A Transactive Network Template for Decentralized Coordination of Electricity Provision and Value

Donald J. Hammerstrom, *Sr. Member, IEEE*, and Hung Ngo, *Member, IEEE*

***Abstract*—While many interesting transactive energy systems have been proposed, few fully decentralize price discovery and price-responsive control. Even fewer plan for extensibility that will be needed for the ultimate growth of transactive networks and for the inclusion of additional and new objectives and new flexible, responsive assets. This paper introduces a transactive network template from which an agent may be configured at its network node to negotiate for electricity with other neighboring network agents, manage a set of locally-owned supply and demand assets, and induce local power balance through price discovery. A set of base object classes defines the template and may be extended. Three important basic computational responsibilities are allocated among the base object classes—scheduling, balancing, and network coordination. These several basic classes and responsibilities may be used to represent the perspective of large and small devices and regions anywhere in the electric power grid because they are based on the self-similarity of these objects and responsibilities. Agents in a transactive network are not at all unique in these basic responsibilities. The template is designed to be further extended to address and monetize still other valuable objectives.**

***Index Terms*—Multi-agent system (MAS), Transactive energy**

I. Introduction

This paper introduces a transactive network template, from which implementers can hasten the implementation and uptake of transactive energy networks.

Transactive energy systems allocate electricity based on dynamically discovered values or prices. We defer to the GridWise Architecture Council for its authoritative definition of *transactive energy*, which is drafted to be inclusive of alternative evolving understandings. [1] Transactive energy systems have been extensively studied, but their uptake and use have been slow. Transactive energy systems might be classified among price-based demand response, but unlike time-of-use pricing, for example, transactive energy systems necessarily include feedback between price and quantity.

The transactive network template addresses electricity supply and demand quite symmetrically. Supply and demand resources are given equal access and opportunity to offer their flexibilities into the network. The chief difference between supply and demand is their sign. Supply is assigned positive value; demand, negative.

Today, wholesale locational marginal electricity prices are calculated centrally, but the transactive network template strives for fully distributed and decentralized calculations of effective locational marginal prices by each agent in a transactive energy network.

So, the facilitation of future transactive networks is a balancing act between allowing implementers leeway to realize their desired operational objectives, assets, and innovative asset behaviors while also enforcing enough structure and behavioral expectations to make sure that resource decisions are not corrupted and that the interactions between neighboring agents throughout the network are semantically similar.

Few prior examples of extensible transactive network templates were found. The thesis of Koen Kok, [2] which became the foundation for the PowerMatcher Suite, [3] is close in that it formulates a nested, domain-independent auction mechanism that is suitable for a deep, radial distribution network and accommodates heterogeneous device participation. It falls short, however, in that it hard-codes a conventional auction mechanism and single-interval look-ahead forecast horizon. It was not intended for extension of its market mechanism and does not invite incentives apart from those representing energy scarcity. Reports from the Pacific Northwest Smart Grid Demonstration also influenced the transactive network template formulation. [4] That work introduced concepts for "toolkit functions" by which libraries of device responses could be cataloged. [5] A large network of transactive agents was implemented, but the formulation also hard coded its price discovery approach.

The contents of this paper are based on the more detailed description of the transactive network template in report [6]. The transactive network template proof of concept has been demonstrated by its management of a small transactive network that is comprised of a distribution utility, an institutional campus, and multiple campus buildings that have novel, responsive building control assets. These results should be reported later in a follow-on companion paper.

This research makes the following contributions: A template metamodel is offered to accelerate the design and uptake of future transactive energy networks. The objects and

Submitted April 4, 2019 to the IEEE Transactive Energy Systems Conference. This work was jointly funded by the U.S. Department of Energy Office of Electricity and Energy Efficiency and Renewable Energy Building Technologies Office. Pacific Northwest National Laboratory (PNNL) is operated for the U.S. Department of Energy by Battelle Memorial Institute under Contract DE-AC05-76RL01830.

D. J. Hammerstrom (e-mail: donald.hammerstrom@pnnl.gov) and H. Ngo (hung.ngo@pnnl.gov) are with Battelle Memorial Institute at the Pacific Northwest National Laboratory, Richland, WA 99352 USA.

responsibilities of the various agents in a transactive energy network are argued to be remarkably self-similar, not unique. The existence of a template metamodel encourages future implementers of transactive energy networks to therefore collaborate and extend existing design content instead of designing every implementation and feature of their transactive networks from scratch.

II. Transactive Network Template Objects

The transactive network template is a metamodel intended to facilitate extensibility. As an object-oriented design, it can be understood from its base object classes and their behaviors.

The most important object classes of the transactive network template are the transactive agent, market, local asset model, and transactive neighbor model. These classes are introduced by Fig. 1.

A. *Transactive Agent Object*

A transactive network template is instantiated and configured once for each transactive agent object. The transactive agent is one of many transactive agents in a transactive network. It represents the unique perspective of a business entity that owns a specific circuit region, circuit element, or generating or consuming device. The object has few responsibilities other than to keep track of its unique sets of local assets, local market or markets, and the neighboring transactive agents with which it must negotiate and exchange signals.

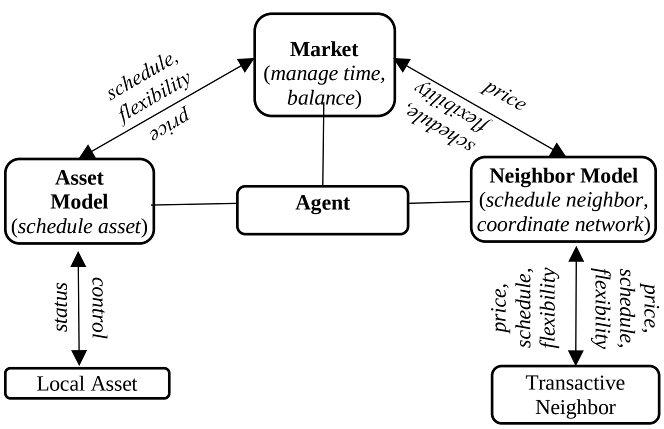

Fig. 1. Four of the most important object classes of the transactive network template are the transactive agent, market, neighbor model, and asset model.

B. *Market Object*

A market object manages the local balancing of supply and demand for its transactive agent. In a distributed system, this means that the sum of generated, imported, consumed, and exported electricity must be zero in every forecast time interval. A price-discovery mechanism determines the price at which balance will be induced. There may be multiple market objects for each transactive agent if there are sequential correction markets—*e.g.,* a real-time hourly correction to a day-ahead market clearing—or if the transactive agent participates in markets for commodities other than electricity.

Market time intervals that must be aligned with local scheduling and coordination processes, so the market is responsible to spawn the forward time interval objects that it will need. The time interval objects have a lifetime within the given market.

C. *Local Asset Model Object*

A local asset model represents its local supply, demand, or energy storage asset for its transactive agent. It is responsible to schedule the local asset's power generation or consumption in every forward time interval. Local assets are "owned" by their transactive agent. There exist no secrets between a local asset and its transactive agent. Therefore, the scheduling of a local asset may be made quite complex. The local asset model receives status information from the physical asset and ultimately controls it to take its scheduled actions.

Unique engineering is almost always required to integrate physical assets. The transactive network template should remain indifferent concerning choices of low-level control and communication protocols.

D. *Transactive Neighbor Model Object*

Transactive neighbor models must be instantiated to represent each of the local agent's transactive neighbor agents. Two network agents are *neighbors* if they are committed to transact with one another. However, market solutions may only be approximated if the circuit transport elements between neighbors are not explicitly modeled. Unstructured pooled markets and some aggregator models violate this requirement, require centralized assessment of transport feasibility, and are therefore not amenable to fully decentralized calculations.

A transactive neighbor model manages signals to and from the neighboring transactive agent. This transaction is kept simple at the interface between coordinating agents. The representations of price, schedule, and flexibility should be standardized at this interface, and exceptions should be discouraged. Communication processes at this interface must remain somewhat flexible, however, because the neighbor agent is fully independent from the local agent. In fact, the neighbor might have used different software languages, different computational platforms, and different standards and protocols. The local agent may choose when to send updated information, but it cannot force the neighbor agent to reply. Flexible, event-based communications are therefore preferred at this interface.

The transactive neighbor model is responsible to schedule the power to be imported or exported from the neighbor circuit location. In this way, it represents the neighbor during the local balancing process. The products of this responsibility are the same as for the scheduling of local assets. However, unlike the local asset model, the transactive neighbor model represents remote agent entities whose motives and methods are unknown to the local agent. Therefore, while scheduling, the transactive neighbor model simply represents what it has learned through coordination signals with the neighbor agent and nothing more. The neighbor model is a useful construct because it allows local scheduling and balancing computations to proceed without expecting or requiring remote communications to neighbor agents for every iteration.

III. Representing Flexibility using the Transactive Network Template

As a prelude to addressing the behavioral responsibilities of the transactive network template, we introduce some insights that prove advantageous toward the representation of flexibility in a transactive system.

A. Time Interval Object

A time interval object is created by a market object and then exists to serve the market during its lifetime. Once created, it remains affixed to its starting and ending dates and times. The time interval object is foundational to careful management of a transactive energy network that includes multiple time intervals, as is the case when a market defines a forward planning horizon of future time intervals. A time interval is relevant from the time it is created until its market delivery period expires and all its associated transactions have been resolved. Anticipating rich, sequential transactive markets, the time interval also keeps track of the market object that created it and its state within that market, which enumeration is extended from a smaller set of market states defined in the OASIS EMIX specification. [7]

B. Interval Value Object

The interval value class helps the transactive agent manage a large, extensible set of measurement types and values while keeping them clearly associated with a time interval object. Every piece of information, whether predicted, measured, quantifiable, or qualitative, is instantiated by the transactive network template as an interval value.

An interval value object binds a value with a specific time interval object. These bound items are frequently needed for predicted information. Values must not move in time. There should never be any error-prone assumption that a series of values are correlated to a series of time interval objects. An interval value object must store precisely one piece of data; however, nothing prevents the data value from itself being a complex struct, as is the case for active vertices (to be introduced in the next section). Metaproperties of an interval value keep track of provenance of the data value and the nature of the stored value.

C. Active Vertices

If production cost functions and utility functions are defined piecewise quadratic, then supply and demand curves are piecewise linear. An *active vertex* represents an inflection point in such a piecewise linear supply or demand curve. The transactive network template requires that vertices be used to represent both scheduled power points and flexibility. Both scheduled powers and flexibility are represented and stored by the transactive network template as vertices.

The principles of this practice are demonstrated in Fig. 2. A single lone vertex (Fig. 2a) represents a constant, inelastic generation or consumption of electricity. The marginal price of any lone vertex is treated as infinity because an inflexible device will not change its production or consumption at any price. Two or more vertices may be used to represent a monotonically increasing, signed power with increasing marginal price. Vertices should span the range of flexibility that local asset or neighbor can offer. The local asset or neighbor is presumed to be indifferent concerning its operations within this span in a given time interval. Arrow tails are always assumed from the highest priced vertex to infinite price and from the lowest priced vertex to negative infinite price. The implied vertices at plus and minus infinity need never be explicitly stated.

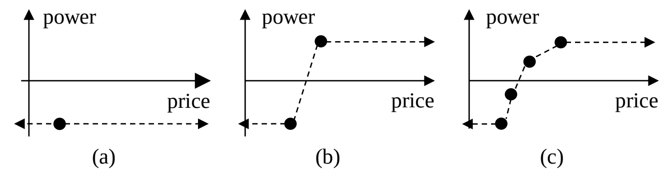

Fig. 2. Flexibility representations using active vertices. (a) An inflexible electric load. (b) A flexible energy storage represented simply by two active vertices. (c) A more complex, flexible energy storage that requires multiple active vertices to represent its flexibility.

The summing of supply and demand thus reduces to the summing of average interval powers at all the defined vertex price locations. If all local assets and neighbors are required to state their flexibilities in this manner by using vertices, then the clearing balancing price may be found by linear interpolation of the net supply and demand curve at the point where net power equals zero, often greatly reducing the numbers of iterations needed for price discovery.

Eq. (1) helps clarify the relationship between scheduled power values and a larger power flexibility range. Scheduled power values should always lie within the range of powers that could be induced by prices. Consequently, the vertex that represents a scheduled power also lies between vertices that represent price flexibility. The vertices in the price flexibility range are described as *active* vertices. The active vertices may change over time and according to local conditions.

There is still a wider range of powers that are defined by physical capacity limits. These are hard, physical limits and therefore are not associated with any price or vertex, but they place a hard constraint on the powers that may be represented by active vertices.

$$\overbrace{\underline{p}_{\text{capacity}} \leq \underbrace{\underline{p}_{\text{flexibility}} \leq p_{\text{scheduled}} \leq \overline{p}_{\text{flexibility}}}_{\text{price flexibility range}} \leq \overline{p}_{\text{capacity}}}^{\text{power capacity range}} \quad (1)$$

D. Transactive Records and Transactive Signals

Transactive signals are sent from the local agent to its neighboring transactive agents as part of the coordination process, much as described in [8]. A transactive signal is comprised of a plurality of transactive records. The transactive record binds a time interval object with real electrical power, electricity pricing, and potentially other qualities. Each transactive record represents either a scheduled price-power pair or a point on the respective agent's flexibility curve.

It may seem unnecessary to distinguish a transactive record from a vertex and interval value, but vertices and interval values are classes that exist only within each transactive agent's chosen computational platform and reference code implementation. A transactive neighbor is hosted on a remote computational platform; uses a separate, independent instantiation of the transactive network template; and cannot be presumed to even have used the same computer language. A transactive record (and composite transactive signal) must therefore be a flat record, like a text file, that can be read correctly by any computer program.

The independence of the transactive agents pose additional challenges for the coordination process. Because a transactive agent cannot insist or coerce a neighbor agent to promptly respond with its own transactive signal, event-based communications are preferred. Each agent must be given means to correctly infer whether neighboring agents have converged to an agreeable power and price schedule. Consensus is asserted in an active time interval if 1) the recently calculated transactive records are acceptably close to those last sent, and 2) the sent transactive records are acceptably close to those that have been received. A convergence flag is set if convergence has been confirmed in all active time intervals. If consensus cannot be asserted, or if local conditions change, then updated transactive records must be sent to the neighboring transactive agent again. A minimum wait time is recommended between the sending transactive signals to avoid having the computations race.

IV. Transactive Network Template Behaviors

A key to facilitating extensibility within the transactive network template is to offer and enforce a small set of required computational responsibilities (see Fig. 3). These computations must produce semantically similar outputs while providing a future implementer great flexibility concerning which local objectives are to be included during the calculations. The three most important computational responsibilities required by the transactive network template are 1) the balancing of supply and demand, 2) the scheduling of power in forward market time intervals, and 3) coordination of the price and quantity of electricity to be exchanged between neighboring network agents.

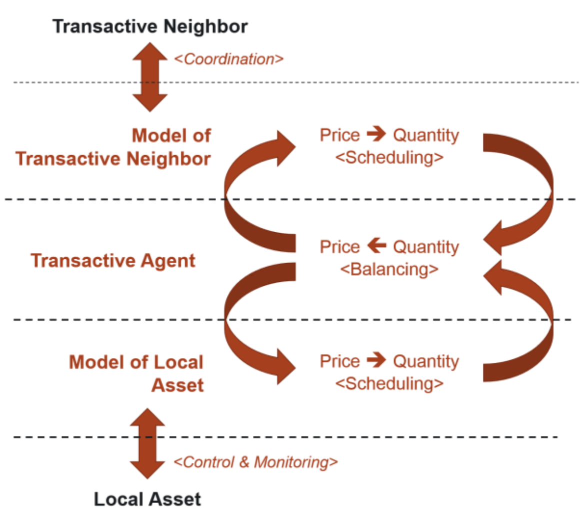

Fig. 3. Scope of the transactive network template and its behavioral responsibilities

The formulation of the scheduling and balancing computations was inspired by [9], which clearly explains the separability of these two calculations. Furthermore, it suggests how the formulations may be augmented to include unit resource commitment, fixed variable costs, and resource reserves might be integrated, all of which are of interest for future transactive network template versions. While [9] was intended to solve conventional wholesale generation supply commitment and dispatch, the transactive network template extends its applicability to include flexible, dynamic demand.

A. *Scheduling Responsibility*

We begin with the scheduling responsibility because this calculation is quite independently performed for each local asset and transactive neighbor model. The basic premise of the scheduling responsibility is: given forward electricity prices, how much electricity would a local asset or neighbor consume or generate in each forward time interval? Additionally, what flexibility does the local asset or neighbor offer to change its generation or consumption if given alternative prices?

All local assets and neighbors must schedule themselves, regardless whether they can offer any flexibility. Local asset models are given considerable leeway concerning the rigor with which they tackle this challenge. The transactive network template is indifferent in this matter. Some implementers will chose purely heuristic methods, as was done in the Olympic Peninsula Project [10] and PowerMatcher [3]. Others may use rigorous optimization methods. Transactive neighbor models, however, simply represent electricity availability, need, and flexibility that has been reported by the respective neighboring agent via transactive signals. The multiple scheduling sub-problems converge upon each transactive neighbor model or local asset model independently calculating a satisfactory power schedule and flexibility for itself.

A basic form of the scheduling problem is (2). Each local asset or neighbor model must find average powers p_τ that maximize sum utility U_τ and energy revenue product $\lambda_\tau \cdot p_\tau$, less production costs C_τ, over a set of forward market time intervals τ and subject to constraints. Note that the optimization is defined over an entire set of forward time intervals, which is particularly important if there exist intertemporal effects and constraints, as is often the case.

$$\max_{p_\tau} \sum_\tau (U_\tau(p_\tau) + \lambda_\tau \cdot p_\tau - C_\tau(p_\tau) + \cdots) \tag{2}$$

The basic form (2) is equally applicable to either supply or demand assets. In practice, demand assets typically use only the revenue and utility terms, where the utility term monetizes the value of things like personal comfort and other preferences. Heuristic decision curves are often formulated for demand-side assets directly in the marginal price plane, the derivative form of (2).

Supply assets normally use only the revenue and cost function terms of (2). Requiring quadratic production cost function of the form (3) enables the use of vertices and piecewise linear marginal price curves. The supply curves of conventional fueled supply resources are usually defined by production costs, and representations like those in Fig. 2 must then be found by taking a mathematical derivative of the constrained production cost function (3).

$$C_\tau(p_\tau) \triangleq a + bp_\tau + \frac{1}{2}cp_\tau^2 \tag{3}$$

B. *Balancing Responsibility*

Every transactive agent must balance local electricity supply and demand for its circuit region in one or more forward market intervals. *Local electricity demand* includes summed electric demand for local assets and also electricity that must be exported to neighbors. *Local electricity supply* includes summed generation by local assets and also electricity that is to be imported from neighbors. The balancing responsibility is the process of price discovery. An effective marginal price is calculated by the transactive agent for each market time interval. The clearing price induces balance of demand and supply, at which point the balancing sub-problem has converged.

The balancing responsibility is treated as a *local* sub-problem by the transactive network template. Many alternative price discovery formulations are feasible. The problem formulation (4) is used in the current version. It represents a minimization of the sum net revenue, production costs, and utility (disutility) function values among all local assets and neighbors that are indexed by i in a single forward time interval.

$$\min_{\lambda} \sum_{i} \left(\lambda \cdot p_i(\lambda) + C_i\big(p_i(\lambda)\big) - U_i\big(p_i(\lambda)\big) + \cdots\right) \qquad (4)$$

This formulation requires minimal iteration if the terms of (4) are linear, or nearly linear, functions of the effective locational marginal price λ. The price may then be found efficiently by interpolation on the net supply and demand curve, which includes the effects of all the terms of (4). More precisely, the balance price occurs where the net supply and demand curve intersects the line $\sum p_i = 0.$ If the terms of (4) are nonlinear in λ, or if any local asset or neighbor model is unable or unwilling to provide the terms as functions of λ (*i.e.*, asset and neighbor price flexibilities are unknown), then a more general and iterative solution method like subgradient search must be used instead to discover price λ.

If the balancing responsibility were fully independent, then a single, common formulation might be applied by all transactive agents. Unfortunately, the inputs and outputs of the balancing sub-problem must align with those of the coordination sub-problem. The balancing process might need to be tweaked depending upon whether neighboring transactive agents will share their flexibilities and whether they elect to plan multiple forward market intervals, for example.

C. Coordination Responsibility

The coordination sub-problem does not require any separate computations. Coordination is the process by which neighboring transactive agents procure the information that they need from the network in order to calibrate their local balancing calculations and confirm price and quantity expectations between agents throughout the network. Each agent of the network must negotiate an exchange of electricity with at least one other network agent. Other network agents' electricity prices are also dynamic functions of their own local electricity supply and demand.

The transactive network template facilitates the exchange of transactive signals between neighboring transactive agents. Preferably, these signals indicate the net scheduled electricity to be exchanged, the price of the electricity, and the agent's flexibility to alter its electricity supply or consumption if the price is changed from the currently scheduled electricity quantity and price. The coordination sub-problem has converged if transactive neighbors agree (within an error threshold) concerning the scheduled electricity and its price.

The coordination responsibility is assigned to each transactive neighbor model. The current version of the transactive network template defines transactive records that are similar to the pairing of a vertex (Section III.C) and interval value (Section III.B). The transactive records are therefore conveniently aligned with the representations of schedules and flexibility within the transactive network template.

Another useful feature of the coordination calculation in the current transactive network template is the exchange of *residual* flexibility. The scheduled electricity to be imported from or exported to the neighboring transactive node is excluded from the net supply and demand curve (as is used in the balancing process of Section IV.B) in the signal sent to that neighbor. The agent's flexibility, too, is calculated in respect to the exchange opportunities between only the two negotiating agents. An important benefit of this approach is that neighboring transactive agents can update and exchange their transactive signals without having to wait for many other agents' bids and offers to be aggregated.

REFERENCES

[1] GridWise Architecture Council, "GridWise Transactive Energy Framework, Version 1.0," Pacific Northwest National Laboratory, Richland, WA, 2015.

[2] J. K. Kok, The PowerMatcher: smart coordination for the smart electricity grid, Amsterdam: Vrije Universiteit Amsterdam, 2013.

[3] Flexiblepower Alliance Network, "The Power Matcher Suite: Transactive Smart Energy," 2 April 2019. [Online]. Available: http:/flexiblepower.github.io/. [Accessed 2 April 2019].

[4] D. J. Hammerstrom, D. Johnson, C. Kirkeby, Y. P. Agalgaonkar, S. T. Elbert, O. A. Kuchar, M. C. Marinovici, R. B. Melton, K. Subbarao, Z. T. Taylor, B. Scherer, S. Rowbotham, T. Kain, T. Rayome-Kelly, R. Schneider, R. F. Ambrosio, J. Hosking, S. Ghosh, M. Yao, R. Knori, J. Warren, J. Pusich-Lester, K. Whitener, L. Beckett, C. Mills, R. Bass, M. Osborne and W. Lei, "Pacific Northwest Smart Grid Demonstration Project Technology Performance Report Volume 1: Technology Performance," Battelle-Pacific Northwest Division: Pacific Northwest Smart Grid Demonstration Project, Richland, WA, 2015.

[5] D. J. Hammerstrom and T. Esram, "Transactive Node Framework and Toolkit Library Functions, Version 1.0," Battelle-Pacific Northwest Divition: Pacific Northwest Smart Grid Demonstration Project, Richland, 2015.

[6] D. J. Hammerstrom, "The Transactive Network Template Metamodel," Pacific Northwest National Laboratory, Richland, WA, 2019.

[7] Oasis, "Energy Market Information Exchange (EMIX) Version 1.0," OASIS, 11 January 2012. [Online]. Available: http://docs.oasis-open.org/emix/emix/v1.0/cs02/emix-v1.0-cs02.html#_Toc319594512. [Accessed 2 April 2019].

[8] D. J. Hammerstrom, "Transctive Coordination Signals," Battelle Memorial Institute: Pacific Northwest Smart Grid Demonstration, Richland, WA, 2013.

[9] C. Murillo-Sánchez and R. J. Thomas, "Thermal Unit Commitment Including Optimal AC Power Flow Constraints," in *Hawaii International Conference on System Sciences*, Kona, Hawaii, January 6-9, 1997.

[10] D. J. Hammerstrom, R. Ambrosio, T. A. Carlon, J. G. DeSteese, G. R. Horst, R. Kajfasz, L. L. Kiesling, R. Michie, R. G. Pratt,, J. Brous, D. P. Chassin, R. T. Guttromson, O. M. Järvegren, S. Katipamula, N. T. Le, T. V. Oliver and S. E. Thompson, "Pacific Northwest GridWise™ Testbed Demonstration Projects: Part I. Olympic Peninsula Project," Pacific Northwest National Laboratory, Richland, WA, 2007.