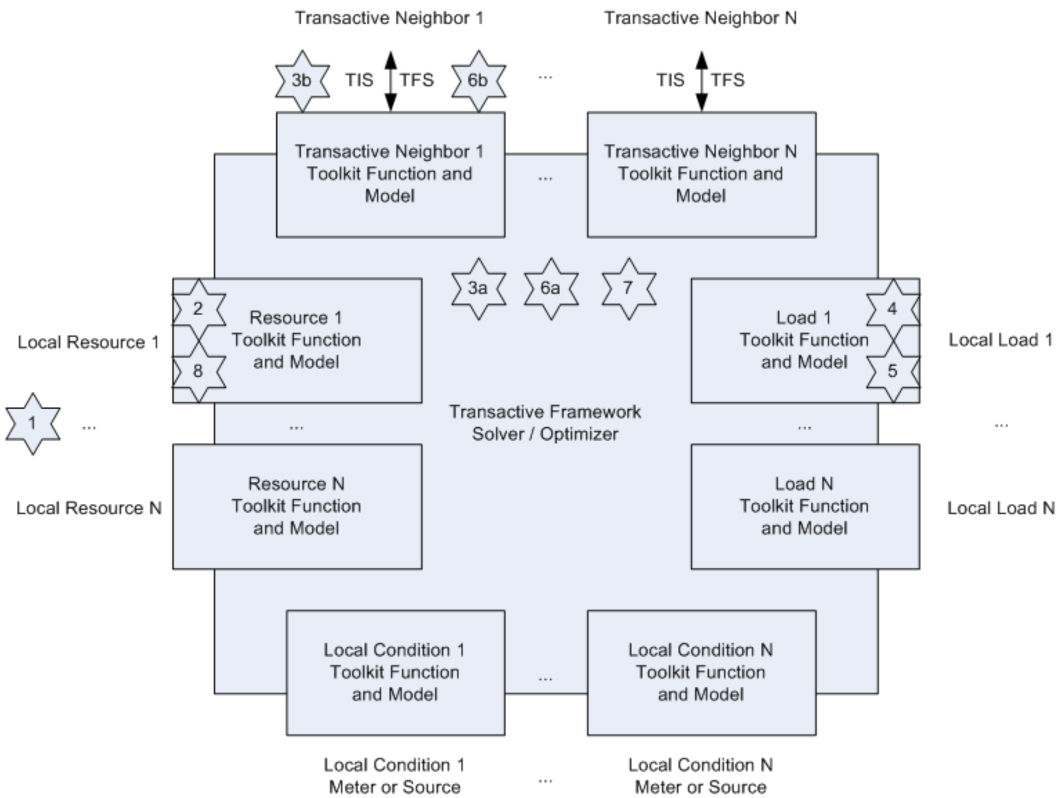

Figure 2.1. Simplified Functional Block Diagram of a Transactive Node. The numbered stars refer to functions that will be referenced as the performance of the transactive system is being discussed in this chapter.

Depending on its sign, the energy exchanged between transactive neighbors is either a resource that is available to, or a load that must be supplied by, the transactive node. A transactive node might also have local electricity resources and an obligation to supply local electric loads.

If a transactive node has its own resident generation resources, then its interface to each generator supply should be represented by a resource toolkit function. These functions are shown at the left side of Figure 2.1. The function represents to the transactive node the energy that is available during each interval and the cost of the available energy. The interface responds to external resources by notifying them if and when they should be dispatched.[2]

[1] The total duration described by the transactive signals varied because of the way the intervals were aligned with 15-minute, hourly, 6-hour, and Pacific Time day boundaries.

[2] The predictive dimension in a decentralized control system like this is perhaps similar to economic unit commitment in today's centralized grid control. The concept of firm future resource commitments may be ccommodated by a decentralized control system, but resources lose some of their value to the system once the commitment becomes finalized.

The theoretical responsibilities and capabilities of resource toolkit function interfaces were greatly simplified for the PNWSGD project. The scale of the demonstration did not allow for the operations of large Pacific Northwest generators to be altered by the project. Instead, the project created a set of informed simulation models that strived to accurately track and predict the dispatch of several of the most important resource types. If the given resource type exists at a transmission zone, a single function represented the aggregate energy from that resource and the calculated wholesale unit cost of that energy.

The following functions were created to model the bulk resources of each type within the transmission zones:

- Hydropower − The Bonneville Power Administration (BPA) helped the project track the dispatch of hydropower generation, and the Dow Jones Mid-Columbia price index was used to emulate the unit cost of the hydropower energy, which closely tracks the costs that are eventually revealed by the recent history of bilateral energy exchanges in the region.
- Wind power − 3TIER, Inc. and BPA helped the project predict and track generated wind power in each transmission zone. The project included the cost of wind power among infrastructure costs, which added a relatively constant offset to the incentive signal at each of the transmission nodes of the transactive system. Thus, the incentive signals, represented as unit costs of electrical energy, decreased when and near where wind power is being produced.
- Thermal power − Alstom Grid and BPA helped the project track the dispatch of thermal resources in the region. The wholesale unit costs were calculated using fuel costs and typical conversion heat rates.
- Transmission power exchanged at the boundaries of the transactive system − BPA helped the project estimate and predict the energy moving across the system's transmission boundaries and the unit costs of this energy. The system was connected to Canada, Montana, Wyoming, Nevada, and California transmissions that were not part of the project's transactive system.

Information about generation dispatch practices, history, and costs was found to be very business sensitive. Access to real-time generation information was sparse and incomplete. Access to accurate historical information from which useful trends might be gleaned was reluctantly made available. The project's knowledge of the region's wind generation was strong, but the project was required to aggregate the information so that accurate information about no single wind site could be gleaned.

The coordinated operations of these resource functions were much more centralized than was hoped for demonstrating a decentralized transactive system. Acting on the project's behalf, Alstom Grid set up and solved economic power flow and economic dispatch for the entire transactive system region. The approach was similar to that used for locational marginal pricing that calculates price differentials over both time and across geographical separation. The solution determined which resources would likely be dispatched in each transmission zone and at what price. The project referred to these aggregate resource functions as an *informed simulation* because they necessarily predicted and emulated the behaviors of the region's generators and system operators from incomplete, dated, and otherwise imperfect available information.

The above discussion addressed the formulation of wholesale energy and its costs. The impacts of incentives that were not directly proportional to energy supply were also represented by functions. The project implemented two such functions. First, because the project strived to represent the TIS as “the delivered cost of energy,” the project applied an infrastructure cost function at transmission zones to represent the remainder of wholesale costs beyond what was already represented by the costs of the generated energy alone. The granularity of the project’s transactive system was too coarse to represent each piece of infrastructure and its cost, but the aggregate impact was estimated from the differences between wholesale electricity prices paid by participating utilities (less than $0.05/kWh) and the aggregate blended cost of energy from the energy resources alone (often less than $0.02/kWh).

Several participating utilities that are supplied by BPA and the University of Washington campus designed and implemented incentive functions to predict and represent the impacts of BPA or Seattle City Light time-of-use price differentials and demand charges on their unit energy costs. These functions effected a price differential on the delivered cost of energy (i.e., the TIS) at the transitions between peak and off-peak hours. They furthermore predicted monthly peak hours and reflected the demand charges that would be incurred as new demand peaks were being encountered.

In summary, the toolkit functions—resource or incentive—have the responsibility to monetize resource costs and incentives (Figure 2.1, marker “2”), and should be responsive to the transactive node’s attempts to balance loads and resources, especially as the system’s loads respond to the TIS (marker “8”). The behaviors of the actual generation resources (marker “1”) must be accurately represented by the toolkit resource functions if the transactive system is to perform well.

Toolkit load functions are shown on the right-hand side of Figure 2.1 at interfaces between the transactive node and its locally served electric loads. Much of the system load is inelastic, unresponsive to any change in the TIS. The inelastic load must be represented and predicted anyway because of its impact on energy balance at the transactive node and throughout the transactive system. At transmission zone transactive nodes, the total forecasted BPA load, which is quite inelastic, was scaled and allocated among the 14 transmission zones. IBM worked with the owners of individual site transactive nodes (i.e., the utility sites) to create and train a function that would accurately predict the inelastic load magnitude at the point where the site electrically connected to the remainder of the transactive system. The only responsibility of an inelastic load function was to accurately predict the energy consumption (Figure 2.1, marker “5”).

More interesting are the elastic loads and their toolkit functions. These functions represented individual, or systems of, electric loads that might change their energy consumption when informed of changes to the TIS. The first responsibility of these toolkit functions is to determine the timing and degree of the elastic loads’ responses based on the TIS and available local conditions (marker “4”). The project found it helpful to categorize the responsive loads as having event-based, daily, or continuous-response capabilities. The differences between the various systems’ responses within each of these categories could often be tailored simply by modifying configuration parameters. The capabilities and limitations of the systems’ responses must be accurately configured if these responses are to also be accurate and meet their owners’ objectives.

Each toolkit load function is also responsible for maintaining a model of its performance from which the energy impact of an elastic response by the load may be estimated and predicted for the transactive system. For example, a toolkit function that represents a thermostatically controlled building might model changes in its consumed heating or cooling energy as a function of thermostat setting, outdoor temperature, building thermal storage, building occupancy, and so on. The modeled change in load must be accurate (marker "5") if its impact is to be recognized and influence its transactive node and the larger transactive system. Asset models were created and implemented for systems of battery energy storage, distributed generators, portals and in-home displays (i.e., voluntary responses), voltage management, thermostatic space conditioning, and electric tank water heaters.

In the prior discussion, the words *resource* and *load* have been used to differentiate the purposes of toolkit resource and load functions. Some may prefer the terms *price-maker* and *price-taker* instead for resource and load functional interfaces, respectively. Indeed, the project modeled distributed generators and renewable generation resources using toolkit *load* functions. The distinction is perhaps that the actions of the systems being represented by toolkit resource functions have their energy production specified during the balance of the system energy, and they compete and influence the system by affecting blended costs in the system. The systems represented by toolkit load functions receive unit cost information and compete based on their flexibility and ability to modify the net electric load at the transactive node. With this understanding, the distinction of *source* (as generation) versus *load* (as energy consumption) becomes less important. It is entirely possible that a more complex asset system may be represented by either or both resource and load toolkit functions, as conditions dictate. The separation of price-maker and price-takers' responsibilities may be an important construct for the architectures of distributed energy systems.

The last interfaces remaining to be introduced are the transactive node's interfaces to local conditions, as shown at the bottom of Figure 2.1. The above-mentioned functional interfaces to loads and resources may be influenced by local conditions. For example, the prediction of inelastic load is usually dependent upon local ambient temperatures. The individual toolkit load or resource functions may individually procure access to such information, but the system may be simplified if frequently needed information, like local ambient temperature predictions, is available from a single interface between the transactive node and the sources of such information. Information sources may be simple meters, systems of meters (e.g., occupancy sensor systems), or accessible Web services, for example.

The following eight markers in Figure 2.1 point to specific functions that were necessarily well implemented and accurate for the transactive system to have achieved and demonstrated useful outcomes. These functions will be referenced as the performance of the project's transactive system is reviewed in the remainder of this chapter:

1. The system must accurately represent the region's strategies for the dispatch of its energy resources.
2. The system must meaningfully monetize resource costs and incentives.
3. Energy costs and incentives must be blended and distributed throughout the transactive system.
4. The responsive loads in the system must be able to allocate their responses and events, based on the incentive signal and local conditions.
5. Responsive loads must accurately predict the energy impacts of their responses.

6. The exchanges of power with the system must be predicted and communicated throughout the transactive system.
7. Plans to exchange energy with the transactive system must be accurate.
8. Supply resources must respond to planned energy exchanges to the degree that the exchanges dynamically affect system balance.

2.2 Step 1: The System Must Accurately Represent the Region's Strategies for the Dispatch of its Energy Resources

Additional section coauthors: SF Joseph and D Watkins − Bonneville Power Administration

The project investigated whether the transactive system accurately reproduced the mix of resources and other grid conditions that actually transpired in the Pacific Northwest region that was modeled by the transactive system. The results of that investigation are reported in this subsection. If grid conditions were accurately represented, then there is a chance that the incentives generated by the transactive system, which were driven by the resource mix and grid conditions (as to be described in Section 2.3), were meaningful and useful. Otherwise, the system diverged from and misrepresented actual power grid conditions. Incentive signals based on erroneous resource mixes and incorrect grid status would unlikely prove to be meaningful or useful.

The transactive system's data-collection layer kept track of its modeled energy resources in several broad categories—hydropower, wind power, thermal power, and power that is either imported into or exported from the region to locations outside the transactive system. A matrix manipulation was devised to also decompose the power being exchanged between the transactive nodes into these four listed categories. Therefore, the project can reproduce a precise accounting of the modeled energy resource mix at each system node. These findings will be reported in this section for a full project year term and by season. The region-level result may be compared against data from BPA, but the comparison is not perfect because the modeled transactive region differs from BPA's.

The project worked closely with BPA to analyze whether not only the static mix of resources, but also the dynamic dispatch of resources and events, matched what BPA reports to have actually happened in the grid. The project determined to conduct this evaluation by comparing the project's transactive system data with data from the BPA's transmission system during exemplary project days. The days were selected by BPA because they represented times when the power grid might have become stressed by extreme weather conditions, generation outages, wind incidents, or transmission incidents. Altogether, seven such scenarios were identified for this investigation,[1] as follows:

[1] The analysis in Section 2.2 was conducted collaboratively by SF Joseph, S Kerns and D Watkins (BPA) and DJ Hammerstrom (Battelle). Much of the text and most of the figures in this section were adapted from unpublished presentation materials that were prepared for and presented at a Project Review Board Meeting, BPA Headquarters, Portland, Oregon on June 5, 2014.

- winter and summer peaks[1]
  - peak winter load event on December 5, 2013
  - peak summer load event on August 5, 2013
- generator outage
  - outage at the Columbia Generating Station on February 5, 2014
- wind incidents
  - rapid wind ramp event on February 15, 2014
  - periods of wind undergeneration and overgeneration on March 5, 2014
- transmission incidents
  - transmission outage on April 1, 2014
  - overloaded flowgate event on April 11, 2014.

Historical BPA loads and resources are reported for hourly intervals by BPA (2015) using substantially the same four resource categories as used in the transactive system data collection. The project did not track any distinction between federal and non-federal hydropower resources, so the federal and non-federal hydropower magnitudes from BPA data were combined in the figures of this section.

The total electric loads will be consistently shown as negative quantities among the diagrams in this section. This practice facilitates visual confirmation that system power is balanced—that all resources and load are being shown. Wherever a visual comparison is being invited between BPA and project data sets, the scales of the figures' power axes were forced to be identical.

The PNWSGD total load was necessarily inferred. The project failed to capture in its data-collection system layer the total system load and allocated node site loads. Therefore, the total system load was necessarily calculated from the sum of all the modeled resources in the region, including the energy imported or exported through the modeled exchange boundaries to entities outside the transactive system.

Finally, the grid region of the PNWSGD transactive system is even larger than that covered by BPA. BPA operates impressive hydropower resource and manages much of the transmission system in the Pacific Northwest. But there are other balancing authorities with generation and transmission assets in the project region. Consequently, total project load should be somewhat greater than that in the BPA data, and the project's resources in the summed categories may be greater as well.

[1] The weeks leading up to winter and summer peak days were also selected for simulation studies (Section 2.10).

2.2.1 Generation Mixes Modeled by the Transactive System

The data-collection layer of the PNWSGD transactive system allowed analysts to reconstruct the mix of energy resource types that were modeled to have been used in the region. The accounting of these resources was quite naturally accomplished in the transactive system because the incentive signal was formulated, in part, by the blended unit costs of these resources. The history of resource usage at each location would not normally be broadly shared between the owners of system nodes, but the information was centrally captured by the project for this research.

Five broad types of energy resources were tracked at each transmission-zone node of the system: hydropower, thermal, wind, imports from outside the transactive region, and the energy received from neighboring transactive nodes. Each of these resources contributed at each node to the node's blended incentive signal in proportion to the magnitude of energy received from that resource. The sum generated and imported energies were then eligible to be consumed or exported at the blended unit cost.

While the power imported from a node's transactive neighbor is a useful magnitude in the distributed system, the component is not itself informative about the resource types that it includes. Fortunately, the project collected complete information of every transactive neighbor's resources. Therefore, a matrix operation was developed to decompose the imported transactive energy components into the remaining four resources—hydropower, thermal, wind, and the imports from the boundaries of the transactive region.

The transactive system had no knowledge of the resources that compose the energies imported at the boundaries of the transactive region, so that component cannot be further decomposed. The project referred to these imports at the region's exchange boundaries as "non-transitive" imported energies.

Figure 2.2 compares the average relative resource mixes of the transactive system before and after the matrix operation that decomposed imported energy from transactive neighbors into its component parts. The data from a complete project year (September 2013 through August 2014) were used. The right-hand figure is the project's best representation of the resource mix that was modeled by the transactive system throughout the last year of the project.

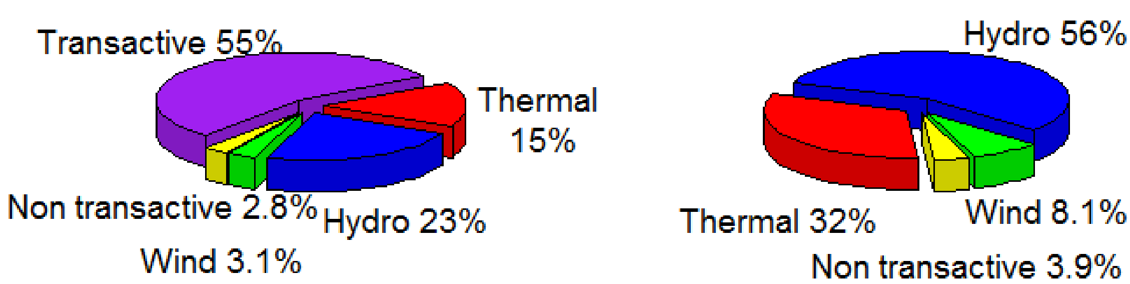

(a) Before Reallocating Exchanged Resources (b) After Reallocating Exchanged Resources

Figure 2.2. Composition of Modeled Resources of the Entire Transactive System in the Last Full Project Year (a) Before and (b) After the Energy that Was Exchanged between Transactive Nodes (the "Transactive" Component) was Reallocated

For comparison and using the same 1-year term, the averaged BPA resource mix (compiled from data on the BPA transmission webpage [2015]) included 67% hydropower, 23% thermal generation, and 9.7% wind. While direct comparison is not possible because the project region differs from that of BPA, the comparison is informative. The project's usage of hydropower is less than BPA's because the transactive region extended outside the Columbia River basin that is the source of the abundant hydropower in the Pacific Northwest. The difference is made up for by using additional thermal resources that become more prevalent toward the south and southeast boundaries of the transactive region. The wind resource percentages are comparable between the transactive and BPA data.

While the non-transactive component should be comparable to BPA's "exchange" component, the term is not accounted for similarly in BPA data and for the transactive system. BPA is almost always a net energy exporter. Project analysts did not have ready access to the individual BPA exchanges, some of which would at times *import* energy. The transactive system, on the other hand, counted imported exchange energy as a resource, even if the entire transactive region might have been a net energy exporter at the time.[1] For these reasons, the BPA data offered for comparisons in this section will not show an exchange resource component, but the transactive system will show a non-transactive exchange component.

The project reformulated the comparison by season in Figure 2.3. The relative resource mixes of the transactive system and BPA system are compared side by side for the four seasons of the last full year of the PNWSGD project (September 2013 through August 2014). All of the limitations of the comparison that were discussed in the prior paragraphs apply to these seasonal comparisons, too.

By season, relative hydropower and thermal energies rose and fell in the transactive system model much as in the BPA system. Every season, the BPA system used a relatively greater percentage of hydropower and relatively smaller percentage of thermal resource than the transactive system did.

[1] Similarly, little or no distinction is made by the transactive system between electric load and the exportation of electrical energy.

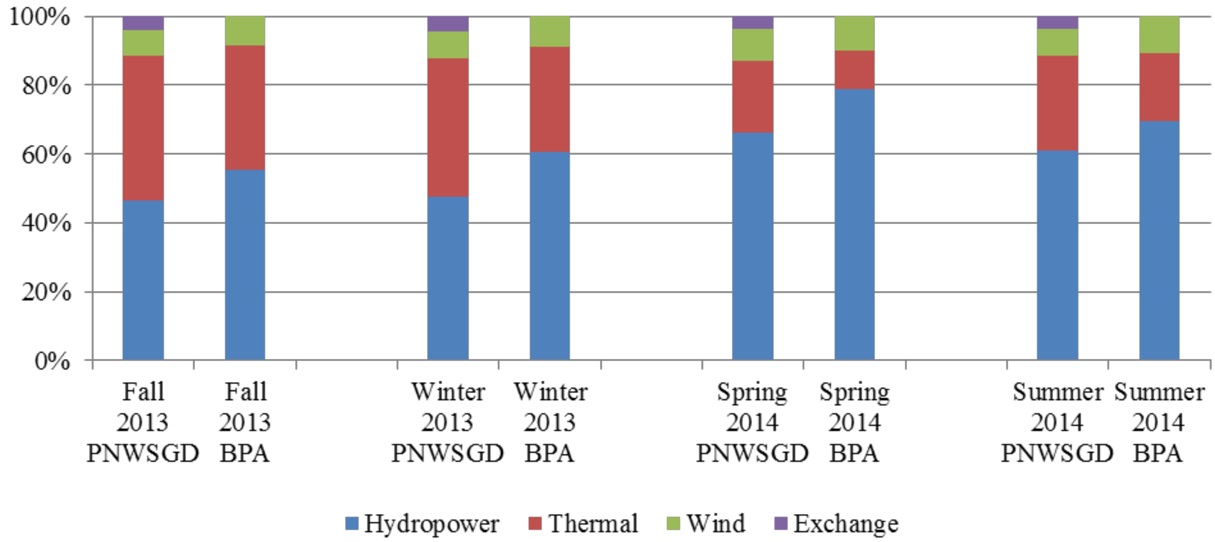

Figure 2.3. Comparison of Average Relative Resource Mix that Was Modeled by the Transactive System and the Mix from BPA Data for the Same Four Seasons

Figure 2.4 compares the modeled resource mix with that of BPA data by hour of day. These data sets both cover the time period from September 2014 through August 2014. As before, the imported exchange energy does not appear in the BPA data because the source for the data did not separate imported energy from exported. This omission will cause the small percentage of imported exchange resources to have been distributed among the other resources in the relative resource mix of the BPA data.

According to this figure, the transactive system relied primarily on thermal resources to balance diurnal load, while the BPA system relied primarily on hydropower resources to do so. This statement follows from the swell of one or another of the resource types especially in late morning when peak load often occurs.

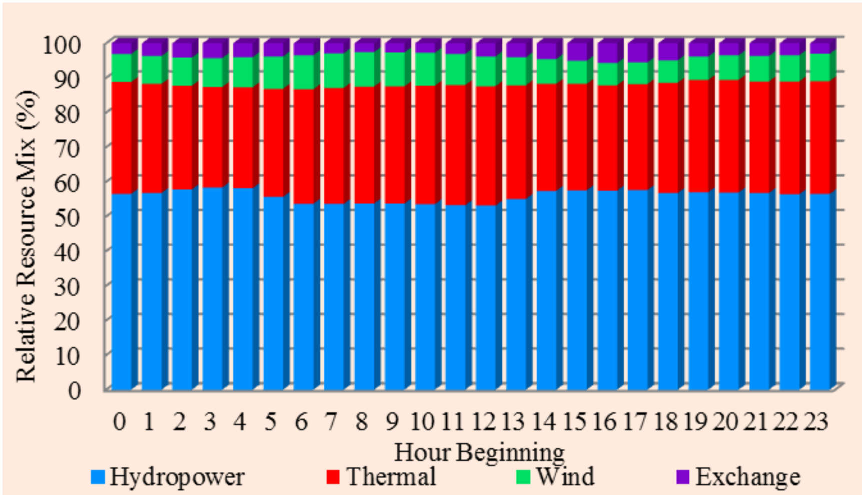

(a) PNWSGD Transactive System

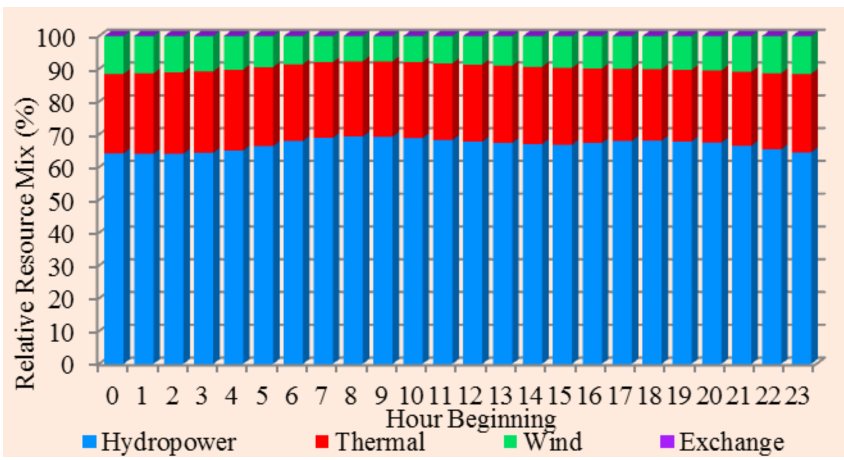

(b) BPA Data

Figure 2.4. Average Relative Resource Mixes (a) Modeled by the PNWSGD Transactive System and (b) According to BPA Data (BPA 2015) from September 2013 through August 2015

Figure 2.5 provides interesting insights into the variability in the transactive system's resource usage according to site locations. Pie charts have been displayed for each of the 14 transmission-zone nodes that were modeled by the project (Appendix B). Data were included from the entire final year of the PNWSGD from September 2013 through August 2014. The pie charts have been approximately positioned at their sites' relative geographical locations among the five Pacific Northwest states that had representation in the PNWSGD. The relative mix of especially hydropower, thermal generation, and wind power are shown to vary according to the local resources at each location. For example, the region's largest hydropower resources reside in the modeled Northcentral Washington zone, where a great fraction of hydropower resource is shown. The Northeast Oregon node includes impressive Columbia Gorge wind resources. The importation of non-transactive exchange energy is most evident in the northernmost and northeast zones that frequently import energy from Canada and eastern Montana.

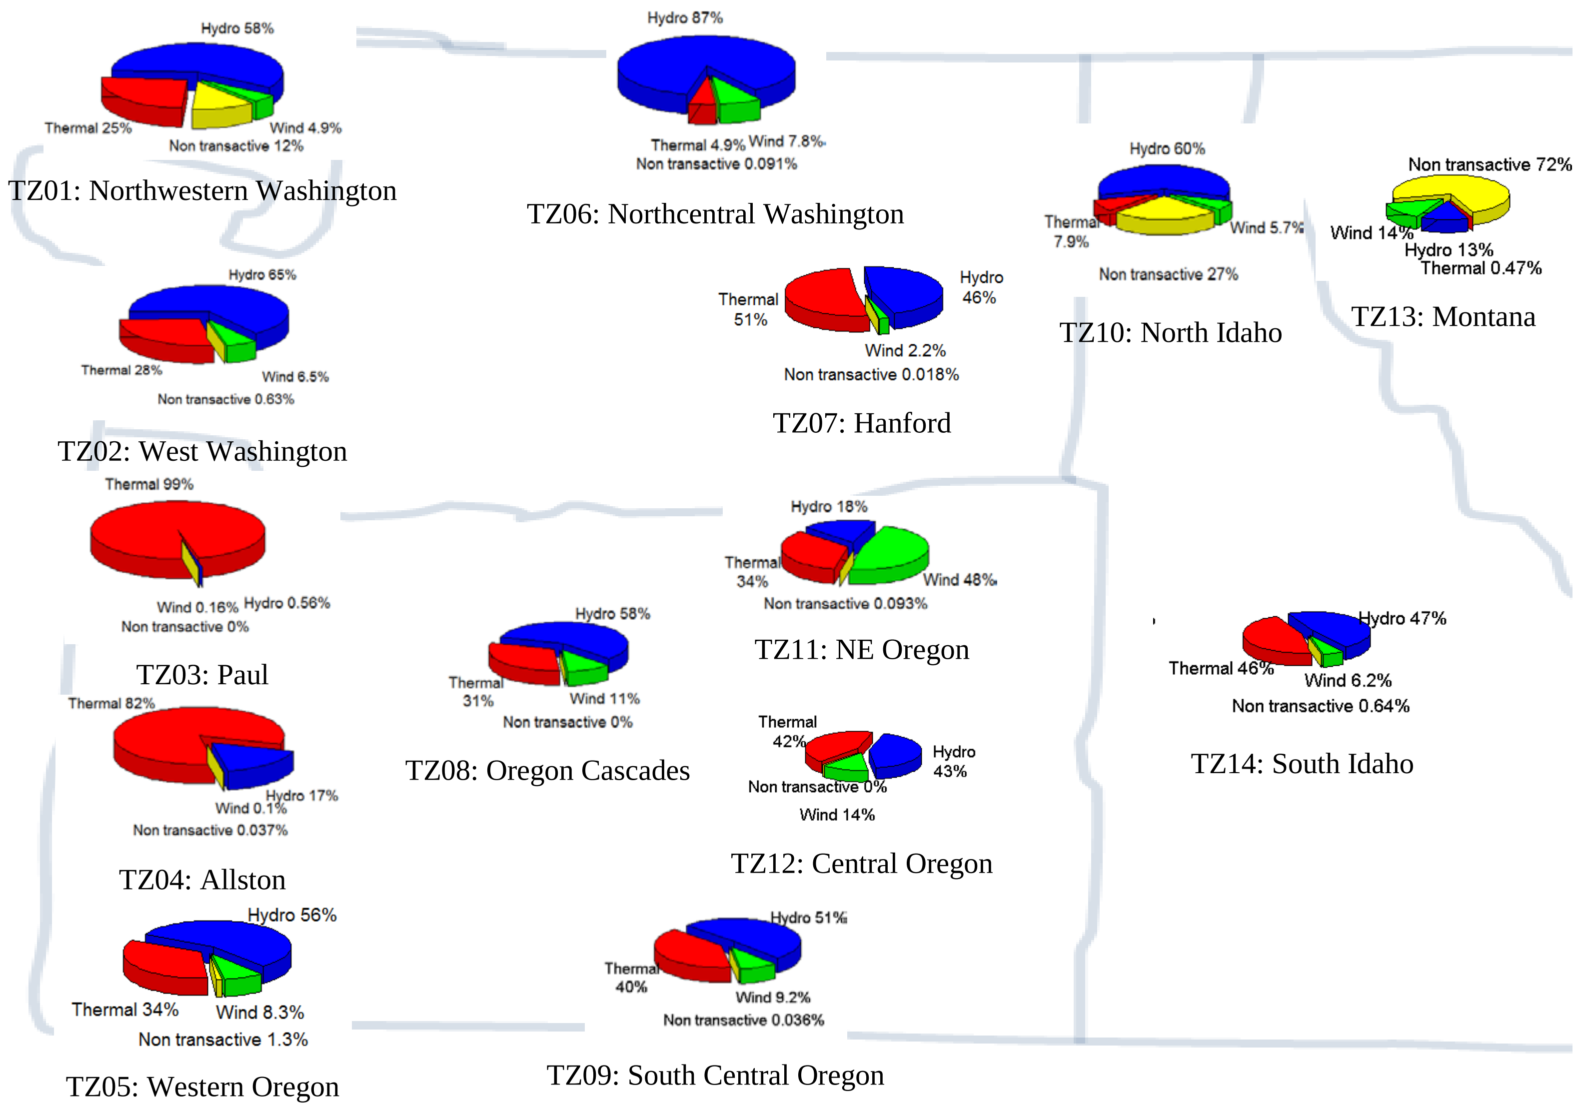

Figure 2.5. Average Relative Mix of Generation Resources Available at Each Transmission-Zone Node during the Last Full Year of the PNWSGD. Each node's pie chart has been placed near its approximate geographical location in the Pacific Northwest.

The sections that follow will investigate the dynamics of transactive systems.

2.2.2 Winter and Summer Peaks

The days of peak winter and summer demand in 2013 were selected by BPA for evaluation, based on the peak total load that it served. These scenarios might be expected to stress the power system as it strives to supply the year's greatest heating and cooling loads.

Peak Winter Load on December 5, 2013. A winter cold snap occurred in the region on December 4–10, 2013. On December 5, morning peak generation by the federal hydropower system that is managed by BPA reached almost 11.5 GW. Because of the cold weather that day, BPA needed to purchase 18 GWh and had little surplus energy to sell. BPA experienced its peak winter load during the hours ending 07:00−09:00.

The BPA and project total generation and load data are compared for this day in Figure 2.6. The components being compared include total hydropower, total thermal generation, total wind generation, and total net exchange power in the BPA and modeled transactive systems. Unlike Section 2.2.2, the total net exchange powers include the sums of all imported and exported exchange and are therefore fairly compared. Exported exchange power is shown as a negative value, as is total system load.

The transactive system modeled a considerably larger peak load than the BPA system on this day. The PNWSGD peak load occurred in the afternoon hours 17:00−19:00, not in the morning. Wind is minimal on this day in both the BPA and project system representations. The magnitudes of net exchange are also similar in the two representations, but there appears to be disagreement concerning the patterns of the increased and decreased exchange power during the day.[1]

Hydropower and other resources are managed differently in the BPA and project representations on this day. First, the transactive system relies more heavily than the BPA system on thermal resources during this cold snap. At midday, there is more than three times as much thermal generation modeled in the transactive system as in the BPA one. Some of the difference may be attributed to the transactive system's extension west beyond the Columbia River hydropower basin. Thermal resources were almost constant in the BPA system, mostly unaffected by BPA system load magnitude, but the transactive system changed its dispatch of its thermal resources during the day.

[1] Exchange power was made available to the project only as typical seasonal trends, leaving the designers of the transactive system model to infer the ways the exchange would be managed. The mismatch between actual and modeled exchange powers means that the strategy was not inferred well in this case.

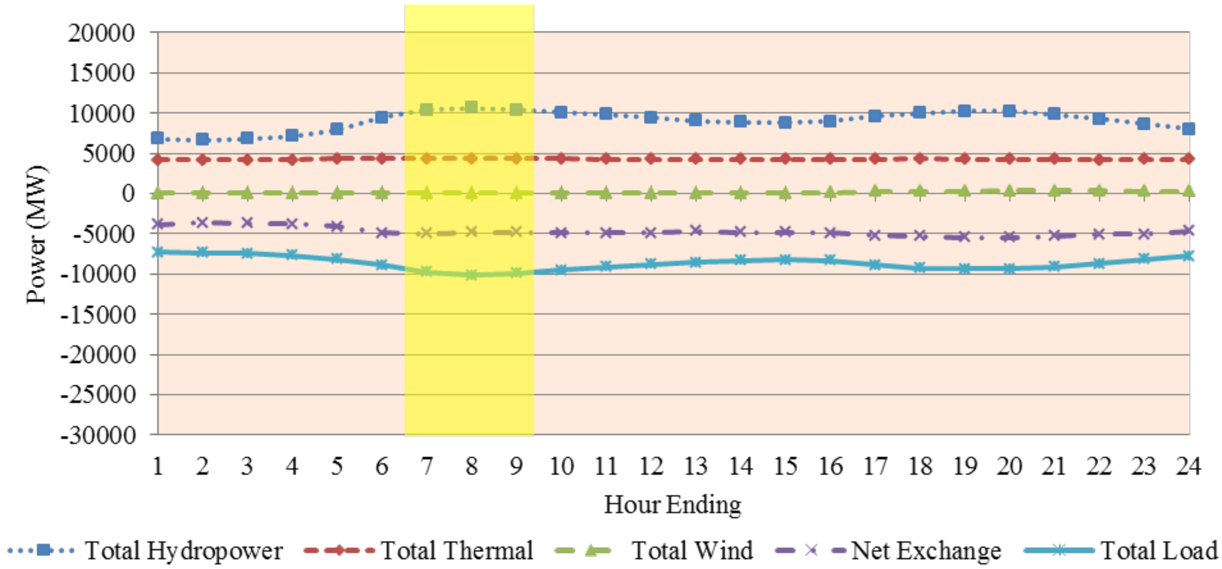

(a) BPA Data (BPA 2015)

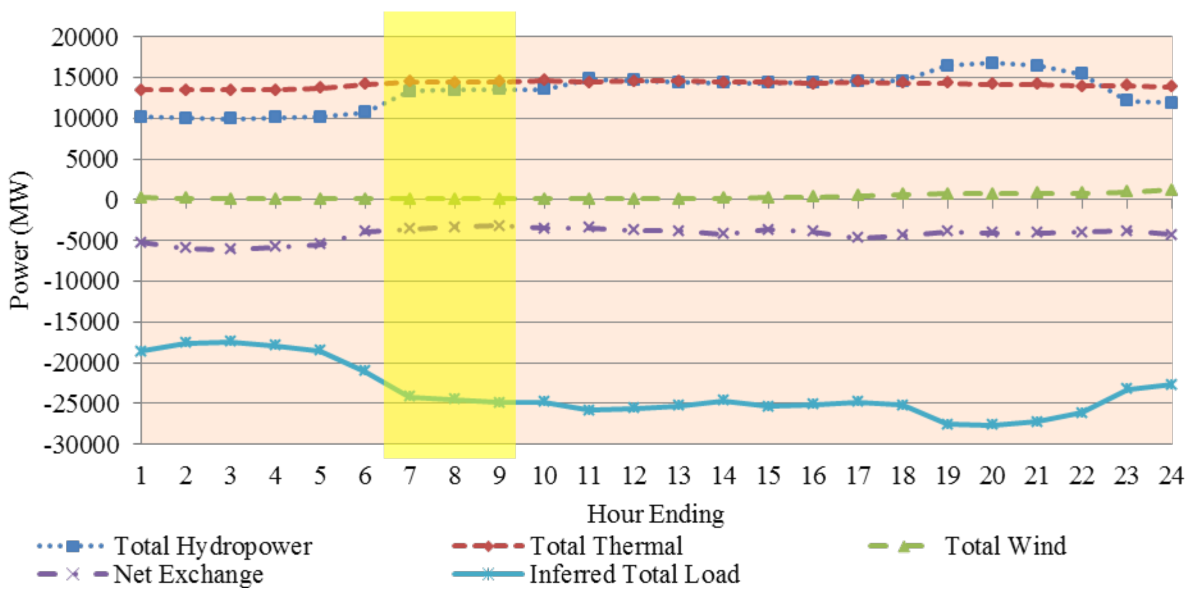

(b) PNWSGD Transactive System Data

Figure 2.6. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System Model on December 5, 2013

A peak summer load event on August 5, 2013. The peak 2013 summer load occurred on the BPA system August 5, 2013, between hours ending 15–19.[1] The total BPA resources and load are compared against those of the project's transactive system on that day in Figure 2.7.

The transactive system data agreed with the BPA data that the peak total load occurred in the afternoon. However, the project modeled about 4 times as much total load as in the BPA data. The transactive system's total load was designed to be a scaled version of BPA system load. The patterns for the various resources were similar through the day for the compared systems. However, the transactive system required more of each resource type to balance the much greater system load.

The biggest difference between the transactive system and BPA data was in the strategies that were followed to dispatch thermal resources. The transactive system used thermal resources more than the BPA system to follow the diurnal load pattern. The transactive system more than doubled its thermal resources to supply the afternoon peak load.

[1] BPA uses "hours ending." Hours ending 15−19 covers the time period 14:00–19:00.

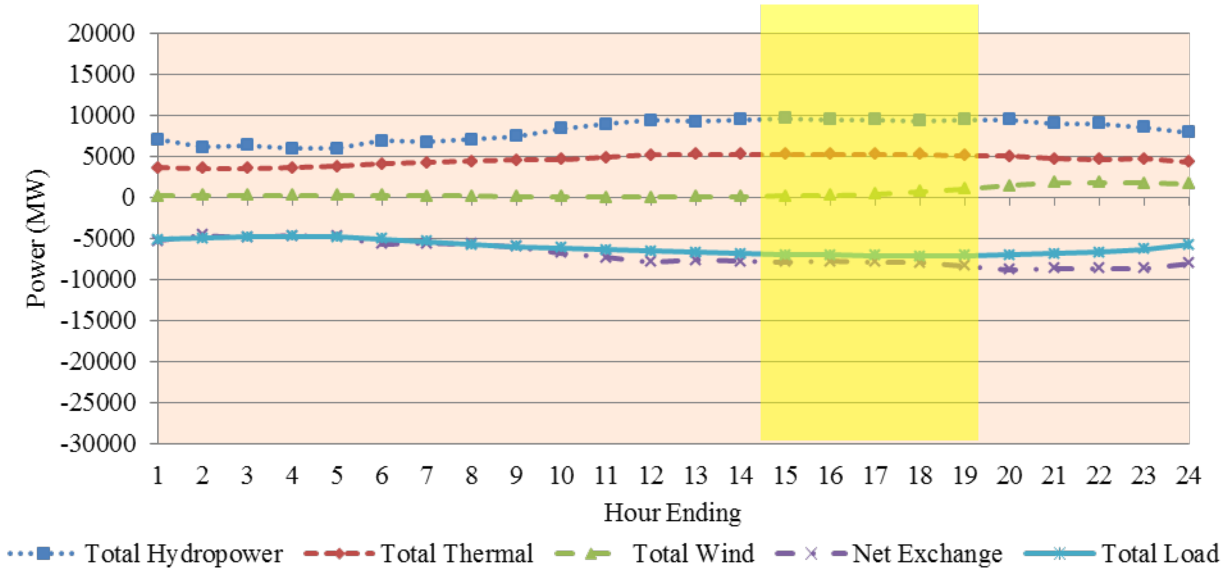

(a) BPA Data (BPA 2015)

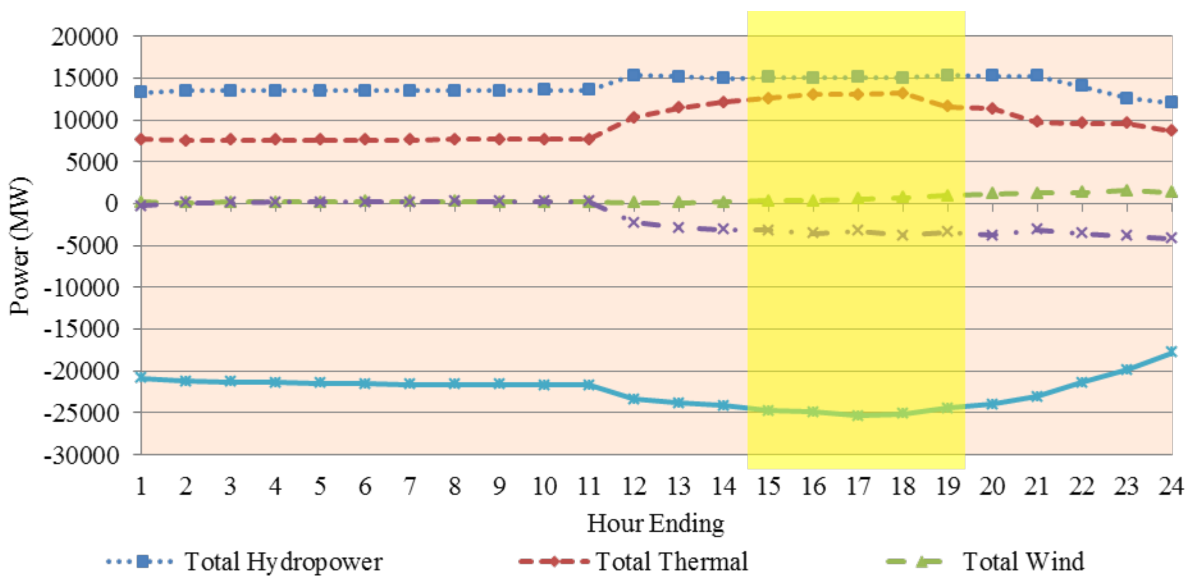

(b) PNWSGD Transactive System Data

Figure 2.7. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System Model on August 5, 2013

2.2.3 Generator Outage

The following event was selected to determine whether the PNWSGD transactive system accurately tracked a generator outage in the BPA system.

Outage at Columbia Generating Station on February 5, 2014. Columbia Generating Station is the Pacific Northwest region's only nuclear power generator. At hour ending 17 (16:00) on February 5, 2014, the Columbia Generating Station went into single-loop operation because of a recirculation pump trip that was traced to an electrical circuit breaker malfunction. Its normal average generation is 1,128 MW, but generation dropped during the outage to 477 MW, less than half of its normal generating capacity.

The BPA and transactive system data from this day are compared in Figure 2.8. Nuclear power generation was grouped with other thermal resources by the transactive system model. While the generator outage in the BPA system (i.e., a loss of ~0.5 GW) was substantial, the change was quite small at this figure's scale, even in among the BPA thermal generation data that was certain to have represented the outage.

The total modeled thermal resources in the transactive system were about 3 times as great as in the BPA data throughout this day.

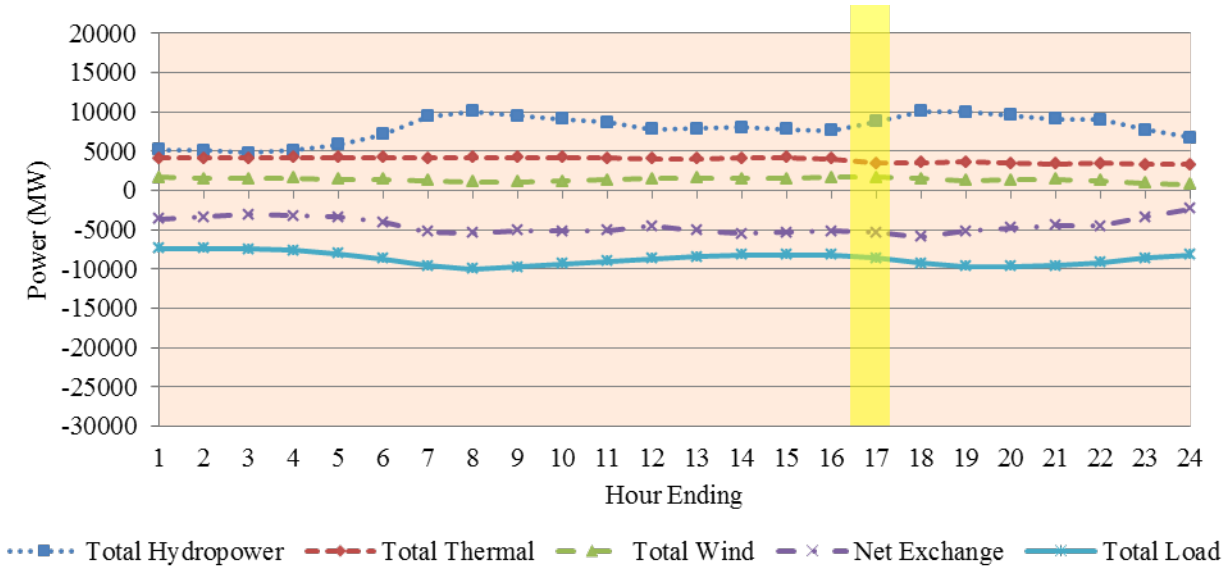

(a) BPA Data (BPA 2015)

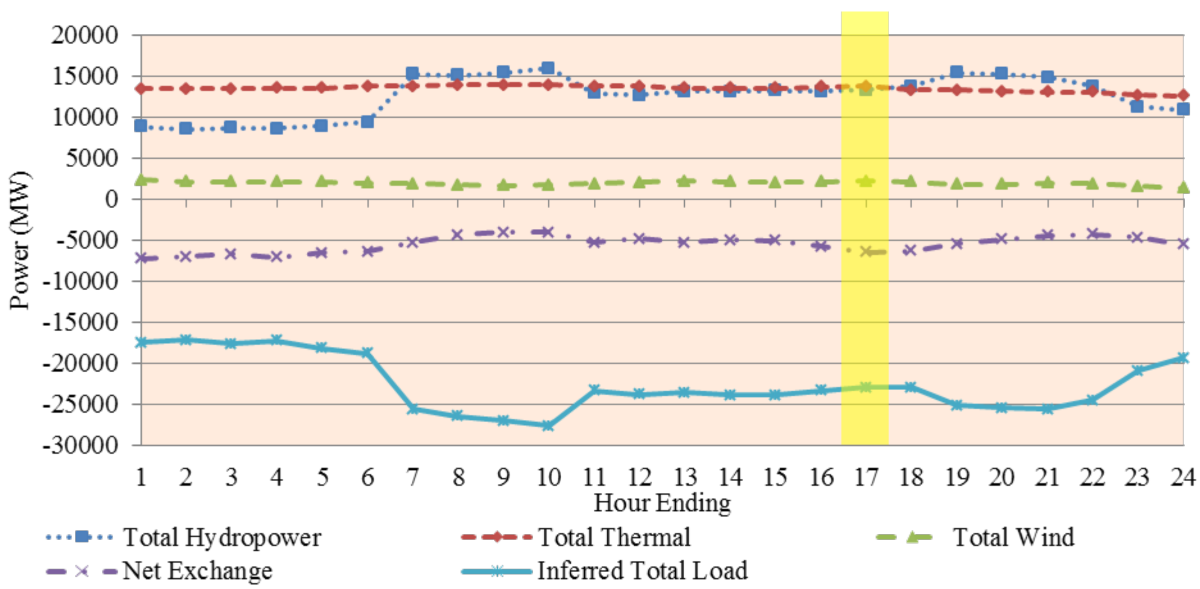

(b) PNWSGD Transactive System Data

Figure 2.8. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System on February 5, 2014

Figure 2.9 focuses on only the thermal generation resources. BPA's thermal generation data in this figure is the same as what was shown in Figure 2.8. At this improved scale, the impact of the outage is evident hour ending 17. The thermal power generation remains reduced at this level for the remainder of the day. The figure also shows both the transactive system's total thermal power generation and the thermal generation in the Hanford transmission zone (TZ07), in which the Columbia Generating Station is located. About 0.6 GW of thermal generation was dropped that hour according to the transactive system model. However, the impact appears after a 1-hour delay in the transactive system data. The source of this delay was not determined.

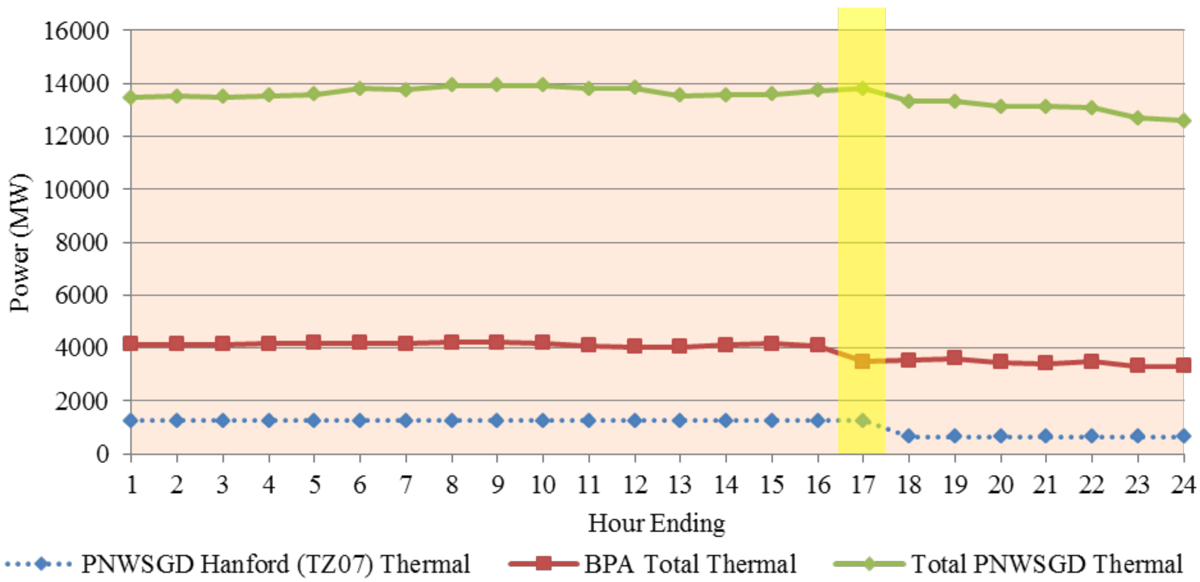

Figure 2.9. Comparison of BPA's and PNWSGD Transactive System's Thermal Generation Data on February 5, 2014, when a Significant Thermal Generator Outage Occurred

With the exception of exchange power, modeled resources in the PNWSGD transactive system are proportional to those in BPA data. The consistency of this proportionality through the day may be seen in Figure 2.10. Here, the resource and load data from Figure 2.9 has been expressed as the modeled transactive system data divided by the BPA data that represents the same resource or load. The transactive system did not curtail as much thermal load as in the BPA system upon the Hour-17 generator outage. Otherwise the dispatch strategies were similar.

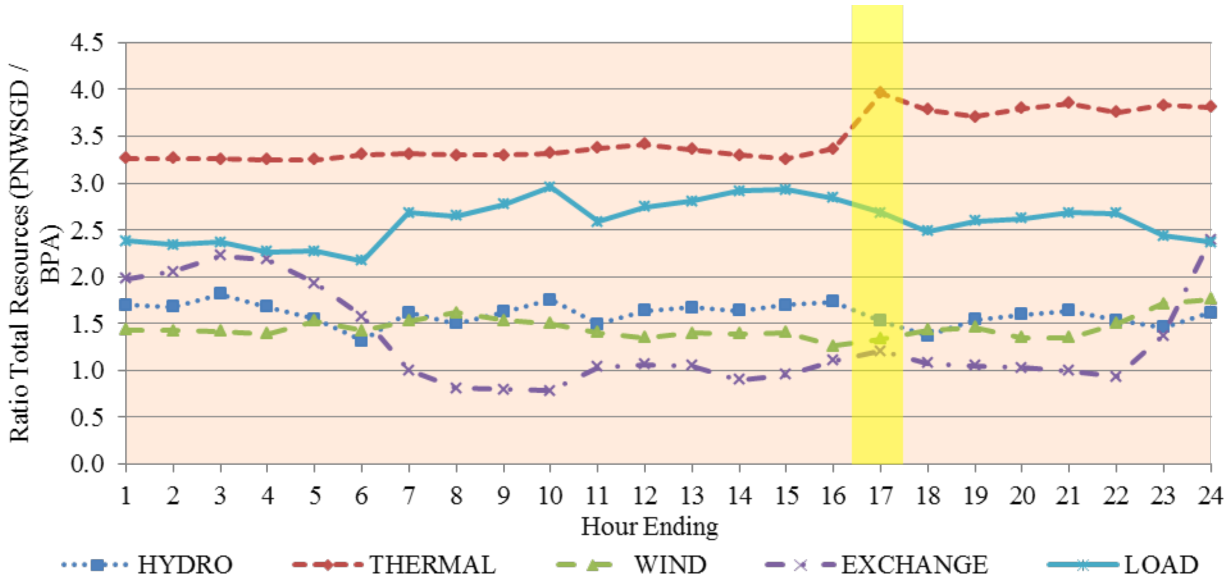

Figure 2.10. Ratio of Transactive System and BPA Total Resource and Load Powers on February 5, 2014. The Columbia Generating Station outage happened during hour ending 17.

BPA described its actions on this day as follows: BPA had forecasted a need for additional power because February 4, 2014, was the first day of what was anticipated to be a 3-day cold snap. BPA purchased over 31 GWh and sold over 3.2 GWh on this day. The purchases were to meet its balancing obligations. The generator outage did not trigger any significant change in its energy purchasing strategy that day.

Figure 2.11 presents the TISs for the entire project region and for select transmission zones at or near the Hanford TZ07 where the affected generator was modeled to reside. The incentive signal should have been affected hour ending 17 when the outage occurred. An increase in the TIS incentive is observed at the TZ07 Hanford zone and at two of the three zones that are attached to the Hanford zone. The change was about \$0.01/kWh. The effect on the overall average regional incentive was quite small, but the changes to the nearby transmission zones' incentive were in a direction that would help mitigate the outage. That is, the neighbors that receive energy from the Hanford transmission zone incurred a cost increase that might have reduced load and helped mitigate the generator outage.

The response in the incentive signal is delayed an hour. The source of this delay was not determined.

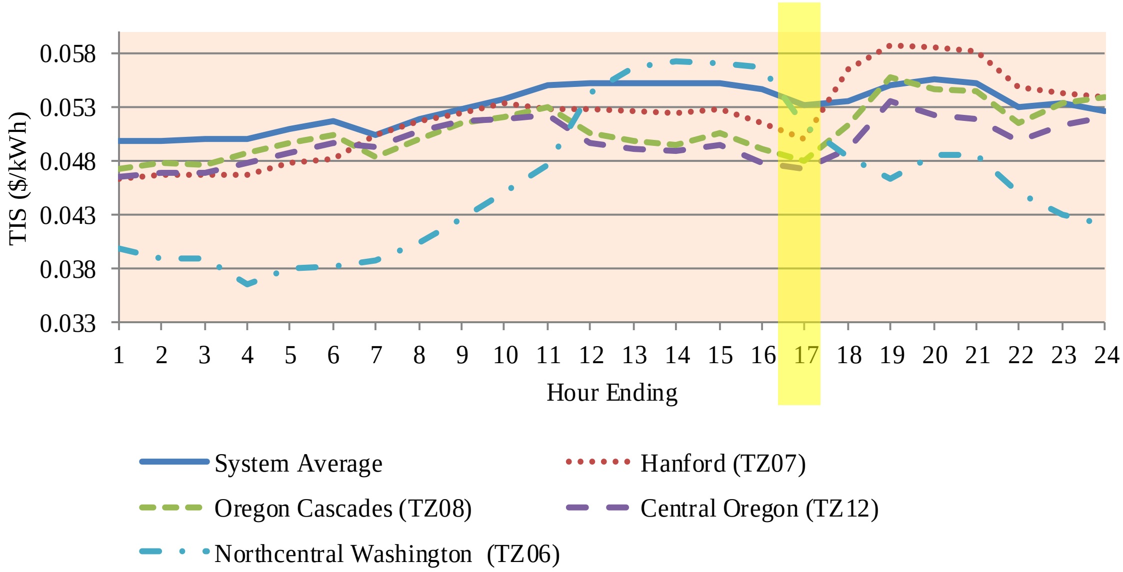

Figure 2.11. Average Transactive System TIS and for Selected TZ Nodes February 5, 2014, when a Significant Thermal Generator Outage Occurred

2.2.4 Wind Incidents

Pacific Northwest renewable wind resources have grown fast. The region is challenged to integrate the growing intermittent resource. This subsection evaluates how accurately the project's transactive system modeled its energy resources, including wind, during rapid ramping of wind energy and at times that BPA reported its wind resource predictions to have been inaccurate.

Rapid wind ramp event on February 15, 2014. Wind plant limitation orders were sent out by BPA between 20:10 and 20:15 Pacific Time on February 15, 2014, when wind generation peaked at 2,884 MW. This peak triggered a fleet level limit order (DSO216) to deploy balancing reserves once the peak pushed balancing reserve levels beyond −995 MW, or 90% of the available "dec" reserves. No further mitigation was needed to recover from the temporary oversupply.

Overall system generation was at a shortfall because the Columbia Generating Station was operating at only 25% capacity, because of scheduled maintenance. At first glance, it appears contradictory that a wind overgeneration incident can occur on a day that there is a generation shortfall.

Refer to Figure 2.12, which compares BPA data and project data for this day. A new data series—total wind forecast—was added to the BPA data to help explain the seemingly contradictory conditions. During hour ending 22, the generated wind exceeded the forecast wind. Because more wind occurred than had been forecasted during these hours of rapidly increasing wind resource, BPA had to call on the types of reserves that can either reduce overall generation or increase system load.[1] As the reserves are

[1] These are referred to as "dec" resources.

dispatched, the pool of remaining reserves of this type decreases. As the available reserves diminish to certain threshold values, the balancing authority must take actions to maintain system balance and reestablish the depleted reserves. Emergency actions can include the curtailment of wind resources in the region.

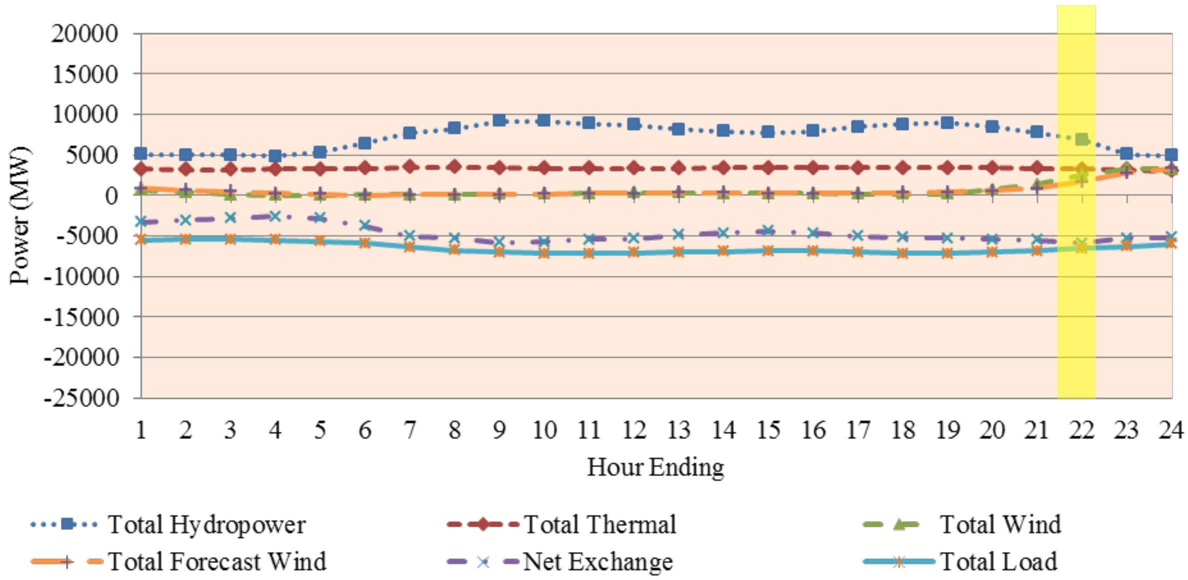

(a) BPA Data (BPA 2015)

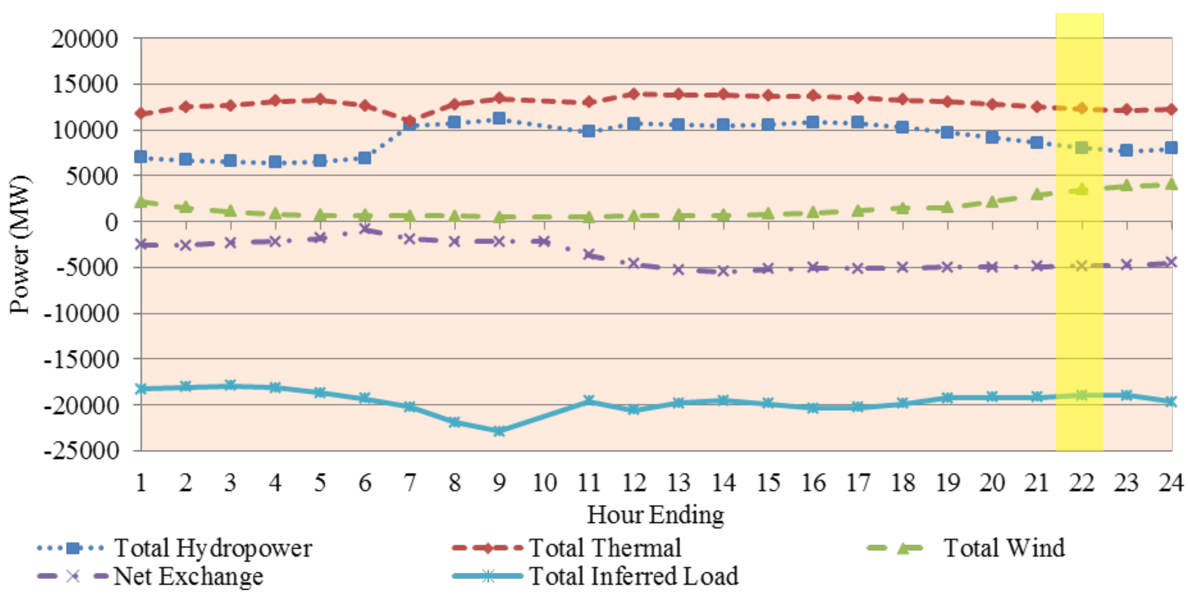

(b) PNWSGD Transactive System Data

Figure 2.12. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System on February 15, 2014

Based on BPA data, hydropower resources were used heavily on this day to both follow system load and to respond to the increase in wind resource late in the day.

The wind resource in the transactive system data is very similar in magnitude and shape to the wind resource shown in the BPA data. This is not terribly surprising given the attention that the PNWSGD project paid to monitoring and predicting wind resources in the region. As for the BPA system, hydropower may be seen to track the impacts of changes in load and wind resource. However, the transactive system's modeled thermal loads were much more dynamically controlled and responsive than the thermal generation in the BPA data. Furthermore, there was more thermal generation resource than hydropower resource in the modeled transactive system, the opposite ordering observed in the BPA data.

The PNWSGD transactive system did not compare wind generation against forecast wind as was described to affect BPA this day. Generating units were not modeled to become committed (*scheduled*) in the project's transactive system implementation. The project predicted wind generation, but the predictions were used on the transactive system's planning horizon without resulting in commitments from the modeled wind farms.[1]

Figure 2.13 features the BPA and transactive system wind data that was shown in Figure 2.12. The source of the BPA data was the BPA transmission webpage (Wind Generation & Total Load in the BPA Balancing Authority, BPA 2015). The discrepancy between forecast and generated wind is easily seen as the wind ramped up. Transactive system data was unavailable for the hour ending 10 on this day. The project had broader visibility of and participation by wind resources than exist within the BPA system, so the magnitude of wind energy was typically greater for the transactive system. The transactive system modeled the rapid wind ramp well. The timing of the wind ramp was similar between the two systems.

[1] Had the impacts of scheduling accuracy and reserve margins been incorporated into the transactive system, the system might have responded to help mitigate over- and under-generation events. Nothing prevents a future transactive system from including the impacts of resource commitments, but committed resources are no longer available and responsive to help mitigate emergent situations thereafter.

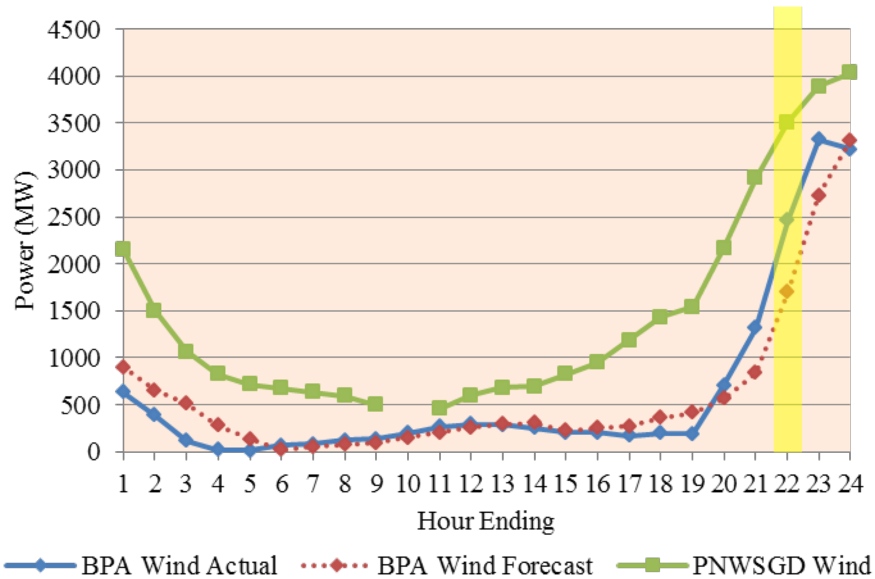

Figure 2.13. Comparison of Transactive System Wind Generation, BPA Wind, and BPA Forecast Wind Data from February 15, 2014

On this day, BPA purchased 1.4 GWh and sold over 25 GWh. Observe in Figure 2.14 that both the averaged transactive system incentive signal and that of the Oregon Cascades transmission zone, from which much of the region's wind emanates, decrease late in the day as wind power increases. Some of this regional impact is a natural diurnal pattern caused by load following, but the transactive system's wind resource functions were designed to make the TIS incentive signal inversely proportional to the magnitude of wind power that is being generated, thus creating a downward pressure on the incentive costs especially at the Oregon Cascades transmission zone.

No significant change occurred in the TIS the hour that BPA observed an overgeneration event. The TIS was not expected to be influenced by forecast errors or by the status of system reserves. These influences were not represented among the inputs to the transactive system. If the system had been responsive to scheduling errors, a wind resource function would have been designed to decrease the TIS in response to the imbalance near hour ending 22. The reduced incentive would encourage consumption and discourage generation until the imbalance was resolved.

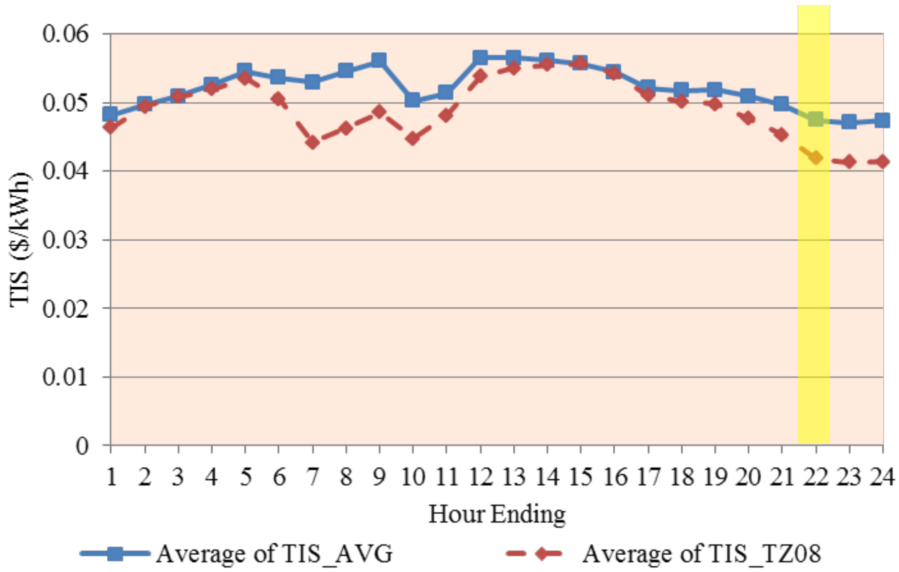

Figure 2.14. Transactive System Average TIS and TIS in the Oregon Cascades TZ08 on February 15, 2014

The inverse relationship between wind generation and the TIS may be seen in Figure 2.15. This figure plots the incentive signal of the Oregon Cascades transmission zone as a function of average hourly wind power generated in this transmission zone. The trend line seems to confirm the inverse relationship. Remember that many inputs, not wind magnitude alone, influence the incentive signal.

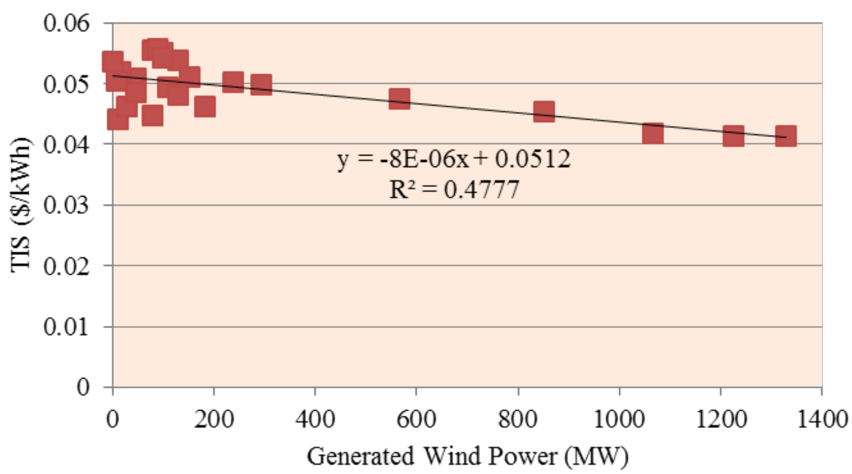

Figure 2.15. TIS as a Function of Wind Power in the Oregon Cascades TZ08 on February 15, 2014. The slope of the line is −$0.008/kWh per GW of generated wind power.

Periods of wind undergeneration and overgeneration on March 5, 2014. Another day of challenging wind conditions occurred on March 5, 2014, when both over- and undergeneration events occurred and were attributed to inaccurate wind forecasts. BPA experienced very heavy wind generation on this day. During Hour 14, wind generation fell short of scheduled wind generation by almost 1 GW. Up to 91% of the available “inc” resources—reserved generation resources—were exhausted to make up the shortfall. Wind states 1 and 2 were issued by BPA as an alert that its “inc” resources were nearing depletion.

The resource and load data from the modeled transactive system and BPA data are compared in Figure 2.16. The most striking observation may be that the transactive system modeled the dispatch of thermal resources to have assisted with load and wind following, whereas thermal resources remained unchanged in the BPA data. The transactive system did not reproduce the resource dispatch strategy in this case.

To make matters even more interesting, starting Hour 21, wind generation then exceeded scheduled wind generation by up to 1.2 GW. Over 90% of the available “dec” resources then became exhausted. Wind states −1 and −2 were issued.

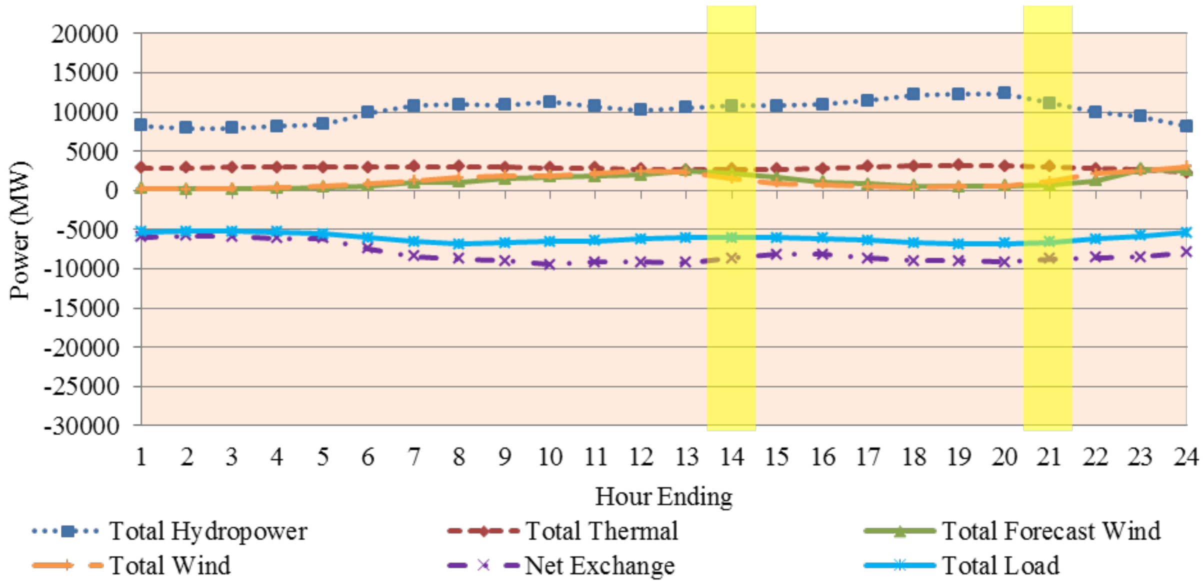

(a) BPA Data (BPA 2015)

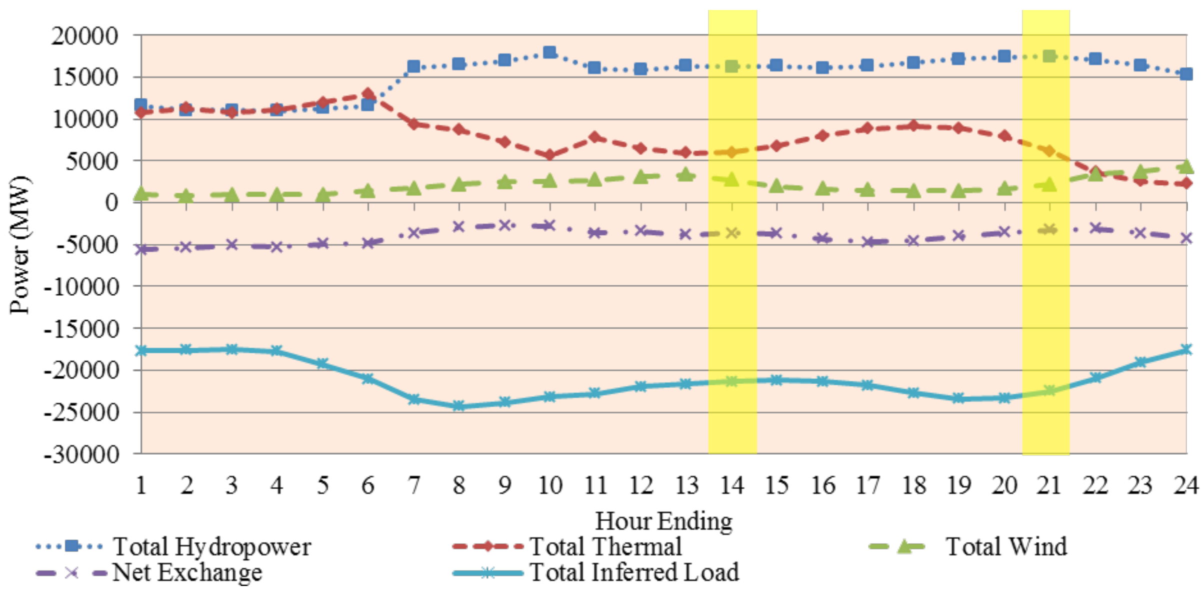

(b)PNWSGD Transactive System Data

Figure 2.16. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System on March 5, 2014

Figure 2.17 focuses in on the wind components that were shown in Figure 2.16. Wind generation in the PNWSGD transactive system closely paralleled that reported by BPA. However, the impacts of differences between scheduled and actual wind generation are not addressed by the transactive system. Wind power may contribute to the need for “inc” resources when the resource falls below the forecast, and “dec” resources when the resource exceeds the forecast.

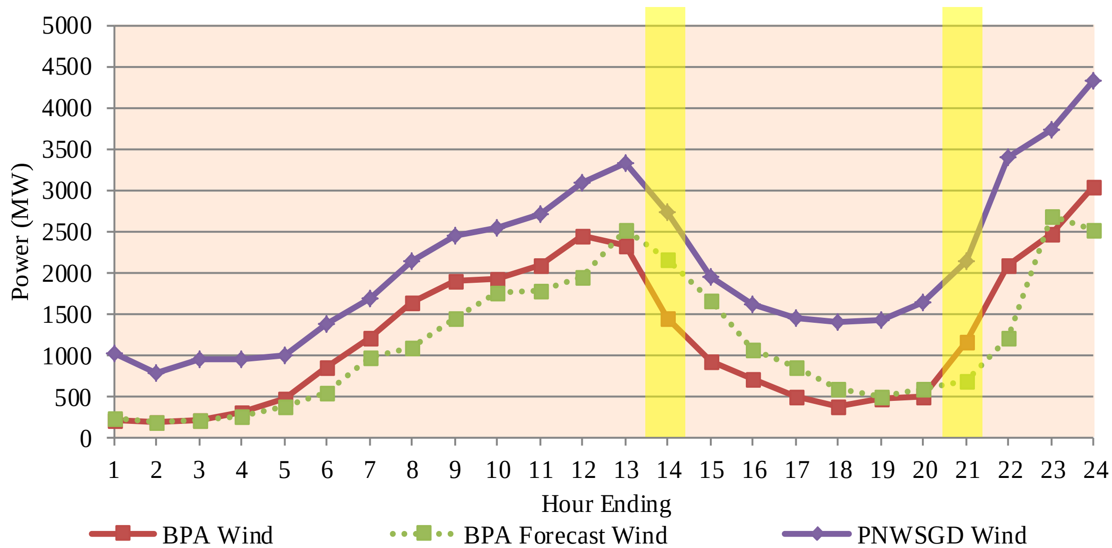

Figure 2.17. Comparison of BPA and PNWSGD Transactive System Wind Generation Data on March 5, 2014

Some influence may be seen in the TIS incentive signal at windy zones like TZ08 − Oregon Cascades (Figure 2.18). The effect is not obvious in the TIS averaged over the entire region. The incentive costs do not reflect the system imbalance, but overall costs were reduced near peak wind power generation near hours ending 12 and 24. The influence is not large because wind remains a relatively small fraction of the total system energy resources.

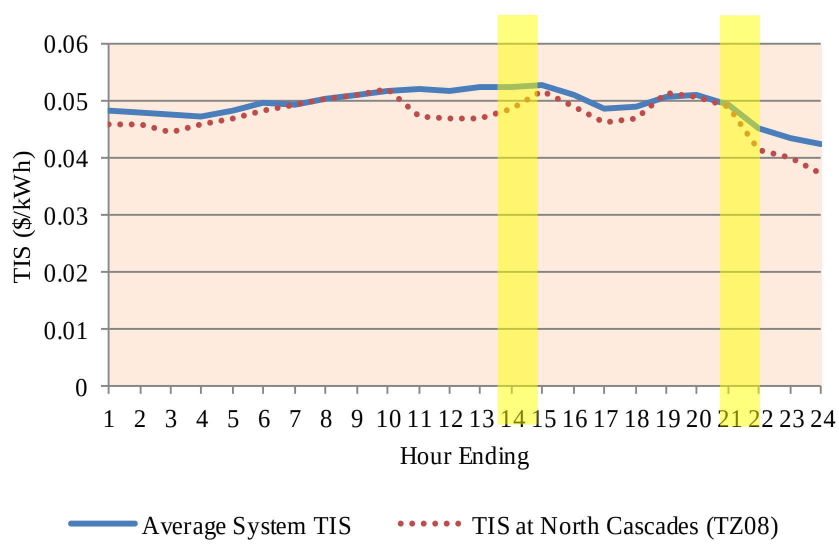

Figure 2.18. Average Transactive System TIS and the TIS at the Oregon Cascades TZ08 on March 5, 2014

Again, the trend is for a zone's TIS to decrease with increasing wind power, as was designed. This trend is demonstrated in Figure 2.19, which plots the incentive signal of the modeled Oregon Cascades TZ08 as a function of wind power that was generated there. The correlation is weak because many conditions influence the blended cost incentive signal.

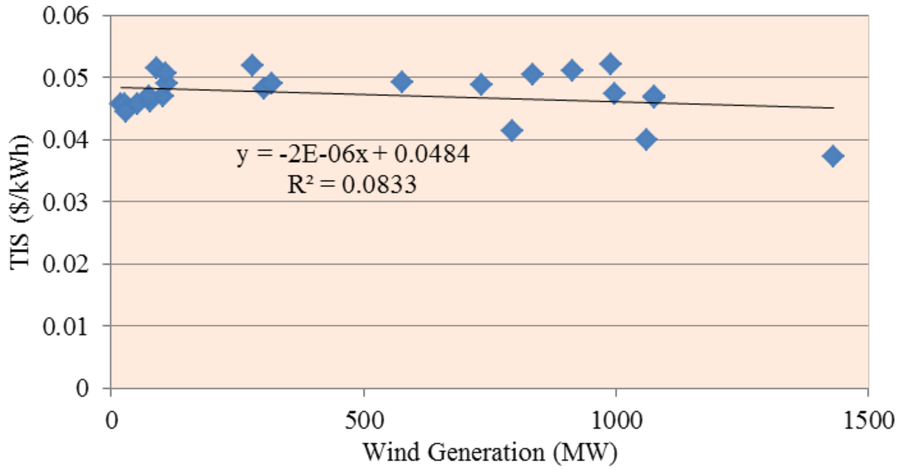

Figure 2.19. TIS as a Function of Generated Wind Power in the Oregon Cascades TZ08 of the Transactive System. The slope of the line is −$0.002/MW. The slope is $0.002/kW per GW of wind generation, but the correlation is poor this day.

2.2.5 Transmission Incidents

This section compares BPA and PNWSGD transactive system data for two BPA transmission system events. At times, this section refers to flowgates, which are transmission lines, or groups of transmission lines, the loading of which are carefully tracked as potential locations of transmission congestion.

A transmission outage on April 1, 2014 (North of Monroe). On April 1, 2014, a transmission curtailment occurred hours ending 10 – 11 on the North-of-Echo-Lake flowgate after the flowgate came within 8 MW of its normal transfer capacity. A planned outage on a nearby transmission line had caused the operating limit to become reduced. BPA reported to the project that load on the flowgate peaked at 1,219 MW at 9:45 Pacific Time, within 7 MW of its maximum normal transfer capability. Between 10:00 and 10:15, the curtailment order decreased load on the line in excess of 90 MW and then an additional 100 MW between 10:15 and 11:00. The curtailment shifted balancing reserves deployed from drawing 260 MW at 09:50 to backfilling 231 MW by 10:20.

Despite the multiple changes made to manage the transmission capacity, BPA purchased no power and sold 24,448 MWh on this day. During the curtailment period, there were no adjustments from the plan for sales on the day-ahead market.

The impact is apparent at the system level in neither BPA nor PNWSGD transactive system data. See Figure 2.20.

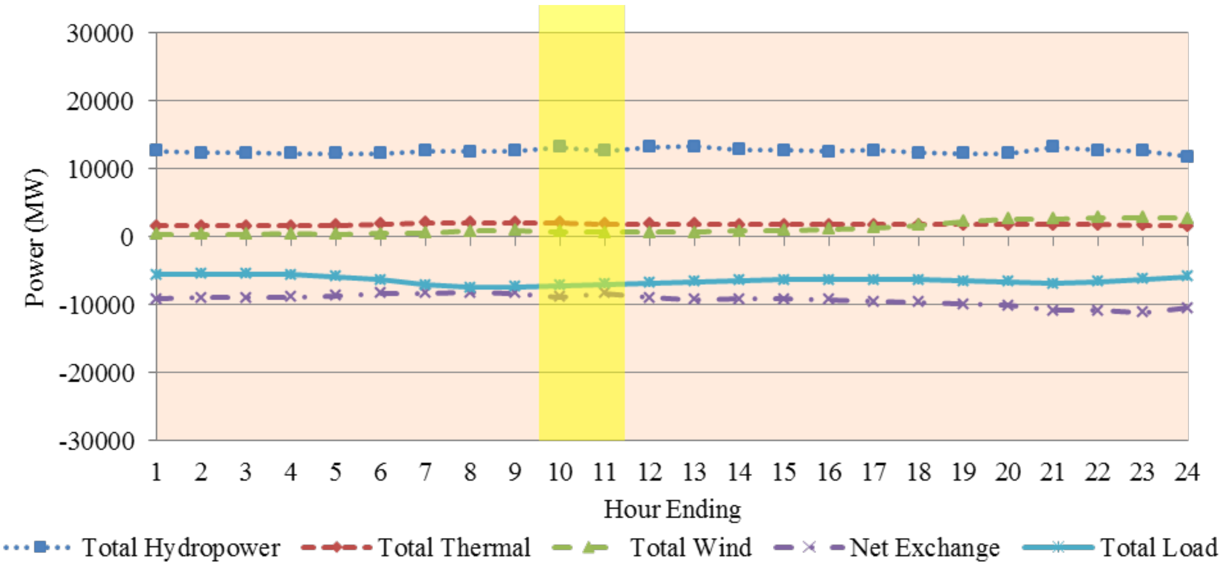

(a) BPA Data (BPA 2015)

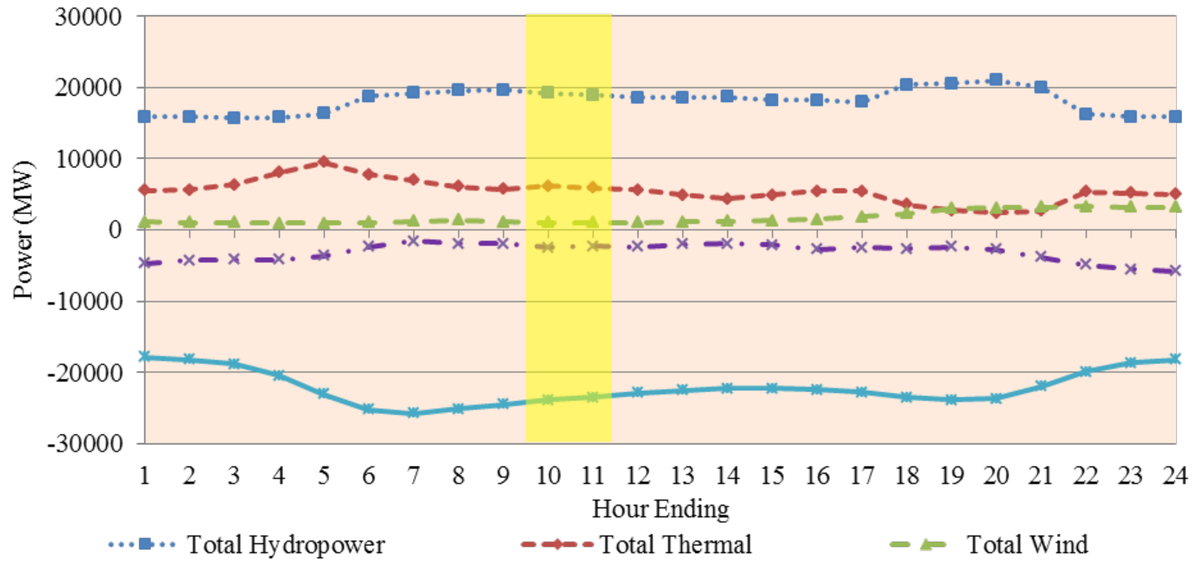

(b) PNWSGD Transactive System Data

Figure 2.20. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System on April 1, 2014

Figure 2.21 more narrowly focuses on the reported power flow in the North of Monroe flowgate on this day and on the way this flowgate was modeled by the transactive system. This flowgate is modeled approximately by the modeled flow between TZ01 and TZ02. Precisely, the flow is therefore modeled by the TFS between these two nodes of the transactive system.

The modeled flow from TZ01 to TZ02 approaches the constraint level at about the time of the event. The flow is much less, however, off peak. The flows are not intended to be identical in the PNWSGD model, but they are found to have comparable magnitudes during the peak period of the day. The transactive system did not implement the construct of a normal transfer capacity during the PNWSGD. The project failed to design and implement a satisfactory function that would have monetized stresses on the transmission system using the incentive signal. The project would not have been able to usefully help BPA avoid this incidence of transmission congestion given the lack of accuracy with which the transactive system emulated the magnitude of power on this flowgate. The magnitudes of individual flowgate power could not be directly mapped to power flows in the project's simplified transmission model.

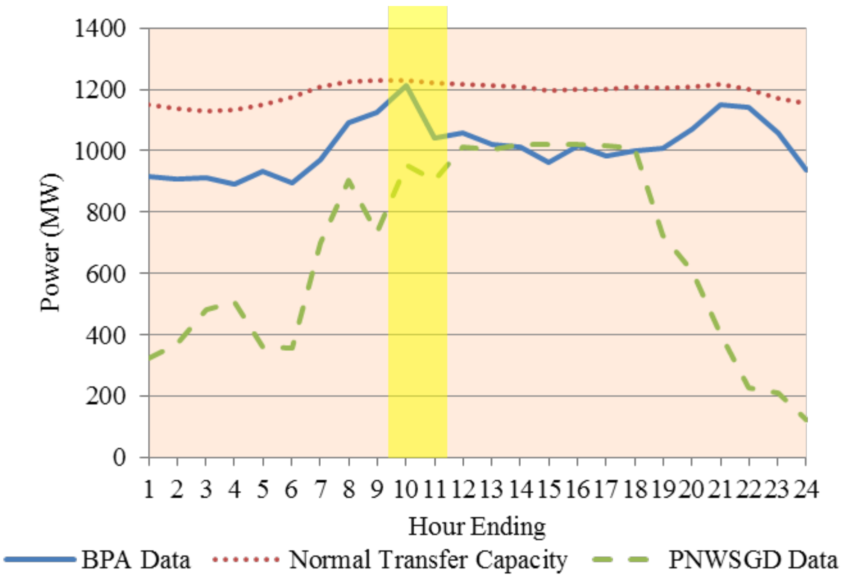

Figure 2.21. TFS Flow North of Monroe on April 1, 2014 According to BPA Data and PNWSGD Transactive System Data

During the event, the magnitudes of the TIS incentive signals at TZ01 and TZ02 do not have the correct relationship that would have helped mitigate the overloaded transmission, as is shown in Figure 2.22. If the transactive system were to help mitigate the overload condition, an incentive difference should appear across the flowgate to discourage consumption (or encourage generation) downstream from the overloaded lines. That is, an incentive would be introduced to make the TIS at TZ02 relatively larger than that at TZ01. The transactive system did not recognize or help mitigate this condition. The incentive signals are found to have the opposite relationship.

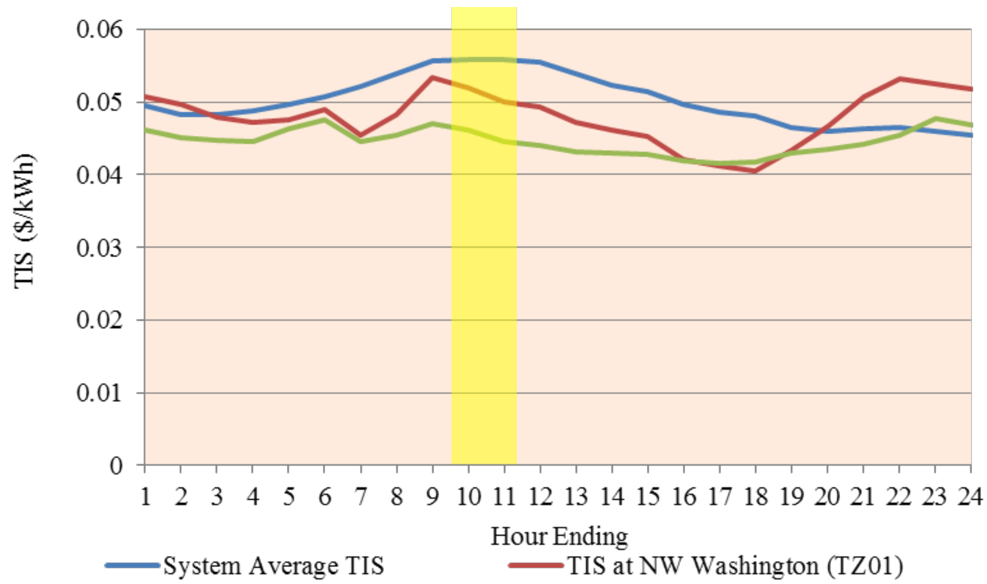

Figure 2.22. TIS Values on both Sides of the Transmission Outage and the Average TIS for the Entire Transactive System

An overloaded flow gate event on April 11, 2014. On April 11, 2014, the flowgate North-of-John-Day system operating limit became surpassed by 120 MW at around 13:30 hours. The system operating limit had been reduced due to a planned outage of two nearby 500 kV transmission lines.

Figure 2.23 compares the resources and loads that day as captured by BPA data and by the transactive system model. No impacts are evident in either the BPA or transactive system data at this level. The traces are similar, and the dynamics are mostly uneventful.

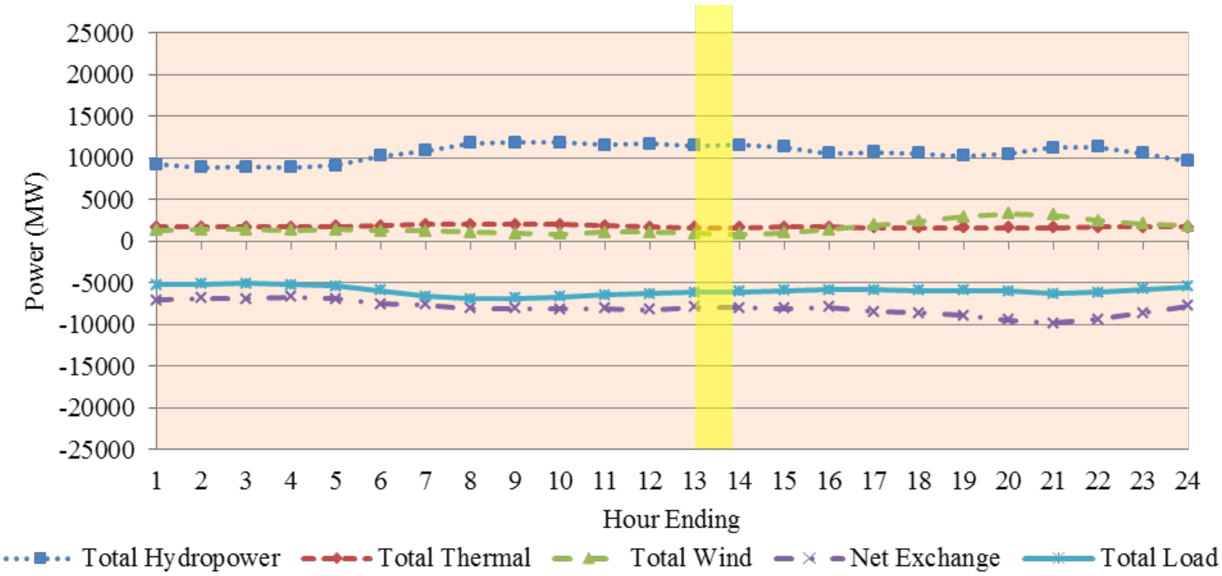

(a) BPA Data (BPA 2015)

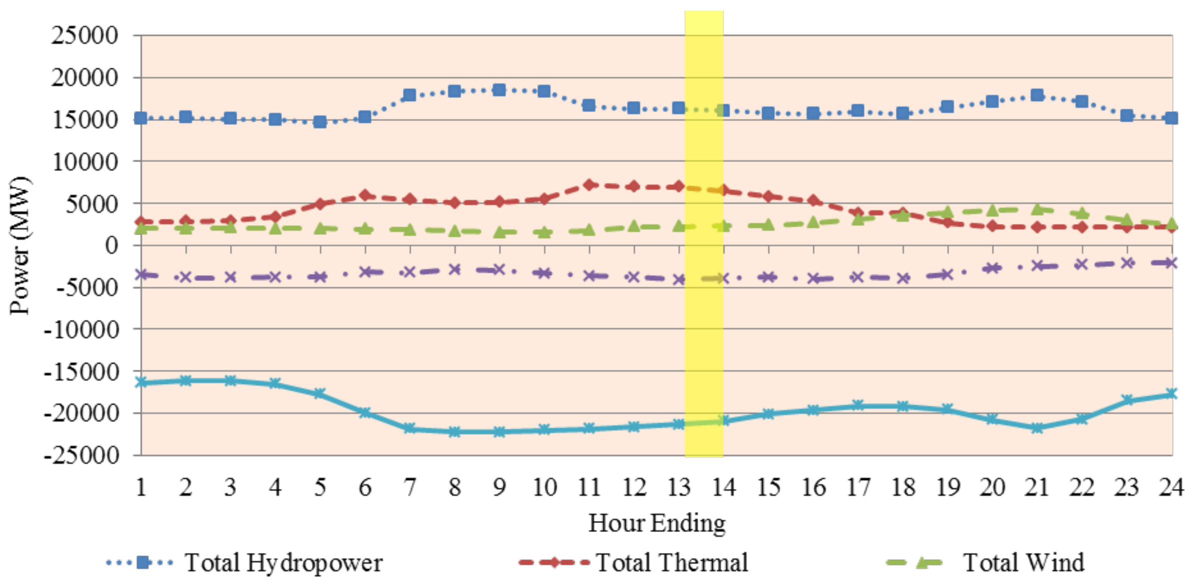

(b) PNWSGD Transactive System Data

Figure 2.23. Comparison of Total Generation Resource and Load Data for BPA and the PNWSGD Transactive System on April 11, 2014

The PNWSGD models the sum of power flowing from TZ07 to TZ08 and TZ12 as being comparable to that of the North-of-John-Day flowgate. The sum of these two TFS flows from zone TZ07 (Hanford) to zones TZ08 (OR Cascades) and TZ12 (Central Oregon) is a pretty accurate representation of the actual flowgate loading on this day, as shown in Figure 2.24. The sum power flow that emulated the flowgate power flow exceeded the net transfer capability of the flowgate at times on this day.

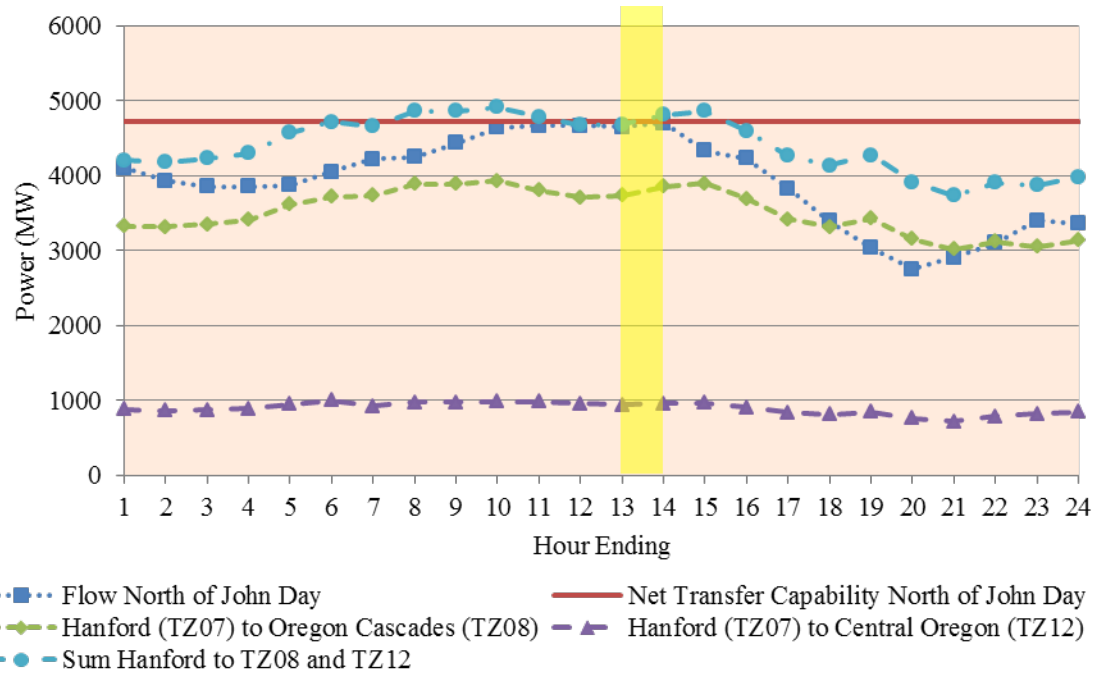

Figure 2.24. Comparison between the Actual Power Flow North of John Day and the Sum of Power Flows between the Transactive System's Hanford TZ07 and Neighboring Transmission Zones Oregon Cascades (TZ08) and Central Oregon (TZ12) on April 11, 2014

2.2.6 Relative Accuracy of Resource Predictions

The PNWSGD transactive system included a predicted future time horizon several days into the future. The predicted future dispatch of resources and incentives were therefore updated every 5 minutes. This proved to be a very challenging innovation for the project implementers. The intention of the future prediction horizon had been to facilitate day-ahead planning, much as is accomplished today by day-ahead and shorter-term markets, but with even greater resource flexibility.

As was discussed in Section 2.3, the value of the TIS follows directly from the unit costs of the energy resources that are being dispatched and perhaps other incentives that follow less directly from the dispatch plan and other grid conditions. If the present or future predicted dispatch and other grid conditions are incorrect, the incentives will also be incorrect and might induce undesired behaviors.

The future predictions would be critical in a truly distributed transactive control system where resources might be viewing both the balance of power and the incentive signal to determine when best to operate and not. The accuracy of the future intervals is especially critical for demand-side elastic loads that often have very few available event periods with which to participate.

Figure 2.25 demonstrates one symptom of an inaccurate resource prediction that badly plagued the transactive system implementation. The horizontal axis represents the difference between the time that predictions are made and the time interval for which conditions are being predicted. The far left position is the nearest-term prediction—the prediction that is being made for the next 5-minute interval. Toward the right, the predictions are being made further into the future until the far right where predictions are being made almost 4 days into the future.

The vertical axis of Figure 2.25 is the average of the summed resources that are being predicted the given time into the future. The window of these calculations progressed 5 minutes each 5-minute interval to include about one-half month from May 18 to June 4, 2013. "Total resources" means the sum of both the power that is modeled to have been generated within the given transmission zone and any power that was imported into the transmission zone from transactive neighbors or non-transactive exchange boundaries. Results are shown for all the transmission zones (see Appendix B).

If total resources are averaged over multiple hours and weeks, the average should represent an accurate average of the total resources; if predicted total resources are averaged over hours and weeks, the same average should be calculated. If not, a bias exists in the predictive calculations. In this case, all of the transmission zones undergo a significant change in average predicted total resource for the predictions that are being made about 3.5 or more hours into the future. Some increase, others decrease. Regardless, no such change should occur in the calculations.

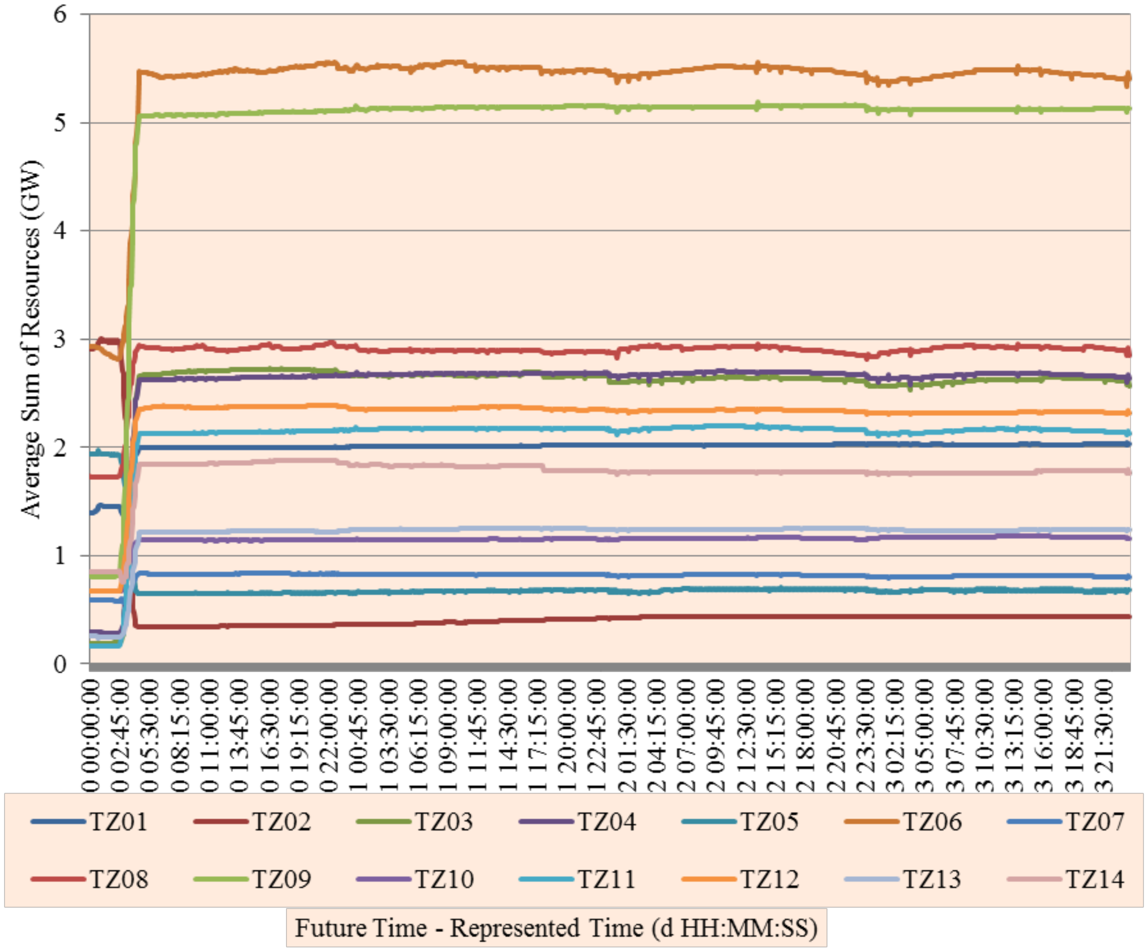

Figure 2.25. Average Total Resource Energies at TZs of the PNWSGD Transactive System Plotted against the Distance into the Future that the Predictions Were Made. This plot includes values from the transactive system production environment between May 18, 2013 and June 4, 2013 before the future predictions were improved.

The project hypothesizes that the change was caused by the use of different calculations in the Alstom Grid-informed simulation for intervals predicted less than and more than about 3.5 hours into the future. After much effort, Alstom Grid managed to decrease the magnitudes of the changes, as is shown in Figure 2.26. They were not able to completely eliminate the discontinuity.

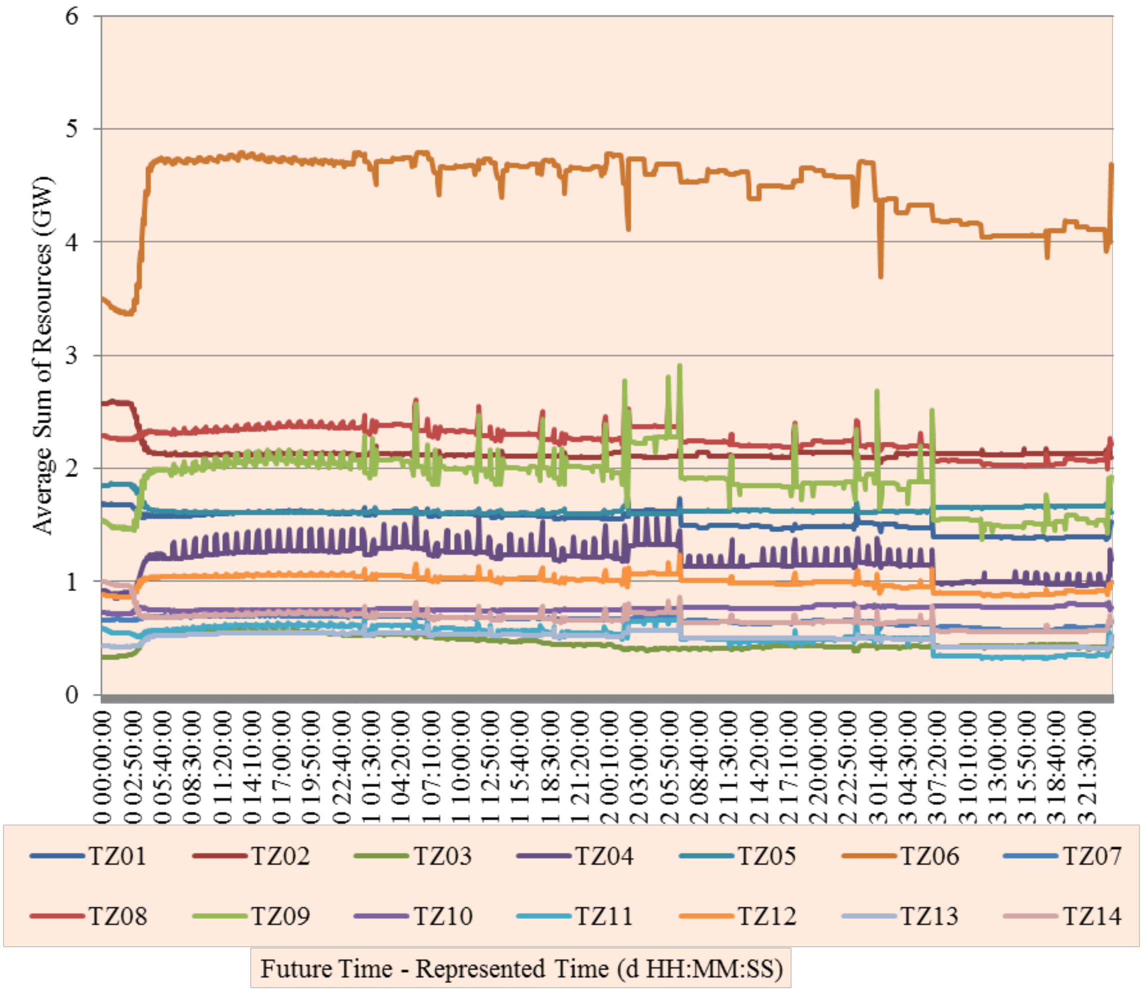

Figure 2.26. Average Total Resource Energies at TZs of the PNWSGD Transactive System Plotted against the Distance into the Future that the Predictions were Made. This plot includes values from the transactive system test environment between May 18, 2013 and June 4, 2013 after the future predictions had been improved.

2.2.7 Step 1 Evaluation Conclusions

In Step 1 of the analysis of the performance of the PNWSGD transactive system, the transactive system was confirmed to have represented the actual statuses of regional generation and transmission where such data was made available to it.

- The mix of generation resource types modeled by the transactive system paralleled those that had been reported by the BPA system. The system separately modeled thermal, hydropower, and wind power, plus the power that was imported into or exported from the region, of which the resource type was unknown. A direct comparison was impossible because the transactive system's region was larger than that of BPA, but the relative resource mixes were credible.
- The transactive system achieved superior visibility of actual and predicted wind power resources throughout the Pacific Northwest. The magnitude of wind resources closely paralleled BPA's wind

power data. The project could therefore anticipate and observe wind power magnitudes and rapid changes in wind power magnitudes—both up and down. However, the project was oblivious to the impact of wind power intermittency on BPA balancing reserves. The status of such reserves was not revealed to the transactive system.

- The transactive system appears to have recognized and represented an unexpected outage at a large power generator. However, the visibility of the outage in the transactive system may have been delayed for an hour. Because the transactive system's model of the Pacific Northwest transmission is coarse, the impact from losing even 1 GW of generation was relatively small at the incident's node.
- The transactive system did not accurately represent and respond to transmission events, including line outages and actions taken to keep loads under capacity limits. The transactive system's transmission model was not formally designed from the actual transmission system in a way that maintained correspondence between individual transmission loading and modeled ones. The status of the system was not explicitly available to the system, so it was expected that the transactive system would not represent such events. The project failed to implement a function that would incentivize transmission loading levels, the purpose of which would have been to assist constrained economic dispatch that is used today.

2.3 Step 2: The System Must Meaningfully Monetize and Predict Resource Costs and Incentives

In this evaluation step, we review the methods by which the transactive system monetized its energy resources and the system objectives to which incentives were applied. In the ideal—a fully distributed system of nodes, where each node independently selects its supply resources and its objectives to be incentivized based on transactive signals and local conditions—the functions might be unknowable and uncountable. What the PNWSGD project was able to implement was instead an "informed simulation" having a small number of defined resources and incentives that were designed by and fully monitored by the project. The project referred to the functions as "toolkit functions" to emphasize that once designed, the functions could be placed in a toolkit library of functions available to be adopted, revised, or reconfigured to suit the needs of future implementers.

The project possesses much unpublished documentation about the workings of the informed simulation that was used to monetize the transactive region's resources and incentives. Figure 2.27 supports an adequate review of how this subsystem worked.

Four dynamic data feeds are shown at the top. These four inputs drove the dynamics of the informed simulation:

- BPA hydropower schedule
- BPA load forecast
- 3TIER wind forecasts
- independent power producer generation schedules.

Other more static system data is also important, including the typical status of the Western Electric Coordinating Council region generators that are needed to emulate the exchange of energies at the transactive system's exchange boundaries, tables of fuel prices and infrastructure costs that especially affect the region's modeled thermal resources, the topology of the transactive system that states the connectivity between the system's nodes, and the mapping of the region's resources into the transactive system's nodal model.

At the bottom left, the region's circuit state was modeled from a limited number of representative historical condition sets ("NETMOM"). These models affected load flow calculations and at times modified the modeled generation and load profiles.

Alstom Grid used its unit-commitment and economic dispatch engines to facilitate the scheduling of modeled resources for the project. From the perspective of Alstom Grid and the informed simulation that is portrayed in Figure 2.27, its resource and incentive toolkit functions (bottom right) were the mappings of the resources' costs and dispatched powers allocated according to the model of the system, grid, and modeled generation resources.

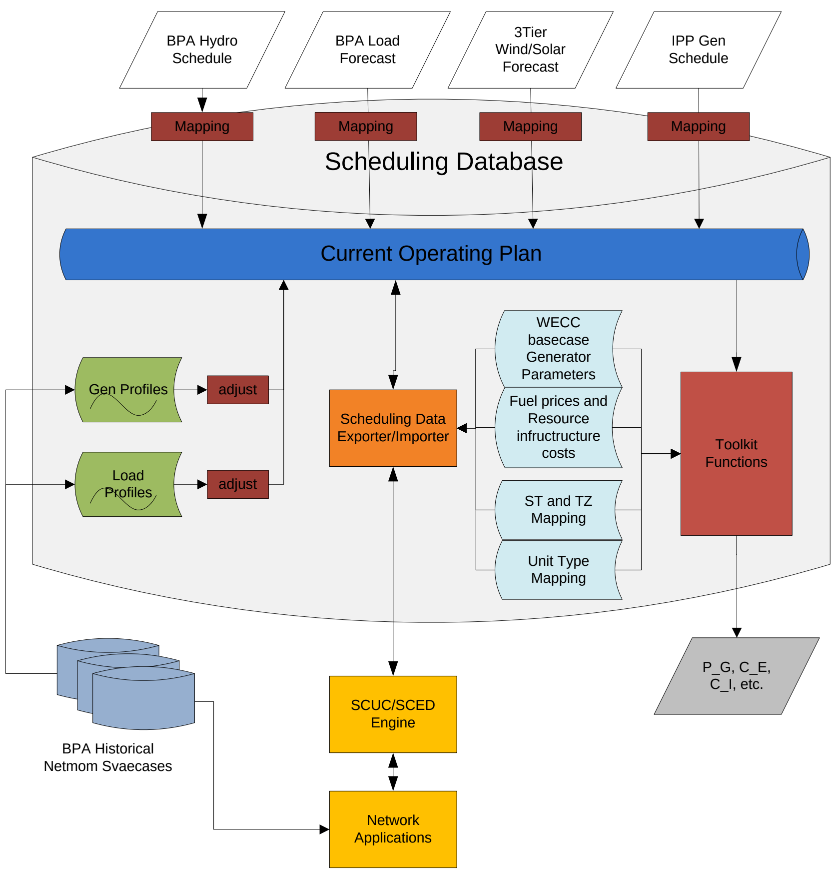

Figure 2.27. Alstom Grid Toolkit Functional Overview[1]

[1] From p. 1 of PNWSGD ALSTOM Toolkit Function Description, Version 0.3. Alstom Grid, 10865 Willows Road NE, Redmond, WA 98033, September 9, 2014, unpublished.

2.3.1 Toolkit Resource and Incentive Functions

Each resource or Incentive Toolkit Function was specified by the project as a black box having defined inputs and outputs. Many of the inputs were different from one function to the next, but the set of output coefficients were specified to facilitate calculations of the blended TISs, as discussed in Section 2.4.1. Only the limited set of coefficients could be assigned values. The limited set of allowed output coefficients fosters interoperability at this interface.

The purpose of each function's monetization is to influence the delivered cost of energy, but the dynamics with which the influence becomes applied over time is a free design variable that is available to the resource's owner (in this case, the project acted on behalf of resource owners) to further incentivize desired energy behaviors. This concept was challenging for the project's utility participants to grasp and accept. It is not the way electricity costs are charged today. Today's regulatory environment would need to be changed to allow this approach while still enforcing fairness.

The sum costs represented by the toolkit functions should sum, at least over long periods of time, to the actual cost of electrical energy at its location in the transactive system. With this understanding, the cost of infrastructure had to be modeled to represent any discrepancy between the transactive systems energy costs and the energy costs that are eventually borne by the region's distribution utilities. The TIS must be equivalent to the price of energy over time if, in the future, the transactive system is ever to be accepted as a basis for energy billing.

These following resources and incentives were monetized by the transactive system. The parenthetical numbers reference the project's numbering convention for its toolkit resource and incentive functions. The project generated as-built design documents for each (see Appendix A).

- Non-transactive imported energy (1.1) − This function emulated the impact from the exchange of power across the region's exchange boundaries. The energy that was imported through these boundaries was treated as a resource to the importing transactive node. The unit cost of the imported energy was based on recent trends in the Dow Jones Mid-Columbia price index.
- Transactive imported energy (1.2) − This is a trivial function, but it is included for completeness. No new calculations were required in the informed simulation. This function is accomplished by the correct blending of neighbor nodes' transactive signals such that the quantity and costs of energy imported from these transactive neighbors influences this node's TIS.
- Hydropower (2.1) − Scheduled hydropower generation was assigned costs according to the recent history and trends of the Mid-Columbia Dow Jones Price Index that was subscribed to and used by the project. Most of the modeled hydropower inflexibly followed schedules. Two large hydropower generators were modeled to be responsive to changes in system power balance.
- Wind power (2.3) − Total wind power in the region was reported and predicted by BPA and 3TIER. No cost was applied for the energy itself, but the infrastructure costs of wind farms were included among the general infrastructure costs (function TKRS 4.0). This approach might encourage consumption of wind power when and near where it is available. There was a downward pressure on the incentive signal magnitude as wind was blowing.

- Thermal generation (3.0) − Scheduled thermal generation was assigned costs according to heat rate curves and fuel costs for the corresponding generator types.
- General infrastructure costs (4.0) − At each transmission-zone node, a cost offset was assigned to represent the costs of infrastructure that had not been otherwise represented in the system. The target costs were based on typical wholesale prices paid by utilities near the given transmission zones. The coefficient slowly tracked that target price with a response time of about 1 month.
- Transmission congestion (5.1) − This function was implemented and used prior to May 2013, but it was turned off at all system locations after it was found to create undesirable, rapid changes in TIS values. The intention had been to disincentivize consumption downstream and incentivize production upstream from any congested flowgate.
- Demand-charges functions (7.x) − These incentive functions were implemented at utility sites, not at the transmission-zone nodes, so they were not part of the informed simulation. However, they fit into the present discussion because these functions modify the effective TIS at utility locations to reflect the impact of demand charges that are imposed by the energy supplier at the site. The University of Washington campus function also included time-of-use impacts that are part of the campus's contract with supplier Seattle City Light.

2.3.2 The Accuracy of TIS Predictions

The TIS represents the blended costs of resources and incentives. This section reviews the relative accuracy of TIS predictions as a surrogate for the aggregated resource and incentive influences. Every TIS included a series of predictions for 56 sequential time intervals. The nearest-term prediction was for the imminent 5-minute interval (interval start time IST0). Successive prediction intervals represented a series of 5-minute, 15-minute, 1-hour, 6-hour-long, and day-long future time intervals.

Figure 2.28 shows the relative prediction error at one of the transactive sites for the 8 months that the PNWSGD was operating in 2014. This refers to the Fox Island site (field site node ST01) in Washington, just one location in the transactive system. By the time the system is acting on the information from a specific
5-minute interval, the TIS of the interval has been predicted many, many times. The relative prediction error here is defined as the average difference between those predictions and the final, best calculation of the TIS that occurs just prior to the interval time, divided by the final calculated TIS value.

A positive result means that the predictions, on average, were greater than the final TIS calculation. Negative results, of course, then mean that the TIS tended to be under predicted. The results are averaged for each of the future 56 time intervals. The averages were further separated out by project month, as is indicated by the figure's legend.

Only the values that were calculated by the TIS are being compared. The true representation of the delivered cost of energy that the TIS emulated was nowhere dynamically available for comparison. Today, one would need to compare long-term averages of the TIS to average energy costs to complete a meaningful comparison against actual energy costs. And comparable costs are rarely available.

The principal observation from Figure 2.28 is that the predictions were mostly accurate before prediction interval 20, but the accuracy became worse further into the future. This boundary between prediction intervals corresponds roughly to the transition between 15-minute and 1-hour prediction intervals that occurs about 3 hours into the future.

The prediction accuracies became progressively worse through winter 2013 and spring 2014. A persistent negative bias is observed, meaning the transactive system tended to under predict the TIS at this site. The prediction bias was between 0 and −4.5%.

A surprise was that the day-long intervals used to predict multiple days into the future exhibited among the worst biases. While the challenge is increased by the distance into the future that the predictions are being made, these are also the coarsest averages of TIS intervals.

Future implementation should eliminate biases like these, tracking and correcting them over time.

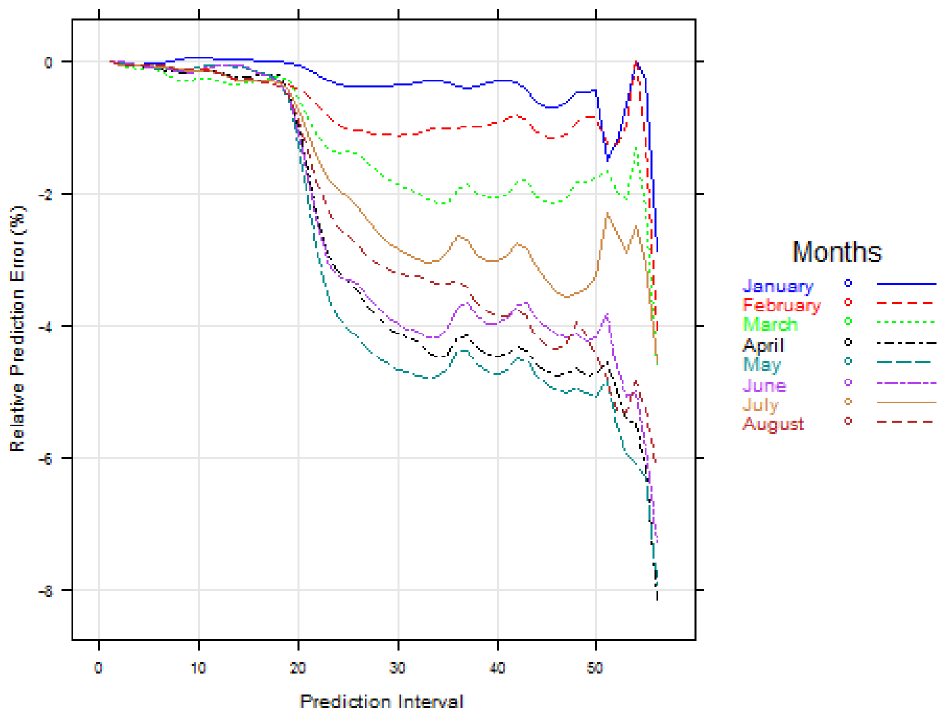

Figure 2.28. Average Monthly Relative Prediction Errors of the TIS Prediction Intervals throughout the Project Months of 2014 at the Fox Island Site (ST01)

Figure 2.29 exhibits standard deviations of the same relative errors that had been shown in Figure 2.28. As should be expected, the standard deviations of the relative errors increase with the predictions' distances into the future. The standard deviations ranged from 0 to about 8%, and the results were similar from month to month.

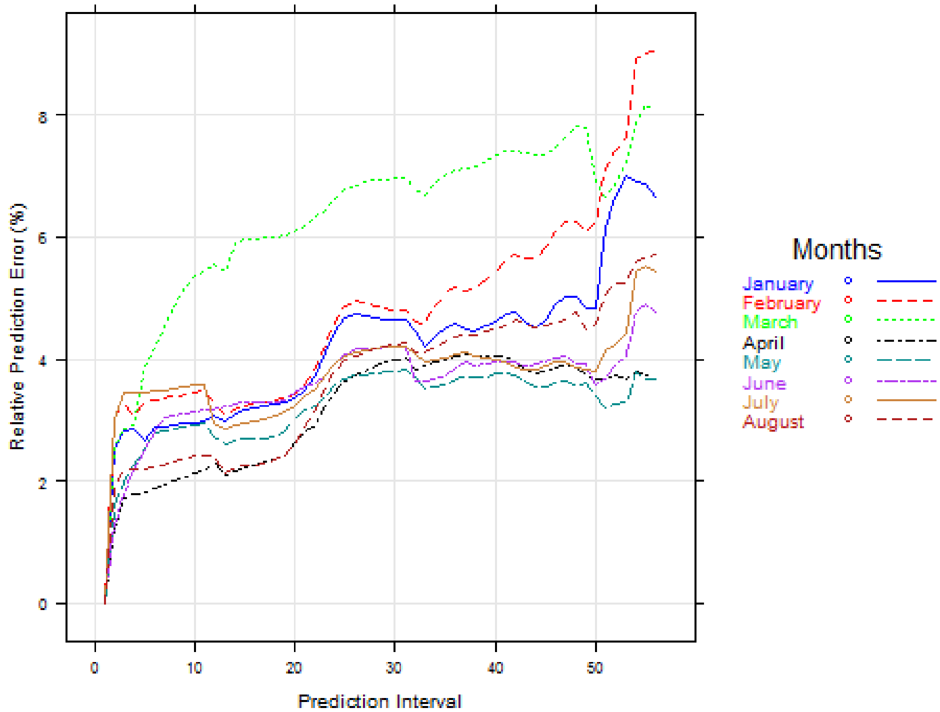

Figure 2.29. Standard Deviations of the Monthly Relative Prediction Errors Eight Months of 2014 at the Fox Island Site (ST01)

Figure 2.30 shows the calculated relative prediction errors again, but this time the results have been parsed by the local hour of day in which the predicted interval fell. Panel (a) graphs heavy load hours (HLHs), and panel (b) plots light load hours (LLHs). The HLHs were generally predicted with less bias than the LLHs. Some of the artifacts in the last intervals (e.g., 55, 56) were predicable. A coarse, day-long interval will tend to under predict the TIS during HLHs and over predict it during LLHs.

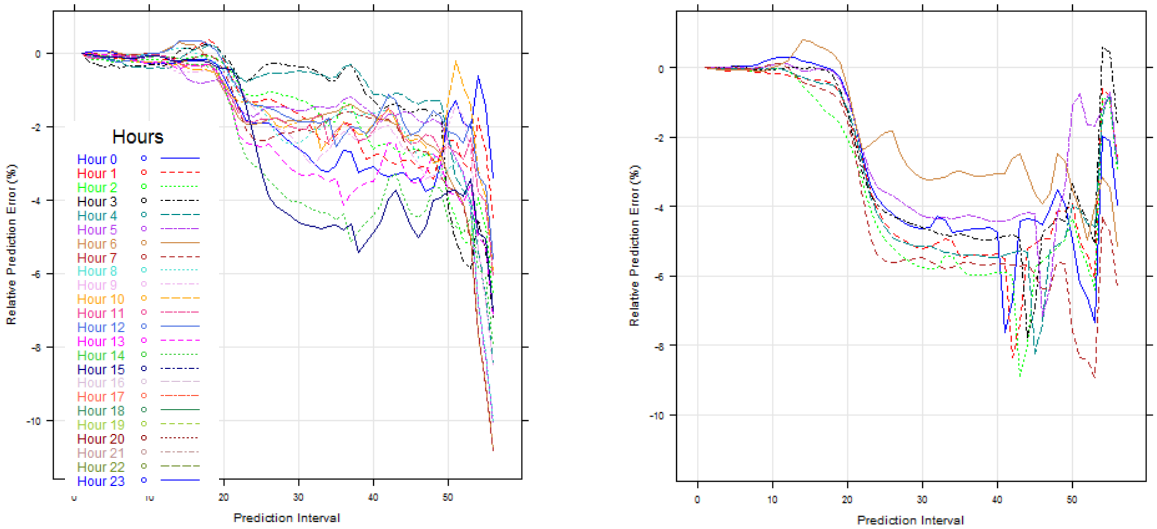

(a) Heavy Load Hours

(b) Light Load Hours

Figure 2.30. Average Relative Prediction Errors of (a) Heavy Load Hours and (b) Light Load Hours at the Fox Island Site (ST01) from January through August 2014

The comparison was also made by interval minute. There were 12 5-minute intervals each hour. The results were similar for all the sub-hourly intervals. As for the monthly and hourly assessments, the accuracy diminished rapidly near interval 20.

The relative prediction errors for another 10 of the other transactive system sites over the 8 project months of 2014 are shown in the panels of Figure 2.31. The prediction error biases were found to be pretty evenly split among those that over and under predicted the TIS. The relative prediction errors were generally small before interval 20, about 3 hours into the future. Predictions of TIS were probably most accurate in panel (b) for the Salem, Oregon site (ST03). Predictions at the Teton-Palisades Interconnect site (Lower Valley Energy ST12) were probably least accurate, approaching 30% error at least one of the months.

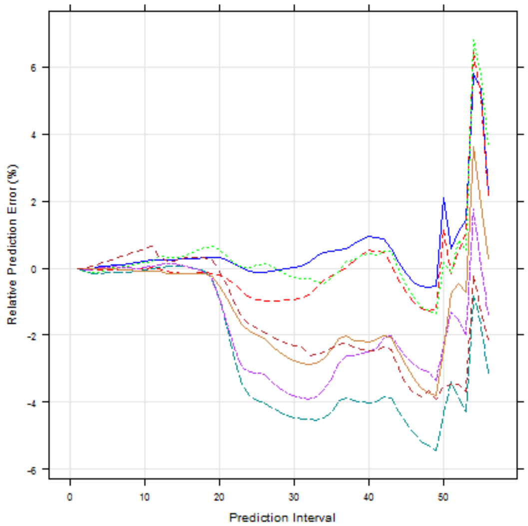

(a) U. Washington Campus Site (ST02)

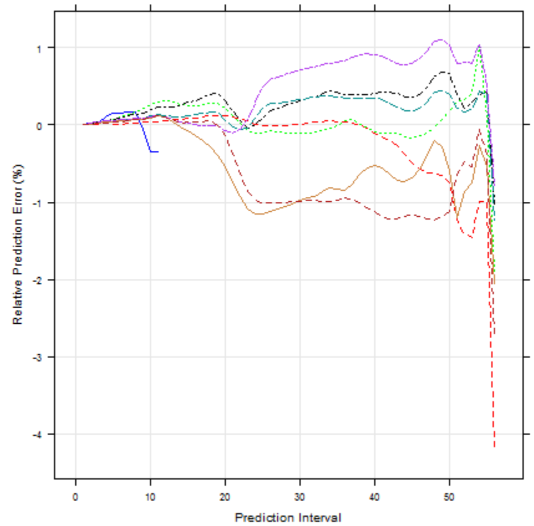

(b) Salem, Oregon Site (ST03)

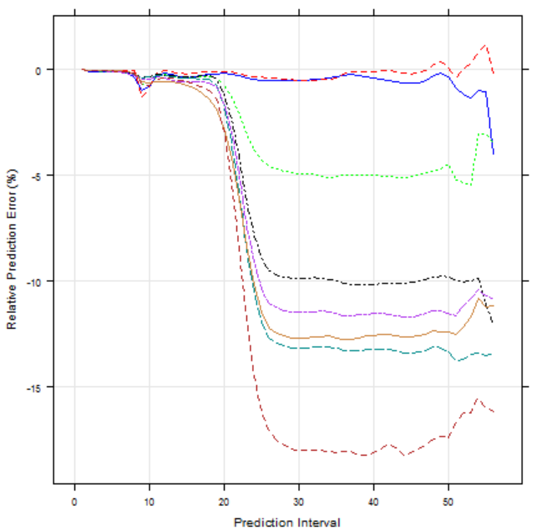

(c) Reata Site (ST06)

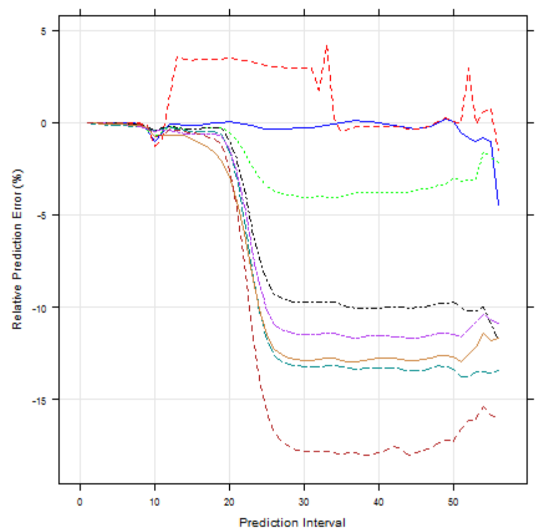

(d) Libby, Montana Site (ST07)

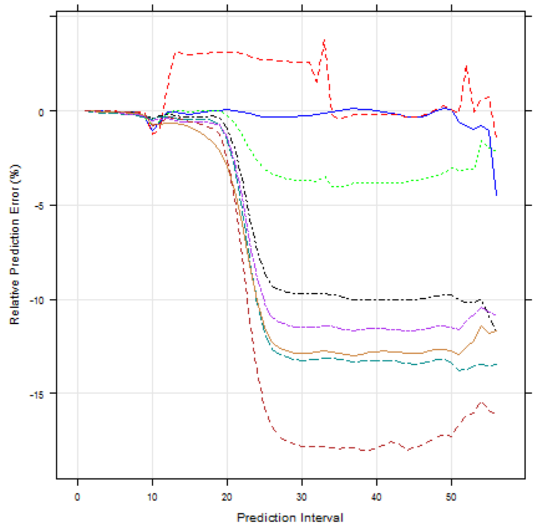

(e) Marion/Kila, Montana Site (ST08)

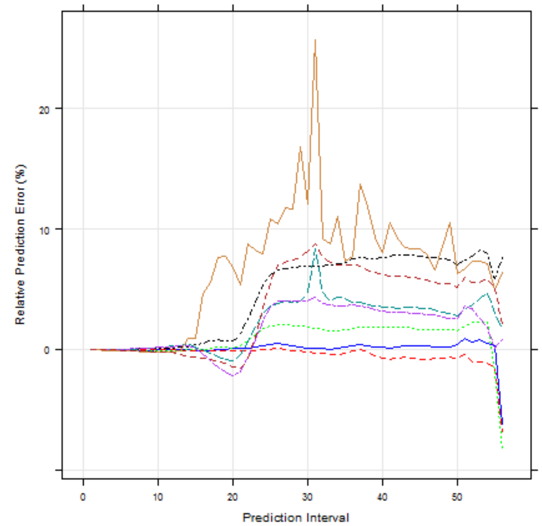

(f) Milton-Freewater Site (ST09)

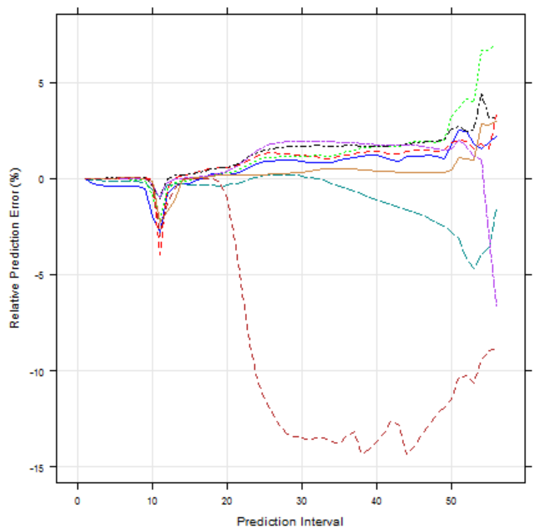

(g) Helena, Montana Site (ST10)

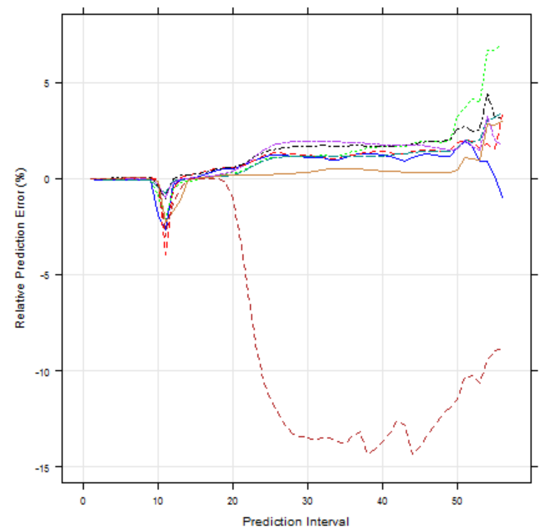

(h) Philipsburg, Montana Site (ST11)

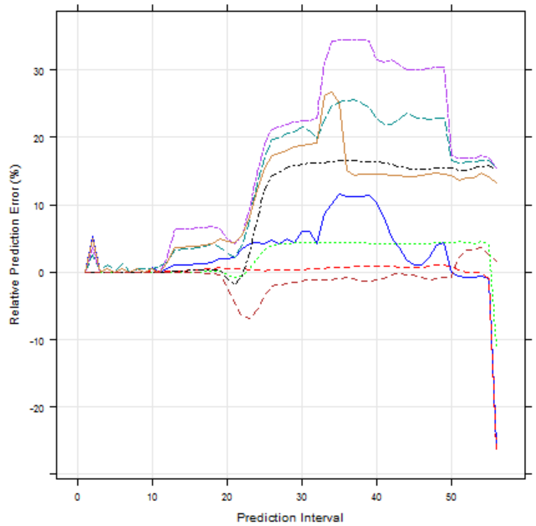

(i) Teton-Palisades Interconnect Site (ST12)

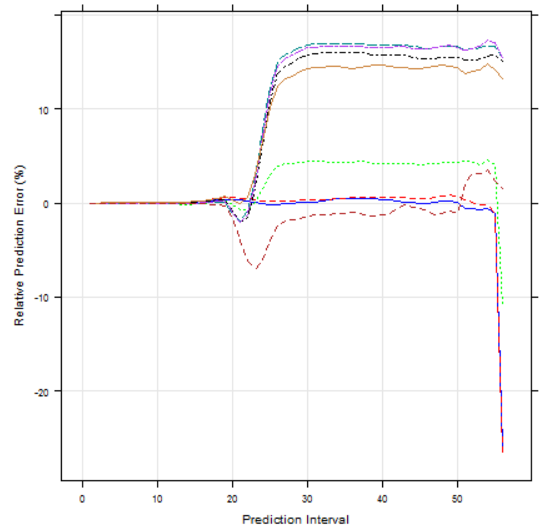

(j) Idaho Falls, Idaho Site (ST14)

Figure 2.31. Relative TIS Prediction Errors for the First Eight Months of 2014 at Ten Transactive System Sites. The month legend from panel (a) works for all the 10 panels.

2.3.3 Changes in Monetized Incentive Mix over Time

Figure 2.32 compares relative resource power mix and relative resource cost mix at one of the transmission-zone nodes over time. This example happens to be for the North Idaho TZ10 during August 2013, but any other transmission zone or month might have supported the comparison equally well or better. The raw resources are shown, just as they are accounted for by this node as it calculates its TIS values. This node occasionally imports non-transactive power from Canada. It typically imports about one-third of its power from its transactive neighbors. Hydropower at this location is not a significant part of the raw resource mix, but there is much wind resource. Thermal resources at times make up one-quarter of the raw resource power mix.

Looking now at panel (b), nearly half the energy cost is allocated to infrastructure costs. Wind energy has been assigned only a small cost. The relative fractions of power and costs are similar for the power and cost panels because the neighbors' costs have already been blended at the neighbors' locations.

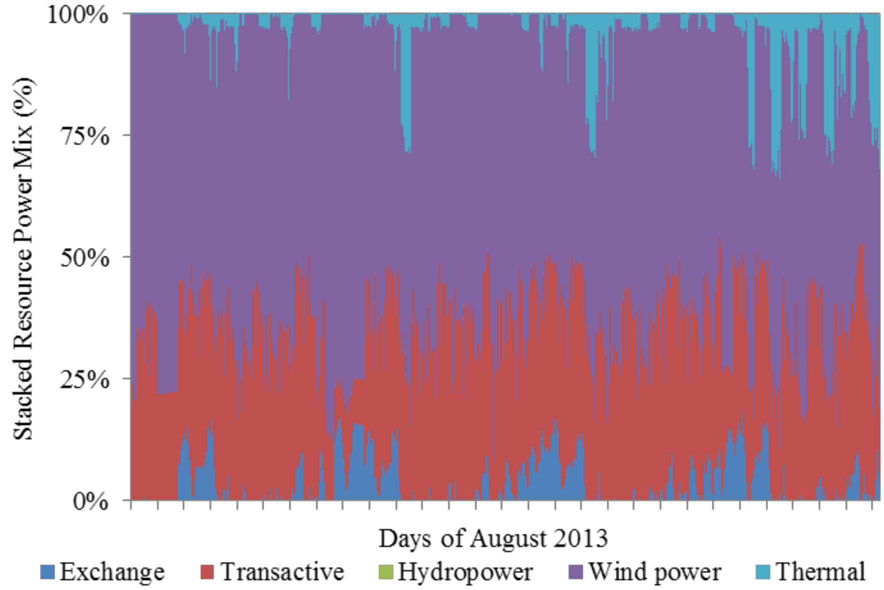

(a) Relative Resource Power Mix

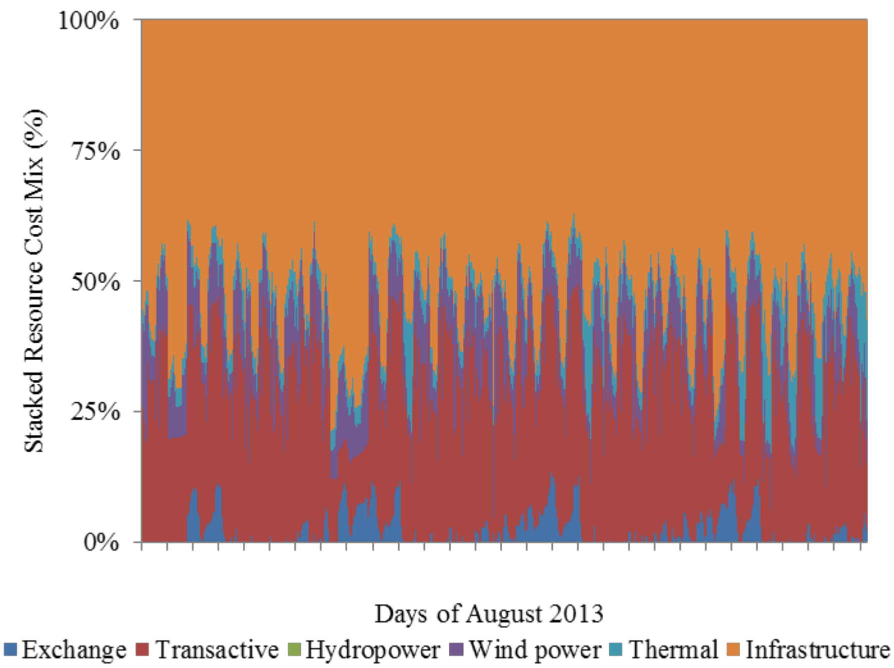

(b) Relative Resource Cost Mix

Figure 2.32. Comparison of (a) Relative Resource Power Mix and (b) Relative Resource Cost Mix for the North Idaho TZ during August 2013

A general observation is that the dispatch of resources may be wild and discontinuous over time. Changes in the dispatch of bulk generation in the transactive system model necessarily created step changes in both the resource mix and the corresponding mix of costs. The next dispatched resource might be distant from the prior one, potentially even causing reversals of power flow in the meshed transmission system. Some of the wild behaviors in the transactive system were caused by the coarseness of the system's transmission model. Some may have been caused by oversimplification or incorrect understandings of the region's resource dispatch strategies. If the discontinuities are real and immutable, this may have adverse implications for the viability of automated distributed control systems like the transactive system.

2.3.4 Lessons Learned Concerning Monetization and Prediction of Resource Costs and Incentives

Need improved prediction tools. The informed simulation that emulated the dispatch of the transactive region's resources persistently predicted values that were either greater than or less than the final resolved value. Some of these biases may be attributed to having used different calculation methods for the long-term and nearer-term predictions. Regardless, the biased predictions of resources also produced biased cost predictions. This was found to be a serious issue. Elastic responsive assets reviewed the predicted incentives to plan their responses. The prediction biases caused these assets to either respond too soon or incorrectly defer their responses depending on whether predicted incentives were always greater than or less than the unbiased value. Much more work must be done to improve the accuracy of predictions.

Integrating wind. The project learned too late during the project term that the toolkit function chosen to monetize wind energy in the transactive system did not adequately address the project's stated objective to integrate wind energy resources. While the project's approach did indeed incentivize consumption of wind energy, it did not monitor and help correct the occasional depletion of balancing reserves that is attributed to wind intermittency. The state of BPA balancing reserves was unavailable to the transactive system.

Demand charges. The demand-charges functions were moderately successful, but fundamental challenges were revealed during the project. The main objective of imposing demand charges is to incentivize a flatter, more consistent electrical load. BPA and other energy suppliers design workable metrics that indicate the overall "peakiness" of the loads. Often only the worst hour of the month is monetarily penalized by the metric. Distribution utilities currently have few resources with which they can truly flatten their load shapes, so they carefully aim their few resources at the one or two worst monthly peaks. They sometimes miss. Regardless, the impact on the transactive system was that its demand-charges functions also behaved this way and applied the modeled monetary impacts (according to the actual incentives) at the peak hour. And the hour was not identified as accurately as we would have liked. The actual cost impacts, when applied to the few peak intervals, were overwhelming and created cost discontinuities in the TIS. In future implementations of demand-charges functions, implementers should do the following:

- Improve the accuracy of the predicted distribution system load.

- Smooth the function's disincentive over more time intervals, perhaps including statistical functions to apply the disincentives in line with the likelihood the peak will be occurring during a given hour.
- Employ enough responsive smart grid resources to truly flatten distribution system load.

In addition, BPA should consider revising its demand-charges metric to incorporate data from many hours, not just the peak hour each month.

Infrastructure cost impacts. A problem was encountered early in the project with the general infrastructure costs function. Its impacts were initially incorporated in a way that unintentionally disincentivized the flattening of system load. This issue is discussed more in Section 2.4.

2.3.5 Step 2 Evaluation Conclusions

- The PNWSGD used a centralized "informed simulation" to emulate the dispatch of generation resources and their impacts on the delivered costs of electricity.
- Toolkit functions, working in conjunction with the informed simulation, specified how much of each type of dispatched energy was to be modeled in the system and this resource's impact on the delivered costs of energy in the transactive system. The project was able to reproduce the power and costs introduced by each resource through its data-collection system. Consumers' energy behaviors may be influenced by the way that resource costs are monetized by the functions.
- Toolkit functions may have merit as a template for distributed calculations. A defined set of output coefficients from the functions served as an interoperability interface in the transactive system. The system must be tested using more distributed nodes to fully confirm the value of the construct.
- The project did not correctly understand and respond to BPA's objective for improved integration of wind power. Future implementations must track and disincentivize the depletion of balancing reserves, which turned out to be the real challenge of wind integration for BPA.
- Incentive functions were similar to resource functions. The project failed in its attempt to design and implement an incentive function for the mitigation of transmission congestion. The demand-response incentive functions were more successful, but further improvements are needed.
- The costs of infrastructure were included in the incentive signal. Unlike locational marginal pricing, the transactive system strived to represent all costs of the delivered energy, not just the marginal costs.
- The dispatch of the transactive system's modeled resources was discrete and at times created discontinuities in the incentive signals. The project hypothesizes that some smoothing might occur in richer transactive systems that have more, and more independently acting, nodes and resources.

2.4 Step 3: Costs and Incentives Must Be Meaningfully Blended and Distributed through the System

Having reviewed the way that the transactive system emulated the dispatch of energy resources and monetized the various resource components, the next analysis step evaluates how those influences were blended and distributed throughout the transactive system. There are certainly many ways that the signals and their conveyance through the system could have been designed and accomplished, and the best method remains debatable. This section will simply remind the reader of the system's specification and affirm that the project's design was adhered to.

2.4.1 The Transactive Incentive Signal Is a Blended Cost

Each nodal location "owned" a unique TIS time series that represented the blended costs of all of its available resources at each time interval. Its TIS represented the unit cost of the energy that was either consumed there or was exported from there to another transactive neighbor node or through an exchange boundary.

The TIS equation—Equation (2.1)—is from the Transactive Coordination Signals report (Battelle Memorial Institute 2013, p. 2.8). The TIS was calculated—blended—at each node by summing all energy-related costs at the nodal location and dividing that total cost by the total energy resources that were available to the node during the interval. The costs may include the costs of generated energies, cost impacts of power capacity during an interval (demand charges, for example), or pure monetary impacts (resource startup costs, for example). In addition, offset costs shown in Equation (2.1) proved useful to represent bulk infrastructure costs. Total energy resources refer to all of the generated and imported energy that is available to be consumed or exported from the nodal location. The resulting units of measure for the TIS are dollars per energy (e.g., $/kWh).

$$TIS = \frac{energy\ cost + capacity\ cost + other\ costs}{total\ energy\ resources} + offset\ costs \quad (2.1)$$

The project collected all component costs and energy quantities that had been used at each nodal location and for each 5-minute data interval, and the project affirms that its calculations adhered to Equation (2.1). A TIS can be recalculated to confirm its value at any system location and time.

2.4.2 Distribution of Paired Energy Quantity and Unit Price Confirmed

Equation (2.1) is recursive in that the costs of energy from transactive neighbor nodes were also necessarily represented in the calculation. Neighboring transactive nodes are required to share their TIS (i.e., a unit energy cost) with one another. The two neighbors must also negotiate and resolve, through iteration, the power that is to be exchanged between them—the TFS.

The TIS of the node that *receives* power from the other is affected by the transactive node that supplies the power. As the recipient node uses Equation (2.1) to calculate its TIS, the supplying neighbor's TIS is among the energy costs, and the quantity of supplied energy—the TFS—is included among the summed total resources.

Therefore, the influences of energy costs and incentives were distributed through the system in the direction of power flow and in proportion to the magnitudes of energy that will flow between the system's nodes. This distribution of influence is confirmed again by the fact that the project can accurately recalculate a TIS at any system location and time.

2.4.3 Lessons Learned Concerning the Blending and Distribution of Incentive Signals

The TIS of the node that *receives* power from the other is affected by the transactive node that supplies the power. The term worked as intended, evenly allocating a constant dollar cost at each transmission zone where the toolkit function had been implemented. The unintended consequence was that, when this term was divided through by total resource energy (see Equation [(2.1]), an undesirable inverse relationship was created for the TIS, which is expressed as a unit cost of energy. That is, the unit cost of energy became smaller when the node had large total resource energy and greater when it had little. This had the unintended consequence of disincentivizing energy consumption when less energy was being generated and consumed. The preferred impact would have disincentivized energy consumption during peak load, thus helping flatten the system load.

The undesirable inverse relationship was fixed by moving the constant cost to the cost offset term, which is unaffected by the magnitudes of resource energy (and load). This correction also demonstrated the flexibility of the transactive incentive calculation. The same infrastructure costs were represented before and after the correction, but the dynamics of the costs could be changed to incentivize preferred energy consumption.

<u>Iterative solution required</u>. As the transactive system was being formulated, there was much debate about whether the impacts of changes in the system would adequately permeate throughout the system. The system's electrical connectivity (Appendix B) defined a network of peer-to-peer communication pathways. Must the timing of communications then be ordered and controlled to ensure that the impacts extend beyond the nearest neighbors? In the end, simplicity won out, and the timing of most communication events became scheduled at 5-minute intervals. Simple logic was adopted to receive anticipated signals from neighboring nodes. The simple timing approach worked for now because the topology was small and shallow, and risks could be managed. In addition, influence within the system was found to fall off quite quickly with distance in the transactive system.

A compromise was the design of relaxation[1] logic that would instigate further rounds of signal exchanges if received signals were found to have modified output signals by more than a configurable threshold. This approach worked. Iterations were occasionally found to have happened, but they were infrequent. The system converged quickly because, in part, the balancing responsibilities of the nodes were deferred (as discussed in Section 2.9), and the incentives of the transmission zones were centrally calculated by the informed simulation, not calculated in a distributed fashion.

Consensus is growing among the project's system implementers that the exchange of signals in the transactive system should become more event-driven and less timed. Iterative calculations and exchanges of the transactive signals will be needed.

[1] The word *relaxation* is borrowed from a well-known simulation solution technique. The convergence of the system is beginning to resemble that of other problems that are set up to solve by relaxation methods. Relaxation methods often employ criteria that cause more iterations to occur where they are most needed.

2.4.4 Step 3 Evaluation Conclusions

- Equation (2.1) guided the blending and distribution of energy cost influences in the transactive system. The equation, while similarly implemented at each system location, provides a great deal of flexibility for implementers to represent the costs of energy resources while also incentivizing desirable dynamic energy behaviors.
- An undesirable inverse relationship was at first created by the way the costs of infrastructure were modeled in the system. The influence was corrected using an alternative representation of costs in Equation (2.1), but the correction did not break the system design.
- The demonstration topology was probably not rich enough to confirm the validity of the combination of time-based intervals and event-based iterative calculations. Future systems should probably become more event-driven, as was exemplified by the systems relaxation criterion. These issues might be resolved by simulation.

2.5 Step 4: Responsive Loads in the System Must Be Able to Allocate Appropriate Responses Using the Incentive Signal

Presuming that the incentive signal is meaningful and is a representation of the actual cost of energy at a location in the transactive system, is a responsive load system able to discern response event periods using the incentive signal and other locally available information? For example, can a system of responsive water heaters select the no more than five useful curtailment periods each month that the customers had been promised? This question refers specifically to the responses determined by toolkit load functions, a module of the transactive system at a transactive node. These toolkit functions are diverse, but the following categorizations proved useful:

- Event-driven − the challenge is to allocate a limited number of allowed events and limited event durations over a relatively long period like a month or year. The event-driven function therefore anticipates and responds to monthly peaks, for example, in the TIS.
- Daily − the challenge is to allocate a limited number of events and limited event duration each day. The daily function therefore anticipates and responds to daily peaks in the TIS. Often, these functions are configured to respond differently (or not at all) on weekends, weekdays, and holidays.
- Continuous − a continuum of allowed responses is possible based on the real-time assessment of the relative magnitude of the transactive system’s transactive signal. Battery energy storage was the only type of asset in the PNWSGD that responded this way. The continuous functions may be configured to constrain the responses of the asset for their given power and energy capacities and other of the owners’ operational preferences.

The output signal from the toolkit load functions was called an advisory control signal (ACS), a signed byte advising assets when to respond and by what relative magnitude. Most responsive asset systems responded in a binary way and had the capacity to either curtail their loads (ACS = 127), or not (ACS = 0). Battery systems, in principle, could be advised to discharge at full power (ACS = 127), charge at full power (ACS = −127), or respond at any charge or discharge power level between these two extremes.

The project preferred that the loads be made automatically responsive to the advisory control signals, but not all systems were amenable to automation. The connection between the functions' advice and assets' actual performance was often tenuous. This section reports the functions' output—the status of the advisory control signals—even though many of the asset systems were found to have, in fact, often ignored the advice from the transactive system.

2.5.1 Event-Driven Function Events

Event-driven functions were designed to anticipate and generate a given number of events over relatively long periods (e.g., a month). The allowed numbers of events, the minimum and maximum durations of an event, and the sum of all of the event durations in the longer period were configurable. The occurrences of events could be configured differently and allowed or disallowed by hour, weekday, and holiday. For example, many of the PNWSGD utilities were subject to demand charges during defined HLHs, and the functions at these locations should have therefore been configured to preferentially respond during these HLHs if that was the preference of the asset's owner.

The event-driven functions individually maintained histories of the local incentive signals' statistics; the statistics helped the functions learn which incentive signal values were high, low, or normal. Thresholds were managed by each event-driven function to help it detect the TIS magnitudes at which the function's asset should respond. As the allowed numbers of monthly responses were used up, the threshold was then readjusted to best use the remaining events within the remainder of the present month.

The allowed numbers of events and event durations were often defined according to agreements that utilities or aggregators had made with their customers. Through configuration of the functions, each utility could enforce that its customers not be inconvenienced more than they had been promised.

In the PNWSGD, event-driven functions were frequently combined with asset models that represented electric water heater curtailment, thermostat setback, distributed generator control, dynamic voltage management, or in-home and portal notifications. The role of the asset models in the transactive system was to predict the impact of the assets' responses on load if they indeed responded at the times that become determined by the event-driven toolkit function.

Altogether, the performances of 14 event-driven functions are combined in Figure 2.33. Persons familiar with electricity supply will recognize the similarity of this figure to price duration curves that show an ordering of unit price of energy from the most expensive hours (left side) to least costly (right). This similarity is intentional. In this case, all of the interval values of the TIS during 2014 have been ordered from highest unit cost (left) to least unit cost (right).

For each 5-minute data interval of 2014 and for each event-driven function, a response flag was paired with the local transactive site's TIS value. Because the definition of relatively high and low TIS values may differ from site to site, the TIS values were transformed at each site to numbers of standard deviations above or below the site's average TIS. The interval pairs were then combined from 14 event-driven functions and the pairs were ordered from greatest to smallest relative TIS. Finally, the events of the event-driven functions were summed from smallest relative TIS (right side) to greatest (left side) and were scaled to represent cumulative event hours for the event-driven functions across the entire transactive system.

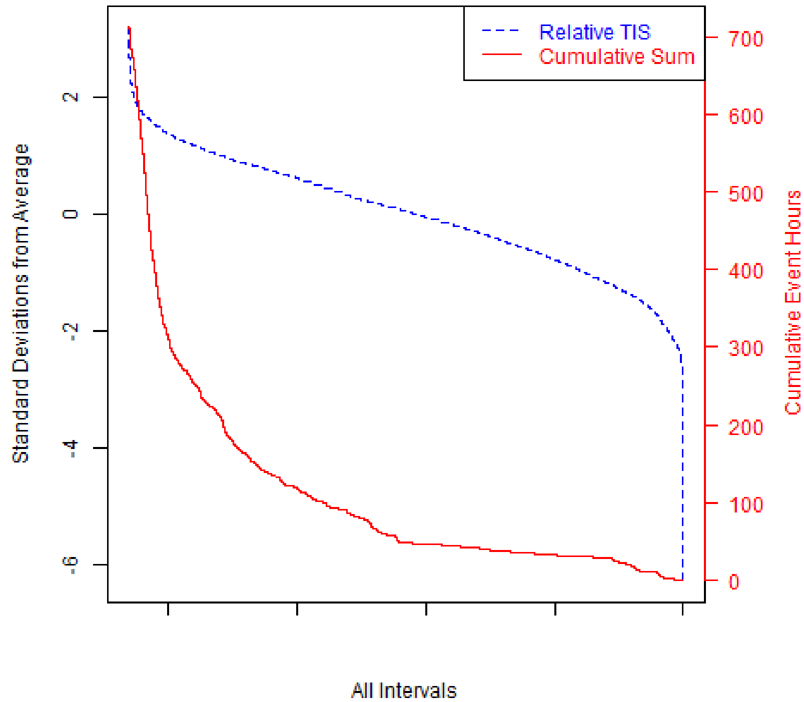

Figure 2.33. Ordered Relative TIS, Stated as Numbers of Standard Deviations from the Average TIS, Paired with the Cumulative Sum of Event Hours from all the Event-Driven Toolkit Functions during 2014

The relative TIS values were quite normally distributed. This figure excluded about 10 high TIS values that had occurred at certain sites. Such high values sometimes occurred at sites that had deployed demand-charges toolkit functions to help them anticipate and lessen their monthly demand charges. The demand-charges functions were found to apply very high costs (disincentives) as demand was peaking, but the additional costs were applied over only several 5-minute intervals, thus causing spikes in the local TIS.

The sites also observed several zero TIS values, which were not removed from the data for Figure 2.33. By 2014, the project had prohibited negative TIS values throughout the system. Earlier in the project, erroneous TIS values, including negative values, had plagued the system.[1] Naturally, the problematic TIS values had caused many of the assets to respond at nonsensical times. Therefore, the figures in this section include only January through August 2014 project data.

Now looking at the cumulative event hours in Figure 2.33, curtailment event periods occurred preferentially at high relative TIS values, as was intended. Overall, the 14 event-driven functions advised about 700 hours of asset responses in 2014. That cumulative duration accounts for about 6-1/4 active hours, on average, per asset per month. Had the events been randomly selected, the cumulative sum of event hours would have been nearly linear. The project did considerably better than that.

[1] Fundamentally, there is no reason to constrain the system from applying negative TIS values, but the project chose to constrain the system to encourage more stable behavior as the system was being designed and tested.

About 50 of the event hours occurred at TIS values that lie below the median relative TIS value. These event periods were undesirable and might have been avoided by better function design and better, more attentive configuration of these toolkit functions. For example, at least one event-driven function had been configured to allow practically unconstrained event durations, and the function therefore advised that events should continue almost indefinitely. On the plus side, about 85% of event periods were advised by event-driven functions while the relative TIS was in its highest quartile.

Panels of all of the individual event-driven functions' performances have been provided in Figure 2.34 to help demonstrate the range of individual performance by these functions. Because each asset was influenced by only one site's TIS, there was no need to normalize the TIS values at the individual sites. Observe that differences exist in both the TIS ranges and patterns of occurrences at the different transactive sites. As was discussed above, very large TIS values were discarded from the set of intervals at several of the sites, even though the values may have meaningfully resulted from a demand-charges function. All sites encountered intervals when the TIS value was zero.

The cumulative response hours for each asset and site need not be individually discussed. Discussion preceding Figure 2.33 should have provided enough background for the interpretation of the panels of Figure 2.34. The best performing event-driven functions correctly identified the TIS intervals representing the greatest delivered unit costs of electrical energy. Ideally, all event-driven responses should have been advised while the TIS was at its very greatest values, far to the left side of these figures. Panel Figure 2.34c exemplifies decent performance; panel Figure 2.34e, not as good.

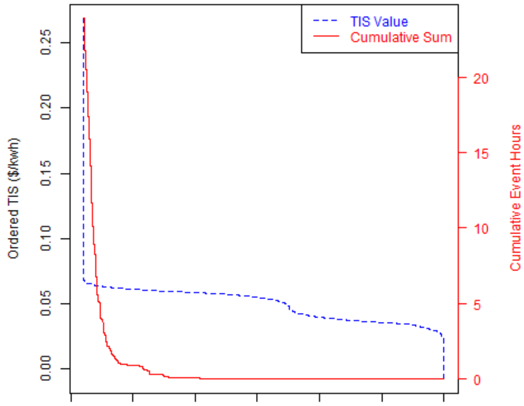

(a) University of Washington: Building HVAC

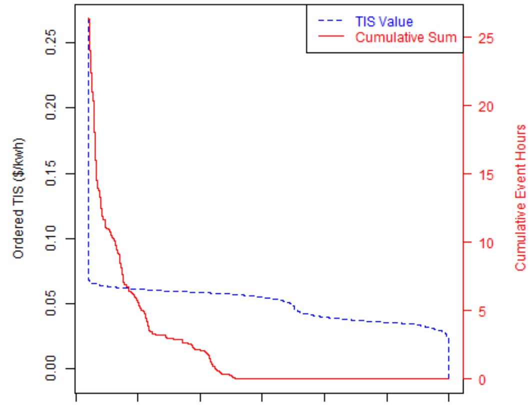

(b) University of Washington: Diesel Generators

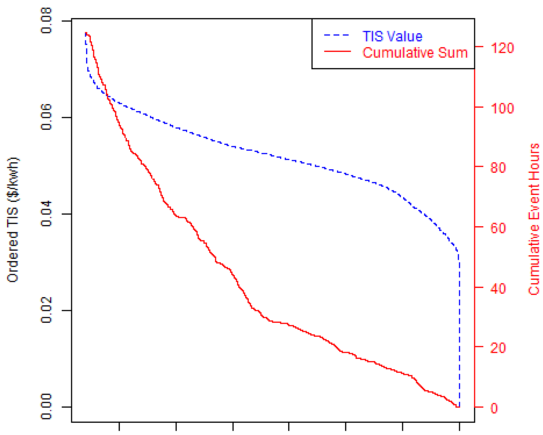

(c) Avista Utilities: Residential Demand Response

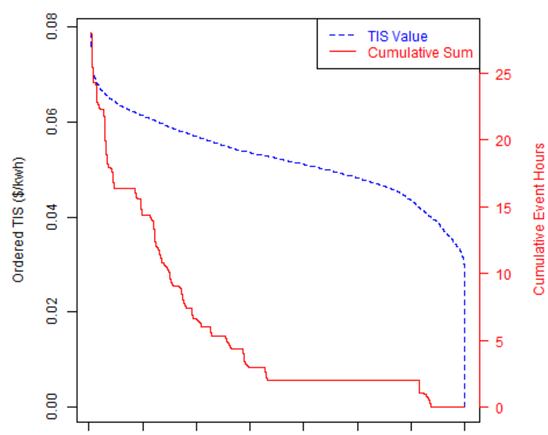

(d) Flathead Electric Coop.: Water Heaters

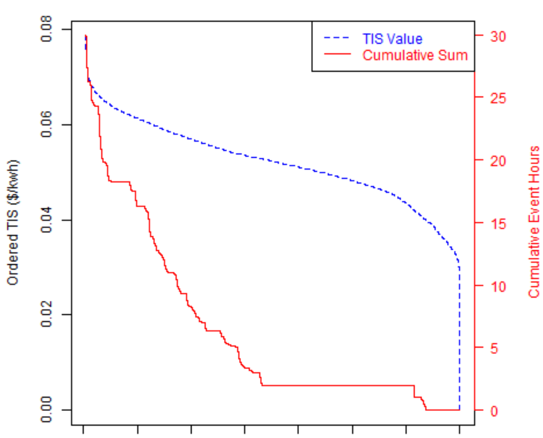

(e) Flathead Electric Coop.: Smart Appliances

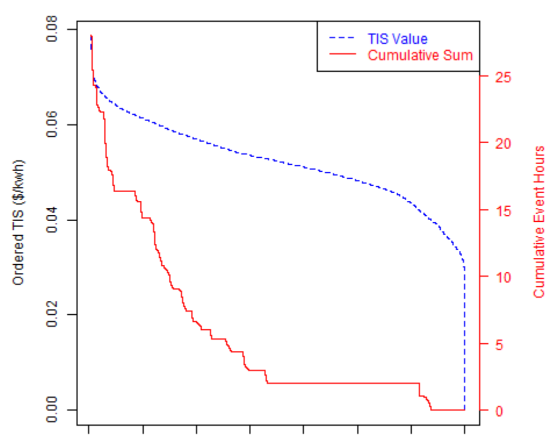

(f) Flathead Electric Coop.: In-Home Displays

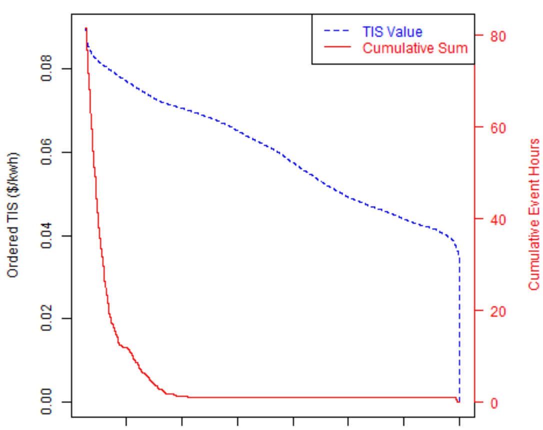

(g) Milton-Freewater: Demand-Response Units

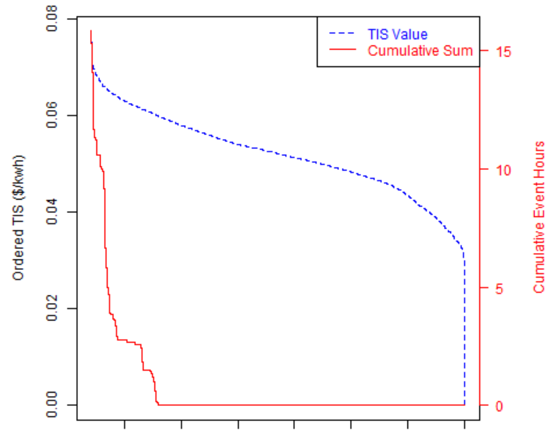

(h) Avista Utilities: WSU Chiller Control

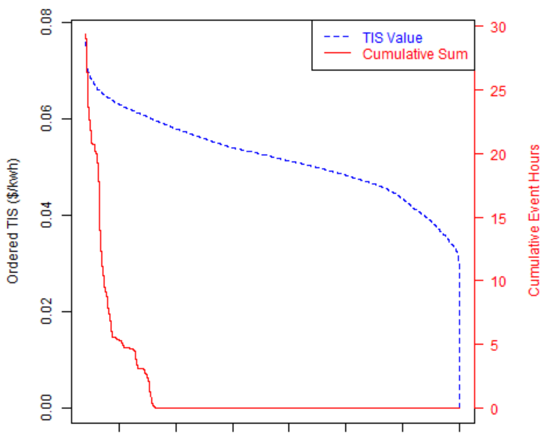

(i) Avista Utilities: WSU HVAC Control

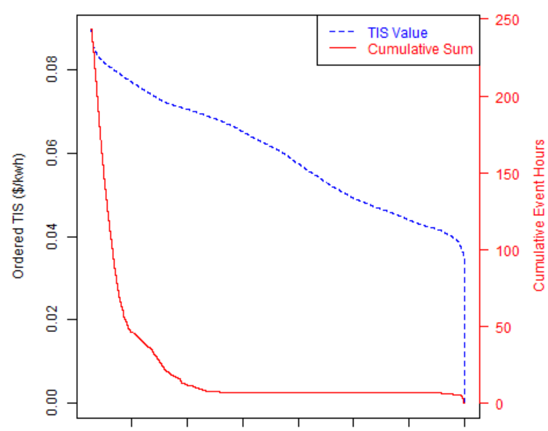

(j) Milton-Freewater: Dynamic Voltage Management

Figure 2.34. Cumulative Responses of Individual Event-Driven Assets to the Transactive System's Incentive Signal. (HVAC = heating, ventilation, and air-conditioning; WSU = Washington State University)

2.5.2 Daily Function Events

The *daily* functions were designed to select one or more event periods in a day. The function reviewed the TIS values that had been predicted more than 24 hours into the future and selected the event period when the delivered costs of energy represented by the TIS would be maximal. Because all of the controllable assets that were selected by PNWSGD utilities targeted *curtailment* of electric load, every project implementation identified the maximum TIS values, when the consumption of the most expensive energy might be avoided. Similar principals would guide the development of functions that would take advantage of minima in the TIS to preferentially consume energy during those periods.

The project had originally called these functions “time-of-use,” but that term did not adequately communicate the dynamic flexibility and economic elasticity that were possible and intended. Utility participants and implementers had preconceived interpretations of what “time-of-use” functions were and how they would behave based on the prior application of the term to demand-response programs. The alternative “daily” was recommended instead.

As for the event-driven functions described above, the daily-event functions could be configured to behave differently on different weekdays, to favor or allow responses certain hours, and to modify the allowed event durations. Some utilities, for example, configured the functions to allow daily events only during HLHs.

The daily-event functions were paired by the PNWSGD utility site owners with various asset systems and their models, including communicating thermostats, dynamic voltage management, networks of in-home displays or Web portals, distributed generators, and water heater control. The role of asset models was to model the impacts on system load when daily events were under way.

Figure 2.35 is a summary of all of the event periods that were designated by nine daily-event functions during 2014. This figure pairs ordered relative TIS interval values with the cumulative hours that the daily-event functions advised events during 2014. These intervals have been ordered from the most expensive intervals (left), when the TIS was at its greatest, to least expensive on the right. Additional details about the representation in Figure 2.35 were provided during the discussion of Figure 2.33 above and will not be repeated here.

The distribution of relative TIS values was quite normal during 2014 after a small number of very large TIS values were removed. The event hours were distributed through much of the range of the relative TIS, but the numbers of event hours increased with increasing relative TIS. About 2,500 total event hours are represented in this figure, which accounts for about 35 event hours per asset system per month, on average. The daily-event function typically designates more events and more active event hours than the event-driven functions. As expected, the cumulative sum of event durations is less steep and gradual than was the case for the event-driven functions (review Figure 2.33). That is expected because daily-event functions could correctly identify the lowest costs in a day, but the day’s energy may have been relatively inexpensive compared to the rest of the month or year. About three-quarters of the event hours were advised while the relative TIS was greater than its median.

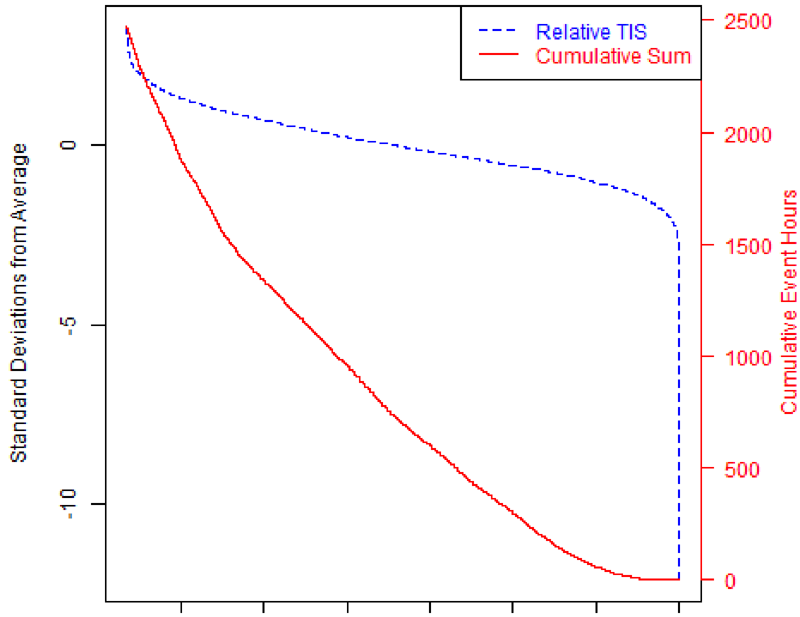

Figure 2.35. Ordered Relative TIS, Stated as Numbers of Standard Deviations from the Average TIS, Paired with the Cumulative Sum of Event Hours from All of the Daily-Event Toolkit Functions during 2014

The range of responses by the individual daily-event toolkit functions to their local TIS values is shown in Figure 2.36. The asset types have also been indicated for each. TIS values greater than about $0.25/kWh were removed from these representations because the several large values disallowed observation of the variability of the TIS values. The large TIS values had been generated at some sites by the demand-charges toolkit functions, which increased the local TIS values to deter consumption during monthly peak demand periods.

Had the daily events been determined randomly, the cumulative event hours would have exhibited a nearly linear relationship with the ordered TIS values. Again, the PNWSGD did considerably better than that. Generally speaking, the best performance of these functions is indicated by having the cumulative event hours pushed close to the left side of these figures.

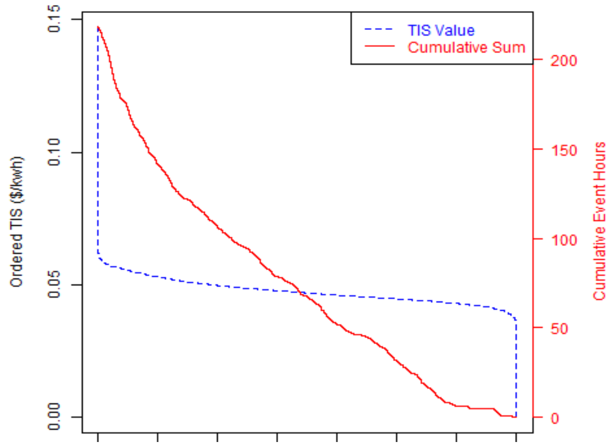

(a) Lower Valley Energy: Water Heaters

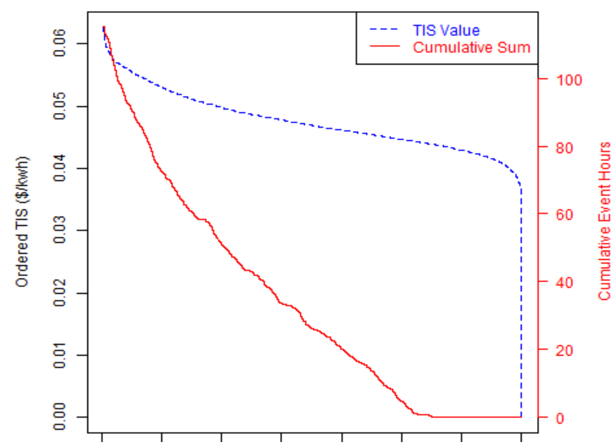

(b) Idaho Falls Power: Water Heaters

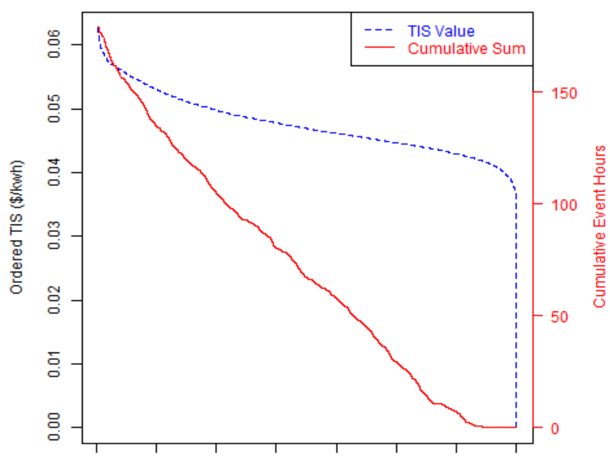

(c) Idaho Falls Power: Thermostats

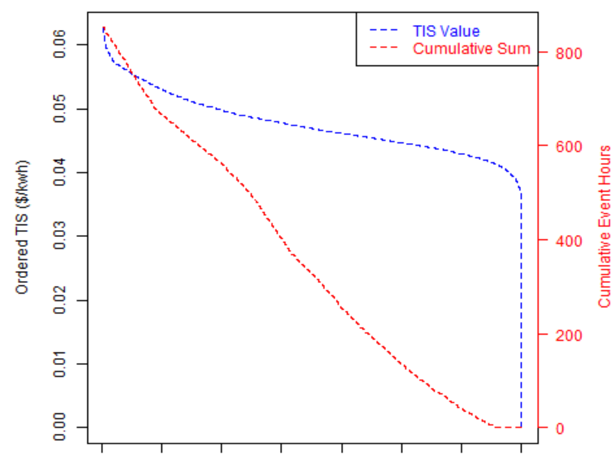

(d) Idaho Falls Power: Voltage Management

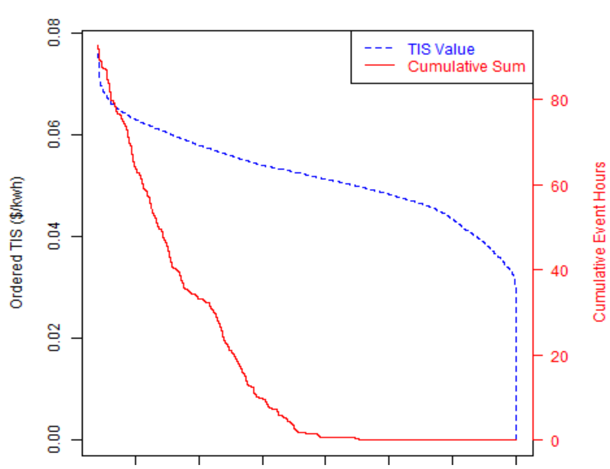

(e) Avista Utilities: WSU Gas Generator #1

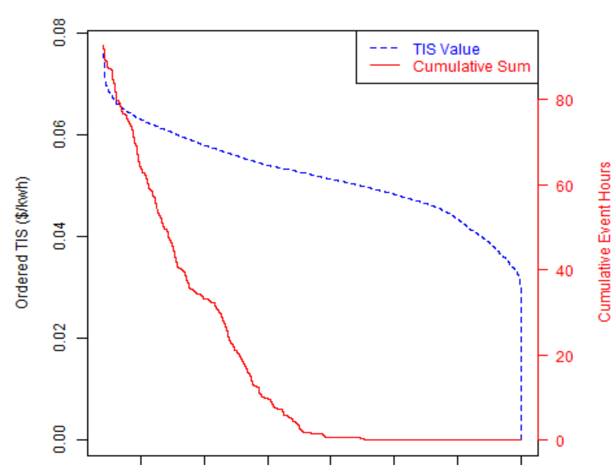

(f) Avista Utilities: WSU Gas Generator #2

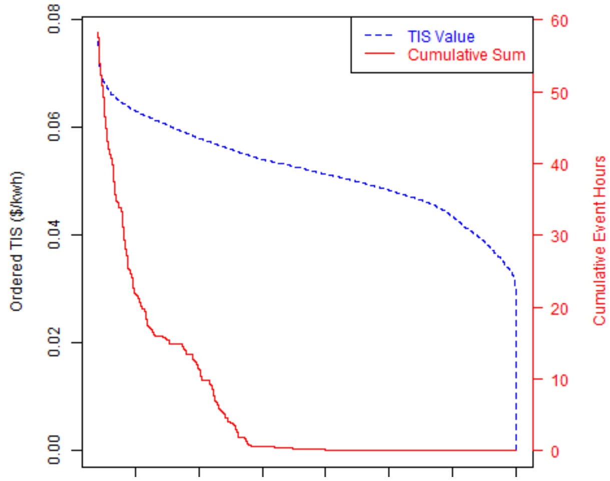

(g) Avista Utilities: WSU Diesel Generator

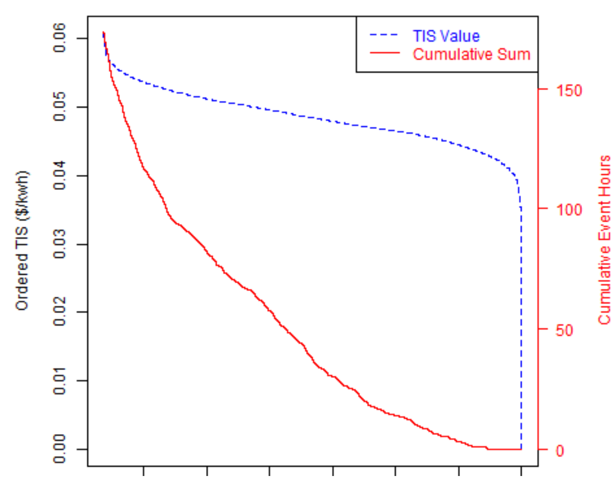

(h) Peninsula Light Company: Water Heaters

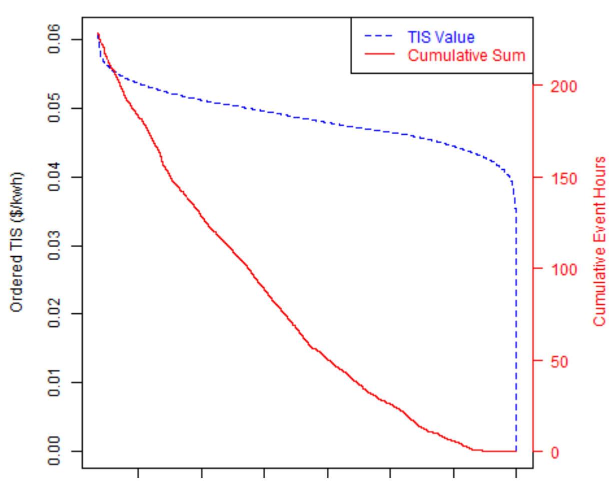

(i) Peninsula Light Co.: Voltage Management

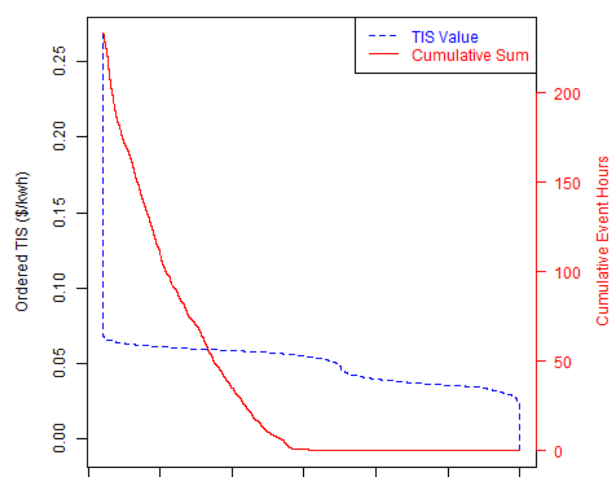

(j) University of Washington: Steam Generator

Figure 2.36. Cumulative Responses of Individual Daily-Event Assets to the Ordered Transactive System's Incentive Signal at Each Transactive Site during 2014. The event hours are accumulated from right (lowest TIS) to left (highest TIS).

2.5.3 Continuous Function Events

Continuous functions were designed to generate a dynamic range of responses—more or less power consumption each interval—based on the predicted TIS values. In the PNWSGD, continuous functions were applied to only battery energy storage systems. The function managed the battery systems' states of charge while optimizing an arbitrage of energy value. The function strove to recharge batteries with the cheapest energy and discharge the most expensive energy, according to the delivered costs of energy represented by the time-dependent TIS. The system was constrained, through configuration, by the batteries' allowed states of charge and by the power ratings of the batteries' converters. The utilities and battery venders constrained the systems even more according to their preferences and concerns about affecting battery life.

An interesting feature of the function's formulation was a dissipation term that tempered the responsiveness of the system. Without the term, the function advised rapid changes between charge and discharge modes nearly every time interval. Because the function was based on an optimization, the system might be advised to switch from full charge mode to full discharge and back after each 5-minute interval in response to small changes in the TIS. Battery vendors typically advocate much more gentle treatment of their systems. The dissipation term, once introduced, moderated the numbers of times that the system alternated between charge and discharge modes. For example, the term could be tuned to advise no more than one charge and discharge event per day.

A classical optimization solver was used by the function that represented the Lower Valley Energy battery energy storage system. Its cumulative charge and discharge durations are shown along with the ordered TIS at this site during 2014 in Figure 2.37. The ordering of TIS magnitudes for figures like this has been discussed already in conjunction with Figure 2.35 and Figure 2.33. The cumulative charge and discharge durations have been shown separately here. Discharge events should correspond to high delivered costs of energy and were accumulated from lowest costs (right side) to highest (left side). The function advised discharge mode using positive advisory control signals.

Charging events were advised at lower costs using negative signals. The charging event durations were accumulated from lowest costs (left side) to highest (right side).

Because a continuum of responses between full discharge to full charge was allowed, the right-side axis has been stated as "cumulative capacity hours," defined as the product of the fraction of the converter's full capacity multiplied by the duration over which the response was advised. In practice, the optimization function often requested that whenever the system was active it should charge and discharge at the greatest allowed power.

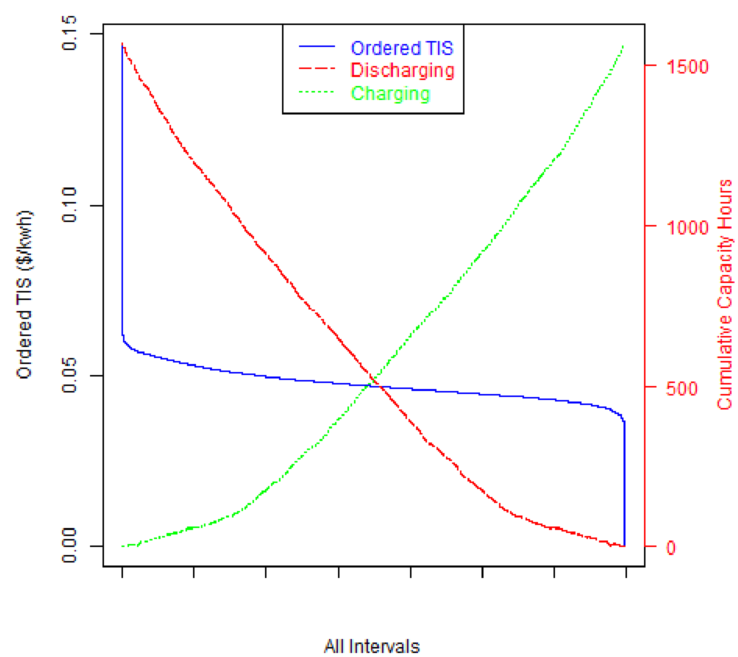

Figure 2.37. Discharging (Red) and Charging (Green) Capacity Hours Advised to the Lower Valley Energy Battery System during 2014

A continuous function was used with the Idaho Falls Power battery system, too, but that function was not updated after the system's vendor stopped supporting the battery system.

The continuous functions that represent battery energy storage necessarily monitor the system's actual state of charge. State of charge is one example of "other local conditions" that must be known by a toolkit function. If the state of charge is unknown to the toolkit function, the function's state will diverge from reality, and the function will incorrectly recommend charge and discharge events. When the function was misinformed about the actual system state, performance was even worse. For example, if the function was advised that the battery was fully charged when it was not, the function would never then advise the system to charge, even if that should have been the preferred control action.

2.5.4 Step 4 Analysis Conclusions

Three different classes of toolkit function were designed by the project and represented assets having event-driven, daily, and continuous event types. These three methods of determining events were found to be applicable to a wide variety of assets. Furthermore, the selections of events by these functions could be configured by the assets' owners to address the assets' specific capabilities and the asset owners' preferences.

The PNWSGD transactive system was, in fact, able to determine event periods based on the TIS and other local conditions. Toolkit functions—event-driven, daily, and continuous—were flexible and effective tools for accomplishing this objective. The functions must be well designed and configured if they are to perform well.

2.6 Step 5: Responsive Assets Must Accurately Predict the Impacts of Their Reponses

Presuming that event periods are being selected well by the toolkit functions and further presuming that the assets do indeed respond to the events, do the asset models accurately predict total load and the impact of the events on elastic load? This section evaluates the asset model algorithms and also tests the care with which the asset models were configured.

The discussion in this section has been divided into two components. At most of the project sites, the project collected data about the total site electric load that should be directly comparable to the transactive feedback signal, or TFS, that was intended to represent and predict that total load. Section 2.6.1 discusses how well the actual load was modeled by the TFS and whether the predictions were self-consistent. Section 2.6.2, discusses the accuracy of the individual models relative to whether the predicted changes in load were meaningful.

2.6.1 The Utility Sites' Demonstrated Abilities to Predict Their Total Load

Each of the project's 11 utility sites was asked to model and predict its total electric load. Total electric load, in this case, referred to either a defined subset of the site owner's distribution system that was participating in the project or the site owner's entire load.

The total electric load was predominantly inelastic—not affected by the price-like incentive signals of the transactive system. IBM worked with site owners to create and calibrate models of the bulk inelastic load at the project's sites. Using time of day, day of week, temperature, and a history of prior electric load at the site, IBM used regression to predict the sites' load curves, represented by a set of periodic smoothing-spline basis functions as described by Harvey and Coopman (1993). The predictions were updated every 5 minutes for each of the 56 time intervals of the project's transactive signals. This modeling approach was found to be computationally efficient.

Each elastic, demand-side asset system further modeled the *change* in load that would occur as it responded to changes in TIS. What is important here is that the responses from those assets changed the total predicted load at the site.

Based on an analysis of the TFS signals between the project's Fox Island, Washington site (Peninsula Light Company site ST01) and the West Washington TZ02 for 8 months of 2014, Figure 2.38 depicts, by month, the average relative error between TFS magnitudes as they were predicted for each of the 56 transactive future intervals and the final prediction of that interval (i.e., the first 5-minute interval at IST_0). The Fox Island site is used to support some general observations about the project's predictions by month, hour type, and minute. Then the available prediction errors at all the project's sites are shown.

The transactive prediction intervals shown in Figure 2.38 on the horizontal axis are displayed such that the prediction closest to the observation (e.g., the prediction 5-minutes preceding the observation) is displayed on the left-hand side of the plot, and the prediction with the greatest lead time (i.e., the estimate of this observation that was generated approximately 4 days prior) is displayed on the right-hand side of the plot. On the vertical axis, the relative prediction error from the observation value was calculated. Relative prediction error was calculated as shown in Equation (2.2).

$$\text{Average Relative Prediction Error for Interval } i = 100\% * \frac{1}{N}\sum_{n=1}^{N}\frac{x(n,i)-x(n,0)}{x(n,0)}, \qquad (2.2)$$

where N is the total number of predictions of this interval's value in the month (e.g., $N = 1$ for the other 5-minute intervals, but $N = 12$ for the hour-long intervals because the interval was predicted 12 times during those hours), n is one of the N intervals, i is one of the transactive signal's intervals from 0 to 56, and x is the predicted TFS value in kilowatts.

Graphing the relationship between the average prediction interval and the relative prediction error allows prediction quality to be assessed relative to the amount of time in advance that the prediction was made. If the prediction methods had been successful at eliminating all but random errors, the plot would be mostly flat. The plot shows data series for 8 months in 2014 (January–August) to compare prediction quality between these months.

The prediction error being addressed here references the nearest-term prediction that was produced just prior to an interval's final 5-minute prediction. That is why the relative prediction error of the first, far left prediction interval is always zero.

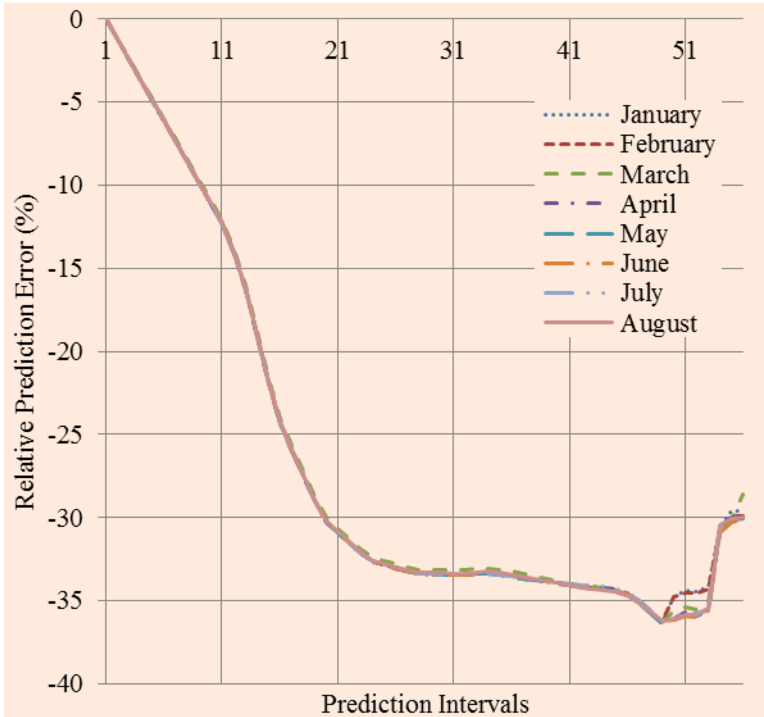

Figure 2.38. Average of the Transactive System's Relative Load Prediction Errors at the Fox Island Site (ST01) Site for the Eight Project Months of 2014

The magnitude of relative prediction errors increased the further into the future that the predictions were made, dropping to about a 35% underestimation of what would eventually become the final prediction. This is a bias in the prediction horizon. The load was found to be persistently under predicted. A prediction error greater than about 10% is probably harmful to resource planning. Only the prediction errors of the first nine intervals fall within 10% error or less, mapping to a successful look-ahead prediction horizon of about 45 minutes.

Every month, the relative prediction error between intervals 1 and 49 was the same. The transactive node stubbornly applied the same training set and methods each month, failing to learn from its prediction errors and adapt.

Figure 2.38 described an average relative prediction error over time, but the variability of the individual prediction errors over time must also be addressed. Figure 2.39 presents the standard deviation of those same relative prediction errors over the same 8 months of 2014.

The magnitudes of the standard deviations increased into the future. At the last prediction intervals, the standard deviation was about 25%. Magnitudes of the standard errors were almost as large as the average prediction error biases. The standard errors in March were a little greater than for other months. The standard errors the other 7 months were indistinguishable at most of the prediction intervals.

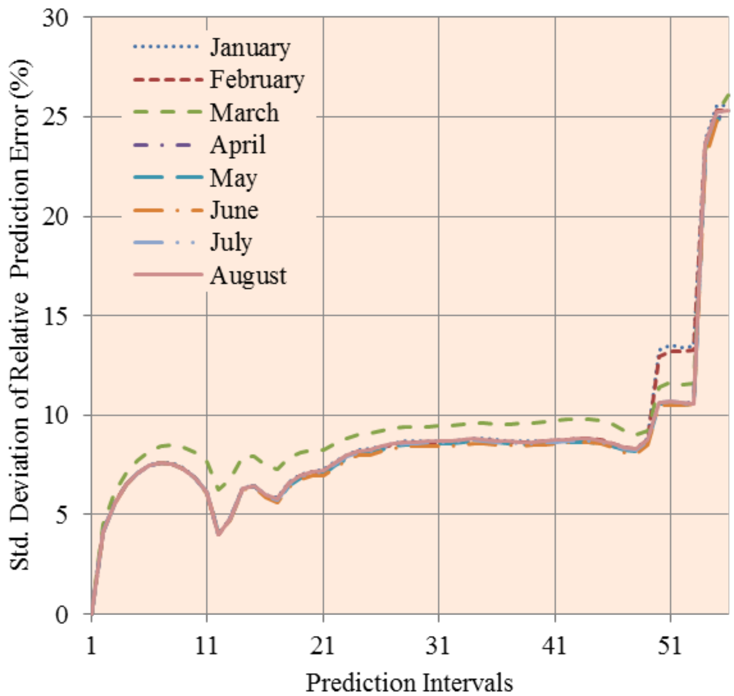

Figure 2.39. Standard Deviation of the Transactive System's Relative Load Prediction Errors at the Fox Island Site (ST01) for Eight Project Months of 2014

Figure 2.40 drives home the magnitudes of these standard deviations. The relative prediction error is shown for only March 2014. The error bars represent the standard deviations of the relative prediction errors during March 2014. It was shown that the monthly average errors and their standard deviations were similar for all months of 2014 at this, the Fox Island, Washington site.

The predictions at this site proved to be persistently under predicted, almost always lower than the final prediction. The ramification of this bias is that mistakes would be made while planning. This site underrepresented its future load. Had the transactive system and this region accepted and acted on the under predicted load, too few energy resources would have been scheduled and dispatched, possibly resulting in the purchase of more costly real-time resources than might have otherwise been procured.

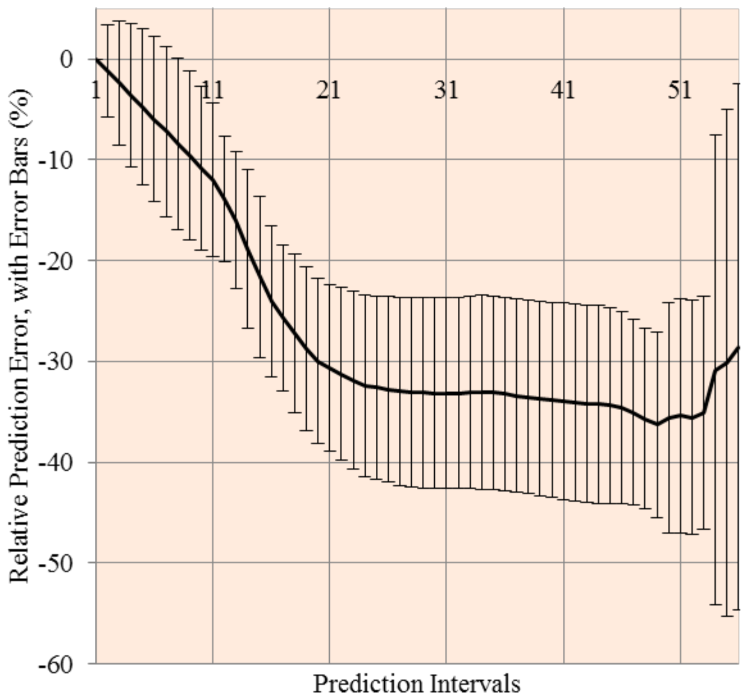

Figure 2.40. Average of the Transactive System's Relative Load Prediction Errors and Standard Deviations of the Average Prediction Errors (Error Bars) at the Peninsula Light Company Site for March 2014

Figure 2.41 depicts the same data as Figure 2.39, but the data is summarized by hour, removing day and month variability. For example, all data between 01:00 and 02:00 local time (i.e., 01:00, 01:05, 01:10, etc.) is denoted as Hour 1 in Figure 2.41. The data set included all 5-minute intervals from January 1 to September 1, 2014. The predictions during Hour 0 (midnight to 01:00) provide the best load prediction (an under prediction of 20% or better) and Hour 3 (03:00 to 04:00) the worst at the Fox Island site this year. Clearly, there is room for improvement.

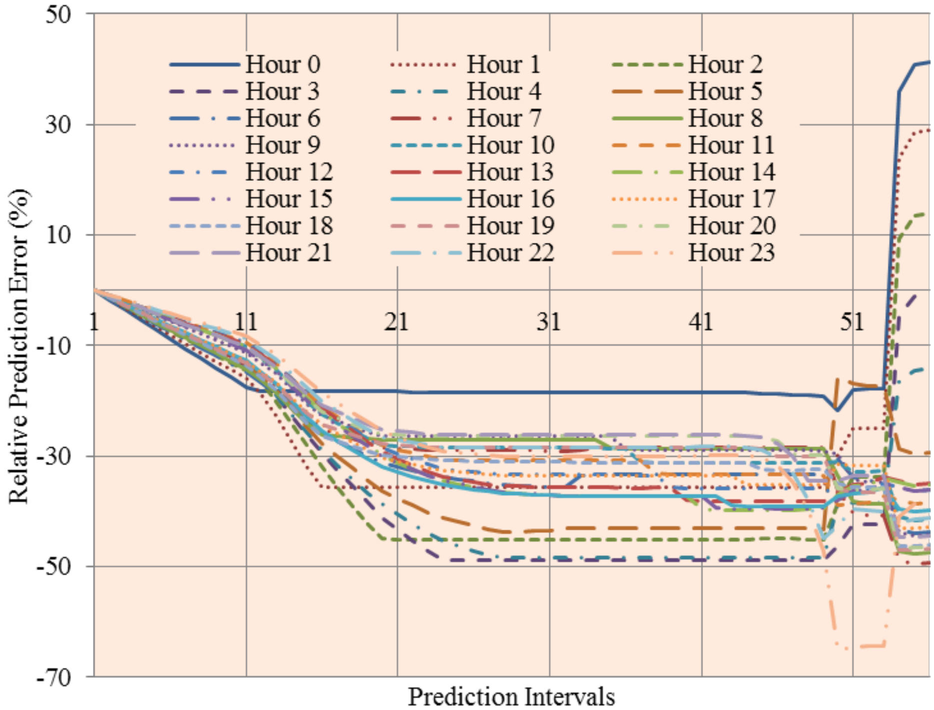

Figure 2.41. Average of the Transactive System's Relative Load Prediction Errors at the Fox Island Site (ST01) by Local Starting Hour. All the 5-minute intervals from 2014 were used.

Figure 2.42 depicts a further breakdown of the hourly summary for HLHs (left) and LLHs (right) at the Fox Island site in 2014. The axes are the same as those defined for Figure 2.41 above. BPA defines LLHs between 10:00 to 06:00 Pacific Time Monday through Saturday and all day Sunday; HLHs are the remainder of the hours. Consider HLHs at Interval 12 (future prediction of load ~1 hour prior to its occurrence). At this point, the prediction value has a percent difference of ~10% load underestimation. The relative load predictions for both hour types quickly become under predicted. The biggest difference between the predictions for the two hour types occurs in the last intervals that predict load multiple days into the future. As might be reasonably expected, the HLHs are under predicted by the day-long intervals and the long intervals over predict the LLH loads.

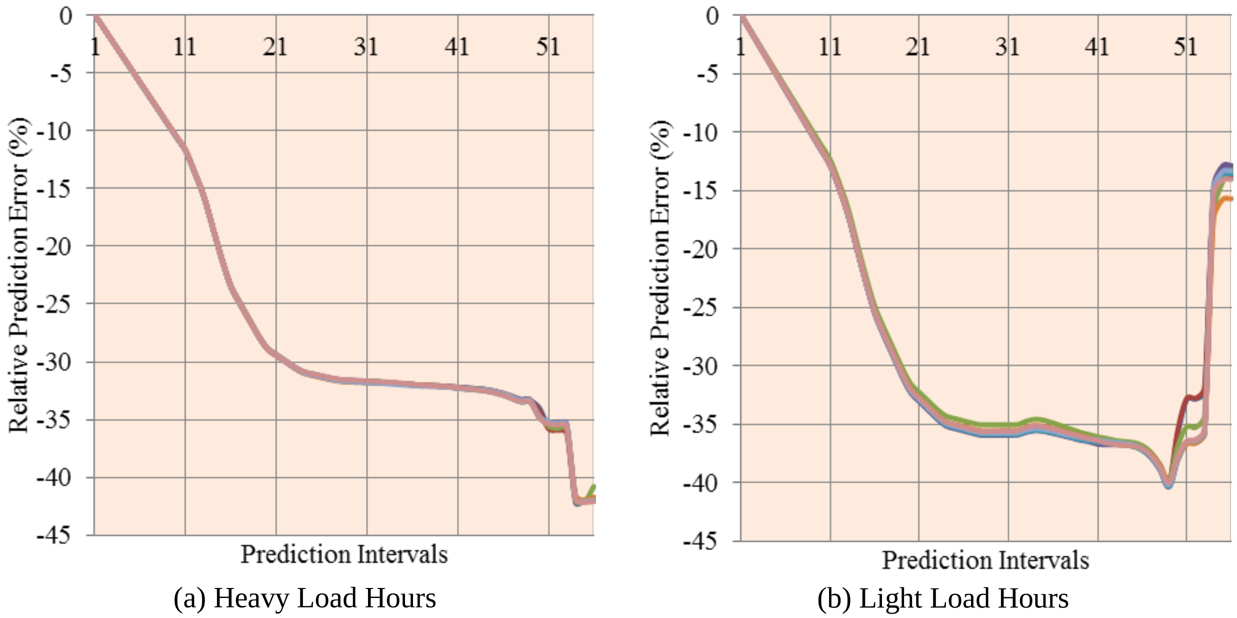

Figure 2.42. Average of the Transactive System's Relative Interval Prediction Errors during 2014 for the Fox Island, Washington Site (ST01) for (a) HLH Hours and (b) LLH Hours. Results from 8 months are shown. Small distinctions between the months' results are not critical to discussion.

Analysis of weekday versus weekend and holiday versus non-holiday periods did not reveal any particularly interesting distinctions in the relative prediction intervals' accuracy.

Figure 2.43 depicts the same data as Figure 2.39, but the data is summarized based on an interval's starting minute, removing hour, day, and month variability. For example, all data that start 5 minutes past an hour (i.e., 01:05, 02:05, 09:05, etc.) are denoted as Minute 5 in Figure 2.43. The data set used is all 5-minute intervals from January 1 to September 1, 2014.

Overall, data predicted for time intervals that begin on an hour or 5 minutes past an hour are predicted more accurately than others. The "waterfall" effect seen in the first 12 intervals is an interesting phenomenon and shows some weaknesses in the predictions' implementations. This effect is probably caused by the dynamic representation of 5-minute data in a system where transmission dynamics were being updated instead no more often than hourly.

Minute 55 was accurate throughout the first 10 intervals, almost an hour into the future. As the clock minutes decreased (with the exception of Minute 0 and Minute 5) the predictions more quickly diverged from the final predicted values. This finding was unexpected. Clearly, more work is needed on predictive algorithms.

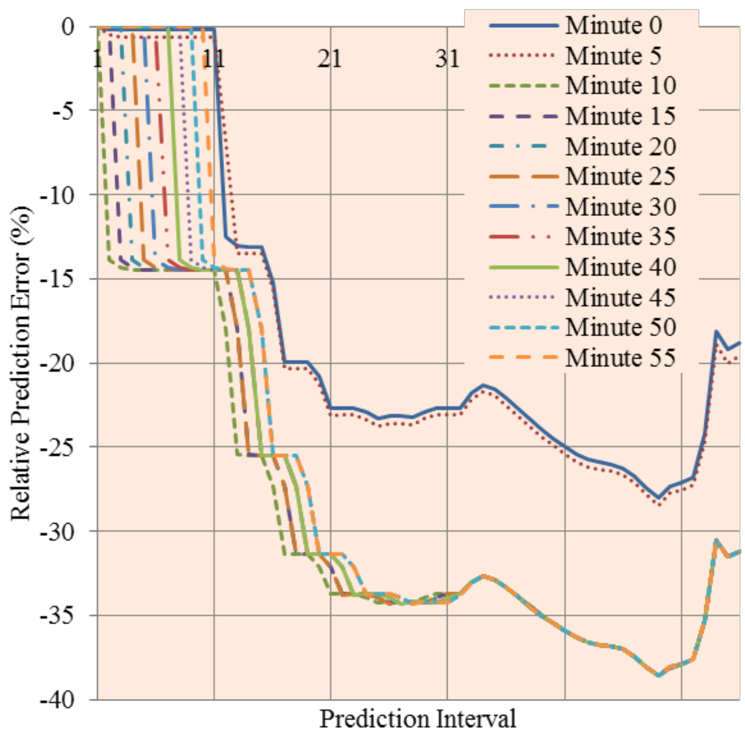

Figure 2.43. Average of the Transactive System's Relative Load Prediction Errors at the Peninsula Light Company Site by Starting Minute of the Hour. Data includes all the intervals of January through August 2014.

Having examined the relative accuracy of prediction errors at one project site according to the data interval's month, hour, and minute, we now sample the remaining sites to understand the variability of this relative prediction accuracy across the set of project sites of the transactive system. Figure 2.44 presents such a sampling across the sites for which transactive data became available to the project. Two utility sites therefore have been omitted—Ellensburg, Washington (ST04) and the Benton Public Utility District Reata Substation (ST05)—because these two did not become active transactive system sites during the PNWSGD.

For each site, a pair of figures shows (left) the average relative prediction error and (right) the standard deviation of the relative prediction error. These two paired plots were introduced using the Fox Island site (ST01) in Figure 2.38 and Figure 2.39, respectively. The left-side plots show bias errors, where the transactive system tended to persistently under or over predict the final predicted value at times into the future. The standard deviations refer more to the dynamic variability with which those predictions were made.

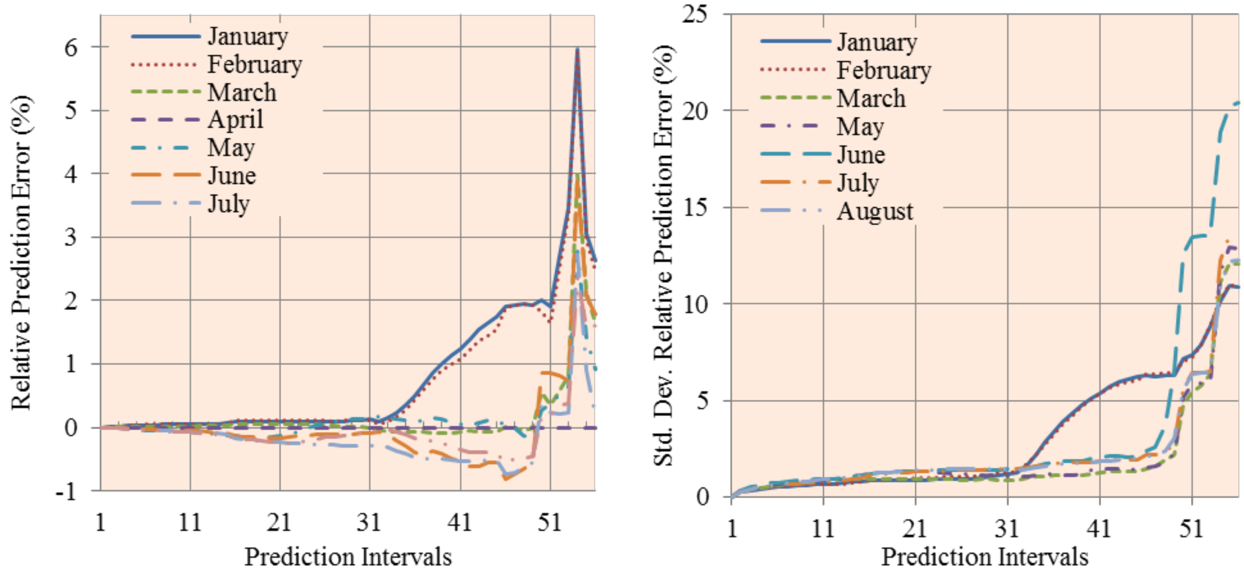

(a) University of Washington Campus Site (ST02). This system was down during April 2014.

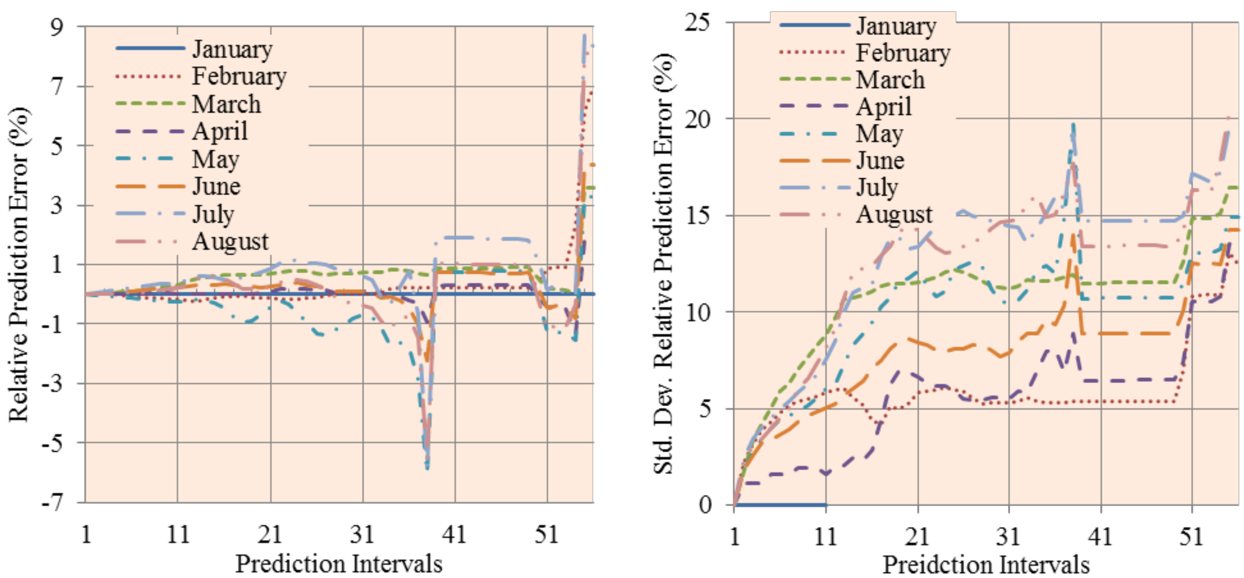

(b) Salem, Oregon Site (ST03)

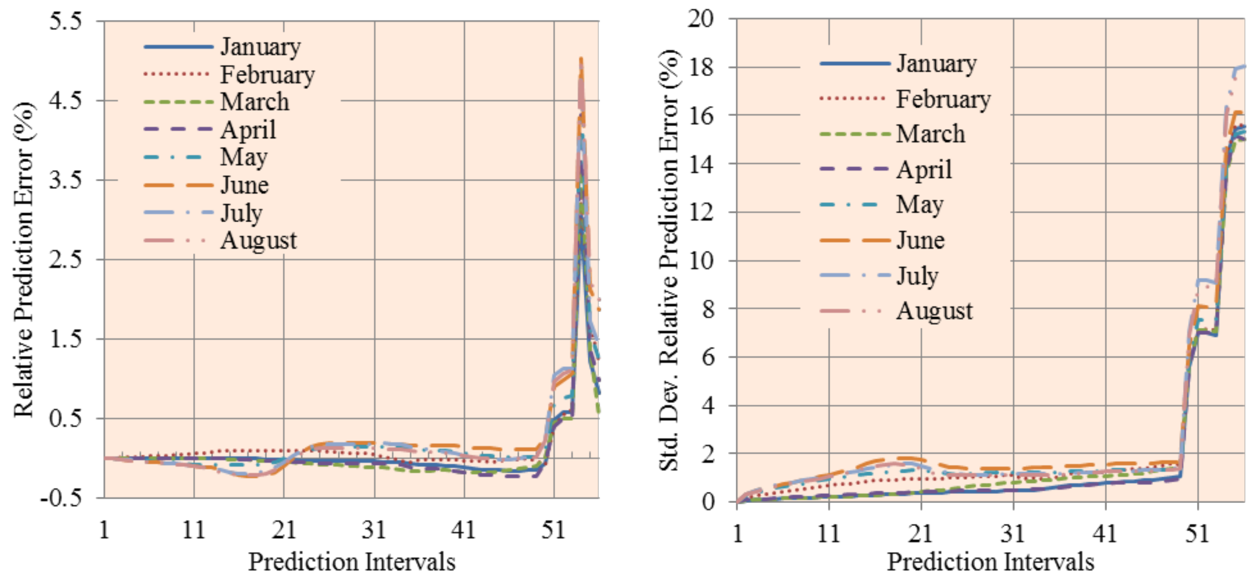

(c) Pullman, Washington Site (ST06)

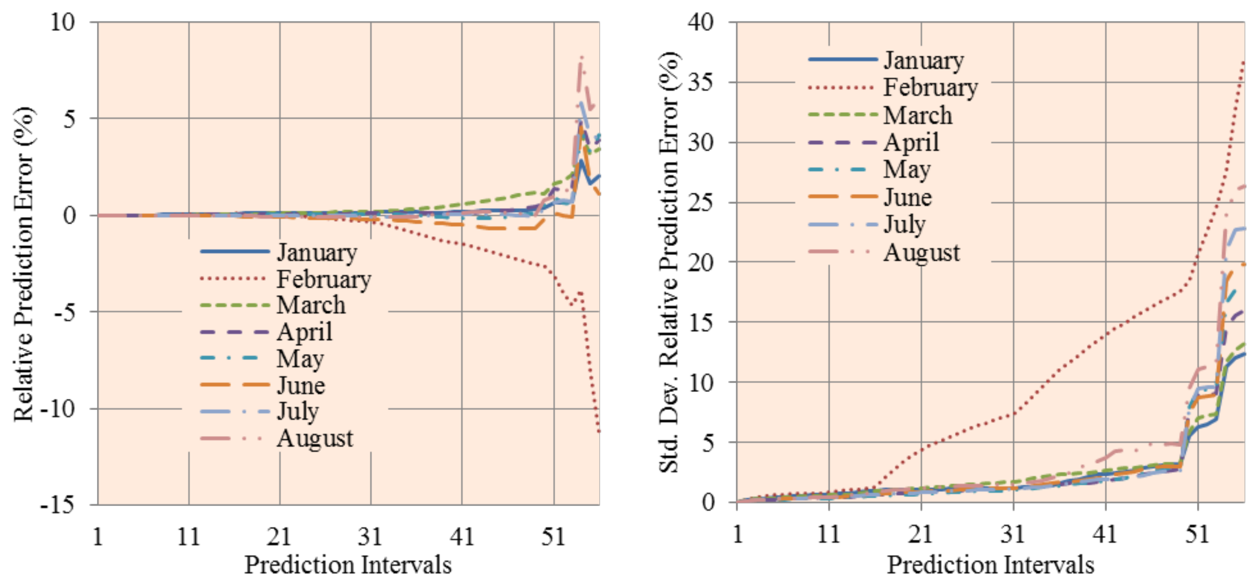

(d) Libby, Montana Site (ST07)

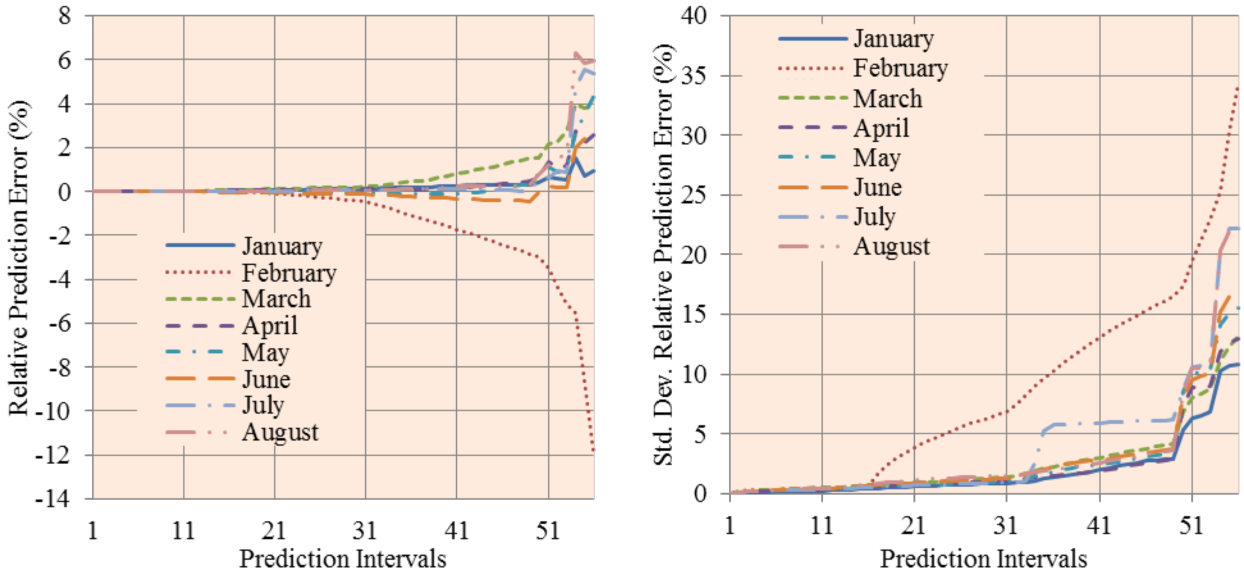

(e) Marion-Kila, Montana Site (ST08)

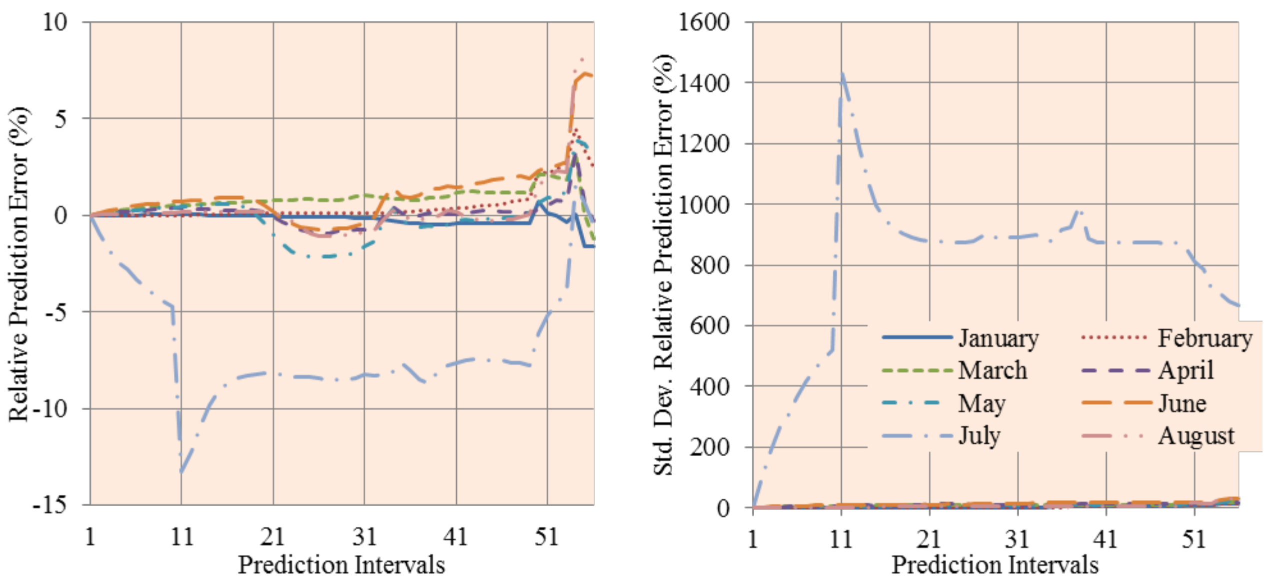

(f) Milton-Freewater, Oregon Site (ST09)

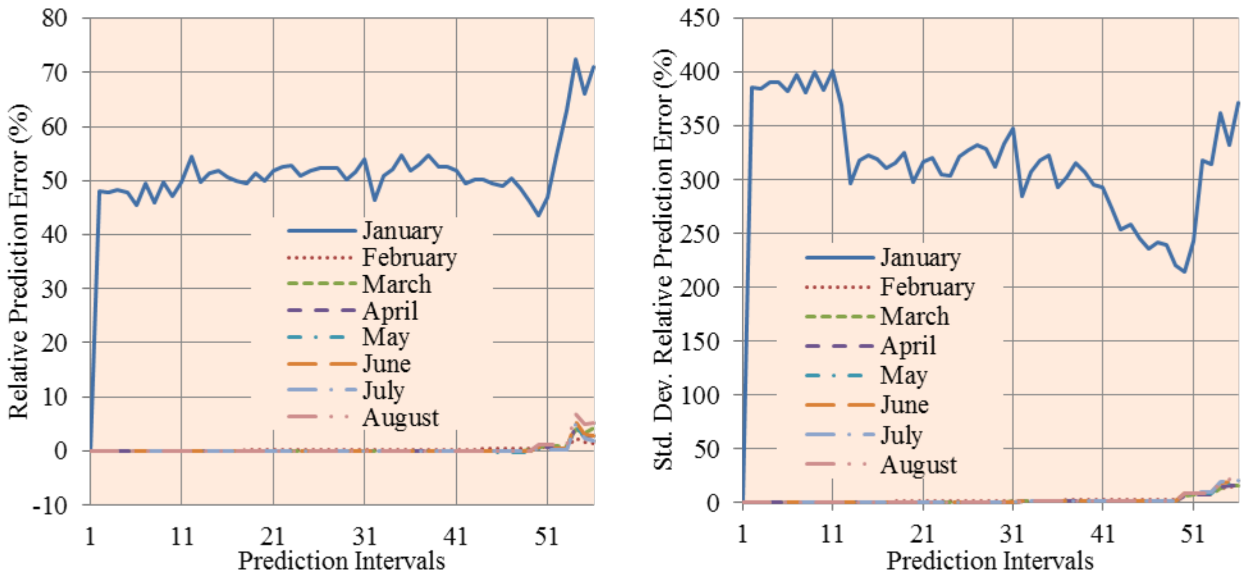

(g) Helena, Montana Site (ST10)

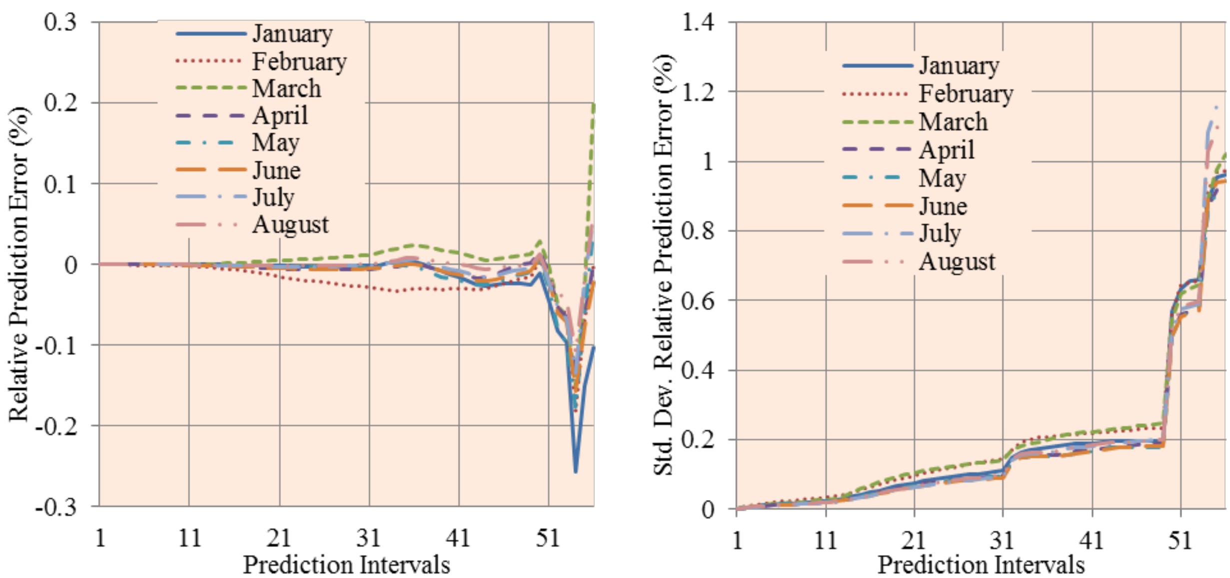

(h) Philipsburg, Montana Site (ST11)

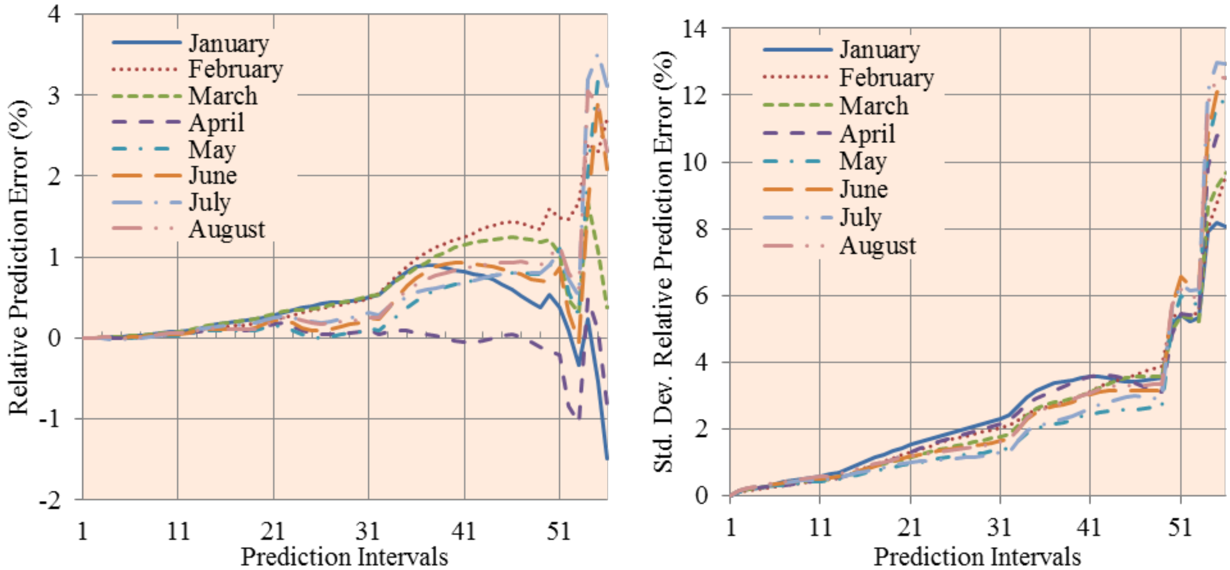

(i) Teton-Palisades, Wyoming Site (ST12)

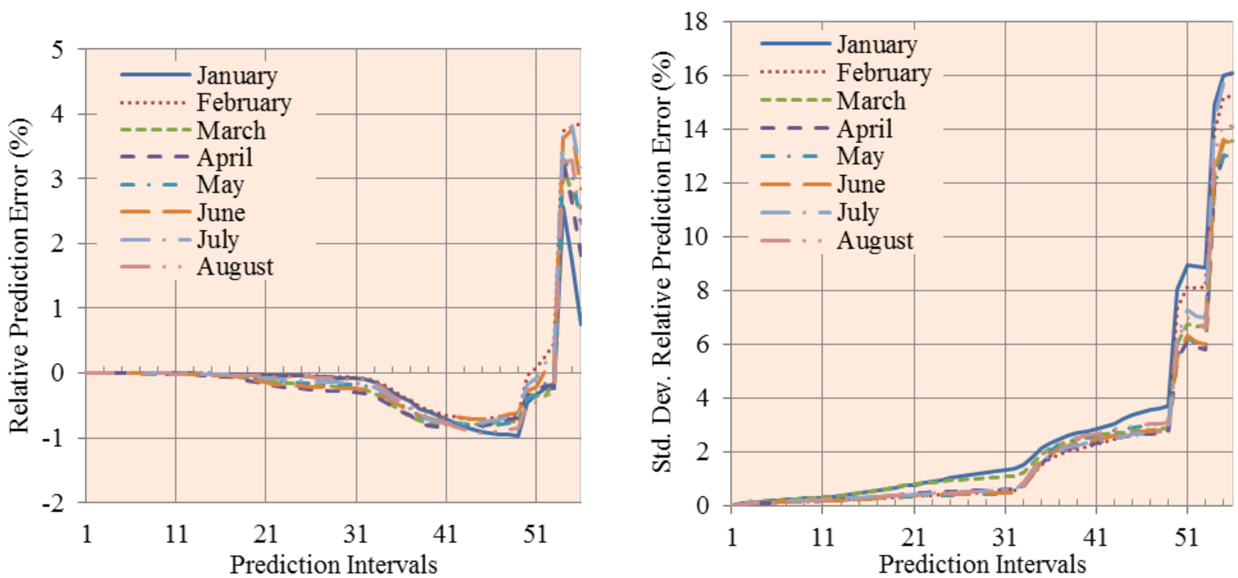

(j) Idaho Falls, Idaho Site (ST14)

Figure 2.44. Paired Average Relative Load Prediction Errors (Left) and Standard Deviations of those Errors at PNWSGD Utility Sites January–August 2014

2.6.2 Asset Models' Modeled Load

Presuming that each function that represents a transactive asset system at a utility site chose time periods when responses from the assets would be useful, each transactive site must then model the asset system to predict the change in load that would accompany the response during those periods. The resulting modeled change in load modifies the TFS, representing the elasticity of the site in light of the TIS to which the site functions are responding. If the response periods were selected during times that the TIS was maximal, the asset system should automatically act to either curtail load or generate more power. Either of these responses reduces the net load that must be supplied by relatively costly energy at this location. The asset model strives to accurately represent the magnitude by which net system load will be reduced by its response.

The functional asset models predicted these impacts during the project for both the impending interval (i.e., at IST0, the next 5 minutes) and for future intervals. The ranges of modeled magnitudes of these power changes are listed in Table 2.1. These are the magnitudes that were automatically calculated at utility sites and reported to project data collection. The table bins the responses rounded to the nearest 10, 100, or 1,000 kW.

The far left column states the site's owner. In a transactive system, the site owner would normally take responsibility for the responsive asset system, would represent the asset in the transactive system, and would benefit from the responses made by the asset. In the transactive system, the asset is operated from the perspective of the site owner. The responsibility might be contractually delegated, as happens today when an aggregator controls assets on behalf of a distribution utility. The degree to which the project's site owners owned this responsibility varied. Utilities today lean heavily on aggregators and other vendors to provide demand-response services for them.

The site, column 2 in Table 2.1, is the defined part of the transactive system topology (Appendix B) that would benefit from the response. A curtailment by a responsive asset system reduces the net load that must be supplied to the site.

Several types of systems have been listed in the column "Asset Description." Similar systems would normally be modeled similarly, but different system capabilities and site owner preferences may be accommodated by functions and their configurations. The functional asset models' details would normally not be shared or revealed between different site owners. The project has considered, however, libraries of functions from which site owners might select. Vendors could compete in this market to offer the most accurate, interoperable, and easily configured functions and asset models. No such libraries of examples existed prior to the project, so the project unilaterally developed examples for the utilities. Of the 11 participating utilities, only one—Portland General Electric—wrote its own functions and asset models from scratch. The others accepted the ones that the project offered them.

Table 2.1. Range of the Modeled Changes in Load by the Various Elastic Transactive Assets at the PNWSGD Project Sites

Site Owner	Site	Asset Description	Asset ID[a]	Load Function[a]	Modeled Change in Power (kW)[a,b]
Peninsula Light Company	Fox Island, WA	Water Heater Control	302	2.4	{−80, …, 0}
			382	3.4	{−350, −310, …, −40, 0}
		Dynamic Voltage Management	383	3.5	{0}
University of Washington	UW Campus, Seattle, WA	Building HVAC Management	303	2.4	{−10, 0, 10}
		Two Diesel Generators	304	2.5	{0, 1500}
		Steam Turbine	305	3.7	{−1500, 0, 10, 1500}
Portland General Electric	Oxford Rural Feeder, Salem, OR	Residential DR	NA	NA	NA
		Commercial DR	NA	NA	NA
		Distributed Generators	NA	NA	NA
City of Ellensburg	Renewable Energy Park, Ellensburg, WA	None	NA	NA	NA
Benton PUD	Reata Feeder, Kennewick, WA	Energy Storage Modules	316	4.1	{−40, …, 10}
Avista Utilities	Pullman, WA	Residential DR	321	2.4	{−10, 0, 10}
		Dynamic Voltage Control	322	3.5	{−430, …, −230, 0}
		WSU Tier 1 HVAC Control	320	2.4	{−7800, −1500, …, −300, 0, …, 800, …, 1200}[c]
			381	2.4	{0}
		WSU Tier 2 Chiller Control	378	2.4	{0}
		WSU Tier 3 Gas Generator Control	379	3.7	{−1000, 0, 1000}
		WSU Tier 4 Gas Generator control	377	3.7	{−1000, 0, 1000}
		WSU Tier 5 Diesel Generator Control	380	3.7	{−1400, 0, 1400}

Table 2.1. (contd)

Site Owner	Site	Asset Description	Asset ID[a]	Load Function[a]	Modeled Change in Power (kW)[a, b]
Flathead Electric Coop.	Libby, MT	Water Heater Control	326	2.4	{−80, …, 0}
		Smart Appliances	327	2.4	{−80, …, 0}
		In-Home Displays	328	2.4	{−190, …, 80, 100, 110, 120, 150, 160}
	Marion / Kila, MT	Water Heater Control	336	2.4	{−20, −10, 0}
		Smart Appliances	337	2.4	{−10, 0}
		In-Home Displays	338	2.4	{−40, …, 20}
City of Milton-Freewater	Milton-Freewater, OR	Water Heater (DRU) Control	344	2.4	{−700, −600, …, −100, 0}[c]
			375	2.4	{−14,000; −13,000; …; 9,000; 10,000}[c]
		Dynamic Voltage Control	345	3.5	{−30, …, 0}
			376	3.5	{−650, …, −240, 0}
			401	2.2	{−440, …, −220, 0}
Northwestern Energy	Helena, MT	Water Heater Control	NA	NA	NA
		Dynamic Voltage Control	NA	NA	NA
	Philipsburg, MT	Water Heater Control	NA	NA	NA
		Dynamic Voltage Control	NA	NA	NA
Lower Valley Energy	Teton-Palisades Interconnect, WY	Water Heater Control	349	3.4	{−410, −390, …, −40, 0}
		Battery Energy Storage	350	4.1	{−130, …, 130}
			402	4.1.1	{−130, …, 130}
Idaho Falls Power	Idaho Falls, ID	Building DR Management	358	2.4	{−5,800, −5,100, −1,100, …, −800, −300, −0, …, 300, …, 700, 3300, 3,700}[c]

Table 2.1. (contd)

Site Owner	Site	Asset Description	Asset ID[a]	Load Function[a]	Modeled Change in Power (kW)[a, b]
		Water Heater Control	359	3.4	{−170, 0}
		Thermostat Control	360	3.4	{−6,000, −5,000, …, 2,000, 4,000}[d]

(a) "NA" in this column means that the asset system was never fully connected to the transactive system or data was never provided for the asset system from the site's transactive system implementation.

(b) The ellipses in this column mean that the series continues incrementally by bins of 10 kW unless otherwise stated. Unless otherwise stated, the bin size is 10 kW.

(c) These modeled changes in load have been rounded to the nearest 100 kW. Ellipses mean that the series continues by bins of size 100 kW.

(d) These modeled changes in load have been rounded to the nearest 1,000 kW. Ellipses mean that the series continues by bins of size 1,000 kW.

DR = demand response
DRU = demand-response unit
PUD = Public Utility District

The column "Load Function" refers to an organization of the project's example functions. The digits shown are the last two digits of the classification of load functions (e.g., "TKLD_2.4"). The first digit indicates the type of response event that would be selected by the function. For example:

- "2.x": Event-driven responses. The asset responds infrequently and for relatively short durations.
- "3.x": Daily responses. Responses may occur each day. Weekends were often excluded.
- "4.x": Continuous. A continuum of responses is possible.

The second digit (e.g., "TKLD_2.4") refers to the asset model. The project was slow to recognize the opportunity to mix and match event determination (i.e., the first digit) and the functional model. That is, nearly any combination of event function and asset model is feasible. The functional responsibilities of the event function and asset model might be separable, resulting in more efficient coding and a more flexible "code library." Because the project was slow to recognize the power of asset modeling, some inconsistencies emerged in the numbering system. However, "x.1" models were for energy storage, "x.4" models were for water heater and thermostatic loads, "x.5" models were for voltage management, and "x.7" ones were for distributed generators.

Battery system model. An optimization was developed to manage a battery system. The battery model strives to optimize its net cost given the predicted unit cost of energy. The charging and discharging may be constrained within allowed states of charge and conversion power capabilities. A parameter was found to moderate the aggressiveness of the charge and discharge, effectively reducing the numbers of charge and discharge cycles. Surprisingly, the project's battery energy storage systems were found to have quite limited allowed numbers of charging cycles, effectively preventing the batteries from being cycled more than once per day.

Water heater model. Based on the principal investigator's prior work with water heaters, the diurnal pattern of water heater average power consumption was scaled to represent the number of controlled water heaters. Therefore, the number of controlled water heaters and an interpolated time of day could be used to estimate the amount of power that was likely being curtailed by the set of water heaters. Future improvements could model the impacts of event duration more accurately. A water heater dynamic model might be used.

Thermostat model. A first-order building model was created for the project. The model could be configured to calibrate it with the aggregate behaviors of a group of buildings. Once calibrated, an event represents a perturbation of the model's operation. Depending how the asset system is controlled, the thermostat set point might be modified, the heating and cooling might be fully curtailed, or the heating and cooling power might be cycled, giving the buildings a fraction of the heating and cooling energy they would otherwise need. Consequently, the modeled interior temperature falls or rises, and the buildings' thermal mass cools or heats. Snapback was modeled at the end of an event as the modeled building worked to recover from the perturbation. This is a rich research area where more accurate, higher order models of thermostatically controlled buildings might be adopted.

Voltage management model. A simple voltage management model was developed. The model was based on conservation voltage reduction (CVR) factor. A CVR factor states the relative reduction in load that should accompany a reduction in the feeder's voltage. The CVR factor is unique to the circuit and may be affected by season, time or day, and other variables. Regardless, a static CVR factor allowed for a relatively simple prediction of the change in feeder power.

Distributed generator model. The distributed generator model was perhaps the simplest of the project's asset models. The model was simply configured to output the full or a fractional nameplate capacity of the distributed generator while events were active. The output was presumed to be constant during the event.

2.6.3 Step 5 Analysis Conclusions

The future predictions of load by the transactive system sites were used as a metric for how well the assets were able to predict their electric load. These predictions were supplied by the project's inelastic load prediction functions. Analysis reviewed the relative prediction errors and the standard deviations of those errors. These errors were always referenced to the system's final prediction.

The relative prediction error analysis revealed multiple prediction biases, where the transactive system was found to have under or over predicted the final load prediction for the given data interval. Most of the sites predicted their loads well up to a day, or so, into the future, but some of the bias errors were significant even for near-term predictions. Had the region used these biased predictions to schedule generation resources, resources too might have been under or over scheduled.

The elastic, responsive assets also predicted how they would change the load when advised to do so by the transactive system. The ranges of these power differences were listed for each responsive transactive asset system. The predictions are affected by the quality of the asset models that represent them. The asset models are also configurable to scale and otherwise tailor the prediction to the unique assets. Some assets were found to have not been configured properly or to have accepted the default configurations without further modification, which would misrepresent the impacts of the asset systems in the transactive system.

2.7 Step 6: The Plans to Exchange Power with the System Must Be Calculated and Communicated throughout the System

The TFS was intended to predict and state the electrical power to be exchanged between nodes in the transactive system. The calculation of the TFS is, in principle, simpler than the blending of unit costs described in Section 2.4 for the other transactive signal—the TIS. In a branched power distribution system, the TFS is calculable from the balance of generated, consumed, and exported powers. The challenge is much greater in a network of transactive nodes, where one node might import power at times from more than one other transactive node.[1] This significant challenge was deferred by the project after it determined that the transmission region was to be represented by a centrally calculated, informed simulation that was run by Alstom Grid for the project.

Two separate methods emerged to calculate the TFS in the transactive system. The utility sites used transactive toolkit functions and asset models to emulate their electric loads. The relative prediction errors of those predictions were discussed in Section 2.6, and the absolute accuracy of the TFS calculations at the utility sites are addressed in Section 2.8. The biggest influence on these predictions was the prediction of inelastic load, which constitute the vast majority of the utility loads. The elastic assets' behaviors modified the total load according to the event-driven, daily, and continuous toolkit functions that determined the events and the asset models that predicted the impact the asset would have on net power.

One issue that emerged was that the inelastic load predictions by the transactive system became inaccurate where the functions had been inadequately trained and where the systems operated in open-loop mode, unaware of the actual power metering. Accurate modeling and prediction of distribution loads is an important, ongoing research area.

Another lesser issue emerged from the modeling of responsive loads. First, more work is needed to make models correspond to actual asset system behaviors. Even then, the functions that model the assets' effect on system load must be carefully configured. The asset owners must assume the responsibility for ensuring that accurate impacts are being predicted. And finally, the connections between the transactive system and the utilities' asset systems were tenuous. The fact that the transactive system had advised that an event should take place did not mean that the assets, in fact, responded to the event.[2]

The transmission-zone nodes within which the utility sites resided in the PNWSGD model possessed no independent means of correcting or negotiating the power needed by the utility sites, as was revealed by their TFSs. A simple fix was made to accommodate this limitation without breaking the system and its expectations that the transactive signal be exchanged. The transmission zones simply parroted back the TFS values stated by utility sites. There was no negotiation.

[1] The point is that while transmission power flows have been centrally calculated for many years, the methods for doing so with distributed calculations, where each node may observe only its own status, are still emerging.
[2] The issue is not so much whether the assets' owners heeded advice from the transactive system as it is that the status of the transactive system was allowed to diverge from reality.

The TFSs between the transmission zones were calculated centrally by the central informed simulation that was run by Alstom Grid. The responses of asset systems did not affect these calculations. This caused a feedback loop to be broken in the transactive system. The project determined that a simulation would be required to implement the feedback loop and test its performance. The resulting simulation activities are described in Section 2.10.

In Section 2.2, the connection between the power flowing between transmission zone nodes in the transactive system were said to be difficult to accurately pair with real-world power flows. Some similarities were observed in the modeled power-flow dynamics and those in the BPA data. Nothing prevents the flow from being more accurately measured as the granularity of the system model improves and as the transactive system becomes better informed about the status of the actual transmission system that it strives to emulate.

In summary, the system reliably exchanged its transactive signals, including the TFSs. The TFS values were calculated as planned at the utility nodes, although the accuracy of those TFS predictions may be further improved. The transmission-zone nodes relied on the Alstom Grid-informed simulation to calculate their TFS values for them, which broke a critical feedback loop in the transactive system. More research is needed to insert distributed power-flow calculations into transactive systems at this grand scale.

2.8 Step 7: The Modeled Exchange of Energy within the Transactive System Must Be Accurate

The TFSs were to have represented the near-term and predicted future power that was being exchanged between connected transactive nodes. This next step assesses whether the TFS at a node accurately represented the power being exchanged by the connected nodes. In the PNWSGD transactive system, the calculations of TFSs were accomplished differently between connected transmission-zone nodes and connections between transmission zones and the site nodes that they served. Therefore, the accuracies of the TFSs for these two connection types are addressed separately in this section.

2.8.1 Accuracy of the TFS between Transmission Zones

The power exchange between transmission zones of the PNWSGD transactive system was calculated within the *informed simulation* that Alstom Grid had designed to emulate the operations of bulk generation and transmission in the region. The region had been divided into 14 transmission zones. The boundaries between those transmission zones had been defined where the region’s transmission could be defined by one transmission line, or by no more than a few transmission lines. Alstom Grid represented this nodal system by allocating the region’s loads and resources among the transmission zones. A power flow was periodically performed to help ensure that the solutions were feasible, but the impacts of resource dispatch decisions were estimated between these calculations using influence factors.

Formulation of the simulation model for emulation of the regional grid behaviors proved very challenging. First, the reduced model was imperfect. Different resource names were used by different entities in the region. When they were available to the project, lists of resources did not always use the same or compatible and interoperable formats. And it was found to be surprisingly difficult to allocate resources to one side of a transmission-zone boundary or the other with the methods deployed by Alstom Grid. The allocation of generation and transmission resources might have proceeded more smoothly if the process were less time constrained and if Alstom Grid had created the reduced-order grid model from the start using rigorous model reduction methods.

Furthermore, access to real-time operational data was quite limited, so the emulated regional grid's behaviors diverged from the actions that were actually taken by grid operators. Today's regulatory environment dissuaded utilities from sharing much time-sensitive information that would have been useful for this exercise.

2.8.2 Accuracy of the Utility Sites' TFS

At the interface between the PNWSGD transactive system's utility site nodes and transmission-zone nodes, the TFS represents the power that is received by the utility site from the transmission zone in which the site resides. The site owners had worked with the project to define this interface, preferably at a well-metered location. At several of the utility sites, a direct comparison is therefore possible between the TFS and the total metered load that the signal was intended to emulate and predict.

At most sites, bulk load was modeled using algorithms that were developed and trained for the project by IBM. These were the main source of the TFS values that were generated at sites to emulate the sites' electric loads. The relative accuracies of the TFS predictions at site locations are addressed in Section 2.6.1. This section addresses only comparisons at the sites between the nearest-term predictions of the TFS[1] and the meter data.

Table 2.2 compares the TFS against the metered power at the Peninsula Light Fox Island site (ST01) during 2014. The monthly averages and standard deviations are shown for the eight project months of 2014. The last column shows the differences between the monthly averages stated as a fraction of the average metered power.

Stepping down the rows of this table by month, a clear trend emerges. Underestimation improved from January to March. From April through August, a trend toward increasing overestimation emerged. The greatest difference between the monthly averages, a relative error of 60.2%, occurred in June 2014. The model failed to track seasonal changes in load. Because March appears to be the month with the least error, and the errors increasingly diverge before and after March, the project suspects that the load predictor was trained using March data. This site did not make use of real-time feedback from the meters that might have improved prediction accuracy over time.

[1] The project refers to the near-term prediction interval as that corresponding to interval start time zero (IST0) that predicted behavior for the next 5 minutes.

Table 2.2. Comparison of Average Metered Power at the Fox Island Site (ST01) and Its Representation by the Transactive Feedback Signal for the Eight Project Months of 2014

	Average Metered Power[(a)] (MW)	Average TFS[(a,b)] (MW)	% Error[(c)]
January	22.0 ± 4.6	17.8 ± 4.4	−18.8
February	21.1 ± 4.7	17.9 ± 4.4	−15.2
March	18.2 ± 4.6	18.0 ± 4.4	−1.48
April	14.8 ± 3.7	17.9 ± 4.4	21.0
May	11.8 ± 2.6	17.9 ± 4.4	52.3
June	11.2 ± 2.1	17.9 ± 4.4	60.2
July	11.8 ± 2.6	17.9 ± 4.4	51.4
August	11.5 ± 2.4	17.9 ± 4.4	56.7

(a) Unless otherwise stated, both the TFS and metered power have been averaged over the period from January 1 to September 1, 2014. The variability is the standard deviation.

(b) For TFS power, only the “Operational” signals were included in the calculations. If relaxation occurred during a given 5-minute interval, the last result for that interval was the only one used in the calculation.

(c) This error is simply the difference between the average TFS and average metered power expressed as a percentage fraction of the metered power.

Table 2.3 shows a similar comparison, but this comparison was for the power received by the University of Washington campus site (ST02) from the West Washington TZ02. The transactive system consistently underestimated the amount of energy that would be, in fact, required by the utility. The site was offline during April 2014 and came back online the next month with some relatively inaccurate calculations in May. This outage of the site’s transactive node was attributed by the university to a server reconfiguration problem that coincided with a cyber security event at a vendor’s location. The University of Washington continued to have meter data-collection issues, which may be the result of the unusually high metered value in May 2014.

Table 2.3. Comparison of Average Metered Power at the University of Washington Site and Its Representation by the Transactive Feedback Signal for the Project Months of 2014

	Average Metered Power[(a)] (MW)	Average TFS[(a,b)] (MW)	% Error[(c)]
January	32.1 ± 4.8	27.9 ± 3.6	−13.1
February	32.8 ± 4.5	27.9 ± 3.6	−14.8
March	31.3 ± 4.0	27.0 ± 3.7	−13.6
April	28.2 ± 7.8	-	-
May[(d)]	61 ± 1070	28.8 ± 4.3	−52.8
June	33.6 ± 4.6	27.1 ± 4.7	−19.5
July	36.8 ± 5.5	30.4 ± 4.6	−17.4
August	37.4 ± 5.0	30.9 ± 4.2	−17.4

(a) Unless otherwise stated, both the TFS and metered power have been averaged over the period from January 1 to September 1, 2014. The variability is the standard deviation.
(b) For TFS power, only the “Operational” signals were included in the calculations. If relaxation occurred during a given 5-minute interval, the last result for that interval was the only one used in the calculation.
(c) This error is simply the difference between the average TFS and average metered power expressed as a percentage fraction of the metered power.
(d) The project believes the unusually high averaged meter power May 2014 was due to persistent data-collection challenges. The comparison is probably not valid this month.

The comparison was repeated for the Portland General Electric demonstration feeder site in Salem, Oregon (ST03). The results are shown in Table 2.4. The project was unable, working with Portland General Electric, to define a meaningful test region and site metering that might confirm the accuracy of the transactive system’s TFS calculations at the Salem, Oregon node. The comparison between these two quantities is not meaningful.

Table 2.4. Comparison of Average Metered Power at the Portland General Electric (Salem, Oregon) Site (ST03) and Its Representation by the Transactive Feedback Signal for the Project Months of 2014

	Average Metered Power (MW)	Average TFS[a,b] (MW)	% Error[c]
January	18,504	-	-
February	175,014	17.4	-
March	17,289	18.7	-
April	17,432	17.8	-
May	19,088	18.0	-
June	11,651	18.1	-
July	20,253	14.4	-
August	21,419	18.2	-

(a) Unless otherwise stated, both the TFS energy and metered energy have been averaged over the period from January 1, 2014 00:00:00 to September 1, 2014 00:00:00 local time.

(b) For TFS energy, only “Operational” signals were included in the calculations. In addition, if relaxation occurred during a given 5-minute interval, the “last” data point was the only one used in the calculation.

(c) The comparison is not valid at the Salem, Oregon site. The TFS clearly was not emulating this metered energy. The result of this calculation would be an extremely large negative percentage.

The comparison was repeated for the Avista Utilities Pullman, Washington site (ST06). The results are shown in Table 2.5. On average, the TFS overestimated the sites power by about 9%. The most inaccurate comparison occurred in June 2014 when the relative error was 62.4% overestimation of the metered value. However, this error appears to be attributable to a problem with the metered quantity, not the calculated TFS values.

Table 2.5. Comparison of Average Metered Power at the Pullman, Washington Site (ST06) and Its Representation by the Transactive Feedback Signal for the Eight Project Months of 2014

	Average Metered Power[a] (MW)	Average TFS[a,b] (MW)	% Error[c]
January	2,215 ± 529	NA	NA
February	2,170 ± 903	2,229 ± 256	2.71
March	2,009 ± 467	2,184 ± 285	8.68
April	2,018 ± 376	2,152 ± 276	6.67
May	2,138 ± 824	2,148 ± 303	0.45
June	1,348 ± 1180	2,190 ± 323	62.4
July	2,269 ± 899	2,124 ± 334	−6.36
August	2,399 ± 645	2,116 ± 328	−11.8

(a) Unless otherwise stated, both the TFS and metered power have been averaged over the period from January 1 to September 1, 2014. The variability is the standard deviation.

(b) For TFS power, only the “Operational” signals were included in the calculations. If relaxation occurred during a given 5-minute interval, the last result for that interval was the only one used in the calculation.

(c) This error is simply the difference between the average TFS and average metered power expressed as a percentage fraction of the metered power.

Table 2.6 compares the calculated TFS at the Philipsburg, Montana site (ST11) and the metered data that it was to represent. Both the average metered load and the averaged TFS representation were consistent from month to month. The differences between metered and TFS values are relatively small. The TFS was no longer being dynamically calculated during the last 3 months of the project. The same average value is reported with no standard deviation.

Table 2.6. Comparison of Average Metered Power at the Philipsburg, Montana Site (ST11) and Its Representation by the Transactive Feedback Signal for the Eight Project Months of 2014

	Metered Energy[a] (MW)	TFS[a,b] (MW)	% Error[c]
January	1,024.1 ± 11.0	1,023.4 ± 7.5	0.07
February	1,030.8 ± 13.5	10,37.4 ± 14.3	−0.64
March	1,022.2 ± 11.0	990.9 ± 14.8	3.15
April	1,018.0 ± 9.8	976.8 ± 9.9	4.22
May	1,011.9 ± 10.5	964.6 ± 11.6	4.90
June	1,008.5 ± 10.1	956.8 ± 0.0	5.40
July	1,004.8 ± 11.9	956.8 ± 0.0	5.01
August	1,005.1 ± 11.5	956.8 ± 0.0	5.04

(a) Unless otherwise stated, both the TFS energy and metered energy have been averaged over the period from January 1, 2014 00:00:00 to September 1, 2014 00:00:00 local time.

(b) For TFS energy, only “Operational” signals were included in the calculations. In addition, if relaxation occurred during a given 5-minute interval, the “last” data point was the only one used in the calculation.

(c) This error is simply the difference between the TFS and metered energy expressed as a percentage fraction of the metered energy.

The calculated TFS values at the Idaho Falls, Idaho site (ST14) were typically twice as great as the metered quantity the TFS was to emulate, or more, as shown in Table 2.7. Implementers must have misunderstood the connection between the TFS and the metered quantity that it was to predict.

Table 2.7. Comparison of Average Metered Power at the Idaho Falls, Idaho Site (ST14) and Its Representation by the Transactive Feedback Signal for the Eight Project Months of 2014

	Metered Energy[a] (MW)	TFS[a,b] (MW)	% Error[c]
January	38.4 ± 4.4	77.6 ± 11.3	102
February	36.3 ± 5.7	75.9 ± 11.2	109
March	31.4 ± 4.1	72.4 ± 9.3	131
April	28.2 ± 4.2	71.1 ± 9.0	152
May	26.2 ± 3.7	69.3 ± 8.6	165
June	25.9 ± 3.8	68.9 ± 8.8	166
July	29.0 ± 5.0	70.1 ± 9.9	142
August	25.2 ± 7.8	68.3 ± 9.0	171

(a) Unless otherwise stated, both the TFS energy and metered energy have been averaged over the period from January 1, 2014 00:00:00 to September 1, 2014 00:00:00 local time.

(b) For TFS energy, only "Operational" signals were included in the calculations. In addition, if relaxation occurred during a given 5-minute interval, the "last" data point was the only one used in the calculation. The stated variability in this case is the standard deviation of the interval values.

(c) This error is simply the difference between the TFS and metered energy expressed as a percentage fraction of the metered energy.

2.8.3 Step 7 Analysis Conclusions

The project's modeling of its electric load was probably not accurate enough for transactive systems of the design used by the PNWSGD. The relative errors between the TFS values at site nodes and the metered power that the TFS values should have modeled were found to be large. The accuracy varied wildly during the months of 2014. If a transactive system is to use feedback from its nodes to inform and plan the dispatch of its resources, the inaccuracy of such feedback (i.e., the TFS) must be small compared to magnitudes of resources being dispatched. Otherwise, the dispatch of resources will also be inaccurate, and the system will not properly plan the balance of resource to load.

Load forecasting is today done by balancing authorities. The project's utilities did not eagerly accept or own a new responsibility to predict their dynamic loads. If nodes are to accurately predict and report their loads, then automated systems must be developed to track and predict such loads. The load must be metered, and the metered data must be made available to the transactive prediction algorithm in real time.

The PNWSGD transactive system neither rewarded accurate predictions nor penalized inaccurate ones. In the future, incentives should be built into the system to reward accuracy and deter inaccuracy.

2.9 Step 8: Resources Must Respond to Dynamic System Load Predictions, Including the Plans from Flexible Loads

To conclude analysis of the complete control loop, this last analysis step should evaluate whether the predicted loads—both the predictions for inelastic and responsive elastic load in the transactive system—affected the actual dispatch of bulk load in the region. There is not much to discuss for this step. The PNWSGD transactive system was not permitted to directly influence bulk generation in the region. Its scale was considered too small to have a substantial influence, and as an experimental system it was not yet trusted to modify dispatch schedules.

It might be argued that the behaviors of responsive transactive loads did, in fact, change system balance and therefore affected the region's resources and resource mix. If so, this was a passive benefit. The dispatch and scheduling of the region's energy resources were accomplished entirely by existing mechanisms that the region's balancing authorities rely upon.

Having recognized early during the design of the PNWSGD that this step would not be successful in the field system, the project planned to simulate the system, including a more direct influence of the transactive system's actions on the dispatch and scheduling of the region's resources. This simulation is described in the next section.

2.10 Simulation Analysis of the Pacific Northwest Smart Grid Demonstration Transactive System

Additional section coauthors:
S Ghosh, M Yao, R Ambrosio, G Jensen, A Koc,
D Phan, H Wang, and K Warren −
IBM TJ Watson Research Center

IBM designed and built a simulation platform to closely mimic the operations of the PNW power grid. The PNWSGD project has designed transactive response assets to dynamically respond to extreme stresses in the system, and balance the cost of electricity over time. The aim of this simulation effort has been to study the behavior of the grid when the presence of distributed transactive response assets is high, consisting of up to 30% of all load being transactive in nature.

The simulation results show that the transactive control mechanisms designed by this PNWSGD project respond in the expected fashion to reduce the peak total system costs. In addition, certain types of responsive assets are effective in balancing the cost of electricity throughout a day, by consuming energy when system-wide costs are low and reducing load when the costs are high. The strength and the consistency of the response were estimated. The magnitude of the response depends on the number of transactive assets in the system, and can lead to up to about an 8% reduction in total peak costs in the PNWSGD region when the presence of load that is transactive in nature is high.

A second important goal of this simulation study is to analyze the interactions of the high transactive penetration system with the presence of renewable generation as a large part (up to 30%) of the total generation portfolio. Renewable generation has a complex interrelation with the transactive system.

Renewables are considered the cheapest per-unit source of power in the simulation. However, output from renewables is subject to various weather factors and is hard to schedule or predict. So the effect of renewables depends on the periods when high output is realized. If high renewables output coincides with the low-cost periods in the day, the cost-balancing assets that take advantage of lowest-cost electricity increase their interaction. On the other hand, high renewables in otherwise peak-cost periods have the effect of tamping the peak by themselves, thus not requiring further action from the peak-shaving parts of the transactive system. Overall, the summer data set shows wind to have a weakening effect on the transactive response, while in the winter and shoulder data sets, transactive response is strengthened by the presence of renewables.

2.10.1 Introduction to the Simulation and its Objectives

This simulation study seeks to understand the effect of the PNWSGD transactive system in its capability of withstanding extreme events and ensuring grid reliability for now and the future, as the transactive solutions are scaled up within the PNW power grid. The premise of introducing transactive assets into the grid is that they will help mitigate the effect of extreme events, as manifested in peak systemic costs, and the effect of uncertainty in predicting and scheduling renewable generation. Our study aims to find out if its promise is borne out by the design of the transactive control mechanisms introduced and studied in the PNWSGD. More importantly, the dynamic, interactive nature of transactive systems must be fully understood, evaluated and tested before the technology can be deployed at a large region-wide scale. Since the production environment cannot be risked for such a study, simulations are the only method that can be used to fully study and understand these systemic behaviors. In addition, simulations allow the controlled study of the effects of unpredictable, sudden, and fleeting stresses on the system.

Another important motivation for creating a simulation environment is to predict the effects of scaling the level of renewable power generation far higher than the current level of penetration. We study the complex interactions between being highly responsive and having a higher penetration of renewables in the electricity grid.

This IBM transactive system simulation is based on a combination of simplified grid network topology and a simplified model of some existing functionalities (like the ones implemented by Alstom Grid for the PNWSGD project) at regional transmission and bulk generation levels, while using real control toolkit functions (like the ones deployed by the project's participating utilities) at the distribution asset level. A scaled-down model that would allow for fast execution was envisioned. While being fast enough to allow for a large number of simulations to be carried out, this model must be detailed enough to allow for the thorough exploration of various inputs and alternative grid conditions. The inputs to the model should be controllable in order to simulate the effects of interesting scenarios (e.g., simulating extreme weather conditions) and the grid model should be configurable in order to study different network topologies and differing numbers of modeled resources.

2.10.2 The PNWSGD Transactive System

To facilitate the timely intervention of the transactive assets in ameliorating highly constrained situations that might arise in an electricity grid, this demonstration project suggests that certain pieces of information be exchanged in 5-minute intervals between all interconnected electricity assets. The information sent by any asset to its neighbors consists of two values:

- a value representing the average cost of the power required by each node to meet its local demand and export targets, measured in dollars per kilowatt-hour. The cost consists of components that measure the per-unit cost of generating power locally (cost of fuel, etc.), infrastructure costs that capture the amortized cost of installing any infrastructure in the node's local area of control, and the cost of importing power from its neighbors.
- a neighbor-specific value that represents the expected interchange of power (in kilowatts) between itself and the neighbor.

Predictions for these two values are published by each electrical asset in the transactive system over a forecast horizon of up to 5 days, with the information being broken down into fine intervals for the first day and coarser intervals for the rest. The published data allow each node to understand the impact of its local decisions on its own average cost as well as that of its neighbors, and is expected to help make decisions on transactive-load management that are to the benefit of the overall system.[1]

The transactive-load systems in the simulation respond to the forecast average cost of power in the node where they reside. Three types of transactive asset loads are defined by the demonstration project and modeled in the simulated system, distinguished by the nature of the control logic:

- daily-event – these loads typically activate up to once per day, trying to match their load reduction to the period with the highest predicted average cost of electricity within the day at its connecting node. The asset classes behind this control type can be residential appliances such as water-heaters, air-conditioning units, washer-dryers, etc.
- event-driven – these loads activate up to a set number of times within a given rolling or fixed time horizon, again trying to match their activation to the period with the highest predicted average cost of electricity within the time horizon at its connecting node. Unlike the daily-event assets, these could allow a time horizon of any length and could be activated multiple times in the time horizon. This simulation study models event-driven assets that act three or four times within a rolling period of a week. The asset classes that provide this response type are similar to the daily-event type.
- continuous-response – these loads continually try to identify an opportunity to use both low-cost and high-cost periods to strike a beneficial tradeoff between electricity usage and load reduction. A

[1] This design of transactive information has a key limitation that will affect the simulation results when high renewable penetration is being studied. Note that the infrastructure cost component of the average cost is applied to *all* power generated in or imported into the node. In particular, this makes renewable generators an unusually low-cost method of power generation under this scheme, in that its unit cost of production (fuel costs, etc.) is zero, and the cost of installing the infrastructure is applied uniformly to *all* power generation, i.e., not just the renewables but also thermal and hydro plants, and to imported energy. This limitation in the design, namely not being able to attribute structural costs to each source of power separately, will affect some of the effectiveness observed in integrating renewables in this simulation, but we fully expect the broad trends observed here to remain true even when the cost accounting for renewables is changed.

typical example is an electricity storage device such as a large battery installation. Unlike the first two control types that typically model assets that can only drop load when activated (e.g., usage of a residential appliance is postponed or rescheduled), the continuous-response assets can either charge from the grid (increase net load) or discharge to the grid (decrease net load).

2.10.3 Advantages of the Simulation Platform

This work leverages IBM Research's expertise and experience on platform integration, simulation and optimization to construct such a simulation system. The simulation addresses several challenges that are very difficult to resolve in the actual demonstration project, in order to study interesting PNWSGD behavior:

Feedback to bulk generation dispatch. In the field demonstration, Alstom Grid provided static, immutable inputs to the transactive nodes that then were used to calculate the node's average cost of electricity. The input information is calculated by Alstom from grid-level information it obtains from BPA (network load, conventional generation etc.) and 3TIER (renewables output). However, the demonstration project did not leave a pathway for the feedback obtained from the transactive system to be provided to Alstom Grid to modify its calculation. In other words, the predicted changes in system-wide loads due to the presence of transactive assets did not inform the calculations. In essence, this breaks a key feedback mechanism of the transactive system, and results in an open-loop system. This was a sensible choice that limited the real-world impacts that an experimental system such as the transactive system could have, but limited the effectiveness of the demonstrated transactive system. The simulation model closes this loop in order to be able to model various scenarios, such as modeling a system with a higher participation rate of renewable generation resources or to model increased loads. In order for fully closed loop feedback to be enabled, Alstom's proprietary management platform was replaced with an IBM-built unit-commitment and economic dispatch module.

Simulation speed (real-time simulation vs. speed-up simulation). Rapid simulation times are needed for tens of thousands of scenarios to be evaluated quickly to allow for a thorough evaluation of the different possible grid configurations with different generation and consumption patterns. This can only be accomplished by speeding up the simulations—by allowing simulated time to be accelerated more rapidly than real (wall clock) time.

Multiple scenario simulation for transactive system optimization. To design a good transactive system, a number of design parameters need to be optimized. A thorough evaluation of the solution space requires easy scenario specification. Such a flexible, parameterized configuration mechanism that could be used to specify the scenarios under consideration does not exist today. Subsequently, optimization techniques could possibly be used to select the best solution that would meet the objectives of the system under consideration.

2.10.4 Core Design Components of the Simulation Platform

The logic functional block diagram for the IBM simulation system is shown in Figure 2.45. It captures the key system components (light blue blocks), data flows (solid black arrow lines), and important configuration inputs required to run the simulation. In summary, the regional transactive system simulation is achieved through interaction between a collection of simulated transactive nodes (modified and based on field model of PNWSGD) and simulated regional balancing components (unit commitment [UC] and economic dispatch [ED]) that represent bulk generations and transmission. The setup enables simulation of distributed, end-to-end transactive feedback control loops within each modeled network zone and over whole simulated regions. The construct provides a simple mechanism to simulate and evaluate how regional generation resources are dispatched and used under the influence of the transactive system. The simulation itself was carried out in a distributed fashion over multiple compute nodes; the use of a distributed architecture for the original PNWSGD project allowed for an easy porting of the distributed nature of computation to the simulation. This greatly aided in being able to pursue a large, complex set of simulation scenarios for the analysis.

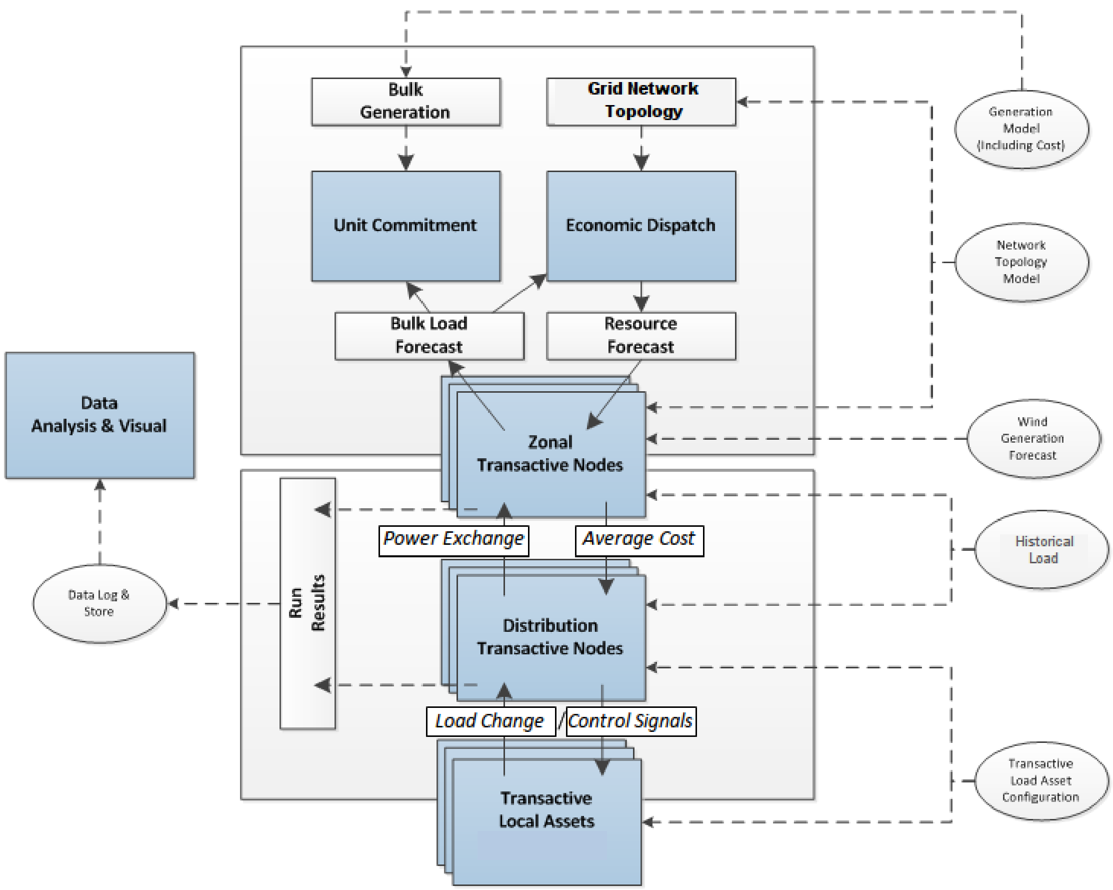

Figure 2.45. Schematic of the IBM Simulation Platform Built to Simulate the PNW Electricity Grid

The simulator is configured and controlled by seven key inputs described below.

Transmission network model. To reduce complexity and computation resource, Alstom Grid developed a reduced model for the bulk generation and transmission network. The reduced regional network model includes 14 transactive bus nodes, representing the transactive system of the region and 4 non-transactive bus nodes representing boundary energy exchange between the simulated transactive system and outside. The 18 node buses and links between them represent the simplified, reduced network model of the simulated regional grid. Figure 2.46 provides a diagram of the simplified network used in this simulation study.

Bulk generation. At each bus node, three types of bulk generation are modeled and provided by Alstom Grid: thermal, hydropower, and wind. The generation model includes generation characteristics, such as minimum and maximum capacities, minimum and maximum on and off times, and cost functions. Both UC and ED simulation components share the same configurational inputs from the bulk generation and transmission network model.

Unit commitment. A day-ahead hourly generation schedule is created by the UC block. This schedule (how much power should each generator generate each hour) is computed using information about load forecasts (how much is needed) and conventional generation characteristics (how much can the generator provide, how long does it take to start up, etc.), how much renewable energy is available, network topology, and cost models corresponding to the cost of generation.

Economic dispatch. The simulation uses the ED module to determine all three cost components of the average cost of electricity at each node. The hourly schedule produced by UC is fed into ED block, which will be run every 5 minutes (in simulated time) to generate actual generation schedules that contain dispatch values (how much power to generate) for the generators so that the overall system load will be met at the lowest cost. For each ED execution, power-flow simulations are also triggered to determine the power flowing through the grid, ensuring that the power flow is within the operating limits. The power-flow simulations are also used to determine power exchanges between different transactive nodes. Each ED execution event produces the output of multiple computations, each run based on bulk-load forecast periods in the power-interchange signal. Correspondingly, the output ED is a forecast, with same time series as the input, of all bus generation capacity (kW), cost ($/h), and power flows of the regional network. These outputs are passed to a transactive control node system for average-cost computation to influence the behavior of responsive local assets across the simulated region.

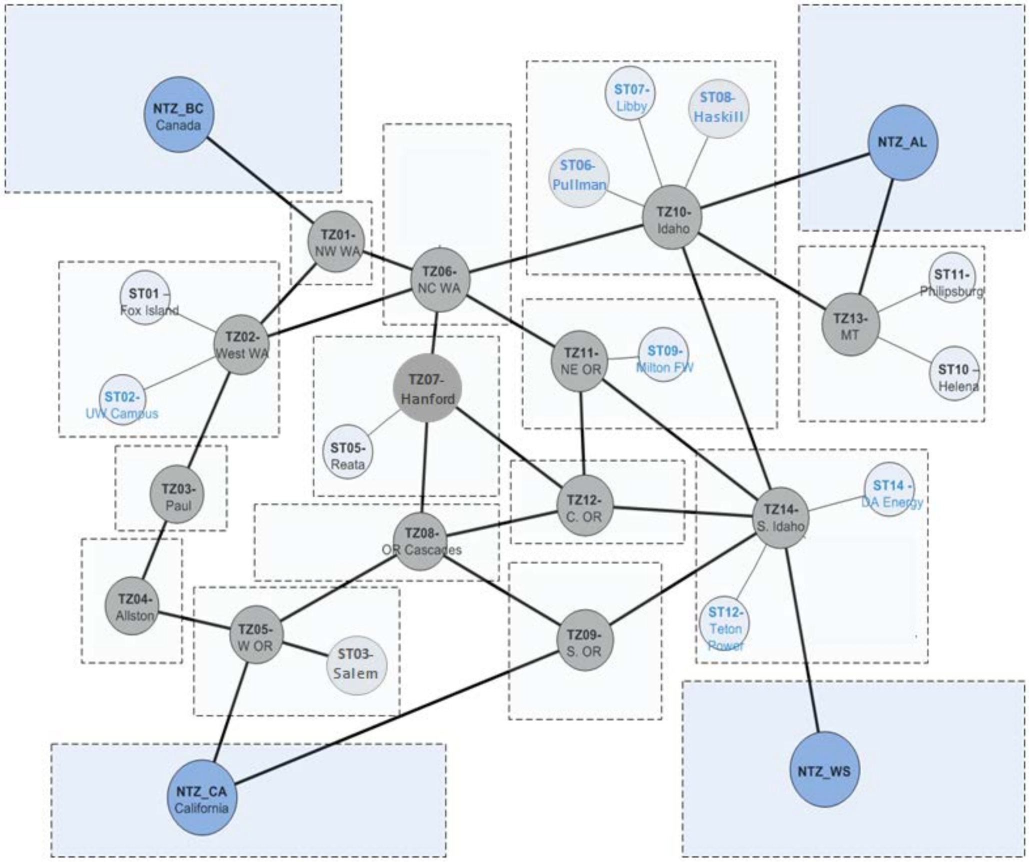

Figure 2.46. Simplified Model of the Pacific Northwest Electricity Grid

Bulk inelastic load. Based on BPA data, bulk inelastic load files for each transactive zone bus were provided for the three simulation seasons. These historical load files provide baseline inelastic load for each transactive node bus and are used to compute and calibrate transactive and renewable wind-penetration levels for the simulation. These are described further in the next section.

Renewable wind generation. For each simulated season, the historical wind power is provided as a renewable wind forecast. To simplify the simulation, the wind power resource was taken into account as negative load on the demand side instead of as dispatchable resource on the supply side. Combined with bulk inelastic load, wind power contributes as base-load for overall net-load forecast computation by a transactive node and is submitted to ED as load forecast. A wind power multiplier is implemented as a configuration parameter to scale the simulated wind-penetration level.

Transactive local assets configuration. All local responsive assets are created and configured by transactive local asset configuration files. These configuration files specify number, type, and characteristics of transactive responsive assets created for each simulation run. These configuration parameters in combination with other controllable input parameters determine the transactive penetration level of the simulation.

The results of simulators are collected as output files for data analysis and visualization. The key output files include the following:

- average-cost predictions – JavaScript Object Notation (Json)-based data collection, each published forecast recorded containing data for each forecast interval, and a breakdown of the cost factors for each type of generation resources and additional infrastructure costs.
- interchange-of-power predictions – Json-based data collection, each forecast published by each node for each of its neighbors containing data for each forecast interval. Also recorded were the local inputs of inelastic load forecast, elastic load change forecast due to transactive assets and control signals for each local responsive asset.
- inputs provided to the ED and UC module, including the net-load predictions for each node as submitted to ED/UC modules.

2.10.5 Simulation Scenarios and Experiment Run Setup

Simulation scenarios were defined and controlled by various configuration inputs and parameters:

Distinct seasons of the year (from 2013). This is configured and controlled by feeding the simulator with different base-load and wind power data corresponding to different seasons of the year. Three season periods, each lasting 1 week and ending in that season's observed peak load for 2013, are defined and targeted, as described in Table 2.8. The shoulder period was selected as a fall season week approximately halfway between the summer and winter peaks.

Table 2.8. Seasonal Data Sets Used in Simulation

Season Data Set	Start Time	End Time
Summer	2013-07-30 08:00:00 UTC	2013-08-06 08:00:00 UTC
Winter	2013-11-29 08:00:00 UTC	2013-12-06 08:00:00 UTC
Shoulder	2013-09-28 08:00:00 UTC	2013-10-05 08:00:00 UTC

UTC = Coordinated Universal Time.

Figure 2.47, Figure 2.50, and Figure 2.53 plot the total system-wide load under the three season data sets. The days in the summer data set have a single flat peak through the 09:00−17:00 (local) period, indicating a likely correlation with cooling load incurred because of the day time temperatures, while the other two data sets have pronounced morning and evening peaks, with the morning peak usually being higher than that in the evening.

Penetration Level of Wind Generation. Wind-penetration level is defined by wind peak power generation capacity divided by total peak base-load power in the region. A calibrated wind power multiplier is applied to the wind power generation forecast input data, consisting of recordings of forecasts for the present from each node in the network, to control wind-penetration level. Three different levels of wind penetration were planned and simulated, as listed in Table 2.9. Figure 2.48, Figure 2.51, and Figure 2.54 plot the total system-wide generation under the medium wind case for each data set. No clear pattern is discernible in the wind output, which serves to underlie its variability and dependence on local weather phenomena.

Table 2.9. Wind-Generation Cases

Wind Generation Penetration	Wind as Percentage of Peak Total Bulk Load
No Wind	0%
Medium Wind	10%
High Wind	30%

Penetration Level of Transactive Control. Three different levels of transactive penetration are planned and simulated, as described in Table 2.10 below.

Table 2.10. Transactive-Load Penetration Cases

Transactive Penetration	Peak Transactive Response as Percentage of Peak Bulk Load
No Transactive Load	0%
Medium Transactive Load	10%
High Transactive Load	30%

To achieve these levels of transactive load, the following steps were applied:

- The total number of transactive-load assets of the three types (event-driven, daily-event, and continuous-response capabilities) was scaled such that the sum of their total peak load reduction equaled the chosen percent of the peak total bulk load observed in the system, for each seasonal data set.
- The relative proportions of the three transactive asset types were always maintained at 20% continuous-response loads, 40% daily-event loads, and 40% event-driven loads.
- The control logic for the event-driven and daily-event type of assets respond to predicted peak average cost of electricity. Having a high number of assets of the same type throughout the network may lead to synchronous large changes in load due to transactive response. In order to simulate a scenario that more closely represents the likely future of uncoordinated asynchronous responsive assets, and also to prevent adverse effects from the simplification of the PNW transmission grid into a 14-node network, we apply a randomization factor to the number of assets that may respond at any given time. In the simulated system, the randomization factor is sampled afresh every 5 minutes uniformly from the range 50−100%, and represents the number of the event-driven and daily-event assets that may be active in that 5-minute period. So, an average of 75% assets of these types is expected to be reacting to incentives at any period.

In addition to the above control parameters, the following configuration parameters were controllable inputs:

- time acceleration factor (defaulted to 50)
- simulation start and end times (to match the start and end times of simulated scenario).

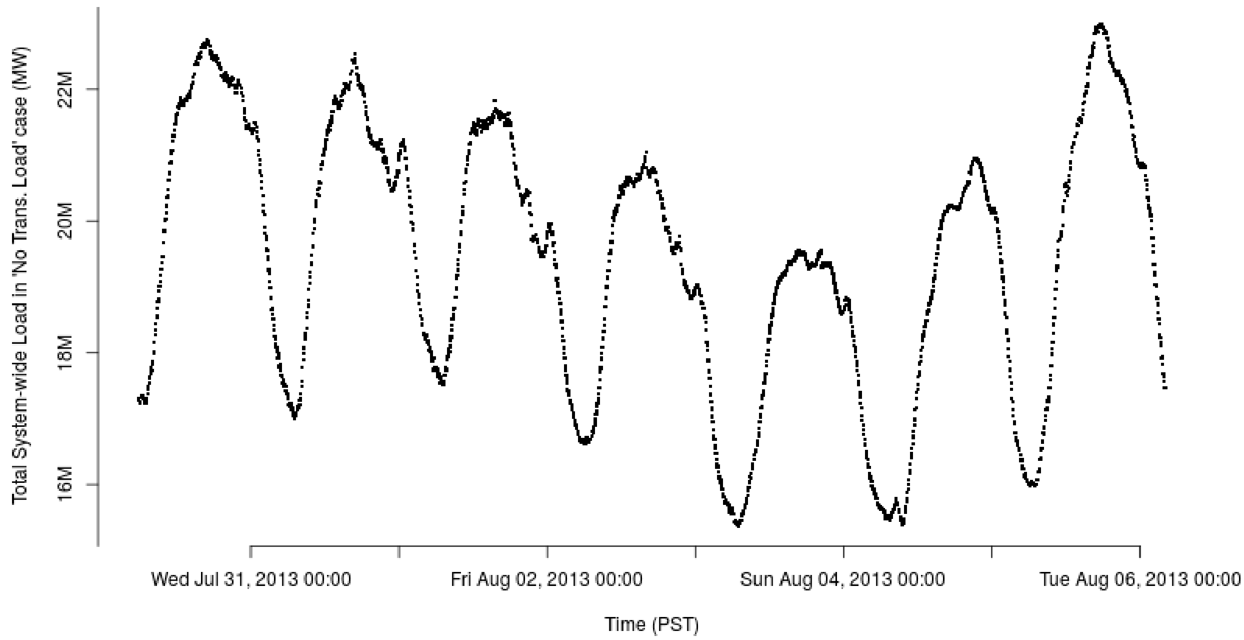

Figure 2.47. Total System-Wide Load in the Summer Data Set in the No-Transactive-Load Case

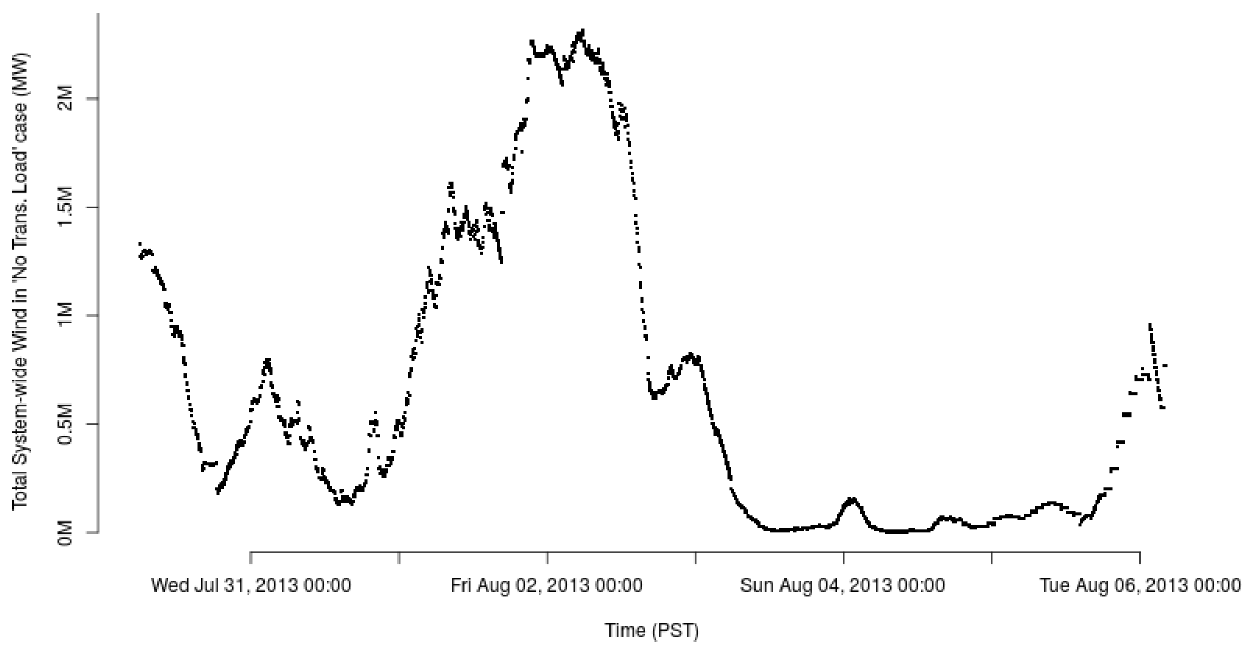

Figure 2.48. Total System-Wide Wind Generation in the Summer Data Set for the Medium Wind-Penetration Case

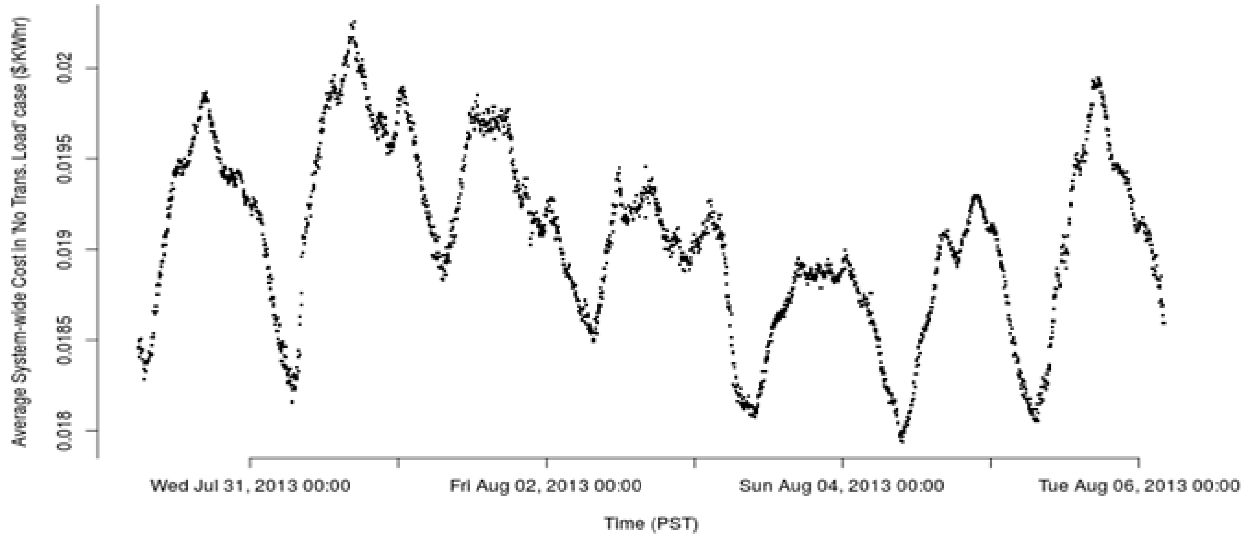

Figure 2.49. Average System-Wide Energy Cost of Electricity in the Summer Data Set under the No-Transactive-Load and No-Wind Cases

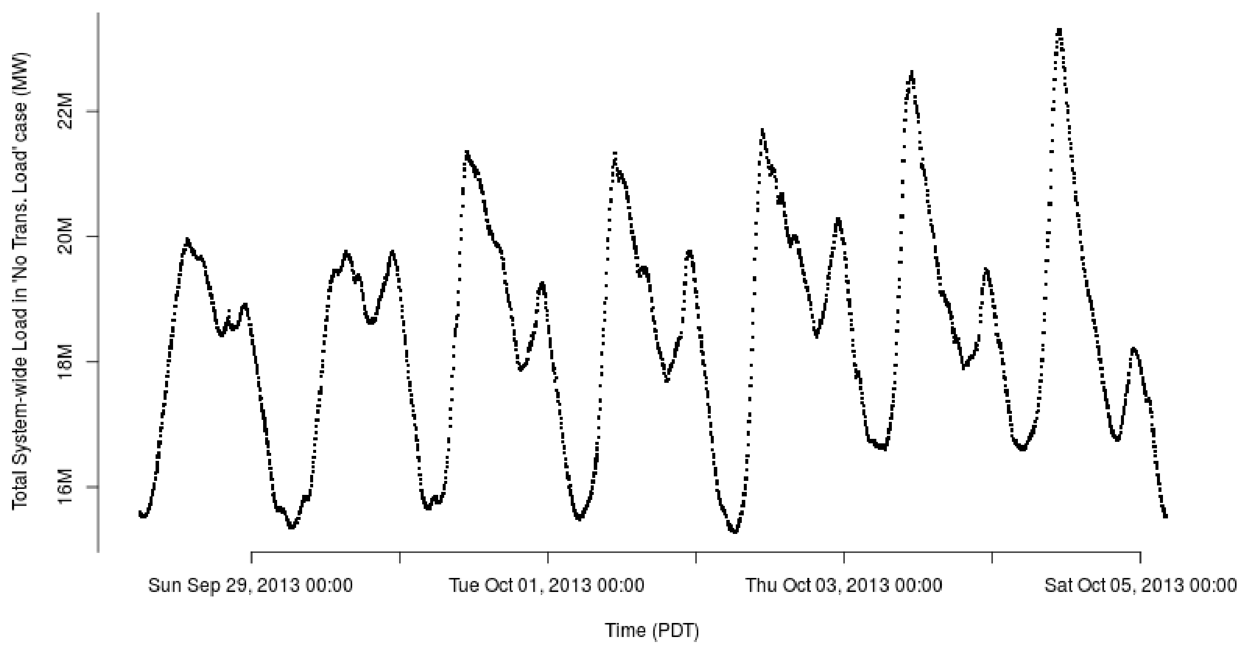

Figure 2.50. Total System-Wide Load in the Shoulder Data Set in the No-Transactive-Load Case

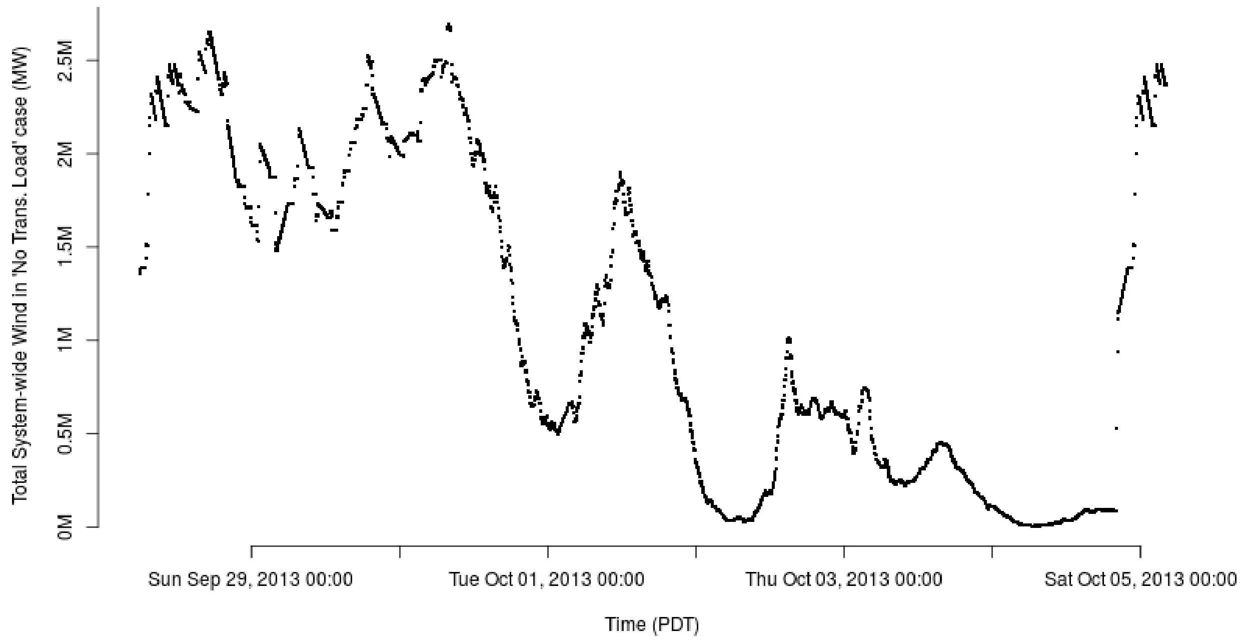

Figure 2.51. Total System-Wide Wind Generation in the Shoulder Data Set for the Medium Wind-Penetration Case

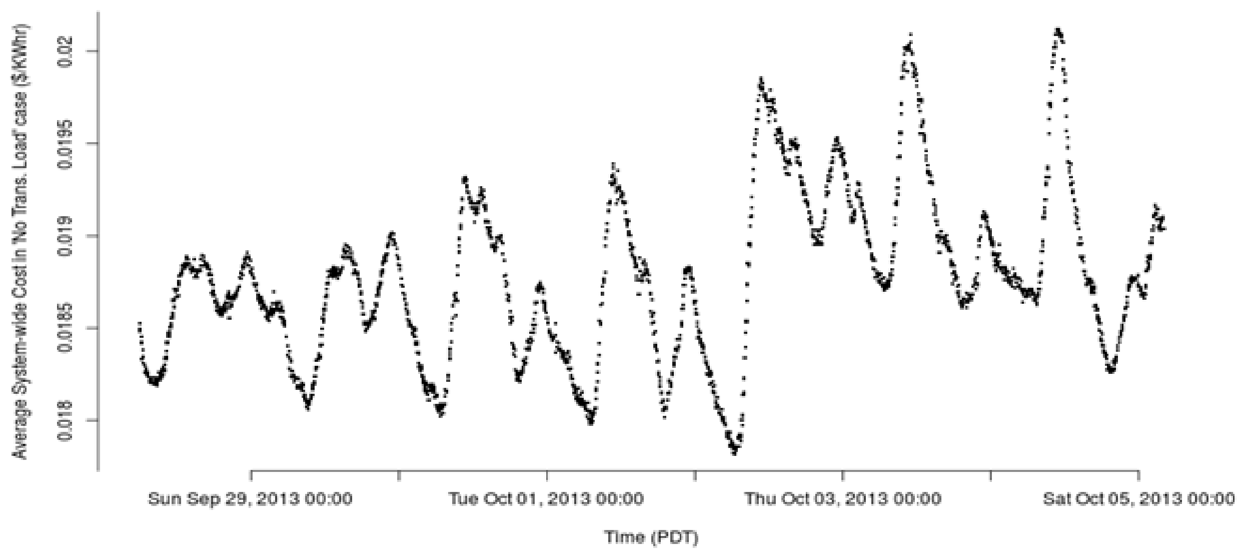

Figure 2.52. Average System-Wide Cost of Electric Energy in the Shoulder Data Set under the No-Transactive-Load and No-Wind Cases

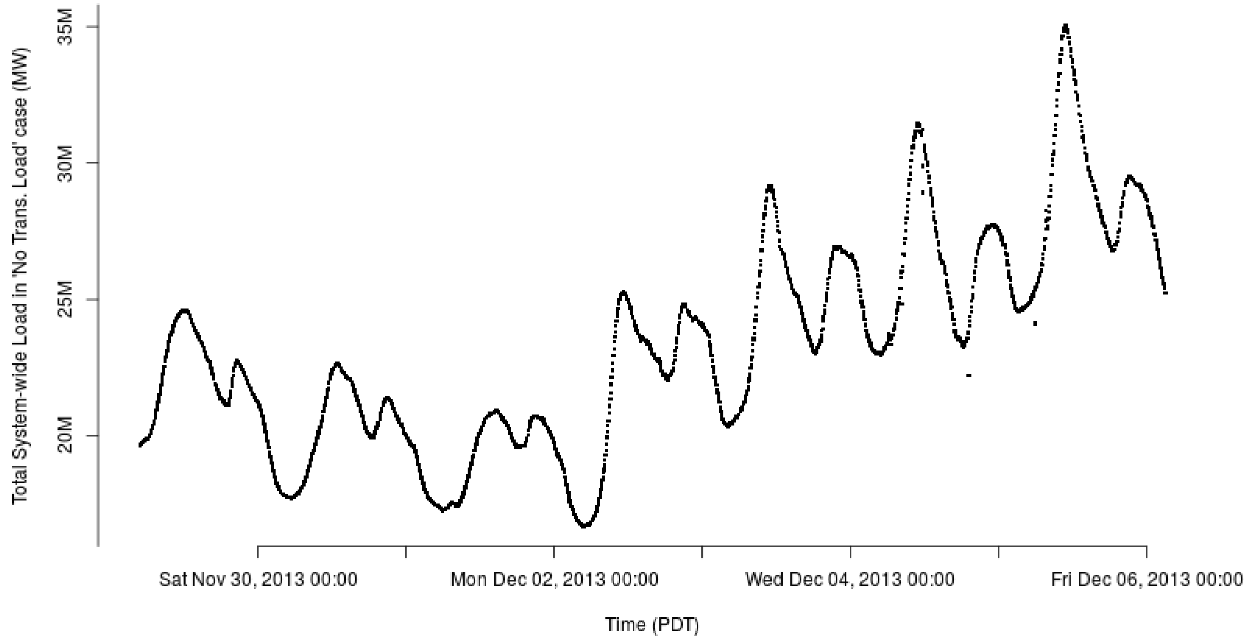

Figure 2.53. Total System-Wide Load in the Winter Data Set under the Case Having No Transactive Assets

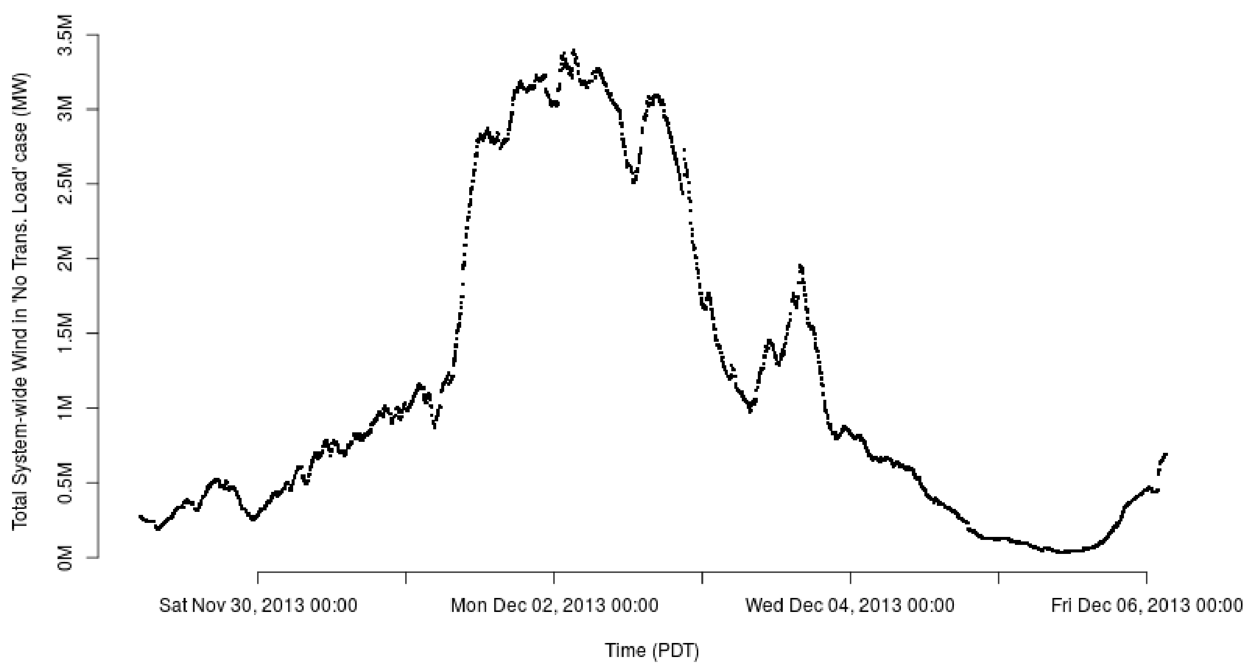

Figure 2.54. Total System-Wide Wind Generation in the Winter Data Set for the Medium Wind-Penetration Case

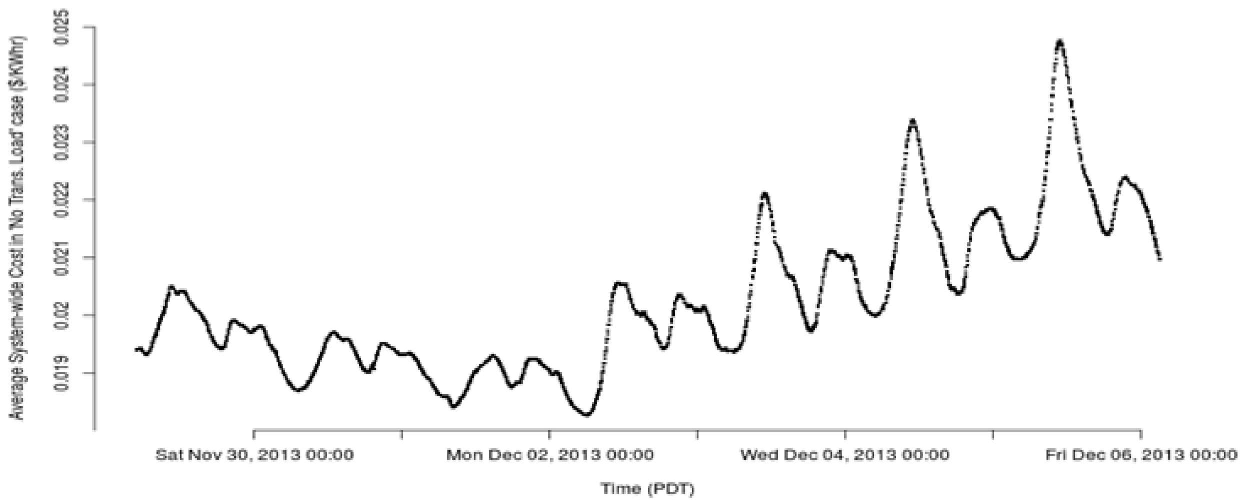

Figure 2.55. Average System-Wide Cost of Electric Energy in the Winter Data Set under the No-Transactive and No-Wind Cases

2.10.6 Output Analysis

This section will analyze the network-wide effects of the transactive system. Recall that all transactive asset control mechanisms are designed to respond to the average cost of electricity at their connecting nodes. Figure 2.49, Figure 2.52, and Figure 2.55 provided the system-wide average cost of electricity for the three data sets when no transactive load or wind is allowed to affect in the system. Overall, the average cost seems to follow the patterns observed in the total system-wide load. Both the winter and shoulder data sets show marked peaks in the day, which will be the times chosen by the event-driven and daily-event responsive assets. The summer data set exhibits flatter high system-wide average costs through the middle of the day, and in this instance the time chosen by the assets to respond will depend more on the average costs in each node.

Key metrics used to elucidate the performance of the transactive system are the total load (W) measured throughout the system at any time period, and the corresponding total hourly cost ($/h) borne by the system to meet this demand.

The first few analysis steps tease out the characteristics of the transactive system independent of the presence of wind in the system, and so use only the no-wind scenarios. The effect of wind over the transactive system is then analyzed.

The simulation scenario that models no transactive load and no wind will often serve as a benchmark for comparison between the other combinations of cases, and so will be referred to as the “base-case” scenario.

A note on the plots displayed in this section: The axes display three quantities with units, the time of day, total system-wide costs in units of dollars per hour ($/h), and the total load incurred in the system in watts (W). On occasion, the displayed units may scale up by a multiple of 10^3, 10^6 or 10^9 to k, M, or G units (e.g., k$/h or MW), as will be indicated. Three dimensionless quantities, scaled total cost and percent relative change in load or cost are also used, and will be introduced prior to first use.

2.10.7 Understanding Transactive Systems

It is instructive to start by taking a deeper look at how the transactive system responds to high-stress situations for the electricity grid in an attempt to alleviate the situation. Toward this goal, we take a closer look at two days from the data sets and the response under all three transactive scenarios, but with the modeled wind-penetration modeling no wind penetration.

The total system-wide costs for the fourth day of the summer data set are plotted in Figure 2.56. The no-transactive-load case has a relatively small, flat peak, in that the peak hours extend from 08:00 to 14:00. The control logic for all of the transactive assets are designed to respond to peak costs, and hence the high transactive load case is seen to significantly respond by reducing total cost. There is no sharp reduction in a single period, but response seems to be spread out in the 10:00 to 14:00 period. This is seen more clearly in Figure 2.57, which plots the change in total system-wide load through the day for the two transactive cases compared to the “No Trans. Load” case. Note the load-reduction response at peak periods being spread over the morning peak. The transactive loads respond to peaks in average costs at their connecting nodes, and these are given in Figure 2.58. The average costs at the nodes are seen to generally peak in the morning but each node’s peak occurs during different periods spread over the 08:00 to 14:00 range. This leads to the spread-out response in the total-load views in Figure 2.56 and Figure 2.57.

Also of note in Figure 2.57 are the slight increases in total load during periods of low average or total costs. This is due to the continuous-response units, which constitute 20% of the total responsive load, consuming extra energy while costs are advantageous.

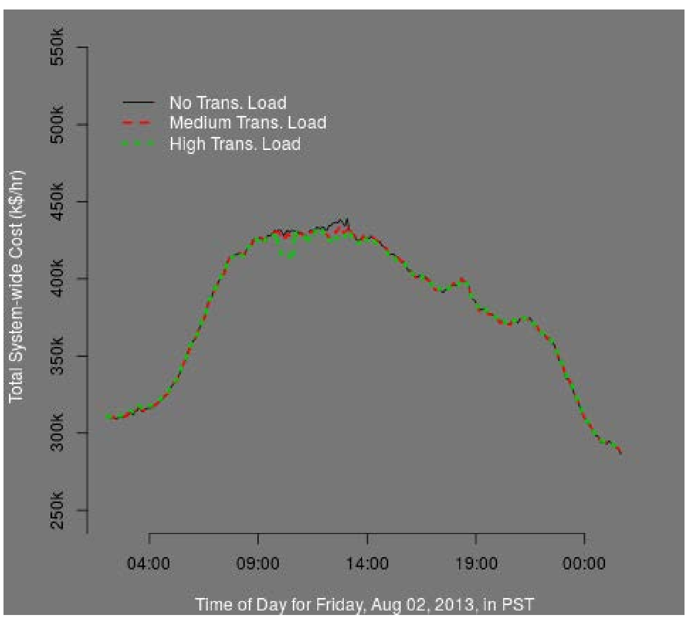

Figure 2.56. Total System-Wide Cost (day 4, summer)

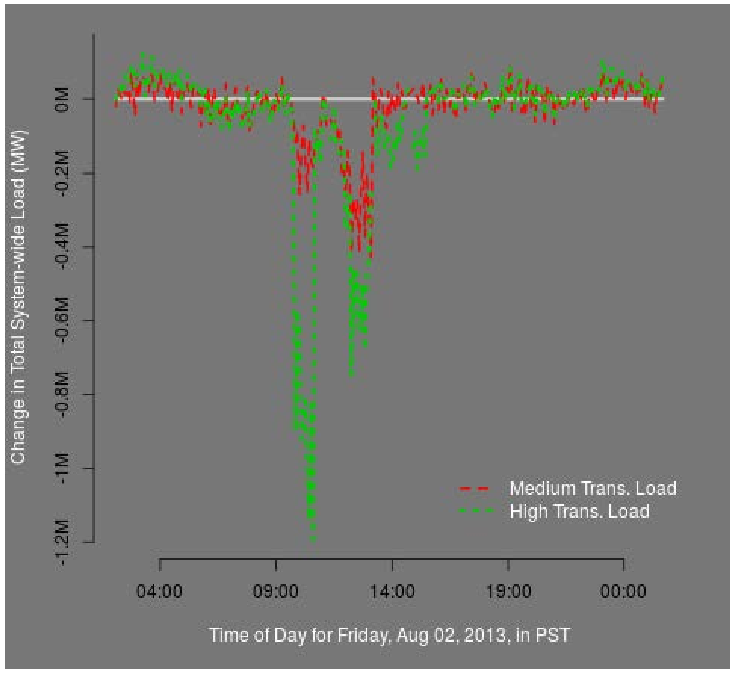

Figure 2.57. Difference in Total System-Wide Load (day 4, summer)

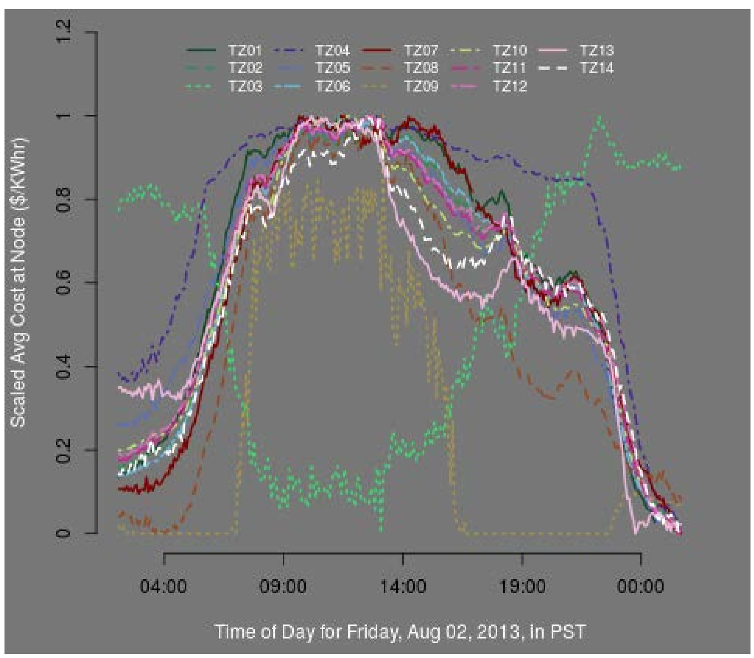

Figure 2.58. Average Cost of Electricity at Nodes (day 4, summer)

The average system-wide cost of electricity in winter (Figure 2.55) and shoulder (Figure 2.52) data sets display two peaks per day as opposed to the single flat peak in summer (Figure 2.49). Figure 2.59 shows the fourth day in the winter data set, which is the most interesting of the twin-peak days because the peak total system-wide cost values are similar for both intra-day peaks. The corresponding drop in total system-wide load in Figure 2.60 displays significant load reductions during both the two short, sharp morning and evening peaks. The individual responses of each asset depends on the average costs at each node (Figure 2.61), which show that each node experiences a slightly different individual morning or evening peak, leading to the responsive loads choosing the corresponding peak for load-reduction activation.

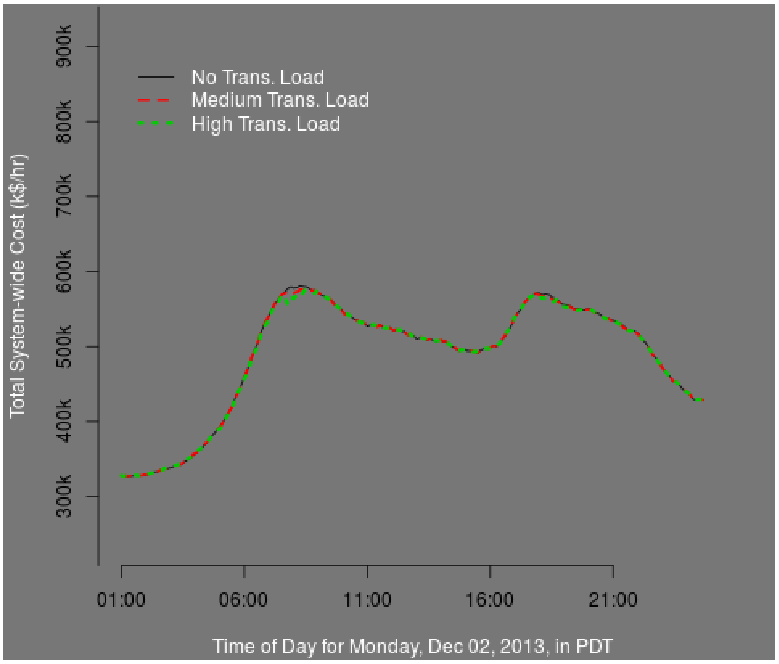

Figure 2.59. Total System-Wide Cost (day 4, winter)

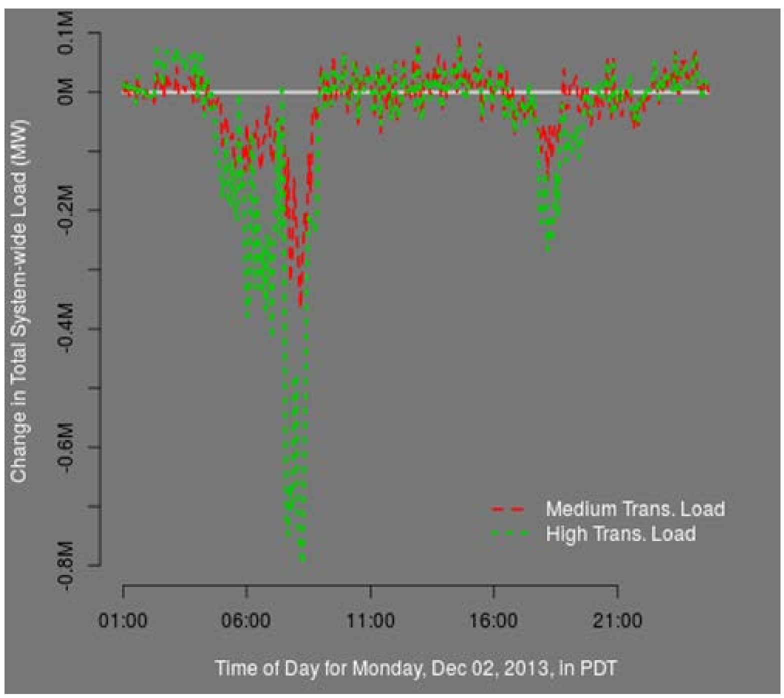

Figure 2.60. Difference in Total System-Wide Load (day 4, winter)

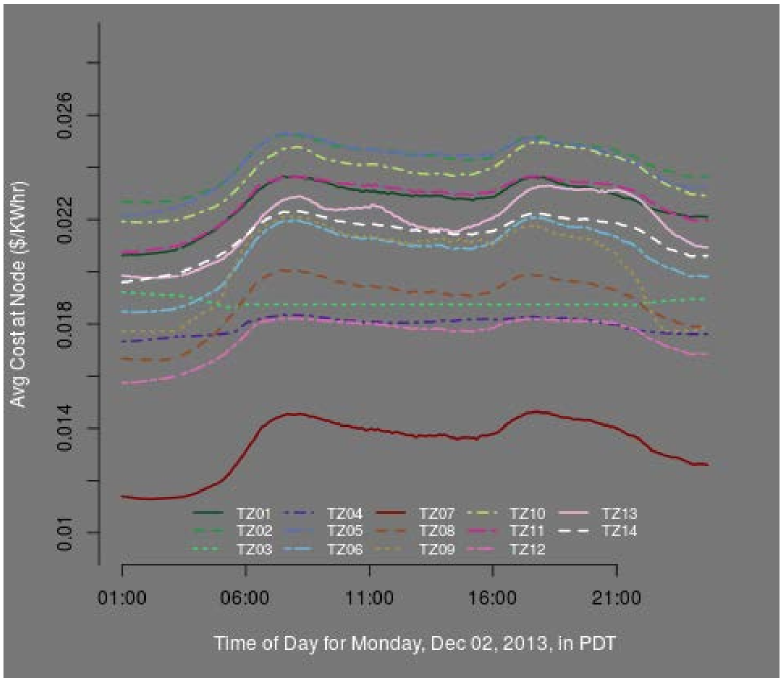

Figure 2.61. Average Cost of Electric Energy at Nodes (day 4, winter)

2.10.8 System-Wide Effects of Transactive Assets

Figure 2.62 graphs the total system-wide cost and total load in system for the base-case scenario. Each point represents a 5-minute interval in the 7-day simulation over the summer data set. Further, each day's data is represented by a different color and point-type. From Figure 2.47, the peak load is observed on the day of Monday August 5, 2013, and the corresponding total system-wide costs (in yellow squares) reach the top-rightmost part of the graph. A key observation is the almost linear relation between the total cost and total load, especially on a per-day basis. This is a result of the ED module, which is able to maintain an almost constant average cost of generation over the entire range of total load in system given the input models of the costs and ranges of thermal and hydro generation in the PNW power grid. The average unit costs vary slightly between days as shown in Figure 2.49. The cost variation for the same total load seems to be due to the slightly different breakdown of the same total load over the nodes of the network in different periods.

Figure 2.63 plots the same two quantities for the high transactive-load case. A general pattern emerges that matches the observations made earlier from the individual day plots. While most of each day's series remains the same as that under the no-transactive-load case, the top-right-most parts of each series are affected by the presence of transactive load. The points on those corners can be visually identified to have moved either to the left or the bottom (or both) of their original location on the left figure, indicating that the total cost and/or load have been reduced as desired.

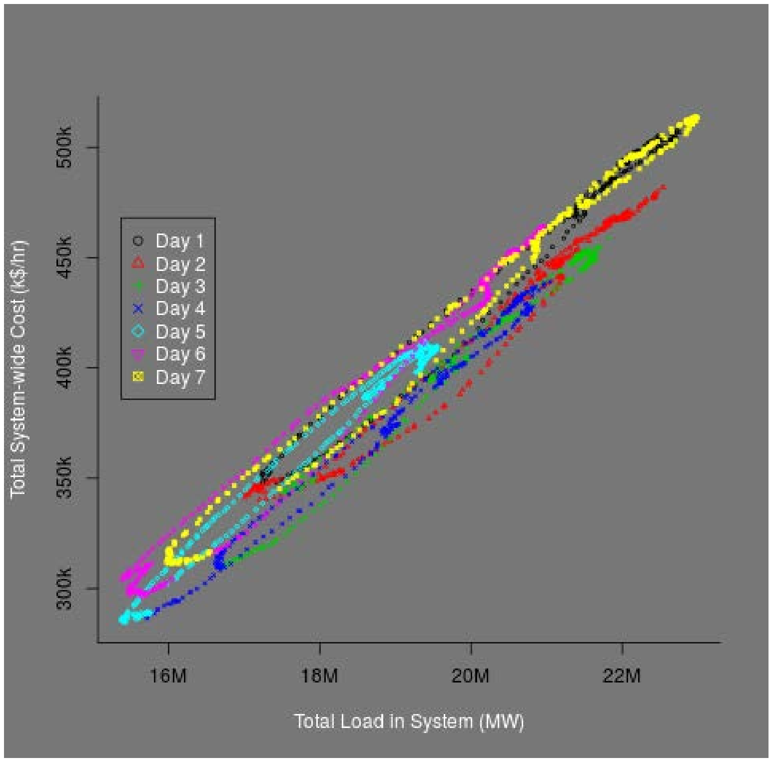

Figure 2.62. Total System-Wide Cost vs. Total Load for the Base-Case

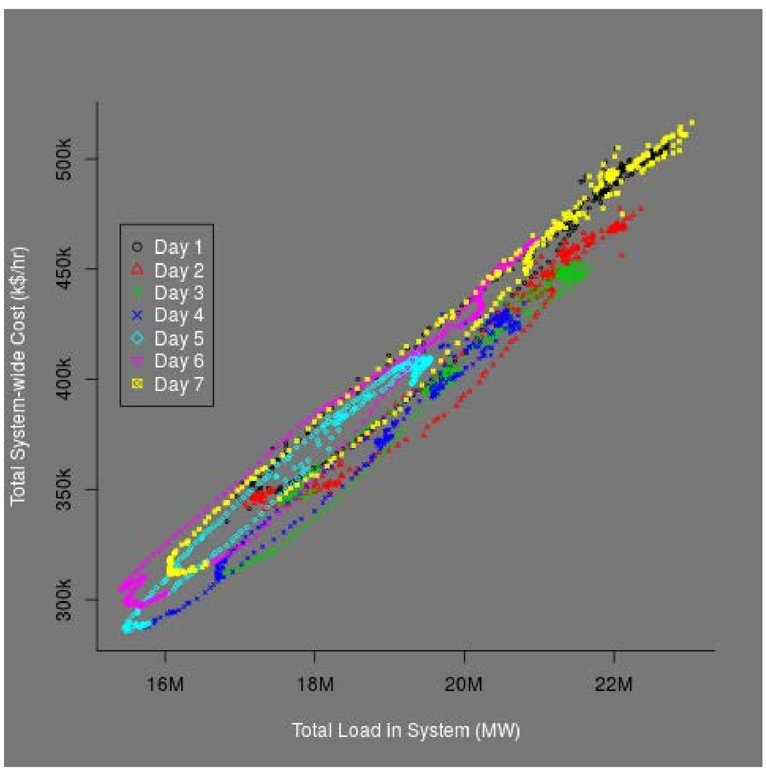

Figure 2.63. Total System-Wide Cost vs. Total Load for High Transactive-Load Case

The response is seen within each day's peak period, and in particular, a transactive response may be observed even in days with peaks that are lower than (to the lower left of) days with higher peaks. This is because of the presence of a significant percentage (40%) of daily-event responsive loads. Further, the event-driven responsive loads may choose to respond in the lower peaks if sufficient event activation budget is available. Moreover, continuous responsive loads may be able to realize a successful tradeoff between relative costs within a day. Indeed, these assets may respond to tradeoff costs in smaller intervals, which we will observe when the effect of wind is taken into account.

Table 2.11 provides the maximum observed changes in total system-wide load between the medium and high transactive-load cases and the base-case. Note that the drops in load generally happen during peak average-cost periods in the day when all responsive assets act by dropping some load, while increases in load occur during low average-cost periods when the continuous-response type assets charge up on cheaper energy. The maximum response in dropping load depends on the season, and can be as high as 7.8% during peak periods.

The next subsection studies the nature of the responsiveness of the transactive system in more detail.

Table 2.11. Maximum Observed Changes in Total System-Wide Load with Respect to the Base-Case Scenario

	Medium Transactive Load		High Transactive Load	
Season	Maximum Drop in Total Load (%)	Maximum Increase in Total Load	Maximum Drop in Total Load (%)	Maximum Increase in Total Load (%)
Summer	−4.34	0.61	−5.78%	1.51%
Shoulder	−7.79	0.65	−7.26%	8.94%
Winter	−4.26	7.86	−4.99%	7.77%

2.10.9 Responsiveness of the Transactive System

An analysis of the transactive responses first needs to adjust for local peaks and troughs in the time series of total system-wide costs. This is achieved in a straightforward manner by defining a new dimensionless quantity that is calculated by scaling the total system-wide cost or load by the corresponding peak and trough within the day. Figure 2.64 plots the result of this scaling, which produces values within the interval [0, 1] for the base-case scenario. While the inter-day variation in the peaks and troughs observed in Figure 2.49 are eliminated here, the intra-day variation in total cost is retained by this scaling.

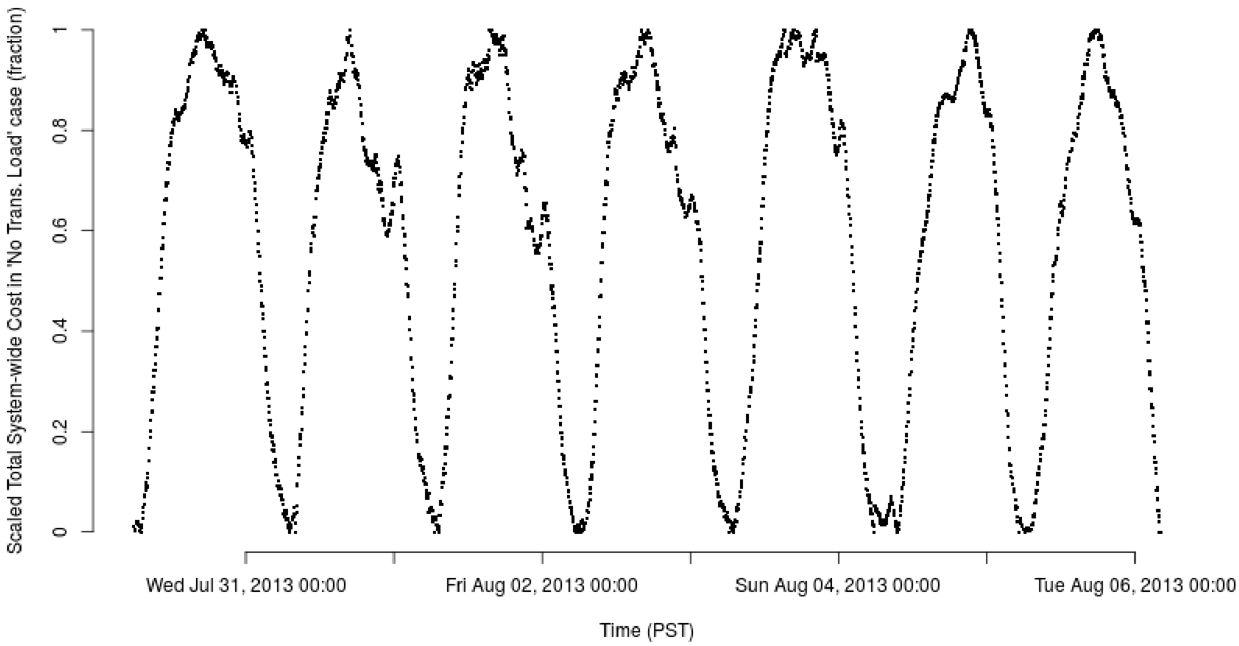

Figure 2.64. Total System-Wide Cost Expressed as a Scaled Dimensionless Quantity, for the Case that Had No Transactive Load in the Summer Dataset

Figure 2.65 provides scatter-plots of the percentage change in the total system-wide cost for each 5-minute interval in the medium and high transactive penetration cases, respectively, as a function of the scaled load value of the total system cost in the base-case. A striking observation is immediately evident: when the system-wide cost is at its highest within a day, the transactive assets generally have the effect of reducing the total cost by reducing their load. The more responsive high transactive penetration case is able to achieve larger relative drops. On the other hand, in the lowest cost periods within a day, a reverse behavior is observed, where the transactive assets might *increase* the total system cost by imposing additional load in the system. This is due to the continuous-response assets using the lower costs to charge up for a successful arbitrage during higher cost periods. The magnitude of the increase of relative total cost near zero-scaled-cost is more modest than the reduction near the high-cost end, which is a reflection of the smaller relative proportion of continuous-response assets (20% of total) against event-based assets (80% of total).

A vertical line on each plot in Figure 2.65 indicates the switch-over from dropping load to increasing load. In both cases, this seems to happen around when the scaled load represents the 35% percentile. The maximum responses observed in these plots are tabulated in Table 2.12.

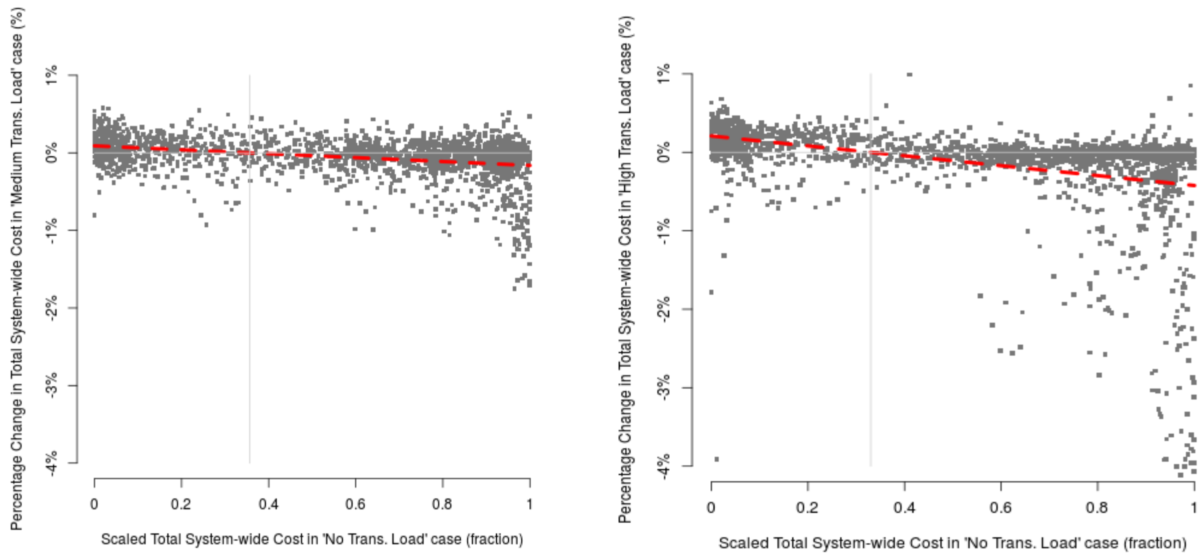

(a) Medium Transactive-Load Penetration (b) High Transactive-Load Penetration

Figure 2.65. Scaled Total System-Wide Load vs. Percentage Change in Total Load under the (a) Medium and (b) High Transactive-Load Penetration Cases

Table 2.12. Maximum Observed Changes in Total System-Wide Load with Respect to a Base-Case Scenario that Had No Transactive Load

	Medium Transactive Load		High Transactive Load	
Season	Maximum Drop in Total Load (%)	Maximum Increase in Total Load (%)	Maximum Drop in Total Load (%)	Maximum Increase in Total Load (%)
Summer	−4.34	0.61	−5.78	1.51
Shoulder	−7.79	0.65	−7.26	8.94
Winter	−4.26	7.86	−4.99	7.77

Table 2.13 provides the results of a linear-regression model fit to the scatter-plots in Figure 2.65. The results for all seasonal data sets are provided. The slope of the modeled linear response is stronger under high transactive load for all data sets. Moreover, the intercept value at zero is always positive and at one is always negative, indicating an appropriate response from the continuous-response assets. The intercepts for the high transactive-load cases are about double those under the medium transactive-load case.

The linear-regression model indicates a direction of response due to the transactive load in the system. A good measure of the strength or determinacy of the response in this single-factor regression model is the correlation value between the percentage change in total load and the scaled load value in the no-transactive-load case. In the summer and winter data sets, we see a fairly high negative correlation value of around −0.30. This indicates that the decreasing relation between the two is significant. The correlation is weaker at −0.19 in the shoulder data set.

Table 2.13. Linear-Regression and Correlation Coefficients for No-Wind Cases for All Seasons and Medium and High Transactive Penetration Levels

Season	Transactive Penetration[(a)]	Slope (% Change/ Change in Scaled Cost)	Intercepts at 0 (% Change)	Intercepts at 1 (% Change)	Correlation
Summer	Medium	−0.25	0.09	−0.16	−0.32
	High	−0.63	0.21	−0.42	−0.34
Winter	Medium	−0.26	0.06	−0.20	−0.28
	High	−0.62	0.12	−0.49	−0.30
Shoulder	Medium	−0.22	0.04	−0.18	−0.19
	High	−0.65	0.16	−0.49	−0.18

(a) In this column, “medium” refers to the 10% transactive penetration case and “high” refers to the 30% transactive penetration case.

2.10.10 Direction of Transactive Response

This section delves deeper into the directions in which the transactive loads move the system-wide costs and loads simultaneously. To do this, Figure 2.66 divides the points in Figure 2.64 into three groups that we will call “terciles” using the two percentile values 33.3% and 66.7% of the scaled load data values. Each tercile is given a distinct point-type and color gradation.

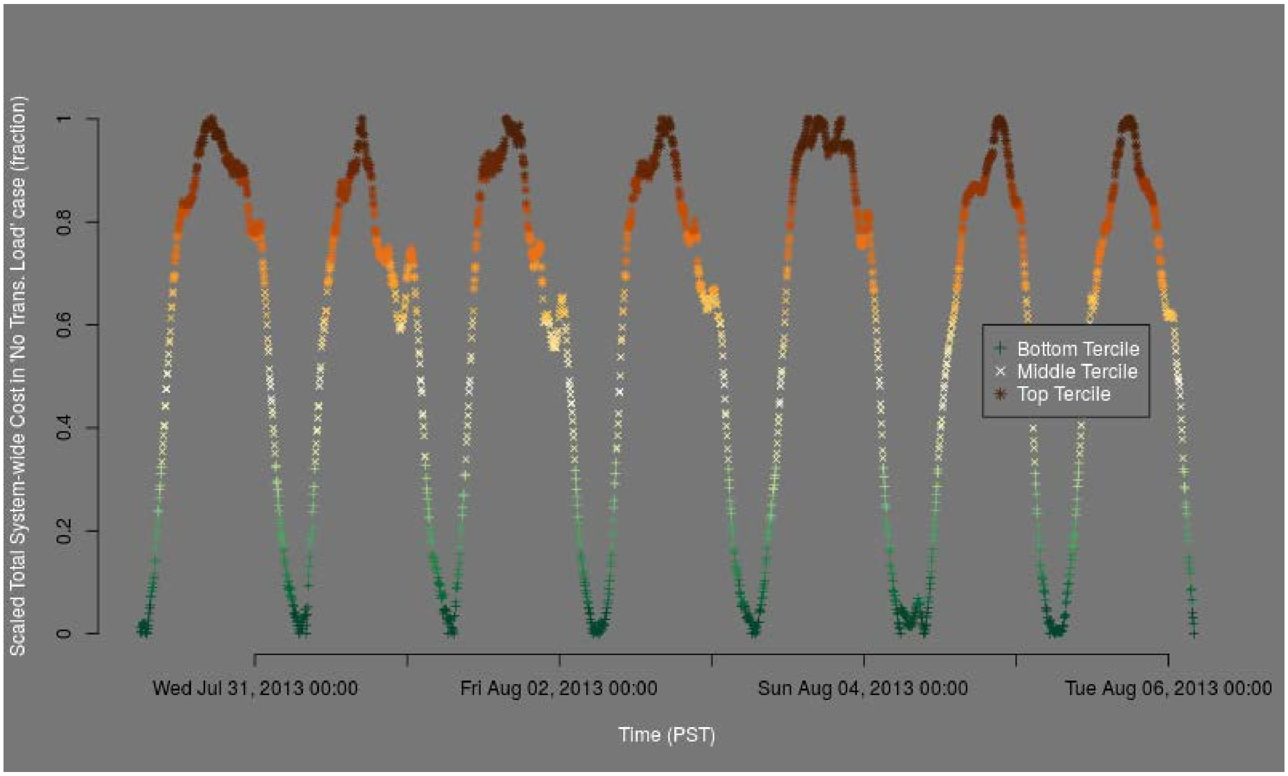

Figure 2.66. Total System-Wide Cost Expressed as a Scaled Dimensionless Quantity, for the No-Transactive-Load Case, using Distinct Colors and Point Types for Each Tercile of the Scaled Total Cost

Figure 2.67a and Figure 2.67b plot the relative change in total system-wide cost and load for each 5-minute simulation period using the colors and point types introduced in Figure 2.66. On the left, the relative change in the medium transactive-load case shows that most of the response has a positive, almost equal, relation between load and cost. In other words, a decrease (increase) in total load leads to the ED engine being able to calculate a corresponding decrease (increase) in total system-wide cost. Also, the larger decreases happen when system loads and costs are high, as was observed in Figure 2.65(a). The plot on the right shows the high transactive-load case. In addition to the observations on the left being strengthened on the right, the periods when load and cost increase concurrently are seen to be predominantly from the bottom tercile, when cost and load are both low. This confirms that most of these responses are due to the continuous-response assets leveraging cheaper energy to attenuate the higher cost periods.

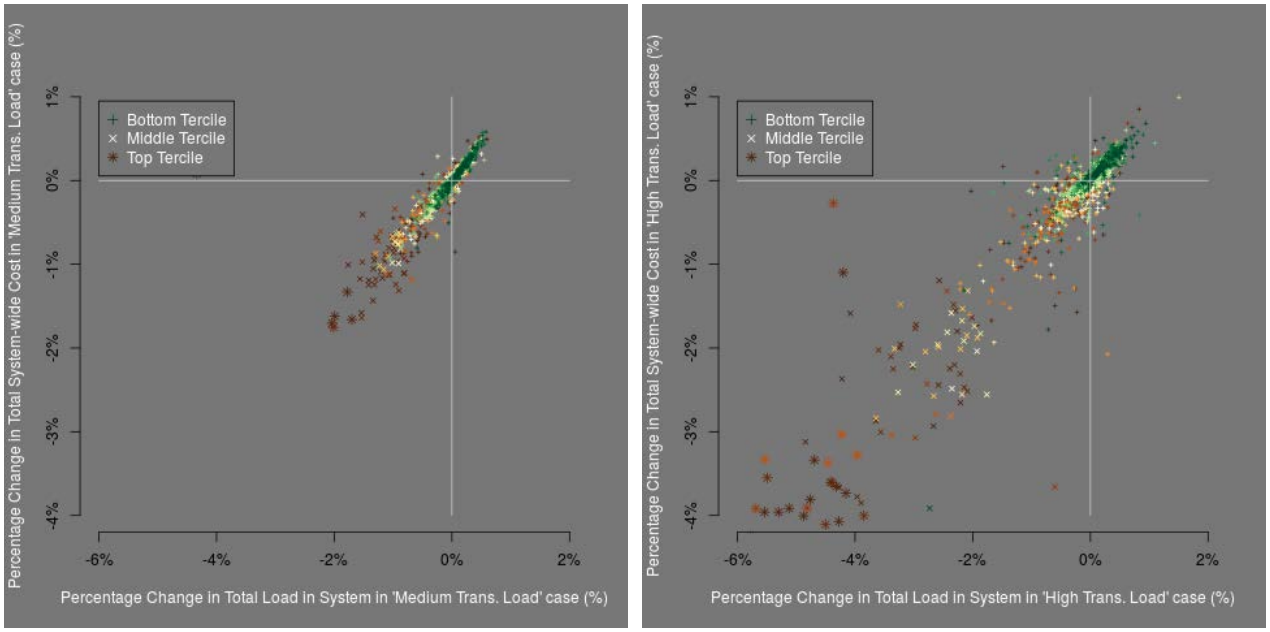

(a) Medium Transactive Load (b) High Transactive Load

Figure 2.67. Percentage Change in Total Cost vs. Load for (a) Medium and (b) High Transactive-Load Cases

2.10.11 Effect of Wind on Total System-Wide Costs

Figure 2.68 plots the variation in total system load vs cost for each 5-minute simulation period as the wind penetration in the portfolio is ramped up from no wind to medium wind and even high-wind cases, keeping the amount of transactive load at zero in all three cases. The points take a lighter shade of green if the total wind in the system is high at that 5-minute period. Wind, having been modeled as a zero-cost quantity, clearly has the effect of decreasing the cost of meeting the rest of the total system-wide load that is not served by the wind generation output.

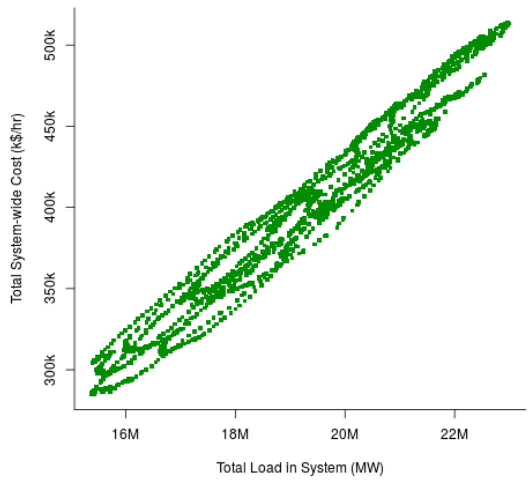

(a) No Wind Penetration

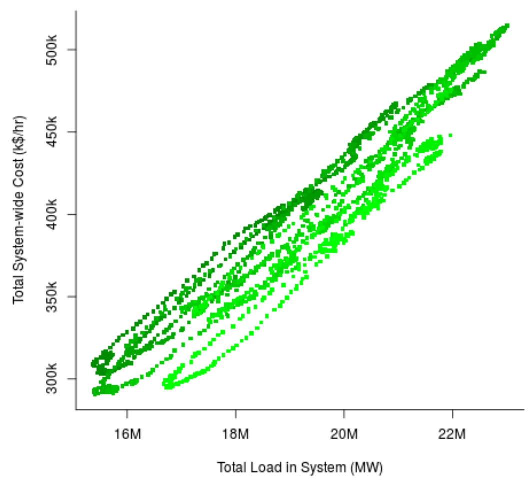

(b) Medium Wind Penetration

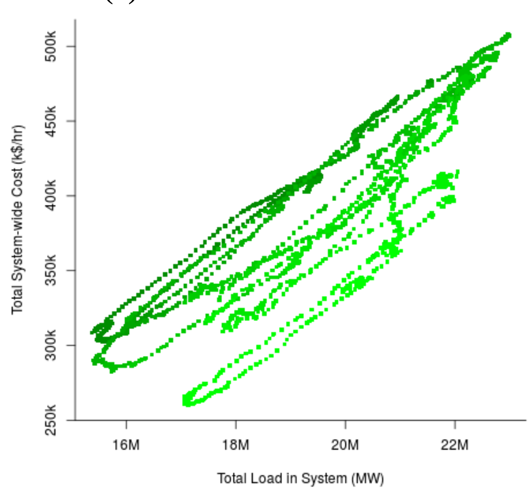

(c) High Wind Penetration

Figure 2.68. Total Load in System vs. Total System-Wide Cost for the No-Transactive-Load Case, with (a) No, (b) Medium, and (c) High Wind-Penetration Levels

A clear correlation is noticed between the reduction in system-wide total cost and the strength of the wind generation output. This correlation is more clearly observed in Figure 2.69, which plots the net change in total system-wide costs between the no-wind case and the medium- or high-wind cases against the corresponding wind generation output. Table 2.14 provides the coefficients of the linear-regression model fit over the data in Figure 2.69. The parameters from each seasonal data set and keeping the transactive load constant show a remarkable consistency in the direction and strength of the linear response. The direction of response (slopes of the linear model) is noticeably weaker in the winter data set. However, a weak positive relation is noticed in the slopes for this data set with the amount of transactive load in the system. This indicates that the two attributes, transactive load and wind-penetration level, may interact in a complex manner. This is the next subject of our analysis.

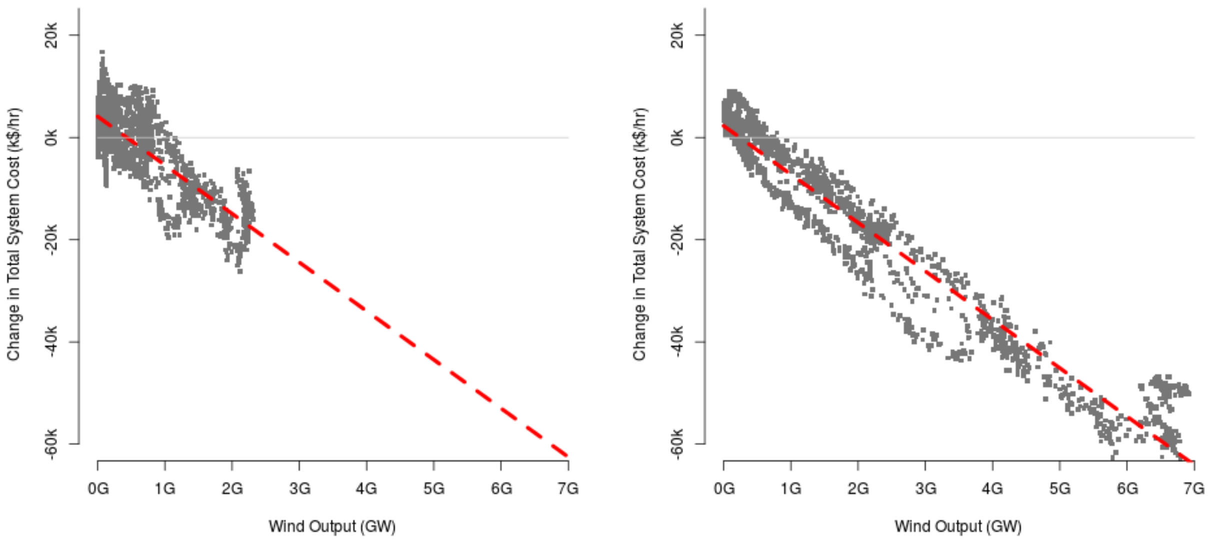

(a) Medium Wind-Penetration Case (b) High Wind-Penetration Case

Figure 2.69. Total Wind Output vs. Change in Total System-Wide Costs from the No-Wind Case, with No Transactive Load for (a) Medium and (b) High Wind-Penetration Cases

Table 2.14. Linear-Regression and Correlation Coefficients between Change in Total System-Wide Costs from the No-Wind Case and Medium and High Wind-Generation Outputs, for Various Transactive Load Penetration Levels

		Medium Wind			High Wind		
Season	Transactive Penetration	Slope (k$/GWh)	Intercept (k$/h)	Corr.	Slope (k$/(GWh)	Intercept (k$/h)	Corr.
Summer	None	−15.45	9.74	−0.95	−10.97	4.42	−0.97
	Medium	−15.44	9.51	−0.92	−10.97	4.29	−0.97
	High	−15.62	9.90	−0.78	−10.91	4.12	−0.97
Winter	None	−9.52	4.12	−0.82	−9.48	2.28	−0.97
	Medium	−9.65	4.26	−0.91	−9.51	2.37	−0.97
	High	−9.70	4.32	−0.70	−9.55	2.57	−0.96
Shoulder	None	−12.44	6.43	−0.95	−9.98	3.98	−0.97
	Medium	−12.63	6.64	−0.92	−9.99	4.01	−0.97
	High	−12.37	6.27	−0.93	−9.98	3.98	−0.97

Recall that the transactive control algorithms are designed to target the most stressed periods in the day. Table 2.15 takes a look at the interaction between wind power availability and the peak periods in a day. Correlations are provided between the observed total system-wide load and cost (as depicted in Figure 2.66 using distinct colors and point types for each tercile of the scaled total cost) for the no wind and the wind output in the medium and high wind cases. In all scenarios, the transactive load is maintained at the zero transactive-load case. Also provided are the correlations between the wind output and the observed average system-wide cost of electricity. The correlations in the summer data set are the weakest across the board, with no clear strong inference possible. On the other hand, the winter data set exhibits strong negative correlations. In other words, the wind output in the winter data set often occurs when the grid tends to have tighter constraints and higher costs in the system. This is also true in a weaker sense in the shoulder data set. Note that the correlations under the medium and high wind-penetration data sets are identical because the correlation metric is insensitive to linear scaling, and each wind data set is derived by linearly scaling a currently observed wind generation to the desired level.

Table 2.15. Correlations between Medium and High Wind-Penetration Cases and Observed Values from the No-Wind Case

Season	Wind Penetration	Total System-Wide Load	Total System-Wide Cost	System-Wide Average Cost
Summer	Medium	0.03	−0.08	0.15
	High	0.03	−0.08	0.16
Winter	Medium	−0.45	−0.38	−0.52
	High	−0.45	−0.38	−0.52
Shoulder	Medium	−0.10	0.01	−0.24
	High	−0.10	0.01	−0.24

2.10.12 Interaction of Wind Output and Transactive Response

Transactive assets are most active at the time periods in the day when the total system-wide costs are at their highest. Wind generation depends solely on weather conditions. Thus, complex interactions may be observed in cases where the high wind period of the day also coincides with periods when the total system costs are traditionally higher. Figure 2.70 graphs such an interaction observed in a day of simulation using the summer data set. On the left is the total system-wide cost under different cases of wind output while no transactive load is imposed on the system. The peak wind period is seen to significantly affect costs, changing a flat peak from 08:00 to 13:00 in the no-wind scenario into a pronounced peak at 14:00 when the effect of high wind penetration is included. For these same wind outcomes, when the system additionally allows high transactive loads, the response from the transactive loads show a marked interaction with the wind outcomes. It is apparent that the peak period of response from these assets changes with respect to the amount of wind available, with a peak response in the no-wind case at about 10:00, which moves to 13:00 and 12:00 for the medium and high wind-penetration cases, respectively.

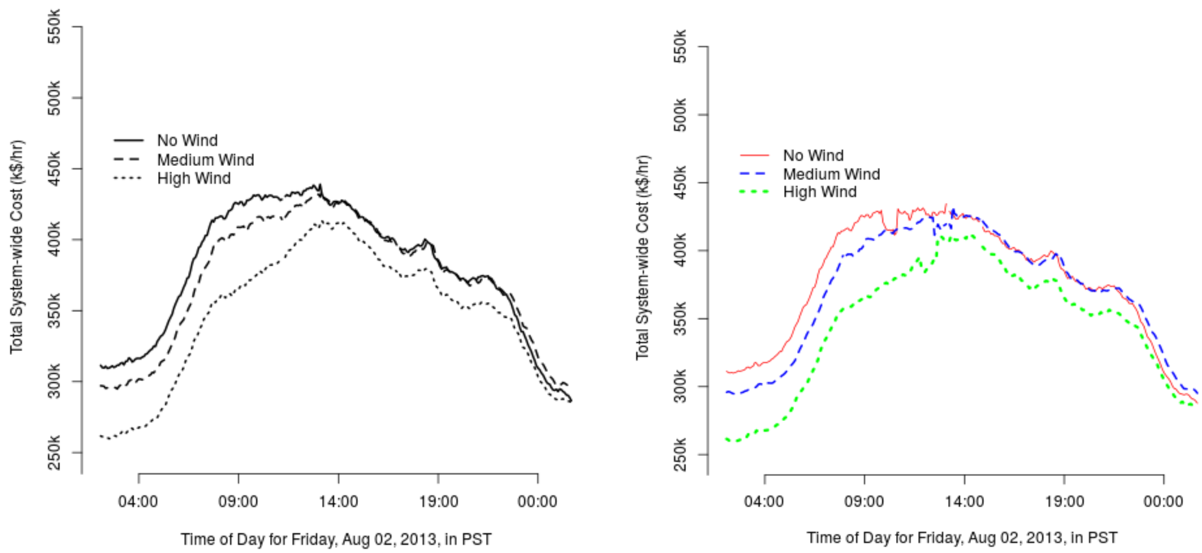

(a) No Transactive Load

(b) 30% Transactive Load

Figure 2.70. Total System-Wide Cost for a Day in the Summer Data Set, with (a) No Transactive Load and (b) 30% Transactive Load Penetration

Figure 2.71 provides some very interesting data on the interactions between transactive load and wind generation. On the left is plotted the difference in total load in the system in a case that combines high transactive load with no, medium, and high wind penetration against a base total load measured in the no-transactive-load and no-wind cases. The plot on the right gives the corresponding differences in total system-wide costs. These plots show that when no zero-cost source such as wind is available, the appropriate response from the transactive loads is to reduce load at appropriate peaks as was observed in Figure 2.70, realizing concomitant reduction in total systemic costs. However, when wind is available, the nature of the local peak in average cost to the system changes significantly. This is most clearly observed in the high wind-penetration case. As observed earlier from Figure 2.70a, the wind generation is high on this particular day during the peak hours, and changes the flat long peak to a short sharp peak. In terms of the average cost of power, some of this formerly high-average period is actually converted to a low average valley in the high wind-penetration case. In Figure 2.71a, we also notice that in the high wind-penetration case, the amount of wind output is high enough for the transactive assets, principally the continuous-response assets, to take advantage of inexpensive wind power by actually increasing their load on the system. This increases the total load in the system during a period that was formerly part of the long flat peak in the absence of wind. However, this increase in load does not affect the overall cost in the system, which continues to be below the no-wind case, as observed on the right. So, wind output being high during peak systemic load can be very beneficial when a significant portion of the load can respond quickly in the transactive system.

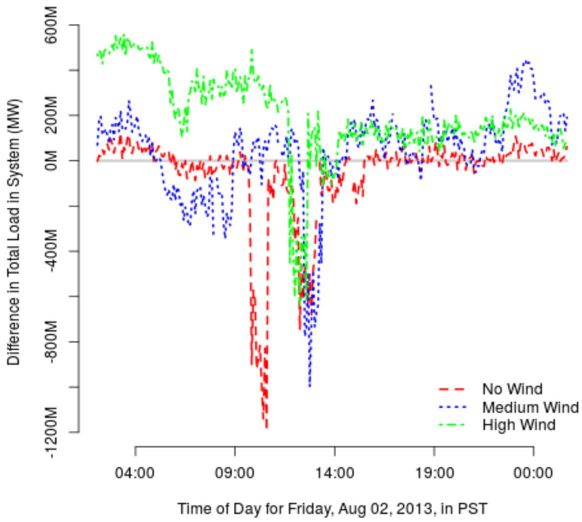

(a) Effect on Total Load in System

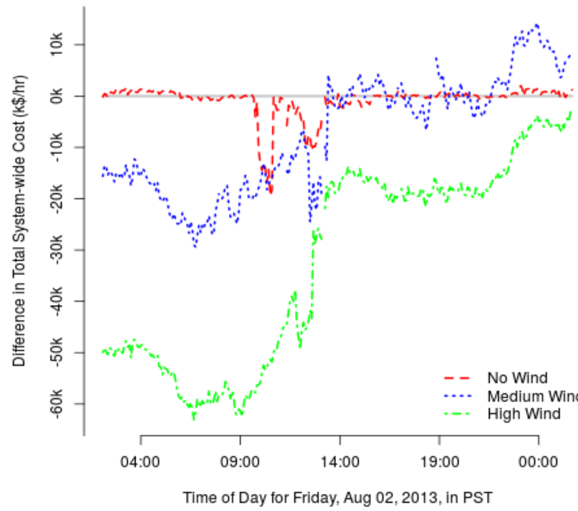

(b) Effect on Total System-Wide Cost

Figure 2.71. Effect of Interaction of the High Transactive-Load Penetration with Wind Output. Each trace shows a difference between high and no-transactive-load cases, with separate traces for scenarios having no, medium, or high wind penetration.

Table 2.16 provides the coefficients of linear-regression models fit to the percent change in total systemic costs observed in medium or high transactive-load cases compared to the no-transactive case costs as a function of the scaled total cost in the no-transactive case. These coefficients are similar to the values provided in Table 2.14 except that they treat the cases where wind is also present in the simulation, either at high or medium penetration levels.

The interaction of transactive assets and wind output shows a markedly different nature in the summer and winter and shoulder data sets. Presence of higher wind output seems to moderate the impact of having transactive load in the system for the summer results, as seen in the declining slope values of the linear-regression models. This may be attributed to the pattern of wind output being slightly in sync with periods when transactive loads typically act (i.e., high systemic cost periods), as discussed in the preceding example in Figure 2.71. On the other hand, wind output has the impact of strengthening the transactive response in the winter months, as evidenced by the similar or higher slope values in this case. This is possible when the pattern of wind output does not match or is antithetical to the high-cost periods, as observed in Table 2.15. However, in all cases, the quality or determinacy of the fit, as measured by the correlation value, decline compared to the no-wind case, indicating that the pattern of wind output and the peaks and troughs of total systemic costs are only very weakly related, as can be expected from the fact that wind output is driven only by weather phenomena.

Note that a majority of the transactive assets studied in this simulation can only intervene by dropping pre-existing residential loads during periods of high average cost of electricity, but are not designed to increase or shift consumption to periods of low average costs. Thus, the intersection of wind output and transactive response is mostly observed when high wind output reduces the peak average costs and obviates the need for a transactive response. A more sophisticated responsive asset that could also respond to high-wind presence facilitating drops in average costs is suited to extract more systemic benefit by providing mitigation to both high and low costs, as they are affected by renewable generation.

Table 2.16. Linear-Regression and Correlation Coefficients Modeling Percent Change in Total Cost in 10% or 30% Transactive-Load Cases as a Function of the Scaled Total Cost in the No-Transactive-Load Case for 10% and 30% Wind-Penetration Scenarios

		No Wind				Medium Wind				High Wind			
			Intercepts				Intercepts				Intercepts		
			at 0	at 1			at 0	at 1			at 0	at 1	
Season	Transactive Penetration[b]	Slope[a]	(% change)		Corr.	Slope[a]	(% change)		Corr.	Slope[a]	(% change)		Corr.
Summer	Medium	−0.25	0.09	−0.16	−0.32	−0.169	0.050	−0.119	−0.084	−0.049	−0.033	−0.082	−0.043
	High	−0.63	0.21	−0.42	−0.34	−0.266	0.000	−0.266	−0.103	−0.290	0.019	−0.271	−0.155
Winter	Medium	−0.26	0.06	−0.20	−0.28	−0.238	0.004	−0.234	−0.101	−0.256	0.030	−0.226	−0.170
	High	−0.62	0.12	−0.49	−0.30	−0.723	0.160	−0.563	−0.113	−0.587	0.063	−0.524	−0.189
Shoulder	Medium	−0.22	0.04	−0.18	−0.19	−0.32	0.09	−0.22	−0.11	−0.29	0.08	−0.21	−0.12
	High	−0.65	0.16	−0.49	−0.18	−0.54	0.09	−0.46	−0.22	−0.74	0.20	−0.54	−0.18

(a) The units for this column are % change/change in scaled cost

(b) In this column, “medium” refers to the 10% transactive penetration case and “high” refers to the 30% penetration case.

2.10.13 Conclusions

IBM's PNWSGD simulation platform demonstrated the complex interrelations that emerge from having a high penetration of transactive assets and renewables in the PNW power grid. Here are some key observations and conclusions derived from the simulation study:

- The transactive assets that use the event-driven or daily-event control logics designed in this project are effective in targeting load reduction to peak-cost periods in a day. The accuracy of the response depends on the sharpness of the peak of the average cost of electricity, as seen in Figure 2.56 through Figure 2.61.
- The cost and output characteristics of the bulk generation capacity in the PNW power grid lead the ED module to maintain a similar average system-wide cost throughout the range of total load (Figure 2.62 and Figure 2.63), but there is enough variation for the transactive assets to behave in the manner expected.
- Transactive assets that are continually responding to peaks and valleys in the average cost of electricity attempt to balance the overall energy usage by buying (charging) more power at low-cost periods (valleys) and selling (discharging) power at high-cost peak periods. This is observed in Figure 2.65 and is more clearly apparent in Figure 2.67.
- The transactive system responds to high average-cost of electricity by reducing the total system-wide load, and also increases the total-load in low average cost periods. As an example, a reduction in total cost of up to about 8% is observed in peak-cost periods when the presence of transactive assets in the system is high (Table 2.12).
- The change in total consumption when transactive assets are introduced in the system exhibits an inverse relationship with total system-wide costs, with a significant negative slope when modeled as a linear function (Figure 2.65).
- The presence of wind in the portfolio of generation has a strong impact in reducing overall system costs in meeting demand (Figure 2.68 and Figure 2.69). Table 2.14 shows that a reliably strong, negative relation exists between renewable production and total system cost. This impact is higher with higher presence of wind.
- The interaction of renewable generation and the transactive system can be complex. Presence of renewable generation can change the periods when transactive assets take action (Figure 2.69a). When high wind output suppresses costs in an otherwise peak-cost period, transactive assets designed only to shave peak load do not show any response. The cost suppression can the high-wind case also create periods of low enough average costs in former peak periods that the continuous response assets that seek the lowest costs of the day exhibit a strong activation (Figure 2.69b).
- The overall effect of wind on the strength and repeatability of the transactive response to the system costs depends on the pattern followed by the wind output.

- Table 2.16 shows that in the summer data set, the presence of wind weakens the transactive response, while in the winter data set its presence does not affect, or even slightly strengthens, the transactive response observed.

2.10.14 Summary and Need for Further Work

The PNWSGD transactive system created an informed simulation to represent much of the region's transmission and generation. It accurately represented wind, thermal, and hydropower resources when timely production data from these resources were explicitly available. The system accurately emulated the scheduling and dispatch order of various load-following resources when those resources had to instead be inferred. The mapping of actual resources, load, and transmission into the transactive system's topology proved to be challenging and the accuracy of the mapping was not convincing.

The TISs were constructed at distributed locations to represent the delivered unit cost of electricity at the nodal location and time. The transactive system's method of monetizing and blending resource energy and incentives was workable. Resistance was encountered to dynamically stating the incentives over time, and much work remains before we should anticipate acceptance of the method by regulators and as the basis for a real energy tariff. A set of parameters were recommended and used to represent dynamic cost components into the TIS formulation, and this approach may be a basis for interoperability in systems like this that must formulate distributed incentives. Caution must be used, however, because the selection of parameters from among this set can create undesirable consequences, as was the case when the project elected to represent infrastructure costs using a parameter with units dollar per hour, which had the undesired consequence of lowering costs during periods of highest energy demand.

The transactive system included predictions over a set of sequential time intervals that extended several days into the future. The TIS predictions and component resource predictions were found to suffer from prediction biases. A step difference was found about 3-1/2 hours into the predicted future that caused some of the transmission nodes to over predict future TIS values and others to under predict them. The project hypothesized that these biases occurred because different calculation methods were used before and after the time. These prediction errors may have serious adverse consequences for assets that rely on the predictions to schedule their operations.

Despite the challenges encountered in formulating the TIS values, the PNWSGD transactive system itself was robust and reliable at communicating its signals throughout the region. IBM created a reference implementation that was eventually adopted by nine of the utility participants. Four utilities attempted to create their own implementations from the project's specifications. Of these, one succeeded, another later accepted the IBM reference implementation, and two were unable to establish a compliant instantiation.

A suite of functions were developed by the project to help the participating utilities identify times that its assets should respond. Different functions were needed for infrequent events, daily events, and continuous responses. The approach proved workable, and the responses were shown to have occurred as they should during the corresponding high and low incentive values. The correlation was, however, strongly influenced by the care with which the functions had been configured. The functions that remained poorly configured or entirely unconfigured did not, of course, perform well.

The responsive asset systems—water heaters, thermostats, dynamic voltage control, etc.—also modeled and predicted their impacts on net load for times that the assets were advised by the transactive system to respond. Simple models were used by the PNWSGD, but these models could be made more accurate by future implementers. The project reported the ranges of the modeled power impacts, but the project lacked means to accurately calibrate or validate the impacts. Again, the approach proved workable, but the accuracy was strongly affected by the attention paid to calibration and configuration of the models.

The changes in load were summed at the utility nodes and were reflected in the predicted net load (TFS) between the utility sites and the respective transmission zones based upon which the sites were modeled to receive their energy supplies. At many sites, the corresponding TFS could be directly compared with the power meter data. The comparison was disappointing at many of the sites. Some sites failed to calibrate their load predictions. Others calibrated the power from only a couple months' data, and the corresponding models then failed to track seasonal variations. All of the sites would have benefited from stronger connections between the models and the real-time metered data that was to be tracked and predicted. The load predictions would need to be accurate (better than about ± 5%) if these predictions are to usefully inform the predictions of the balance between load and supply at the distributed transactive sites. Much future work is needed in this area.

Finally, the PWNSGD transactive field demonstration was not permitted to directly affect the scheduling and dispatch of the region's power resources. Analysis Step 8 (Resources Must Respond to Dynamic System Load Predictions, Including the Plans from Flexible Loads) could not be tested in the field. Instead, simulation was conducted by IBM to help the project scale up the modeled penetration of transactive assets and to close the control loop so that the connection between assets' responses and the dispatch of regional resources could be tested. Total load and total incentive costs were observed in the simulations to have decreased as the daily peak incentive costs were occurring. A smaller increase in load and incentive costs was observed as modeled battery systems reacted when minimum daily incentive costs were occurring. There was a complex interaction between dynamic wind power and these impacts within the transactive system because the wind power dynamically affected incentive values.

Here are some specific recommendations for issues and future work to be addressed in the development of future transactive systems:

- Many more responsive assets are needed. If a truly distributed system is to become viable, the changes in power offered by its responsive assets must be comparable in total magnitude to the changes in power available today from the supply side.
- More flexibility should be available from each asset. Today's demand-response programs and their assets allow for only several events each month for a few hours at a time. These programs might address peak demand, but they are otherwise limited in the services they can provide.
- Even battery systems, which were anticipated to offer great dynamic responsiveness, were found to be limited to no more than about one charge and discharge cycle per day.
- The project instigated the exchange of its transactive signals based mostly on timed 5-minute intervals. There is emerging consensus among project participants that future transactive system implementations should be more event-driven than timed. Communication of transactive signals

should only take place when there is evidence or likelihood that the system has appreciably changed or that predictions have become inaccurate since the time that the information was last communicated.

- In a loosely connected, distributed system, the validity and accuracy of the signals that are being received from neighbors might be in question. The project defined, but did not implement, a confidence attribute that could accompany transactive signals to state the sender's confidence in the calculated quantities. The confidence attribute might then temper the recipient's trust in the signal's values, which might further temper the actions that the recipient takes based on the signal's questionable values.
- Another incentive function had been planned that would have represented impacts of transmission congestion on energy costs, but this function was not successfully implemented by the project. Dispatch opportunities that would stress transmission capacities are outright disallowed today. The project's intention was that as transmission approached a stressful capacity, the cost of the transmission might be smoothly incremented to dissuade consumption of that power. A function having such smooth response was not found. The function's output rapidly changed the unit cost of energy faster than these changes could be responded to by the system. Furthermore, the responses from the transactive system were inadequate to mitigate the congestion and therefore could not stabilize the proposed function.