Standardization of a Hierarchical Transactive Control System

Donald Hammerstrom and	Terry Oliver	Ron Melton	Ron Ambrosio
Pacific Northwest National Laboratory 902 Battelle Blvd., K1-85 Richland, WA 99352	The Bonneville Power Administration 905 NE 11th Ave Portland, OR 97218	Pacific Northwest National Laboratory 902 Battelle Blvd., K1-85 Richland, WA 99352	International Business Machines (IBM) PO Box 218, Yorktown Heights, NY 10598-0218
donald.hammerstrom@pnl.gov	tvoliver@bpa.gov	ron.melton@pnl.gov	rfa@us.ibm.com

Keywords: Transactive control, real-time pricing, distributed energy resources, load control

Abstract

The authors describe work they have conducted toward the generalization and standardization of the transactive control approach that was first demonstrated in the Olympic Peninsula Project for the management of a transmission constraint. The newly generalized approach addresses several potential shortfalls of the prior approach: First, the authors have formalized a hierarchical node structure which defines the nodes and the functional signal pathways between these nodes. Second, by fully generalizing the inputs, outputs, and functional responsibilities of each node, the authors make the approach available to a much wider set of responsive assets and operational objectives. Third, the new, generalized approach defines transactive signals that include the predicted day-ahead future. This predictive feature allows the market-like bids and offers to be resolved iteratively over time, thus allowing the behaviors of responsive assets to be called upon both for the present and as future dispatch decisions are being made. From the resources' perspective, the predictions allow the responsive resources to anticipate and therefore proactively participate in coming peak events, at times taking energy at the current cheaper price on the bet that a future higher price may be avoided.

1. BACKGROUND

In *transactive control*, responsive demand assets bid into and become controlled by a single, shared, price-like value signal, which may be, in turn, influenced by many local and regional operational objectives of the electric power grid. The approach was first demonstrated for the control of a transmission constraint during the Olympic Peninsula GridWise Project that was funded by the U.S. Department of Energy from 2004 - 2007 [Hammerstrom 2008].

The Olympic Peninsula Project's responsive assets included residential thermostats, residential water heaters, residential clothes dryers, commercial HVAC systems, distributed diesel generators, a gas turbine, and municipal water pumps. Algorithms were formulated to automatically generate bids and offers from these responsive demand assets based on user preferences and the degree to which the assets' processes (e.g., room temperature or water level) had been satisfied. Commercially available home energy management system components and additional engineered solutions communicated and acted upon the value signals and bids. Ultimately, the project's long-haul communications were facilitated via the Internet to the project's control center at Pacific Northwest National Laboratory. Coordination of the diverse system components was engineered with IBM using their Internet Scale Control System (iCS), a WebSphere™ based middleware software.

In the Olympic Peninsula GridWise Project, a transmission constraint was imposed on a set of homes and businesses participating in the project, and power supplied by a virtual transmission line was successfully limited to this imposed constraint for over a year. The transactive control approach proved viable and exhibited many useful attributes. For example, the approach demonstrated the practical value of a high degree of automation, by which the responsive assets were called upon only when and to the degree their responses were needed, and this automation resulted in successful operation of multiple complex assets to meet a severe artificial, but probable, constraint. Unexpectedly, because thermostats were configured to track current and future relative price, thermostats behaved quite opportunistically, taking advantage of low-cost opportunities to pre-heat or pre-cool living spaces without requiring any explicit algorithm be applied for that purpose. Customers therefore experienced relatively little discomfort because they were able to select, within some degrees of freedom, how much comfort they were willing to forego in exchange for incentive (price-like values) benefits.

The remainder of this paper suggests improvements to the transactive control approach. The approach has been generalized and formalized to make it practicable for any set of demand assets and many grid objectives. This generalized formulation of transactive control is a worthy foundation for standardizing the practice of real-time price control.

2. IMPORTANT TENETS OF TRANSACTIVE CONTROL

Some consistent tenets have driven the evolution of the transactive control approach:

- Communicate value—communicate the value of the power grid's control benefits via a single value signal at each location. The value signal is not necessarily the monetary energy price that is to be used for revenue and billing purposes. The use of a single value signal forces all benefits and costs to be weighed fairly, in advance, and openly using a common currency, and enables machine response without active daily or real-time occupant action.
- Dynamic signals—demand responds to fluctuations in the value signal and thereby helps moderate the value signal. Valuable grid objectives (e.g., fast frequency regulation) may be achieved if time intervals are short enough to respond to such intervals.
- Facilitate interoperability—allow multiple communication media, protocols, and vendors to coexist and compete
- Multitask—each responsive asset should respond to any operational objectives that it is able to help accomplish. Multiple grid objectives simultaneously influence the value signal.
- Respond 24/7—the control of demand assets can be put to valuable use continuously, not simply for the few critical stressed periods of each year
- User-friendly—if many responsive demand assets participate, valuable responses may be had with little or no inconvenience to customers. Individual customers should always have the right to temporarily override asset responses.
- Distributed control—specific control decisions are best made nearby and by the controlled assets.
- Aggregators are not required—aggregators are not necessarily required if specific responses are decoupled from the communication of a value signal, as is advocated within transactive control.
- Low bandwidth—the use of distributed control and the reduction of communication to a single value signal serve to reduce overall communication bandwidth.

3. CONTROL OF DEMAND ASSETS

The tenets stated above have important implications for the design of responsive demand and distributed resource assets. Presently, demand-response assets are uniquely engineered for specific types of utility programs. In transactive control, many and multiple responsive assets are encouraged and need not be programmatically placed. A responsive asset does not even need to know exactly which objective(s) it is helping to accomplish at any given time. If customer incentives are adequate, populations of responsive assets will grow. The means by which the assets are to respond should be engineered by manufacturers of the assets (or by home energy management system manufacturers) and may further be influenced by how customers configure these assets to respond.

Perhaps the most important requirement placed upon responsive assets is that each responsive asset should reveal its need for or willingness to provide energy. In the prior Olympic Peninsula Project, each asset's bid was an explicit, monetary bid, but any feedback concerning how a device would favor or avoid a given price (i.e., the value signal) is useable. The intelligence of the device can be resident in the device, within a buildings energy manager, or even more centrally.

4. IMPORTANT GRID OPERATIONAL OBJECTIVES

The reformulation of transactive control has been influenced by the operational objectives and benefits that we value most for our power grids. Among the most important are

- Facilitate renewable resources—challenging renewable portfolio standards will require that impressively large amounts of wind and solar resources be accommodated. These resources are imperfectly predictable and have dynamic attributes that necessitate equally dynamic control of supply, demand, and perhaps storage. The consumption of renewable resources can be facilitated if, for example, wind energy is discounted when and near where it is generated.
- Mitigate operational constraints—utilities operate closer to their operational margins. Locations throughout the grid should have means to dissuade the consumption of power delivered through that location if served power threatens to damage or shorten the useful life of equipment. In the Olympic Peninsula Project, the value signal was permitted to rise when the transmission constraint was exceeded, and the higher value signal persuaded demand assets to either use less power from, or provide distributed generation to, the feeder circuit.
- Flatten load—system efficiency improves and less infrastructure and fewer peaking generators are needed when load is moved off peak. Customers would choose to defer responsive loads off peak periods if the costs of supplying such premium power were made transparent to these customers.

5. NEWLY RECOMMENDED ATTRIBUTES FOR TRANSACTIVE CONTROL

The following several improvements should be implemented to increase the applicability of the transactive control approach and to make the approach more amenable to standardization:

1. Enforce a hierarchical communication structure
2. Create an initialization and maturation plan
3. Formalize generalized transactive inputs, outputs, and behaviors
4. Require a forecast time horizon.

6. HIERARCHICAL STRUCTURE ENFORCED

Communication and control within a smart grid should be aligned well with the flow of electrical power. Define a *node* as a physical point anywhere in the electric power grid where demand may be aggregated and predicted. We propose a contiguous hierarchy of nodes from end uses through generation. Demand capacity is to be aggregated through the hierarchy from end uses toward generation (the *upstream* direction); a value signal is to be propagated from generation toward end uses (*downstream*) through the hierarchy. It will be shown that it is the interplay between the demand capacity and value signals that defines transactive control. Control approaches that ignore or jump over points within the hierarchy violate this ideal and will not correctly address local control objectives at nodes.

A great example of such a violation addresses our zeal to electrify our transportation infrastructure and charge the batteries of electric vehicles at our homes. Unfortunately, the pole top distribution transformers that serve several residences were not often sized to simultaneously provide power to multiple vehicle battery chargers. Therefore, unless we specifically include the pole top transformers as nodes in our control hierarchy, the transformers will be unable to help manage the power they provide and thereby protect themselves.

Figure 1 demonstrates the principle of the proposed hierarchical structure. In this figure, the value signals flow downstream toward the left (labeled "operational objectives"), and the corresponding demand capacity signal flows upstream toward the right (labeled "status and opportunities"). Figure 1 does not at all suggest that responsive assets and the formulation of value signals occur only at the extreme upstream and downstream locations. Indeed, just as every node can interject the value of meeting its own operational objectives, responsive assets can reside quite far upstream and even in transmission in the forms of flow control devices, resource dispatch practices, and voltage control devices.

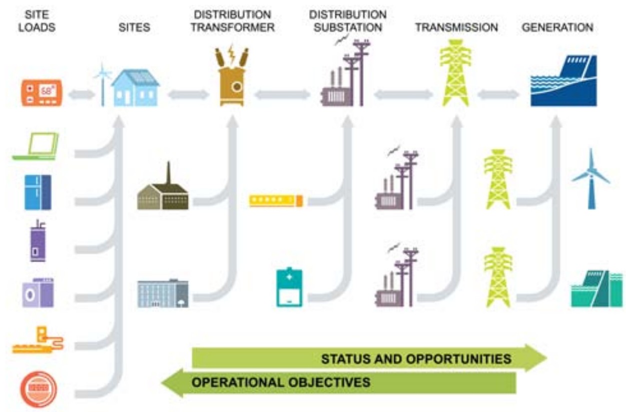

Figure 1. Representation of Proposed Hierarchy [2]

7. INITIALIZATION AND MATURATION PLAN

The hierarchical, transactive control approach varies greatly from the present deterministic way in which the grid is managed. Therefore, a transition plan is suggested that will first introduce hierarchical, transactive control into regions, then will provide ways for the approach to mature and expand.

Initially, a transactive node, or a pair of transactive nodes, is to be assigned at the intersection between the region's transmission system and each utility distribution site (Fig. 2). These initial nodes become an anchor of the hierarchical system from which the hierarchical structure can later become expanded. Initially, some objectives will be imperfectly addressed at the initial nodes. With an incomplete hierarchical node system, neither regional nor local objectives can be accurately connected to resource availability and upstream constraints.

This initial node pair is of interest only to the degree that it will provide for control of assets at that node or downstream of the node. The hierarchical structure is allowed to be temporarily relaxed during this initial stage, allowing some downstream assets to be controlled at these initial nodes and jumping over some passive nodes that will not fulfill their responsibilities to aggregate demand and modify the value signal.

Once this initial installation has been completed and tested, the hierarchical, transactive system of nodes should expand and mature. The hierarchy may expand as adjacent nodes, both upstream and downstream from the initial node, become transactive nodes. Control matures also as transactions at the existing nodes are made richer and more accurate. Ideally, the transactive control system will encourage participation by more and additional types of responsive assets over time.

The maturation plan thus facilitates the introduction of hierarchical, transactive control throughout a region and provides for the evolution of that system into a complete, rich transactive system. This plan necessarily makes compromises during the initial installations. A side-by-side comparison of the initial compromised conditions and ideal final behaviors at nodes has been summarized in Table 1 below, which includes additional maturation indices beyond those that can be addressed here.

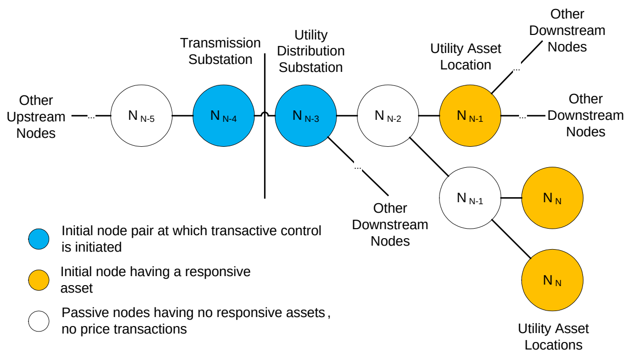

Figure 2. Initial Nodes Anchor the Formation of a Complete Hierarchy

8. STANDARD NODE DEFINITION

A generalized formulation of transactive control includes definition of the inputs, outputs, and functional responsibilities of any node. During this discussion, refer to the simplified functional block diagram of a node's responsibilities in Fig. 3.

The generalized definition of a node includes only two necessary communication pathways. A value signal is communicated downstream through the node, and the demand (or capacity) signal is communicated upstream. Other diverse communication and control signals might be used at a node, but these additional signals are only locally relevant and are not part of the transactive control system.

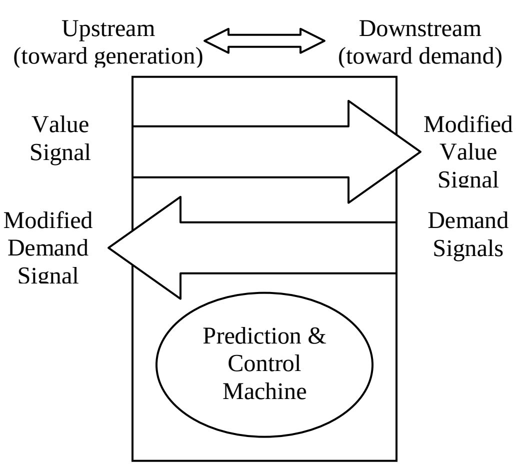

Figure 3. Simplified Functional Block Diagram of a Node

Table 1. Comparison of Initial Compromise and Final Implementations of Hierarchical Transactive Control

	Initial	Improved / Final
Hierarchy	Incomplete, anchored by several initial nodes	Transactive behaviors migrate upstream and downstream Eventual complete hierarchy
Benefits	Small set of grid objectives addressed	Many grid objectives addressed through extensive hierarchy by transactive control
Finesse	Grid objectives addressed crudely	Grid objectives configuration and location dependent
Nodes	Many passive nodes that control no responsive assets	Nodes become more active and transactive
Control	Many responsive assets controlled centrally or by upstream assets	Control becomes more and fully distributed
Intervals	Crude intervals of 15 minutes to 1 hour are acceptable	Time intervals become shorter and event-driven
Predictions	Crude estimates	Accurate predictions and trending
Dynamics	Limited dynamics; emulate traditional demand response and time-of-use	Increasingly dynamic; fast customer-friendly responses
Signal	Value signal not equivalent to expense. Incentive and billing distinct from value signal	Value signal trusted for customer incentives and billing

Value signal input. In most instances, a node will receive a single value time series from one upstream node. Occasionally, due to non-radial circuit configurations or multiple resource inputs, a node might be downstream of two or more upstream nodes and must therefore formulate a single, blended input value time series from the multiple value time series that it receives. A preferred unit for this signal is cents per kWhr. The transactive value signal may be used, but is not necessarily used, to create customer incentives and specify customer billing at that node. The use of actual energy price as the value signal is preferred. In principle, the value time series should represent a predicted energy price over the next 24 hours, or so. The predicted value signal should become increasingly accurate and should have finer intervals near-term.

If a node receives more than one price signal time series from upstream nodes or resources, the node must blend those prices into a single input price series. The recommended approach is to calculate the single input price series as a weighted sum of all input price time series, where each time series component is weighted by the fraction of supply received from each upstream node or resource during an interval.

Value signal output. Often, the received input value time series will be relayed to all the next downstream nodes without modification. However, the node may choose to modify the value time series before relaying it downstream in order to address its own local operational objectives. For example, a node might choose to increase the price signal at a future time interval to avoid an impending constraint at the node.

Demand inputs. The node will receive or measure demand served at the node and by all downstream nodes that are served by the node. The use of the word demand here is not at all intended to preclude cases where distributed resources might actually supply energy, as is the case for distributed or renewable generators. The responsibility of the node will be to aggregate all served demand into a single aggregated demand time series. The preferred unit is kW. The demand time series should represent a predicted demand for time intervals over the next 24 hours, or so. The demand signal should become more accurate and should have finer intervals near term. The time intervals of the price series should be the same as those used for the demand series.

The easiest demand inputs to be aggregated at a node will be those from downstream nodes that have already calculated and provided their demand time series for use by this node. It does not matter whether these downstream nodes are responsive to transactive price signals or not.

More challenging is the demand that is measured at the node but provides no, or an incomplete, time series prediction. In this case, the node is responsible to predict future demand and complete the time series for which it is responsible. The function that creates such a prediction within the node might be called a “prediction machine.” Those devices that are to be controlled at or by the node are additionally responsive to the transactive value signal at the node. Therefore, the price elasticity of controlled assets must be considered by the node’s prediction machine.

Aggregated demand output. A node is responsible to aggregate all present and future demand that it serves and make the demand prediction known to any upstream node. The present demand of a node will usually be verifiable using existing meters. The accuracy of a node’s demand prediction can also be monitored, assessed, and improved over time.

Control machine. Figure 3 uses the word *control machine* to describe a node’s opportunity or responsibility to modify the value signal. A node need not reveal the formula it uses to modify the value signal. The aggregated demand output time series that is calculated at the node is one of the most important formula inputs into the control machine. Several reasons that a node might choose to modify the value signal are limited resources, constrained infrastructure, or efforts toward performance optimization.

There exist no (and might never exist) definitive formulas and practices for the modification of the value signal at a node. However, the author was able to formulate workable initial functions for each of the operational objectives listed in section 4 of this paper. These functions may be improved over time.

Demand prediction machine. A node’s *demand prediction machine* receives measurements and predictions from all responsive and unresponsive demand that is served by the node and aggregates and predicts a future time series aggregate demand. Many types of data can be incorporated by the prediction machine as it strives to produce accurate demand predictions. Simple historical trending is recommended as a first approach, and such predictions can be made increasingly more accurate, if necessary, as additional information becomes available.

Asset participation. The means by which a node affects the energy consumed by devices that are controllable from the node itself will be diverse. By control, we refer only to those devices whose energy consumption (or generation) is managed by the intelligence and transactive behaviors at the node. Again, this control is treated as a black box; there is no need for the node to share its function outside the node. An example of a simple control function might be the curtailment of water heater load whenever node pricing exceeds a price threshold. This simple function could effectively automate most time-of-use responses. Another simple function would be for the node to curtail a water heater’s load when the node’s load (let’s say a home’s load in this example) starts to exceed a threshold capacity. These two simple control functions could run simultaneously using the same controllable water heater. In a more complex example, a home’s thermostat might be controlled from the residential node to move its set points up or down in response to a function of occupancy, price, daily average price, price standard deviation, the home’s temperature, predicted outside temperature, and the predicted home envelope simulated behavior.

9. MARKET RESOLUTION IN THE FORECAST FUTURE

Transactive control will be much more powerful when it includes forecasts of both the value and demand signals. The future intervals of such forecasts have not yet been determined. But the intervals should be course far in the future and should become shorter in the near term. The intervals should align well with the regional dispatch practices.

Several advantages follow from the inclusion of forecasts: First, responsive demand may then be considered at the time of and on a fair playing field with resource dispatch decisions. Customers may provide feedback concerning how their demand resources will respond and potentially avoid using expensive peaking resources that would otherwise become dispatched. Second, when forecasts are used, the importance and urgency of a formal market clearing process is reduced. Instead, the frequent iteration of future forecasts over time will achieve an equivalent resolution—the pairings of value (price) and demand during future intervals.

The inclusion of forecasting is not intended at all to preclude dynamic, real-time control opportunities. The present is simply a special case. The signals should remain dynamic and available for even unforeseen contingency responses in the present.

10. GENERALIZED TRANSACTIVE CONTROL IS AMENABLE TO STANDARDIZATION

This formulation of transactive control is amenable to standardization as a foundation for price-responsive control in a smart grid. The defined inputs, outputs, and responsibilities of a generalized node are scalable throughout a power grid, applicable at any node of the recommended hierarchy.

The inputs and outputs of the generalized node are defined in a way that reduces overall communication bandwidth and facilitates interoperability. The value signals are available to multiple entities that would choose to influence demand at a node and are a simple basis from which responsive demand assets can plan their energy consumption. The feedback of immediate and future demand is concise and benefits from aggregation into a single signal at each node. Because the described approach does not rely on communication of device-specific information, and because decision making is highly distributed, the proposed approach might be less vulnerable to some cyber security threats.

11. REMAINING CHALLENGES

The authors are hopeful that the newly generalized transactive control approach can be tested soon in a Pacific Northwest smart grid demonstration. Admittedly, more work is needed in the following areas:

- The time intervals and future horizon must be selected to accommodate and influence dispatch decisions. The preferred time interval should be short enough to enable innovative ancillary services while supported by existing communication technologies.
- The value and demand signals must be evaluated for ways in which they might augment or supplant existing customer incentive programs. The initial formulation has emphasized grid control while deferring details about how such behaviors can be persistently induced through incentives, regulations, and business cases.
- The methods for predicting the demand of responsive assets are lacking and should be improved. It is acceptable to begin with crude trending and improve the predictive demand models over time.
- New functions for how grids' operational objectives influence the distributed value signals should be formulated and tested. The authors have good confidence in the control of operational constraints, as was demonstrated in the Olympic Peninsula project. A simple formulation for encouraging wind consumption has been developed. But additional formulations will be needed for other important operational objectives, including carbon mitigation, encouraging environmentally preferable generators to become dispatched.
- The value signal itself will be amenable to alarm generation for system operators, but this feature has not yet been defined. We believe this feature will be a fundamental bridge between present operations center practices and inclusion of data from transactive control into operations center toolsets.

12. CONCLUSIONS

The authors have described a generalized formulation of transactive control that is amenable to standardization as a foundation for dynamic price control in a smart grid. The formulation is based on that used in a prior field demonstration, but the new, generalized formulation proposes use of a future time horizon, a hierarchical nodal structural framework, and a generalized functional model of an active node. A plan is offered to first launch a limited version of the control and later incrementally improve the extent and responsiveness of the system as it matures.

References

[1] D. J. Hammerstrom, *et al.*, "Pacific Northwest GridWise Testbed Demonstration Projects: Part 1. Olympic Peninsula Project," Pacific Northwest National Laboratory, Richland, WA, Tech. Rep. PNNL-17167, Oct. 2007. Available at http://gridwise.pnl.gov.

[2] J. R. Salmi Klotz, Bonneville Power Administration, by permission from an unpublished slide presentation, Aug. 2009.

Biographies

Donald Hammerstrom – Don Hammerstom is a Senior Research Engineer at the Pacific Northwest National Laboratory (PNNL) in Richland, Washington, where he has managed smart grid field studies. He is a principal designer of grid-responsive tools including the autonomous Grid Friendly™ appliance controller. He received his Ph.D. in Electrical Engineering from Montana State University in 1994.

Terry Oliver – Terry Oliver has worked globally to advance energy conservation and renewable energy. He has worked for the Bonneville Power Administration (BPA) since 1981 and was appointed in 2005 to become Bonneville Power Administration's first Chief Technology Innovation Officer, responsible for re-energizing, focusing, and managing BPA's research and development activities.

Ronald Melton - Ron Melton currently serves as the administrator of the GridWise® Architecture Council and as a senior technical leader for smart grid research and development projects at Pacific Northwest National Laboratory. He has over 30 years of experience applying computer technology to a variety of engineering and scientific problems. In addition to smart grid related projects recent experience includes research and engineering in cyber security for critical infrastructure protection and process control system security. Dr. Melton is a Senior Member of the Institute of Electrical and Electronics Engineers and a Senior Member of the Association for Computing Machinery.

Ron Ambrosio – Ron Ambrosio oversees IBM's Energy & Utilities Industry activities in its eight world-wide Research Laboratories. Ron joined IBM in 1981 at the T.J. Watson Research Center, working in a variety of areas including embedded operating systems, distributed application frameworks, and pervasive computing environments, ultimately focusing on networked embedded computing with particular emphasis on what he coined "Internet-scale Control Systems."