Spring 2019

The tabu ant colony optimizer and its application in an energy market

David Donald Haynes

Department: Engineering Management and Systems Engineering

THE TABU ANT COLONY OPTIMIZER

AND ITS APPLICATION IN AN ENERGY MARKET

by

DAVID DONALD HAYNES

A DISSERTATION

Presented to the Faculty of the Graduate School of the

MISSOURI UNIVERSITY OF SCIENCE AND TECHNOLOGY

In Partial Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

In

SYSTEMS ENGINEERING

2019

Approved by:

Steven Corns (Advisor)
Mariesa Crow
Cihan Dagli
Ruwen Qin
Donald Wunsch

ABSTRACT

A new ant colony optimizer, the “tabu ant colony optimizer” (TabuACO) is introduced, tested, and applied to a contemporary problem. The TabuACO uses both attractive and repulsive pheromones to speed convergence to a solution. The dual pheromone TabuACO is benchmarked against several other solvers using the traveling salesman problem (TSP), the quadratic assignment problem (QAP), and the Steiner tree problem. In tree-shaped puzzles, the dual pheromone TabuACO was able to demonstrate a significant improvement in performance over a conventional ACO. As the amount of connectedness in the network increased, the dual pheromone TabuACO offered less improvement in performance over the conventional ACO until it was applied to fully-interconnected mesh-shaped puzzles, where it offered no improvement.

The TabuACO is then applied to implement a transactive energy market and tested with published circuit models from IEEE and EPRI. In the IEEE feeder model, the application was able to limit the sale of power through an overloaded transformer and compensate by bringing downstream power online to relieve it. In the EPRI feeder model, rapid voltage changes due to clouds passing over PV arrays caused the PV contribution to outstrip the ability of the substation to compensate. The TabuACO application was able to find a manageable limit to the photovoltaic energy that could be contributed on a cloudy day.

ACKNOWLEDGMENTS

I would like to take this moment to recognize the work of all engineers that make this world a better place.

I would like to thank:

- Dr. Corns, for advising me in the development of this dissertation, and for his many hours of discussion and guidance.
- The faculty committee, for working with me as I developed this dissertation topic and worked to develop it into a finished product.
- Drs. Chandrashekhara and Guardiola, for serving on the faculty committee and for their timely words of encouragement.
- Dr. Daniel Ciardullo, for serving as an off-campus committee member.
- My family (especially my wife), for their support as I spent a decade pursuing advanced degrees at MS&T.
- Aclara, for helping to fund the effort.
- Friends and co-workers, for their encouragement.

TABLE OF CONTENTS

Page

LIST OF ILLUSTRATIONS

LIST OF TABLES

Page

NOMENCLATURE

Symbol	Description
ρ	pheromone deposition or evaporation factor
μ	mean
σ	standard deviation
τ	pheromone concentration
ACO	“ant colony optimization” or “ant colony optimizer”
AMI	advanced metering infrastructure
CDF	cumulative distribution function
DER	distributed energy resource
DG	distributed generation
G&T	generation and transmission
ISO	independent system operator
PDF	probability density function
PV	photovoltaic
RTO	regional transmission organization
QAP	quadradic assignment problem
TE	transactive energy

1. INTRODUCTION

1.1. ORGANIZATION OF THIS DISSERTATION

This dissertation will describe an area of research in computational intelligence in the form of advancements in ant colony optimization. The ant colony optimizer (ACO) is identified in Section 1.2. The subject of this research is called the “tabu ant colony optimizer” (TabuACO). The TabuACO is identified in Section 1.3. A literature review in Section 2 compares the TabuACO to other research efforts. Section 3 describes the several different versions of the TabuACO that have been tested against various solvers using several different benchmark problems. After demonstrating the performance and scalability of the optimizer, this dissertation proceeds to apply the algorithm to a practical problem in Section 4. The optimizer is applied to the energy market. An introduction to the smart grid is provided, along with a background of the operation of the energy market. Transactive energy is explained in Section 4.4. The TabuACO is applied to the transactive energy market in Section 4.7-4.8. Section 5 offers conclusions. Section 6 offers ideas for future work.

1.2. BACKGROUND

Ants solve problems in nature, and their technique serves as a computational paradigm.

1.2.1. Ants in Nature. Ants communicate in nature by laying down pheromones along the trails they travel [1]. The pheromone is an odorless chemical that serves to attract other ants and evaporates over time.

It is believed that at least one species of ants is able to lay a repulsive pheromone to discourage exploration of known unproductive trails [2].

Ants have a scent which is inherited from the queen ant. If an ant encounters a pheromone trail from another nest, it will usually not mistake it for its own. This allows an ant to ignore seemingly “irrelevant” trails if it so chooses. Ants compete with other species for the available food sources, and must contend with other factors such as predators, unfavorable weather conditions, and efficient use of resources over time. Ants do not have a centralized “leader” but instead work instinctively as agents on behalf of the interests of the collective whole. It could be argued however that different ant species employ different methods to reach their objectives.

While foraging for food ants also detect pheromones in the environment. These pheromones are placed by other ants (of the same species) to communicate a particular meaning. Pheromones evaporate over time, butwhile they are still present, they can serve to mark the trail to an important food source. The pheromone, even in small quantities, can influence foraging ants to walk to the find. As they return to the nest with their payload, the successful ants will lay down pheromone to mark the trail. The new pheromone deposit will supplement the evaporating pheromone. Thus, between the repeated deposits and ongoing evaporation, trails are created. These trails enable ants to walk in groups along lines between the nest and the food source.

Ants that find alternative routes also lay pheromone, but the faster route allows more trips and more ants to lay more layers of pheromone than the longer route. This buildup of pheromone allows the optimal solution to emerge. This solution is stored in, and described by, the environment.

1.2.1.1. Ants have various diets and unique pheromones. Ants lack the equipment necessary to chew and digest food. Instead, they squeeze the juices from various foodstuffs. Different species of ants favor different sources of food. Some species prefer sweets, others prefer protein, and others will grow their own fungus. The harvester ants (for example Pogonomyrmex rastratus and P. pronotalis) favor seeds in their diet. Research shows that more than 87% of the items carried to the nests are seeds, and 93% of the seeds are grass seeds [3]. Argentine ants on the other hand prefer to feed on other insects (when available), and then when the supply of insects is depleted, changeover to feed on sweet plant sap [4]. The leaf cutting ant (atta cephalotes) cuts leaves to grow fungus for larvae while adults imbibe liquid directly from the crushed leaf tissue [5].

1.2.1.2. Some species collaborate between nests. A pheromone trail left by a foraging ant from one nest could be worthless to ants from another nest. Yet examples have been found of collaboration among nests on those occasions in which a super colony of the same species is formed over a vast area [4] [6].

1.2.1.3. At least one species of ant can deposit repulsive pheromones. It would be “expensive” for an ant to mark all of the places that are determined to be unproductive while foraging. However, it is not necessary to mark every unproductive location, nor is it necessary to mark the length of every trail to empty locations. It would be sufficient to merely mark the start of unproductive trails. This is precisely where the Pharaoh ant (*monomorium pharaonis*) lays a repulsive pheromone to mark unrewarding trails [2].

1.2.2. The Ant Colony Optimizer. The ant colony optimizer (ACO) is a programming technique which solves problems in a manner inspired by nature.

The ACO is recognized as one of the five nature-inspired computational intelligence (CI) paradigms:

- Artificial Neural Networks
- Fuzzy Systems
- Evolutionary Computation
- Genetic Algorithms
- Swarm Intelligence

Ant colony optimization was developed by Dorigo [7]. Dorigo's original work explored the use of an ACO to develop a closed Hamiltonian graph to solve a traveling salesman problem. ACOs which develop Hamiltonian graphs generally begin by sprinkling ants throughout the search space. The ants lay pheromone in relation to the quality of a find. With successive deposits and evaporations, an optimal solution emerges in the pheromone pattern.

ACOs can also be applied to develop a minimum spanning tree within graphs or trees. ACOs that develop spanning trees generally allow ants to originate from a nest, forage for a "prize," and lay attractive pheromone back to the nest in proportion to the quality of the find. Ants that forage in this way must generally remember the path home. The formation of loops is a possibility when random foraging is allowed. The algorithm must (typically) trim any loops after foraging has detected a food source and prior to pheromone application.

In many cases, the ACO is a particularly good fit in situations where the problem maps nicely to a graph (with nodes and edges). When a network exists that must be optimized according to some cost function the ACO is usually appropriate. The ACO algorithm can be very simple yet effective [8].

1.3. THE TABU ANT COLONY OPTIMIZER

A generalized description of the TabuACO is provided in this section. Customizations of the algorithm follow in subsequent sections, with a particular application of the algorithm at the end.

1.3.1. Dual Pheromones. The tabu ant colony optimizer (TabuACO) is different than a conventional ACO in that it employs two distinct pheromones to speed convergence to a solution. Each pheromone serves a distinct purpose. Much like a conventional ACO, an attractive pheromone is laid in response to a positive find. Unlike a conventional ACO, a repulsive pheromone is laid in response to a negative find. In this sense, the algorithm combines the benefits of a tabu-search algorithm with the benefits of an ACO. The precise rules for deposition, evaporation, and path selection vary by application and are described in Section 3.

1.3.2. Multinest. Some problems lend themselves to a situation in which multiple nests can be activated to simultaneously compete and/or collaborate in the solution of the problem. When applied this way, each ant must know its nest of origin, the travel rules and constraints which govern the implementation of the problem (such as what constitutes a “find,” any cost to travel, or authorization to spend funds), and how to compute the score (the objective function). For a multinest ACO to work, the software must issue at least one ant from each of the nests and allow the ants to collaborate.

Quite often the purpose of multinest activity is to leverage a parallel processing effort to speedily solve a single problem [9]. However, it is possible that each nest could work to satisfy its own objectives. It is possible for each nest to have its own success criteria. The application of the TabuACO for the energy market was coded to support multinest operation. However, a single processor[1] was used as the computing resource in the tests which follow.

1.3.3. Graph Navigation. The TabuACO algorithm requires a graph with reference directions in order for ants to navigate and properly interpret the sign of the pheromone. Figure 1.1 provides an example. The black arrows represent edges with a reference direction. They should not be interpreted as unidirectional edges. The blue arrows represent attractive pheromone. The red arrows represent repulsive pheromone.

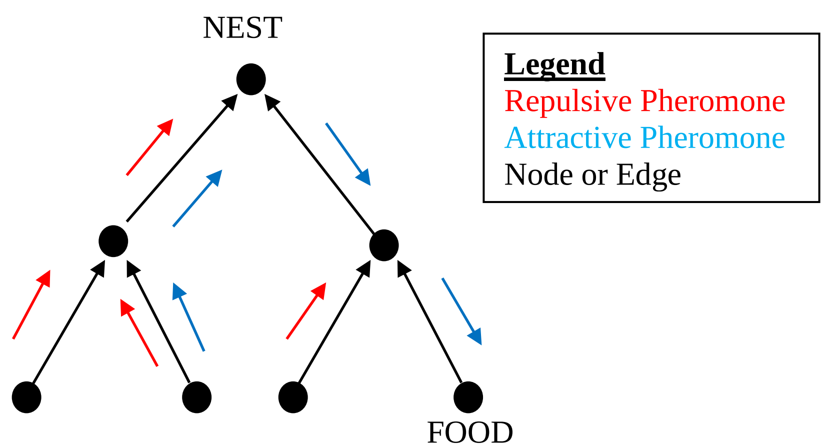

Figure 1.1. Bidirectional Graph with Overlaid Directions, Fully Painted with Pheromones

The proposed ant-inspired method uses a directed graph in order to navigate. This allows a directional pheromone to be laid down by ants. To achieve directionality, and

[1] The processor was usually a multi-core processor, but this property was not exploited. Section 6.1.3 describes the use of parallel processing as a future work.

allow ants to roam freely, a pheromone value between [-1, +1) is used instead of the conventional [0, 1) range. A positive value influences the ant to go in the direction of the arrowhead.[2] A negative value influences the ant to go in the direction of the tail.[1] A pheromone value of zero implies no influence to the stochastic path selection process.[1]

Some optimization problems impose additional constraints beyond a simple travel cost or prize. This additional data can be modeled in the graph and observed by the foraging ant. Nodes and edges can be modeled to have a location, and thus edges can have lengths.

Nodes and edges can have a transit cost represented in the model under multiple costing schemes. There can be a financial cost as well as an efficiency or loss of service cost. A network and its model can have capacity constraints. If a node or edge cannot carry the full payload the ant is authorized to bring back to the nest, then the ant must abide by this restriction and lessen her ambitions to be no more than what the "weakest link" can support along the path she chooses. If the ant should wander back across this weak link in the path, then this constraint imposed by the solution can be lifted (by popping it from a stack) and the new minimum link capacity serve as the payload constraint.

To keep the solution viable from a traffic perspective, additional capacity constraints can be tracked and enforced by the ants as they forage. When an ant traverses an edge, the records pertaining to that assed must first be checked to see if there is sufficient uncommitted capacity to carry the payload the ant has in mind. If the network has made prior commitments to other ants, the ant may have to reduce its expectations

[2] This is the case for both attractive and repulsive pheromones.

regarding the payload that can be carried over a given link. (The network itself may have capacity limits at certain nodes or edges.) The ant will eventually return in the opposite direction across the link. If the ant returns across the edge in the reverse direction empty handed, the temporary reservation originally placed in the prior crossing is discarded. When the ant returns across the edge with food, the temporary reservation is adjusted to the actual payload size the ant is carrying.

Multiple evaluations may occur over a series of time intervals. Each time period may present slightly different circumstances, but within the same network. Pheromone information, developed over the course of one interval, serves as the initial conditions for the successive interval. Ants must deliver food at a certain rate in order to satisfy the requirements of the nest. If food is removed from a source, it is gone for the duration of the evaluation interval, but more food may appear at the same location in a future interval.

To keep the solution viable from the ant's perspective, constraints are imposed by the ant as she travels. Ants start from their nest and memorize the path home by utilizing a stacking mechanism. They are given the authority to acquire a certain amount of food, and expend a limited amount of resources in doing so. The ant may have to settle for a smaller payload than expected for a variety of reasons:

- The food source might have less food than the ant's maximum (or authorized) payload.
- The network may impose constraints on the size of the payload.
- A cost may be associated with the ant's travel and thus reduce the ant's purchasing capacity upon arrival at the food source.

Travel costs might correspond to the fuel expended by the ant in delivering the payload, or to the cost incurred by the ant for using the network. A nonzero travel cost encourages ants to forage near home and not wander infinitely far away.

These constraints correspond to real-world constraints in the ecosystem. The ants themselves may have payload capacity constraints. Some species are larger than others and can simply carry heavier payloads. An ant may have stamina limitations in her ability to reach a distant food source and carry food back to the nest.

Section 4.8 explores the idea of the ant identifying the constraints within the model and pursuing a solution that abides by those constraints. A common model of the electrical grid is studied and the capacity constraints of the distribution asset observed. Yet, despite this capability, it may not be enough. Many complex phenomena can occur which might outstrip a simple model's ability to account for all of the possible capacity constraints. For example, on the electrical grid a shortage of reactive power can often be remedied by closing a switch to engage a capacitor bank. Tap changers can make voltage corrections in response to conditions on the grid. Distributed generation will often cause voltage to rise locally. To model all of these phenomena in addition to the basic power flow, the complexity of the software must grow considerably. At some point it implies that the software architecture should change so that engineering analysis is performed separately from the routing optimization. This is precisely what occurs in section 4.9. A more advanced analysis of a much larger network occurs. An external engineering analysis package is utilized to validate proposed routing solutions.

2. LITERATURE REVIEW

A review of the state of the art is an important aspect of any research. A proper literature review is necessary to ensure that the research does indeed further man's knowledge.

2.1. ACO LITERATURE REVIEW

There are have been numerous developments in the field of Ant Colonly Optimization since its inception.

2.1.1. ACO Origins. Ant colony optimization was first suggested as a metaheuristic in a journal article by Dorigo in 1991, and later in his PhD thesis in 1992 [7]. Dorigo studied the work of entomologists Pierre-Paul Grassé, Goss, and others [1]. Goss et. al. developed an equation (1) based on observations of the Argentine ant for the probability that an ant would choose a path given that she had two choices, and that pheromones of some concentration were present on both paths.

$$p_1 = \frac{(m_1+k)^h}{(m_1+k)^h+(m_2+k)^h}. \tag{1}$$

Where m_1 and m_2 are moments in time since the ant has visited path *1* or path *2* respectively, and h and k are constants with $h \approx 2$ and $k \approx 20$.

Dorgio's first ACO algorithm used a path selection formula that computed the probability of selection based on pheromones that are elevated in an exponential manner as described in (2).

$$\tau_{ij} \leftarrow (1-\rho) \cdot \tau_{ij} + \sum_{k=1}^{m} \Delta\tau_{ij}^{k} \tag{2}$$

Where τ_{ij} is the pheromone along an edge joining cities i and j, ρ is the evaporation rate, m is the number of ants, and $\Delta\tau_{ij}^{k}$ is the quantity of pheromone laid on edge ij by ant k.

Since Dorgio's original work, numerous improvements have been made [9] – many of them by Dorgio. He developed the *Simple Ant Colony Optimizer*, and the *Ant Method*. Other improvements by other researchers have continued to this day.

2.1.2. Repulsive Pheromone Research (a.k.a. "Negative," "Anti-Pheromone," etc.). In 1996, Schoonderwoerd et. al. were pioneers in suggesting that an anti-pheromone could be useful in speeding convergence [10].

In 1999, Roux, Fonlupt, Robilliard, and Talbi reported on a methodology they called the *ANTabu* [11]. Their approach combines a conventional ACO with a conventional Tabu Search. These algorithms are applied repeatedly in sequence to develop an optimal solution. The ACO is used to develop an initial schedule which is then optimized by the Fast Tabu.

In 2002, Montgomery and Randall wrote about three different ACOs that employ an "anti-pheromone" to speed convergence to a solution [12]. Montgomery and Randall described their solvers as variants of an Ant Colony System (ACS).

- Their *Subtractive Anti-Pheromone* algorithm is like the ACS and uses a single pheromone, but it has an additional step during the pheromone evaporation process to favor the evaporation of paths that are shown to be poor performers.
- Their *Preferential Anti-Pheromone* algorithm is like the ACS, but uses two pheromones—one for good solutions and another for bad solutions. Global evaporation is applied at the end of an iteration to all pathways.

- Their *Explorer Ants* algorithm uses a single pheromone like the ACS, but some of the ants are programmed to find their own pheromone repulsive. Thus, these ants explore unexplored territory when given the opportunity.

Montgomery and Randall made these comparisons while studying the Traveling Salesman Problem. They found their *Anti-pheromone* solvers did not offer a statistically significant improvement over the ACS benchmark.

In 2013, Ramos, Rodrigues, and Louca reported a *Second Order Swarm Intelligence* method which was essentially an ACO that used both attractive and repulsive pheromones to solve the TSP [13] [14]. They found that "by using two different sets of pheromones, a second-order co-evolved compromise between positive and negative feedbacks achieves better results than single positive feedback systems." They showed that the algorithm compared favorably against the benchmarks. It is also noteworthy that their use of "no entry signals and negative feedback" allowed "a colony to quickly reallocate the majority of its foragers to superior food patches." They also note that "this (was) the first time an extended ACS algorithm (was) implemented with these successful characteristics."

In 2014, Ezzat and Abdelbar described a *Less Exploitative Variation of the Enhanced Ant Colony System* [15]. It is similar to the TabuACO in the sense that paths that have already been analyzed by ants can be flagged to prevent further (unnecessary) analysis. Their method utilizes a "don't-look bit" to prevent further exploration, and a "don't-push stack" to prevent further participation of the path in the formation of a solution.

In 2015, Haynes and Corns described an *Algorithm for a Tabu Ant Colony Optimizer* [16]. This conference paper introduced the TabuACO algorithm and tested it against a few prize-collecting Steiner tree problems. The TabuACO was compared to a conventional ACO and found to outperform the conventional ACO.

2.1.3. QAP-Related Research. In 1999, Gambardella, Taillard, and Dorigo introduced *Ant Colonies for the Quadratic Assignment Problem* [17]. Taillard also developed a *Robust Taboo Search* (RTS) for the QAP [18]. Taillard was kind enough to post the source code for this algorithm, and the RTS is used later in section 3.3.4 as a key benchmark algorithm. In subsequent years other researchers have conducted experiments in this area. Talbi et. al. combined a parallel ACO with a Tabu search in an effort to solve the QAP. The design alternated between these two search methods to narrow and focus the search [19].

Wiesemann and Stutzle introduced *Iterated Ants: An Experimental Study for the Quadratic Assignment Problem* [20]. Qi presented *A Modified Ant Algorithm for Solving the Quadratic Assignment Problem* [21].

2.1.4. Scalability. In a 2010 journal article by Sameh, Ayman, and Hasan, each ant runs on a separate processor [22]. They describe a parallel ACO algorithm in which pheromones are shared in a common environment. Their research extends parallelization strategies suggested by Stutzle [23].

2.1.5. Multi-objective ACOs. Lopez-Ibenez in his 2004 thesis described an ACO in which he leverages the *Max-Min Ant* system to develop a multi-objective ACO in which a single pheromone is used to represent multiple objectives [24].

In 2005, Pinto, Baran, and Fabregat described a *Multi-Objective Multicast Routing based on Ant Colony Optimization* [25]. The algorithm maintains a population of ants and an external set of Pareto solutions. The algorithm uses a classic evolutionary algorithm to develop solutions. The distribution of the ants within the environment is controlled by the chromosome. Good distributions which lead to good solutions are evolved using crossover and mutation to generate solutions which improve with successive iterations.

2.1.6. Multi-pheromone ACOs. In 2010, Alaya, Solnon, and Ghedira described an *Ant Colony Optimization for Multi-objective Optimization Problems* [26]. Their *m-ACO* algorithm utilizes multiple pheromones – one for each objective. It also supports multiple colonies – one for each objective. Ants randomly choose a pheromone trail corresponding to an objective to optimize.

2.2. DEVELOPMENT OF THE APPLICATION

The improvements to the smart grid present many opportunities for improvement, optimization, and problem solving. This research applied the Tabu ACO solver to implement a planning function which worked to protect distribution assets and organize the delivery of energy through the grid from producers to consumers. A portion of the findings were published in [27]. The results are described more fully in section 4. . The approach taken was rather unique. The distribution network was modeled as a graph. Producers, consumers, and certain other assets were modeled as nodes, while wires were modeled as edges. In one effort the ants themselves were tasked with observing the constraints of the assets and identifying a viable solution in which the interests of the stakeholders were satisfied. The solver was given a scenario from the IEEE 34 node feeder model. The ants, through the network identified a constrained asset and protected it while implementing an energy market function. These results are presented in Section 4.8. This experiment proved the concept but did not test scalability. A larger experiment was conducted in which the test for the fitness of the circuit, and a test for the objective of the solver (implementing an energy market) were performed. Engineering Analysis software from EPRI was used to test proposals that were generated by the Tabu ACO solver. These results were published in [27] and are presented in more detail in Section 4.9. The Tabu ACO was able to form contracts between the stakeholders, while the EPRI OpenDSS software was able to identify a market scenario that limited the PV contribution to manageable levels.

3. THE TABU ANT COLONY OPTIMIZER

In Section 3. the TabuACO will be described and tested against a variety of puzzles. A generalized form of the solver is first presented, then it is applied to the Traveling Salesman Problem (TSP), Quadratic Assignment Problem (QAP), and the Steiner Tree Problem.

3.1. DESCRIPTION OF THE ALGORITHM

This section will describe the TabuACO at a high level. The TabuACO has many of the traits of a conventional ACO. Both the TabuACO and conventional ACO can be described at a high level by Figure 3.1. After initializing the environment, the code implements a loop which implements an ant sortie. If the stopping criteria are not met, the ant is sent out looking for "food." If she finds food (or runs out of time), her sortie ends and she returns home. If food was found, its quality is evaluated and the ant returns home. She applies attractive pheromone in response to the quality of the find. The TabuACO and conventional ACO both incorporate all of these features. The TabuACO is different than the conventional ACO in the way the ant behaves when she finds an empty node, and in how she computes chooses paths. The TabuACO uses two pheromones (attractive and repulsive) while the conventional ACO uses only one pheromone (attractive.) This difference in behavior for the ant is found during the sortie process. This difference is captured in Figure 3.2.

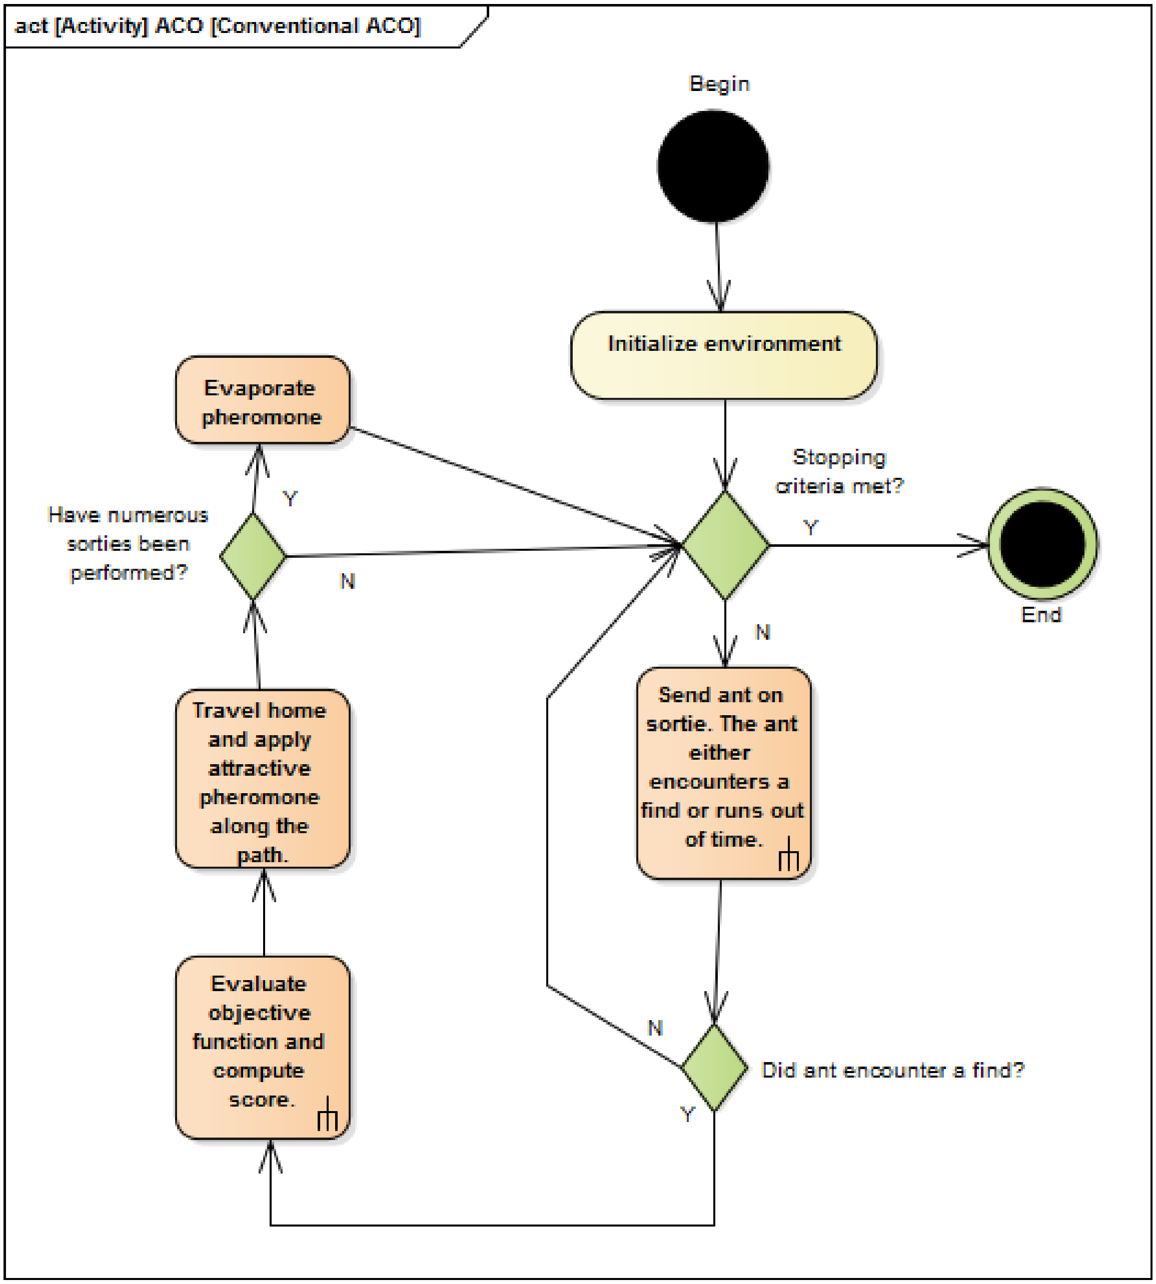

Figure 3.1. ACO Activity Diagram

The activity in red is unique to the TabuACO. This is a distinguishing feature that sets it apart from a conventional ACO. The TabuACO differs from a conventional ACO in three ways:

1.) The ability to apply a rule to identify edges that may be deprecated.

2.) The deposition of repulsive pheromone.

3.) The use of the repulsive pheromone in the ant's decision process.

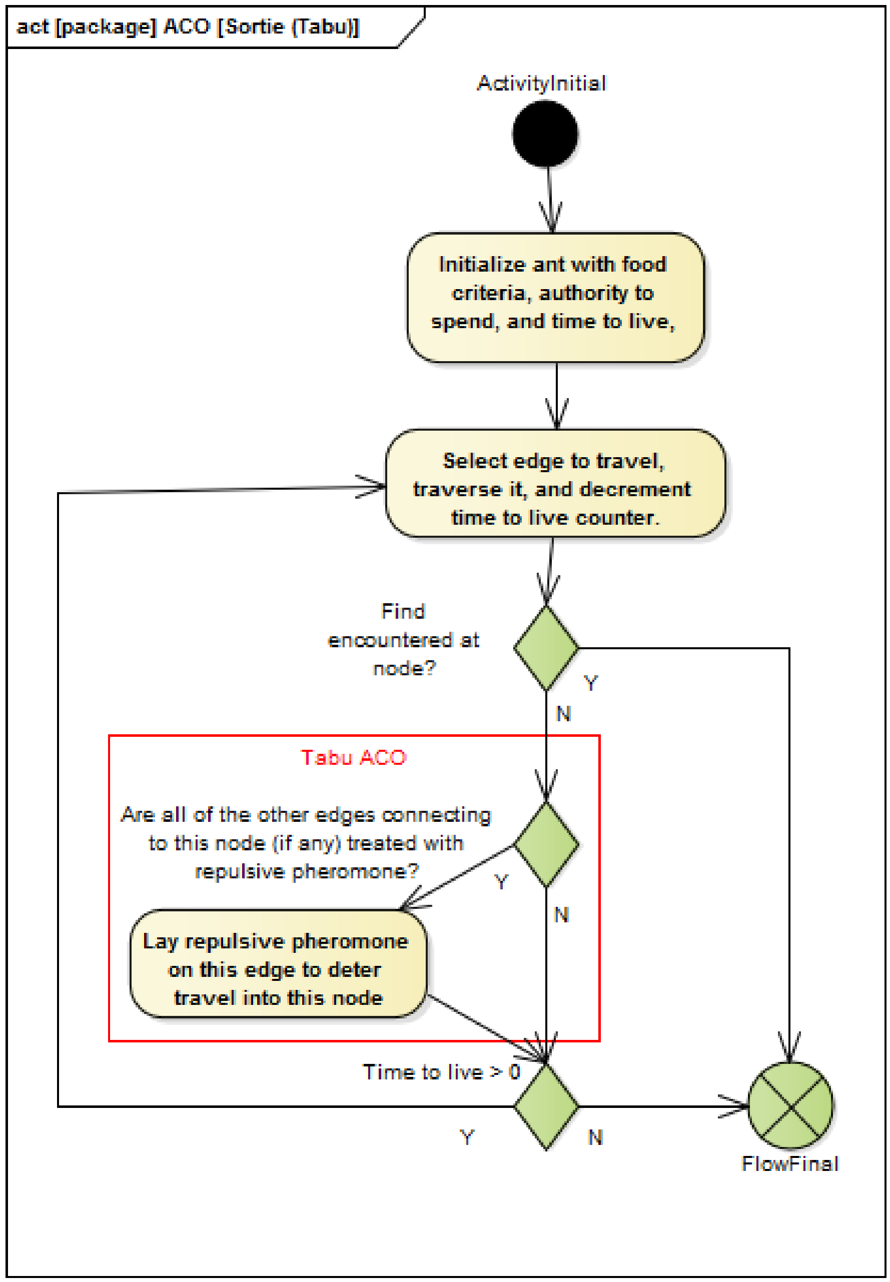

Figure 3.2. Detailed View of Ant Sortie Activity (with TabuACO Activity in Red.)

The TabuACO makes gains over traditional ACOs by effectively “trimming” paths from the graph which are known to be unproductive trails for the target food type. The algorithm requires a sense of direction in the graph. Undirected graphs will need to be overlaid with a reference direction. This is necessary so that pheromone meaning can be interpreted in the graph setting. Attractive pheromones on an edge are interpreted to mean that a recommendation exists from other ants to go in the direction of the arrow. If the arrow is pointing into the node where the ant is located, it represents a recommendation to not go down the edge. Similarly, repulsive pheromones on an edge are interpreted to mean that travel in the direction of the arrow is discouraged. Instead, ants should travel against the direction of the arrow. Attractive pheromones are used to attract ants to a food source. Repulsive pheromones are used to discourage ants from traveling to empty leaf nodes.

When ants are placed at a nest in the graph, they forage until they find food and return home. Ants remember the way home as they forage. Ants can be assigned a specific type of food to find and they lay a pheromone that corresponds to that food type. Different nests or different ants from the same nest may have different agendas and lay pheromones that are ignored when they don’t agree.

If an ant forages, finds a leaf node, and determines that the leaf node lacks the type of food it is looking for, the ant will deprecate the node by marking all of the edges connected to the leaf node with repulsive pheromone.

If an ant forages, encounters an inner node, and determines that it lacks suitable food, she will also consider all of the adjoining edges (except the edge homeward). If none of the adjoining edges have any attractive pheromone, and all of them have some repulsive pheromone, the ant deprecates the node by deprecating all of the adjoining edges.

If the ant finds the food she is looking for, she will lay attractive pheromone all of the way back to the nest. This behavior, with the deposition of attractive pheromone, is just like a conventional ACO.

3.1.1. The Deprecation Rule. If the prize-collecting ant travels an edge to encounter an empty node that has all other edges (insomuch as they exist) deprecated, then the just-travelled edge may also be deprecated.

Rationale:

1.) A deprecated edge is effectively "trimmed" from the solution and temporarily trimmed from the puzzle.

2.) An empty leaf node may be deprecated.

Then, a node containing one unexamined edge, and N-1 deprecated edges, is determined to not contain a prize. All N-1edges may be deprecated per the first rationale, leaving an empty leaf node which in turn may be deprecated per the second rationale.

3.1.2. Pheromone Deposition. Once an ant completes a sortie, the fitness function is computed to determine the goodness of the outcome. The fitness function renders a score, and the score is used to determine the relative goodness of the find (compared to previous finds). The ant then retraces her footsteps back to the nest and applies an attractive pheromone using an exponential moving average function. All edges have a reference direction, and a positive or negative value is deposited to represent attraction depending on the reference direction of the edge.

A repulsive pheromone is deposited by the ant, in full concentration, on deprecated edges, according to the rules described in section 3.1.1 above. All edges have a reference direction, and a positive or negative repulsive value is deposited to represent repulsion depending on the direction of the edge.

The code accounts for attractive and repulsive pheromones separately. One does not merely erase the other. The rules for pheromone deposition and evaporation are different for each type. If the pheromones were to be stored as a single (positive or negative) value, certain path geometries in the puzzle could allow useful information to be lost.

3.1.3. Pheromone Initialization. At the start of a trial, all pheromones in the model, both attractive and repulsive, are zeroed.

3.1.4. Pheromone Evaporation. Attractive pheromones are periodically evaporated using an exponential moving average formula. The formula averages zero into the pheromone values. Both positive and negative values are moved towards zero when this occurs. An attractive pheromone (τ_a) is scaled back by an evaporative constant (ρ_{ae}) between zero and one. A repulsive pheromone (τ_r) can be scaled back in a similar way. However, ρ_{re} is typically equal to 1.

$$\begin{cases} \tau_a = \tau_a \times \rho_{ae}, for\ attractive\ pheromones \\ \tau_r = \tau_r \times \rho_{re},\ \ for\ repulsive\ pheromones \end{cases} \tag{3}$$

3.1.5. Edge Selection. The ant computes a probability density function (PDF) based on pheromone deposits found on each path. Paths that loop back to a node on the path home are assigned a zero probability (unless the path back is the only option). From the PDF a CDF is computed. The CDF also relates to each connected edge just as the PDF. The ant then selects a number between zero and one. The random value selected is then related back to an edge number contained in the CDF, and the process ends with an edge being selected. These details are described in the pseudocode in the sections below.

3.1.6. Problem Setup. The ACO begins with a setup of the problem which relates the process to be optimized to a graph (or tree) structure. The TabuACO requires that the edges have a reference direction. This is so that the attractive and repulsive pheromone clues can be interpreted correctly by the ant. The graph must be organized to represent the problem to be solved. It also provides an environment for ants to store data. Some puzzles naturally have an organization to them which leads to an obvious up/down directional relationship between edges and nodes. Other puzzles do not have an obvious top and bottom. For these puzzles, an arbitrary up/down or left/right directionality can be imposed.

A common stopping criterion is when no improvement has been made in the score after some number of trials.

When the ant forages, the ant selects a path influenced by environmental factors. The edges and nodes that represent the environment may be modeled to reflect the constraints of the puzzle. Certain moves may be illegal due to the puzzle constraints. Other moves may be permissible only if certain runtime criteria are met. Of the moves that are allowed, the ant's selection is influenced by both the attractive and repulsive pheromone present in the environment. The ant sortie is a series of moves in which the ant is allowed to travel until she encounters a "find" or her travel is otherwise exhausted. Ants remember the path home, avoid forming loops, and prefer to not backtrack. Ants are assigned a specific type of food to find and they lay a pheromone that corresponds to that food type. In the TabuACO, if the ant encounters an empty leaf node, the edge leading to the leaf node is deprecated. The ant has just explored a location that will never have a “find” and never needs to be explored again. In a similar manner if a node is encountered that is completely surrounded by deprecated edges, has no attractive pheromones, and the remaining edge has been travelled and fully explored as well, then all of the edges on the node can be marked as deprecated. In this way deprecation can start at a leaf node and spread outwards through the graph as topology and travel allow.

Once an ant has encountered a find, a score is computed, and the ant returns home. During the trip home, the ant lays attractive pheromone as a function of the score. On rare occasions, the environment is updated, and pheromones are globally evaporated.

This general pattern is used for all of the applications of the TabuACO – whether they be known benchmark problems or new applications of the solver which have not been modelled before.

The TabuACO has been tested against several different benchmark problems. In each case the optimizer was modified somewhat to suit the problem being solved. The QAP was modelled as a tree, and unidirectional travel occurred from the root to the leaf. The Steiner tree problem was modelled as a graph, and the ant allowed to roam in any direction. Each variant of the TabuACO offered improvements to accomplish different purposes.

3.2. THE TRAVELING SALESMAN PROBLEM

The Traveling Salesman Problem (TSP) was promoted in the 1800s by William K. Hamilton and Thomas Kirkman. In this problem, there are *n* cities to be visited (numbered one through n). A fully interconnected graph represents the travel between cities, with a cost c_{ij} associated with each edge. The challenge is to find the least expensive path, while visiting all n cities, and returning to the starting point. The TSP objective function can be described by (4).

$$min \sum_{i=1}^{n} \sum_{j=1, i\neq j}^{n} c_{ij} \tag{4}$$

The TSP is considered a classic problem. It has been heavily studied and is described in countless literature. The TSP is considered NP complete. Puzzles with dozens of cities can be challenging. A particularly difficult puzzle containing a million cities have been devised. While considered unsolvable by most methods, some progress has been made with divide and conquer techniques [28] [29].

3.2.1. Proposed Methodology. The TSP puzzles used by [12] will be used here to compare the TabuACO to a reference ACO. The reference ACO will be the TabuACO operated with attractive (and no repulsive) pheromones.

3.2.2. TSP Modeling. The TSP is represented by a fully interconnected graph. Every node is connected by an edge to every other node. Nodes are not connected to themselves. There is a cost of traveling every edge. Puzzles will be drawn from the TSPLIB. The TSPLIB presents data in XML. The cost of traversing each edge is presented in the XML as each edge is declared. Some puzzles are presented in an (older) geographic format. In this format, each row represents a city along with an X-Y coordinate. In this case it is assumed that it is possible to navigate to each city from every other city, and the cost of travel is the distance between the two cities.

Each city has a number. Reference directions are overlaid on the graph by defining nodes (cities) with a lesser number to be upstream from nodes with a higher number. By way of example, Figure 3.3 contains a fully interconnected mesh graph with 4 nodes. Each node has a number. A direction is overlaid on each edge based on the node number so that a node with a lower number is “upstream” of a node with a higher number.

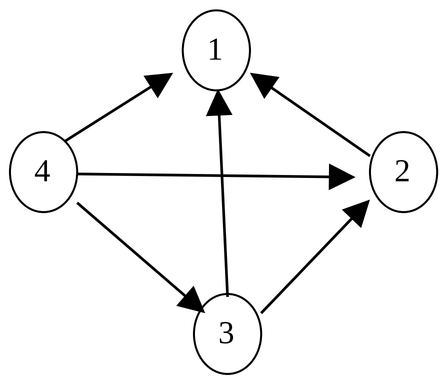

Figure 3.3. Example Imposed Reference Directions for a TSP graph of N=4.

3.2.3. Path Selection. A Probability Density Function (PDF) is computed based upon path history and pheromone density as described in (5).

$$p_{ij} = \begin{cases} 0, & if\ n_j \in P^k_{nest..end-2} \\ \left(\frac{2-\tau_{a_{ij}}-\tau_{r_{ij}}}{4}\right)\frac{3}{4}, & if\ n_j = P^k_{end-1} \\ \frac{2-\tau_{a_{ij}}-\tau_{r_{ij}}}{4}, & \text{otherwise} \end{cases} \quad (5)$$

Where, n_j is node j, p_{ij} is the probability of path selection along the edge from node i to node j, $P^k_{nodeList}$ is the path history for ant k starting at the nest and ending at the current position, $\tau_{a_{ij}}$ is the attractive pheromone along edge between node *i* and *j*, $\tau_{r_{ij}}$ is the repulsive pheromone along the edge between node *i* and *j*.

If an edge selection would cause the ant to loop back and cross into its path history, the probability of selection is zero. If an edge selection would cause the ant to backtrack, the selection is discouraged relative to other choices – being reduced by a factor of ¾. Otherwise, a probability of selection is the sum of the attractive and repulsive pheromones. However, given that the graph is directional, and either pheromone can be positive or negative, the sum of the two lies somewhere between -2 and +2. This is adjusted to become a number between 0 and 1.

3.2.4. Pheromone Deposition and Score Dissemination. For all edges along ant *k*'s path home, deposit positive attractive pheromone along edges where the reference arrow points away from the nest, and negative attractive pheromone along edges where the reference arrow points toward the nest. This is stated mathematically in (6).

For $\forall e_{ij} \in P^{k}_{pathHome}$, (6)

$$\begin{cases} \tau_{a_{ij}} = \tau_{a_{ij}} \times \rho_d + (1 - \rho_d), if\ reference\ arrow\ points\ away\ from\ nest \\ \tau_{a_{ij}} = \tau_{a_{ij}} \times \rho_d - (1 - \rho_d), if\ reference\ arrow\ points\ toward\ the\ nest \end{cases}$$

3.2.5. Expected Results. There is little reason to expect that the TabuACO solver will outperform the ACO solver. Negative pheromone deposition only occurs when the ant can find a leaf node to trim away. There are no leaf nodes in these puzzles.

3.2.6. Experiments. The puzzle is set up at the start of each sortie with fresh food at each node location (each city). City #1 is selected as the nest. An ant will be sent on a sortie, expected to visit each city, and upon claiming the last prize, return home. The code has been set to force the ant to use the remaining single hop home.

If repulsive pheromone application is enabled, the ant will watch for opportunities to lay down repulsive pheromone.

The TabuACO (with both attractive and repulsive pheromones enabled) will be allowed to run for 10^4 sorties, and progress towards completion noted. Twenty trials will be conducted, and the mean and standard deviation calculated.

Similarly, the reference ACO (with attractive pheromones only enabled) will be allowed to run for 10^4 sorties, and progress towards completion noted. Twenty trials will be conducted, and the mean and standard deviation calculated.

The Student's T-test will be used to identify significant performance differences between the two solvers.

3.2.7. Data. The TabuACO solver was tested on a series of TSP problems. The TSP puzzles selected are depicted in Table 3.1.

Table 3.1. Results Comparing TabuACO to Conventional ACO Using TSP Puzzles.

Puzzle Name	Best Known Score	Sorties	Percentage above best-known score	
			ACO	TabuACO
eil51	426	10^4	422	424
eil76	538	10^4	162	162
gr24	1272	10^4	110	111
kroA100	21282	10^4	613	610
d198	15780	10^4	3479	3479
lin318	42029	10^4	1672	1674
pcb442	50778	10^4	1367	1368

3.2.8. Discussion of Results. It is interesting to note that the set of TSP puzzles studied here are the same ones studied by Montgomery in his anti-pheromone research [12]. Montgomery, running a variety of anti-pheromone algorithms, found no discernable improvement over the reference ACO. Our results for the TabuACO also found no discernable difference between the TabuACO and our reference ACO. For the Tabu vs. classic ACO, the outcome is explained by the nature of the problem being solved. The TSP graph has no extraneous nodes. Every node in the TSP must be visited by the ant in order to form the solution. The TabuACO did not find anything it could deprecate. It was not able to effectively reduce the search space. Thus, it offered no advantage over a conventional ACO. Increasing the number of sorties will not make a difference in the outcome because the setup of the problem prevents nodes from being trimmed away. When the TabuACO deprecates an edge, there is always a node somewhere which is also being deprecated. All of the nodes present in the TSP puzzle must be present in the solution. This prevented the TabuACO from identifying any edges to deprecate [30].

3.3. THE QUADRATIC ASSIGNMENT PROBLEM

The Quadratic Assignment Problem (QAP) was described by Koopmans and Beckmann in 1957 [31] [32]. It is NP-hard [33]. It is commonly used as an expression of a facilities planning problem [32].

The QAP can be expressed mathematically as (7-8):

$$\min_{p \epsilon \Pi_{\mathcal{N}}} score(p) \tag{7}$$

$$score(\pi) = \sum_{i=1}^{N} \sum_{j=1}^{N} A_{ij} B_{\pi(i)\pi(j)} \tag{8}$$

Where, $\Pi_{\mathcal{N}}$ represents all of the permutations of $\mathcal{N}$. $\mathcal{N}$ is a set of numbers one to n. A and B are NxN arrays which represent distance or cost. π is a 1xN array which represents the ordering of the solution.

With the solution being drawn from the set of numbers one to N without replacement. The solution size grows as N!. With n=12, a search space of 479 million possibilities is created. If a brute force solver were to test each possible combination at a rate of one every 100 ms, it would take more than 1.5 years to complete the evaluation.

The QAP offers a type of problem which can offer a large search space, and a challenge to those who wish to represent the problem with a model. Memory constraints usually allow only a mere fraction of the model to be retained in memory at any given time. The challenge of the researcher is to identify the salient features that can be used to represent the problem and guide the solver to optimization.

3.3.1. Proposed Methodology. In most forms of Computational Intelligence problems are solved by forming an objective function which is then maximized or minimized. Code reuse is facilitated when the solver can be separated from the objective function. However, for successful convergence of an ACO it is important that as much of the results as possible be retained in the model.

The starting point for the TabuACO is that repulsive pheromones are utilized in addition to attractive pheromones. However, in order for a problem to be solved, it must first be modeled. The ACO is well suited to solving problems that are modelled as a graph. The TabuACO requires that the graph have directional edges. If the puzzle to be solved does not naturally have directionality, it must be assigned as part of the setup of the problem. The TabuACO uses the directionality to help the ant interpret the meaning of the pheromone. In a graph setting, a foraging ant can find herself inclined to choose an edge that has pheromone on it. The TabuACO uses the pheromone sign (positive or negative) along with the direction of the arrow to interpret the direction the ant should travel. A positive pheromone (attractive or repulsive) will guide the ant in the direction of the arrowhead. Similarly, a negative pheromone value is used to guide the ant away from the direction of the arrow. Without directionality defined, an ant can be drawn up a trail she is meant to go down.

A tree is a simplified form of a graph. When foraging occurs in a tree it can be confined to be unidirectional. When this occurs, it simplifies the code. There is only one direction for the ant to follow and pheromones cannot be misinterpreted by the ant. In this special case, the ant decision process does not need to consider the sign of the pheromone with respect to the direction of the edge.

Whenever a foraging ant arrives at a node, she will encounter a series of edges. She must make a decision (as depicted in Figure 3.4), and does so from clues she finds in the environment (if they are available).

The proposed algorithm uses the overlaid graph directionality to interpret a signed pheromone value in a directional sense.

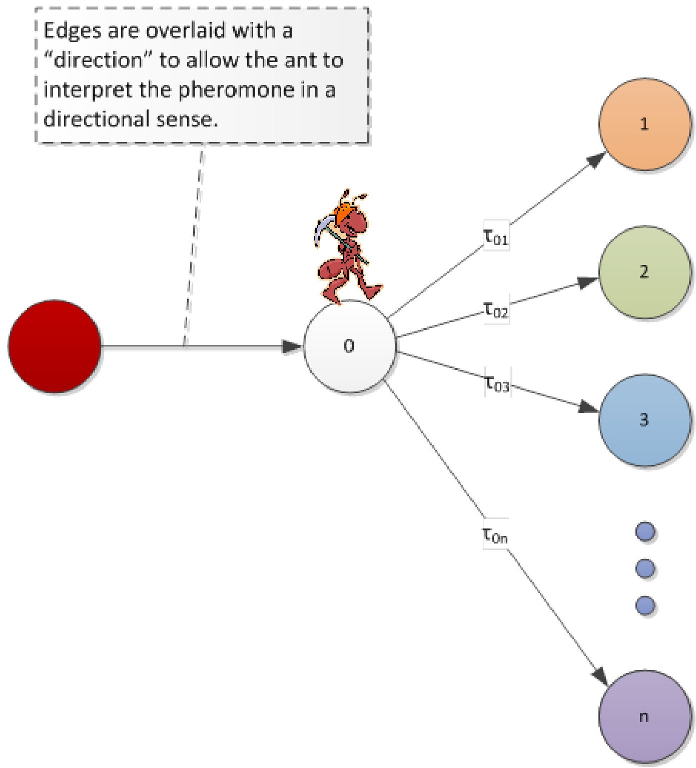

Figure 3.4. The Ant and the Decision Process.

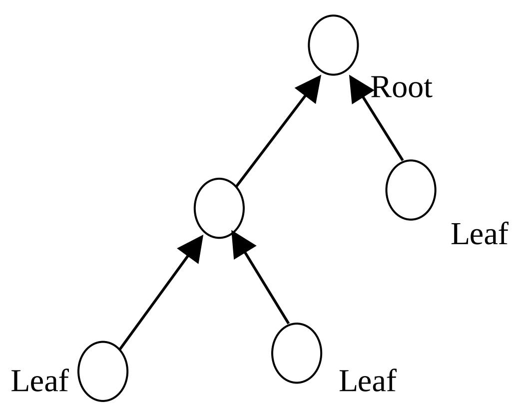

Figure 3.5. Example Tree

In this application of the ACO, ants will start at a root node (at the top as depicted in Figure 3.5) and forage downward. In this particular application, the problem is setup so that treasures are stored in leaf nodes. This application has a zero cost in travelling to a find. The combination of paths used to reach a find represents a potential solution. In a conventional ACO, we would expect that the ant would lay attractive pheromone along paths that yield good results. This approach assumes that good finds are located near each other. If this is not true, and good finds are randomly distributed among all locations, the solver may eventually find the optimum, but it may be no faster than a brute force search.

The TabuACO goes beyond a conventional ACO by adding additional information beyond the attractive pheromone:

- An edge is given a mode which indicates if it is unexplored, partially explored, fully explored, or prohibited by virtue of the puzzle.
- The solver may supplement this with repulsive pheromones. (Note: All pheromone values, attractive and repulsive, must be signed when bidirectional foraging is allowed.)
- The solver may deposit other information in the environment such as an indication of the best score ever discovered while travelling that edge.

The TabuACO will use any combination of the above stigmergy to find the optimum.

3.3.1.1. QAP modeling. A tree is used to model the QAP. An example is shown in Figure 3.6. Black lines represent potential choices. The red lines represent prohibited choices due to prior selections.

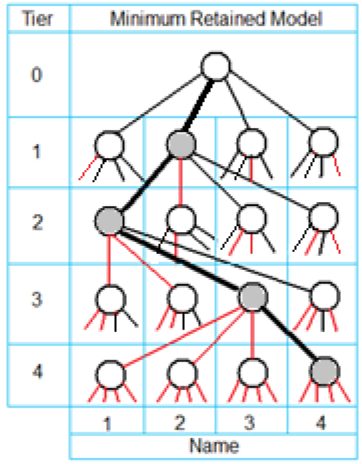

Figure 3.6. Example Encoding of [2 1 3 4] for N=4

A node is modeled as having *N* downstream edges and one upstream edge. A root node is introduced at the top tier. The root node fans out to *N* navigable edges at tier 1. Tier 1 nodes in turn fan out to *N*-1 navigable edges and 1 non-navigable edge. Tier 2 nodes fan out to *N*-2 navigable edges and 2 non-navigable edges. The process continues to Tier *N* where there are no navigable edges and *N* non-navigable edges. This constitutes a leaf node. This creates a structure in which there are conceivably N^N+1 nodes and $N(N^N+1)$ edges.

Figure 3.7 describes the Class Design for the TabuACO experiment used when modeling the QAP. Each node has a “name” one through *N*. The edges take on one of five different edgeModes. The node names can take on one of any of the *N* names in the range.

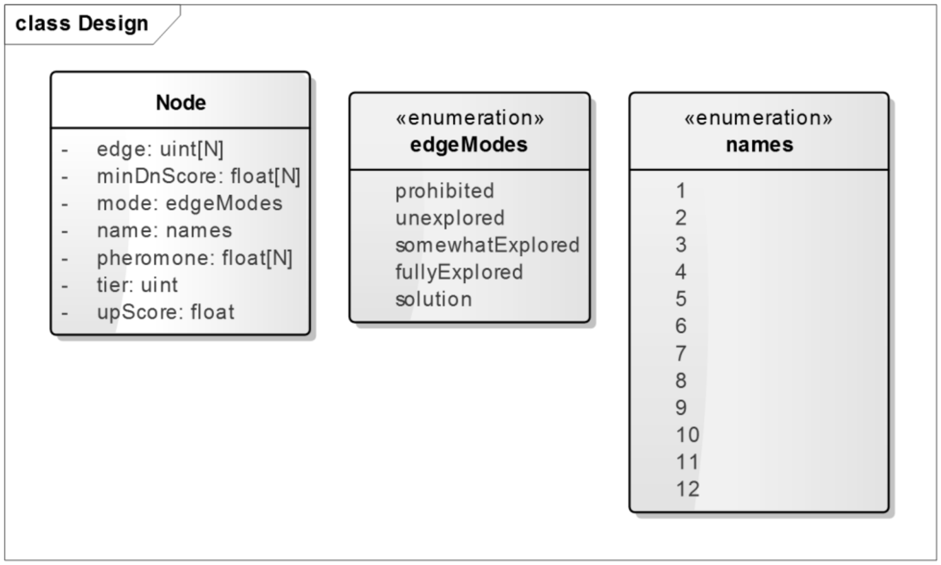

Figure 3.7. Class Design Shown for N=12

3.3.1.2. Pheromone calculation. The research implemented in [16] uncovered a fundamental difficulty with laying attractive and repulsive pheromones in response to the quality of a find. As a solution converges, the quality of the find tends to improve. What was once considered a good find, in retrospect, over the course of time, becomes a comparatively bad find. Most pheromones that are laid on paths indicating the quality of a particular find will become obsolete. A conventional ACO would evaporate all pheromones in an effort to destroy obsolete and misleading information, but not all pheromones are necessarily misleading. The TabuACO attempts to remedy such loss of information by storing non-evaporative "best of" score data along each edge in addition to pheromone information.

The pheromone and score data are then used to compute a PDF of the selection options, the PDF is then used to compute a CDF. The CDF is used to select from the available options.

3.3.1.3. Path selection. When an ant explores and selects an unexplored path, the solver promotes the edge status from "*unexplored*" to "*partiallyExplored*." If every downstream edge subordinate to a downstream node has been explored, then the edge leading to the downstream node may be promoted from "*partiallyExplored*" to "*fullyExplored*." The edge status is used to select the pheromone formula used in the calculation of the PDF.

The TabuACO may be operated in several different modes. Ordinarily both attractive and repulsive pheromones are used. (In this section the notation "TabuACO-A1R1" for convenience.) For test purposes it may be operated in other modes. It can be operated with attractive pheromones only (called "TabuACO-A1R0" in this paper.) It can be operated with negative pheromones only (A0R1), and just for comparison sake, it can be operated with no pheromones at all (A0R0). These options are listed in Table 3.2.

Table 3.2. Probability of Selection Based on Path Mode

Tabu ACO Solver	Attractive Pheromone may be used	Repulsive Pheromone may be used	Path mode			
			prohibited	un-explored	somewhat explored	fully explored
			Probability of selection			
A0R0	No	No	0	$P_{1/N}$	$P_{1/N}$	$P_{1/N}$
A0R1	No	Yes	0	$P_{1/N}$	P_R	0
A1R0	Yes	No	0	$P_{1/N}$	P_A	P_A
A1R1	Yes	Yes	0	$P_{1/N}$	P_{AR}	0

There are four probabilities identified in Table 3.2. The probability of selection for a path that would cause the ant to wander off the model is zero (no formula necessary). Likewise, when a path has been fully explored the repulsive pheromone calculation from the table yields a probability of zero. The probability formula for attractive pheromones P_A is computed by (12), repulsive by (13), and the combination of the two by (11).

$$P_{1/N} = \left(\frac{1}{N}\right)^{\alpha} \quad (9)$$

$$\tau_1 = \frac{localAverage_i - bestOfScore_{ij}}{MaxLocalScore_i + 1} \quad (10)$$

$$P_{AR} = \left(\tau_1 + \frac{1}{N}\right)^{\alpha} \quad (11)$$

$$P_A = \begin{cases} P_{AR}, if\ \tau_1 > 0 \\ P_{1/N}, otherwise \end{cases} \quad (12)$$

$$P_R = \begin{cases} P_{AR}, if\ \tau_1 < 0 \\ P_{1/N}, otherwise \end{cases} \quad (13)$$

Where α is an exponential value used to tune the solver, $bestOfScore_{ij}$ is the best (minimum) score identified along any of the partially explored downstream leaf nodes while traversing edge *ij*, $localAverage_i$ is the average of all path scores among the partially and fully explored edges (prohibited and unexplored paths are excluded) connected downward at node *i*, $maxLocalScore_i$ is the worst score identified among the downstream edges which have been partially or fully explored downward at node *i*, *N* is a puzzle dimension as described by (2).

When an ant is faced with *N* paths, she must choose one of them. Many paths may be *prohibited*. They are given a likelihood of selection of zero. Paths which are *unexplored* are by default given a likelihood of $\left(\frac{1}{N}\right)^{\propto}$. If the score history indicates that traversing the edge has proven advantageous, a contribution proportional to the quality of the find τ_1 is added to the probability function.

Note that the PDF values for the various *N* pathways do not always sum to exactly 1.00 with the above formulae. However, the code contained a normalization function which caused the CDF to provide a value which did indeed start with zero and end with 1.00.

In another design note, it should be mentioned that the formula would have assigned a probability of zero to pathways which loopback onto the ant's path history. However, since the entire problem was modeled as a tree, loops are not possible, and this portion of code was omitted. The ant did however contain a bias to her decision making which prevented her from backtracking. This appeared to speed the solvers to convergence, and was applied to the TabuACO as well as the reference ACO.

3.3.1.4. Pheromone deposition and score dissemination. In this model of the QAP, a tree structure is synthesized where the root is pictured at the top (tier 0), and leaves are at the bottom (tier *N*). The score is computed for a given "path" through the tree which generates a sequence of numbers between the root and a leaf. Whenever the ant completes a sortie by encountering a find, a new score for that combination of pathways is computed. The ant traverses back from the find at the leaf node to the root. The ant updates the "best score" for each edge involved in the solution. If the solution produces a better score than the historical best score for that edge, the new best score replaces the historical best score on that edge. Each edge is considered in turn until the ant arrives home.

The ant uses this score information to compute a local best score, as well as her personal experience of a lifetime best score, to compute pheromones as needed as described in Table 3.2.

3.3.1.5. Pheromone initialization. This version of the ACO begins with an environment in which there is effectively no pheromone concentration in the environment. It is initialized as shown in (3).

Similarly, each edge is set to either *unexplored* or *prohibited*. The "best known score" for each edge is left undefined for edges with a status of *unexplored* or *prohibited*.

3.3.1.6. Pheromone evaporation. Pheromones are derived values in this particular implementation of the TabuACO and ACO. However, the pheromone metadata can be destroyed due to the memory management process. Poor performing pathways may be discarded from memory. This is comparable to rapid evaporation of selected pathways known to be poor performers.

3.3.2. TabuACO Pseudocode to Search Trees. Figure 3.8 describes the main routine used to operate the algorithm. The algorithm begins with some initialization in lines 1-3 and repeats in a loop until the solution score is found. During the trip to the leaf, the ant will choose random paths that are influenced by the pheromone calculation, which in turn is based upon any known scores along previously explored paths. The path chosen represents the QAP solution, and with the ant reaching a leaf node, enough symbols have been selected to complete the formation of a solution. The solution is then scored and the results disseminated along the path traveled. With a large puzzle, it was found that the best use of the ant's time was to perform repeated sorties which explore new territory. Then, armed with some preliminary information, the ant is able to favor certain symbols over others in the first position, second position, and so on, until an optimal solution emerges. As depicted in Figure 3.6, an ant will start at the root of the tree (shown at the top) at level 0. At each junction, she must choose a downward path, and then move downward one level until a leaf node is encountered (reaching the N^{th} symbol.)

```
1  Initialize the root node and an ant;
2  antSortieHopCount ← 0;
3  Reset run timer;
4  Repeat
5  |  createDnCDF(); per Figure 3.9
6  |  decodeCDF(); per
7  |  Acquire a node (new or existing) from the memory manager.
8  |  Move ant down to selected node.
9  |  Increment antSortieHopCount;
10 |  If( ant at leaf node)
11 |  |  Compute score using (8).
12 |  |  Repeat
13 |  |  |  Travel a step homeward. Along the edge homeward update the historical
   |  |  |  best score as appropriate with the newly computed score.
14 |  |  |  Promote the edge status to fullyExplored if all other connected edges are
   |  |  |  fullyExplored or prohibited; otherwise promote the edge status to
   |  |  |  partiallyExplored.
15 |  |  Until( ant home)
16 |  Endif
17 until( stopping criterion met)
18 Report optimal and resulting scores
```

Figure 3.8. Main Routine Pseudocode

Figure 3.8 line 5 has the ant using the “createDnCDF” routine to create a cumulative density function (CDF) in preparation for path selection. Then, by calling “decodeCDF” in line 6, the ant needs merely to pick a random number between 0.0 and 1.0, and determine the corresponding path which has been selected as a result.

```
compute localAverage_i, for local (downward going) edges around node i;
compute MaxLocalScore_i, the high score on local (downward going) edges around node i;
//build PDF from stigmergy data
for (j = 1 to N)
    compute P_1/N using (9)
    compute P_AR using (10);
    compute P_A using (12) and P_R using (13);
    case (the mode of edge j)
        "prohibited": pdf_ij ← 0;
        "unexplored": pdf_ij ← P_1/N;
        "somewhatExplored":
            case (pheromone usage) of
                "no pheromones": pdf_ij ← P_1/N;
                "repulsive only": pdf_ij ← P_R;
                "attractive only": pdf_ij ← P_A;
                "attractive&repulsive": pdf_ij ← P_AR;
            Endcase
        "fullyExplored":
            case (pheromone usage) of
                "no pheromones": pdf_ij ← P_1/N;
                "repulsive only": pdf_ij ← 0;
                "attractive only": pdf_ij ← P_A;
                "attractive & repulsive": pdf_ij ← 0;
        Endcase
        "solution": pdf_ij ← 1;
    Endcase
Endfor
//compute CDF from PDF
for k = 0 to N
    cdf_k ← 0.0;
    for l= 0 to k
        cdf_k ← cdf_k + pdf_l;
    Endfor
Endfor
for k = 0 to N
    cdf_k ← cdf_k / cdf_N;
Endfor
```

Figure 3.9. CreateDnCDF Pseudocode

The acquisition of a node in Figure 3.8 line 7 can be somewhat involved. The memory manager ranks the value of each node to the solution and recycles the node ranked lowest.

Nodes far away from the root, with a poor score, and numerous edges that are fully explored are good candidates for reuse. (There are multiple reasons to retain score information in the environment. One reason is to influence the path selection by the ant in (3), but another reason is to influence the memory manager's ranking and recycling of the limited memory resource in Figure 3.8 line 7.)

Figure 3.9 describes the pseudocode to build the CDF used to select the path the ant will follow. The routine uses path stigmergy as input, and develops a path selection as output.

With the CDF having built with the createDnCDF routine (Figure 3.9,) it must be decoded. This is done with the decodeCDF routine (Figure 3.9.) A random number between zero and one is selected, and this in turn maps to a corresponding range in the CDF array. The array index represents the path number which is chosen.

```
selection ← 0;
rvalue ← random number between 0 and 1;
while((rvalue ≥ cdf_selection) and (selection < N))
    selection ← selection + 1;
endwhile
return selection;
```

Figure 3.10. DecodeCDF Pseudocode

3.3.3. Expected Results. The QAP has been shown to be an NP hard problem. If the entire search space were laid out in memory using a scheme similar to Figure 3.6, we see from Table 3.3 (the N! column) that the memory requirements can easily outstrip the ability of today's computers. Even if this were viable, the time required to solve the puzzle can become inordinate as well. If we were to use a brute force search, and anticipate that on average, the optimal solution would be found by searching half of the search space, and that 125 thousand segments could be processed per second (where a QAP puzzle is *N* segments long), it would take the time described in Table 3.3 (columns 3 and 4) to solve the puzzle.

Table 3.3. Estimated Time Required to Solve a Puzzle
Based on a Brute Force Search Process

QAP length	Search space = N!	Time to search n!/2 @ 125 k segments/s	
12	479×10^6	2×10^3 s	0.5 hr.
15	1.3×10^{12}	5×10^6 s	6 days
18	6.4×10^{15}	3×10^{10} s	812 yrs.
20	2.4×10^{18}	10^{13} s	3×10^5 yrs.
22	1.1×10^{21}	4×10^{15} s	1×10^8 yrs.
25	1.6×10^{25}	6×10^{19} s	2×10^{12} yrs.

An ACO solves its puzzle through information stored in the environment (stigmergy). Large puzzles such as the QAP can be challenging to solve with an ACO. The TabuACO represents the QAP as depicted in Figure 3.6. Since it cannot store the entire model in memory, it will store only a portion of it. This will create a challenge since the ACO requires information from the environment in order to solve the puzzle. A memory manager will attempt to identify nodes that have the least value and recycle them. This could cause a loss of information which slows convergence. All of the forms of pheromone metadata associated with a given node (and connected edges) are destroyed when memory is recycled. When new (possibly recycled) nodes and edges are created by the memory manager, the *unexplored* pathways have a probability of selection of approximately $\left(\frac{1}{N}\right)$. As pathway history matures, scores which prove to be above average will influence the probability in one direction or another away from 1/N.

With the TabuACO it is always possible for deprecation to spread. Given that the problem is modeled as a tree, we would expect deprecation to spread upward towards the root. This particular implementation of the TabuACO applies a path status to deprecate explored pathways. Deprecation does spread upward in the model, and nodes closer to the root are considered more valuable to retain in memory than those far removed from the root. Thus, deprecation that has spread upward has a high probability of retention in the model by upper edges even when the child leaf nodes and edges have been destroyed by the memory manager.

3.3.4. Experiments. Two types of experiments will be conducted: one to show that the use of repulsive pheromones offers an improvement over an attractive-only solver; and another experiment to compare these solvers to other benchmark solvers.

3.3.5. Comparison of Performance Between an Attractive-Only Pheromone Driven ACO to an Attractive-Plus-Repulsive Pheromone Driven ACO. A library of QAP problems is available from Burkard, et al. [34]. The entire series from N. Christofides and E. Benavent [35] was tested. The Student's t-Test [36] was used to measure the significance of an improvement to a 95% confidence level.

3.3.6. Comparison of Performance Between Assorted Benchmark Problems. The benchmark testing by [17] compared the Hybrid Ant System for the QAP (HAS-QAP) with the genetic hybrid method of [37], the reactive tabu search [38], a tabu search of [18], and simulated annealing [39]. Of these tests, Taillard's tabu search [18] will be run alongside the TabuACO solvers on the same machine, matched by iteration count rather than runtime. By comparing Taillard's tabu search to the TabuACO, it provides a benchmark comparison to all of the solvers tested by [17].

3.3.7. Data. Two sets of experiments were performed. One experiment tested the TabuACO against a conventional ACO. The other tested the TabuACO against other solvers.

3.3.7.1. Performance comparison using attractive-only verses attractive+repulsive pheromones. The entire suite from Christofides and Beavent was tested to compare the TabuACO to the reference ACO solver. The results are summarized in Table 3.4.

Table 3.4. Results Comparing the Performance of the TabuACO to the Reference ACO

Puzzle Name	Best known score	Sorties	Percent above best-known score	
			ACO	TabuACO
chr12a	9552	3582	62.7	49.3
chr12b	9742	174	118.2	118.2
		10^4	48.0	44.9
chr12c	11156	1102	58.3	58.3
		$5x10^4$	24.5	22.9
chr15a	9896	1801	132.6	133.3
		10^4	95.0	86.2
chr15b	7990	3081	155.1	142.9

Table 3.4. Results Comparing the Performance of the TabuACO to the Reference ACO (con't)

chr15c	9504	7326	116.9	121.7
		$5x10^4$	89.7	81.8
chr18a	11098	22541	153.1	152.6
		$5x10^4$	140.5	136.3
chr18b	1534	55	116.9	116.8
		$2x10^5$	30.2	29.2
chr20a	2192	220016	99.4	96.5
		10^6	92.1	85.0
chr20b	2298	606779	83.2	67.6
chr20c	14142	28646	167.3	168.9
		$4x10^4$	182.3	167.0
chr22a	6156	452890	31	30
		10^6	29.5	29
chr22b	6194	239838	31	31
		10^6	29.9	26.1
chr25a	3796	440915	166	161
		10^6	156.5	152.1

3.3.7.2. Comparison to other solvers. The puzzles in Table 3.5 were solved by the RTS, TabuACO, and ACO solvers. In most cases the RTS solver drove the puzzle to the best-known minimum. In some cases, the RTS solver could not achieve this minimum. It was allowed to run its course, and the number of iterations used to achieve its best score were recorded. This same number of iterations were then used to allow the TabuACO and reference ACO solvers to achieve the best score they could attain. The percentage above ideal is expressed in the two rightmost columns. Ten trials were conducted and averaged.

Table 3.5. Comparison of RTS to TabuACO and Reference ACO Solvers

Puzzle Name	Best known score	Iterations to solution	Percent above best known		
			RTS	ACO	TabuACO
bur26a	5426670	1042	0.00	4	4
bur26b	3817852	874	0.00	3	3
bur26c	5426795	11032	0.00	3	3
bur26d	3821225	2294	0.00	4	4
bur26e	5386879	2535	0.00	4	4
bur26f	3782044	5014	0.00	4	4
bur26g	10117172	15117	0.00	3	3
bur26h	7098658	5457	0.00	3	3
els19	17212548	2564	0.00	59	59
kra30a	88900	4516	0.00	32	33
kra30b	91420	5467	0.00	29	29
nug20	2570	8135	0.00	13	14
nug30	6124	29810	0.00	17	16
tai20b	122455319	895	0.00	30	30
sko42	15812	90800	0.00	17	16
sko49	23386	245204	0.05	15	15
sko56	34458	264948	0.03	15	15
sko64	48498	18492	0.09	15	15
sko72	66256	53704	0.08	14	14
sko81	90998	915957	0.03	13	13
sko90	115534	5851	0.13	15	15
tai20a	703482	52349	0.00	12	13
tai25a	1167256	84827	0.00	12	12
tai25b	344355646	989	0.00	56	56
tai30a	1818146	38009	0.00	13	13
tai30b	637117113	32393	0.003	36	32
tai35a	2422002	904498	0.00	13	12
tai35b	283315445	27880	0.00	32	32
tai40a	3139370	113095	0.70	14	13
tai40b	637250948	22954	0.00	37	37
tai50a	4941410	13688	1.06	15	15
tai50b	458821517	954177	0.00	33	34
tai60a	7208572	461004	0.89	13	13
tai60b	608215054	639669	0.00	36	36
tai80a	13557864	697395	0.94	12	12

Table 3.5. Comparison of RTS to TabuACO and Reference ACO Solvers (con't)

tai80b	818415043	343559	0.02	34	34
wil50	48816	84827	0.02	9	9

3.3.8. Discussion of Results. Two different experiments are discussed. One pits a conventional ACO against the TabuACO. The other compares the TabuACO to other solvers.

3.3.8.1. Performance comparison using attractive-only pheromones verses attractive+repulsive pheromones. The study shows that (for the puzzles studied) the TabuACO with both attractive and repulsive pheromones outperformed the solver with attractive pheromones only [30]. When a short run was performed, it was difficult to detect a difference in performance between the two. A longer run was often necessary to bring out the differences in performance. It is believed that this is due to the need for a significant amount of information to be present in the environment before the ant's behavior becomes significantly affected. Compare the puzzle with a small number of runs to the same puzzle with a large number of runs. Quite often, a statistically significant difference[3] is not apparent until after ten thousand sorties.

It should be noted that the TabuACO solver is intended to work with computationally constrained applications which are unable to store the entire model in memory. If the important and relevant portions of the model are retained, convergence upon the optimum should be possible. If important portions of data are missing from the

[3] Significance being quantified here by the student's T-test.

model, the solver will have difficulty finding the optimal solution and may malfunction altogether.

With the QAP established as outlined in Figure 3.6, a large search space is created. Only a portion of the travelled edges are retained in memory. An analysis of a completed run typically shows a count of *fullyExplored* edges to be approximately 0.02%. The memory manager recycles *fullyExplored* nodes. Yet, due to their strategic location, these remaining edges served to influence the outcome of the TabuACO solver.

The test results show that in every case which is driven toward a result, there is a statistical advantage to using the repulsive pheromone information in addition to the attractive pheromone information.

3.3.8.2. Comparison to other solvers. The analysis in this paper compared the RTS solver to the TabuACO solver, and by extension, to all of the other solvers tested in [17]. The comparison in this case was based on iteration equivalents rather than computational time. The RTS solver beat the TabuACO solver in every case. The TabuACO and conventional ACO did fairly well on the entire "bur" series of puzzles, but not as well on other series. It was expected that the RTS (as a specialized solver) would outperform the more general-purposed ACO and TabuACO solvers. The TabuACO and conventional ACO both deposited attractive pheromone and attempted to converge to a solution by exploring combinations of numbers that appeared to yield good scores. The TabuACO outperformed the conventional ACO because it prevented previously explored portions of the graph from being reexplored. The Robust Taboo Solver (RTS) however made use of *a prioi* information. It was able to eliminate many combinations of numbers from the search space by exploiting properties of the QAP equation itself.

3.3.8.3. Memory management. The solver retained only 1000 of the most valuable nodes (as determined by the scores leading to them). For any puzzle studied, this is a remarkably small percentage of the solution space. Even for the smallest puzzle, less than 2 part per million of the solution space was retained in memory. For the largest puzzle, 1000 nodes represent a mere $1{:}10^{22}$ portion of the solution space. The results show a statistical difference in performance between single and dual pheromone operation. It is believed the memory manager must be doing an effective job ranking node value, retaining important nodes and discarding less important ones.

3.4. THE STEINER TREE PROBLEM

The Prize Collecting Steiner Tree Problem (PCST) is well known. Given a graph with nodes and edges, where a cost assigned to using every edge, and one or more nodes have “prizes.” The “rooted” version of the problem contains a special node which must be part of the solution. The cost can be expressed as the value of the prize minus the cost of the edges spanning between the required node and the prize node. The objective is to determine the lowest cost path between the required node and the prize node and from this information develop a minimum spanning tree. When multiple prizes exist in the problem, and therefore multiple trees exist in the solution, a “Steiner Forest” is created.

3.4.1. Steiner Tree Literature Review. In 1999, Gendreau and Larochelle reported “A tabu search heuristic for the Steiner Tree problem” [40]. Gendreau and Larochelle leveraged a tabu technique which prohibited certain moves based on a minimum spanning tree analysis of a given tree. The tabu analysis caused certain unproductive moves to be categorically prohibited. This narrowed the search space to solve the tree to a smaller space. The solver alternates between techniques until convergence. The solver was found to perform well compared to other contemporary solvers which were similarly able to leverage unique characteristics of the Steiner Tree.

In her 2004 PhD dissertation, Ljubic described and developed a number of memetic algorithms designed to solve Steiner Tree [41]. Memetic algorithms are evolutionary algorithms which purposefully leverage some aspect of the puzzle being solved in order to gain a computational advantage.

3.4.2. Introduction. The "Second Order Swarm Intelligence" ACO presented by Ramos, Rodrigues, and Louca [14] is well suited to problems such as the TSP in which ants are scattered across a surface, travel, and deposit pheromones to identify a closed Hamiltonian circuit by highlighting the optimal path within the environment. However, other types of problems exist in which the objective is quite different, the cost function is quite different, and the solution is not obtained in the pheromone trails themselves but in other network related information.

The paper presented herein examines a class of problem in which a closed Hamiltonian circuit is not desired. Instead, a Minimum Spanning Tree is desired from a root to every prize in the graph. The type of solution offered needs to support multiple roots competing for the same prizes. This type of problem seems closely aligned with a form of the Rooted Prize Collecting Steiner Tree.

The common Steiner Tree problem allows insertion of "Steiner Points" which are typically representative of modifications to the build out of a network. The research considered in this paper focuses on a variant of the rooted Prize Collecting Steiner Tree problem in which modifications to the original network topology are not allowed.

Table 3.6. PCST - ACO Modeling

Steiner Tree Model	ACO
Root	Nest
Prize	Food source
Objective Function	Ant agent explores network to discover prize valuation minus travel cost.

The ACO seems to be a natural fit to the Steiner Tree model (as described in Table 3.6.)

In order to validate the efficacy of the dual pheromone TabuACO, the proposed ACO will be compared to an attractive-only pheromone ACO using the same code on the same machines.

3.4.3. Proposed Methodology. The TabuACO operates to solve the Steiner tree in a manner similar to the generic algorithm. The following pseudocode describes the algorithm.

```
Initialize network and ants
While stopping criteria are not met
    Set one or more ants to have an antMode of "foraging." Each ant is given a
    time limit to search for food.
    While ants are on sorties
        Case antMode of
            Foraging:
            Select an edge to traverse. As the ant moves, she must maintain a
            path history that shows how to return to the nest.
            Upon arrival at a node, the ant examines the node to see if it contains
            a prize for an unsolved tree. If so, the ant claims the prize and sets
            antMode to be "homeward bound."
            If the new node is a leaf node, and a prize was not found, the ant lays
            down repulsive pheromone on all adjoining edges.
```

Figure 3.11. TabuACO Pseudocode for Solving the Steiner Forest Problem

9	If the new node is not a leaf node, does not contain a prize, all adjoining edges except the homeward edge are considered. If none of the edges (save the homeward edge) contain attractive pheromone, and all (save the homeward edge) contain some repulsive pheromone, then all edges (including the homeward edge) are deprecated with repulsive pheromone.
10	If an ant forages unsuccessfully for an extended period of time so that its time limit expires, then antMode will be set to “homeward bound.”
11	Homeward Bound: As an ant returns home, and an unsolved foodsource was found, ants lay down attractive pheromone along each edge in the path between the foodsource and the nest according to the quality of the find. The antMode is set to “home.” If the ant encounters a solved foodsource, it ignores the find and continues foraging.
12	Home: Ant reports find to nest, all contracts formed and amount spent. The ant sortie is over.
13	EndCase
14	Endwhile
15	Evaporate pheromones
16	Endwhile
17	Report optimal paths to each prize

Figure 3.11. TabuACO Pseudocode for Solving the Steiner Forest Problem (con’t)

In line 1 of Figure 3.11, the network is initialized. All pheromone traces are eliminated (zeroed). All ants are placed at their home nest nodes and not authorized to travel.

In line 2, a stopping criterion is set. A certain number of iterations may be allowed for solutions to be found. In the case of a known solution, the stopping criterion could be for the ants to find the known solution.

In line 3, if multiprocessing is supported, multiple ants may be sent out concurrently.

In line 4, a WHILE loop is set up that allows ants to run sorties in which they forage for food.

In line 6, the ant originated at node ***nest***, now sits at a node ***i*** (n_i), and considers selecting an edge $\boldsymbol{e_{ij}}$ which would allow her to move to node ***j*** (n_j).

First a probabilistic function $\boldsymbol{p_{e_{ij}}}$ is computed for each edge connected to the node the ant is at as described in (15).[4]

$$p_{e_{ij}} = \begin{cases} 0 & , if\ n_j \epsilon P^k_{nest..end-2} \\ \frac{2+\tau_{a_{ij}} \times (1-\rho_{ad}) - \rho_{ad} - \tau_{r_{ij}}}{2} & , if\ n_j = P^k_{end-1} \\ \frac{2+\tau_{a_{ij}} - \tau_{r_{ij}}}{2} & , otherwise \end{cases} \tag{15}$$

If the proposed node ***j*** (n_j) exists in ant ***k***'s path history somewhere between the nest and two hops before node ***i*** (inclusive), it is assigned a probability for selection of zero. This prevents loops from forming.

If the proposed node ***j*** is the node previously visited before node ***i*** it represents backtracking on the part of the ant. The likelihood of selection is discouraged but not prohibited.

Otherwise, the probability of selection $p_{e_{ij}}$ is computed based directly on the positive and repulsive pheromone along the edge ***e*** between ***i*** and ***j***.

A Cumulative Distribution Function (CDF) is then computed as described in (16-17).

---

[4] Of course when the environment is a tree, only one edge may exist between two nodes, but a graph, by definition, allows multiple edges to exist between nodes 'i' and 'j'.

$$c(e) = \sum_{e=1}^{n} p_e \tag{16}$$

$$cdf(e) = \frac{c(e)}{c(n)} \tag{17}$$

The ant will then select a number between 0 and 1, and perform the reverse cdf function to identify which edge has been selected.

Table 3.7 describes an example where an ant sits at node ***i*** and has n=4 paths to choose from. The ant has reached an unexplored portion of the graph, so τ_e=0 for every edge. The ant travelled along edge ***e***=1 to reach node ***i***. The ant's path history shows it must visit node ***j***=2 to return to its nest. Based upon the rules given above, the following table would be built:

Table 3.7. Example CDF Data

e	1	2	3	4
p(e)	0.25	0	0.5	0.5
c(e)	0.25	0.25	0.75	1.25
cdf(e)	0.2	0.2	0.6	1

The reverse cdf function is simply a matter of testing the random value against cdf(e) values until the correct value of ***e*** can be found. The "reverse CDF" process is one of converting a pseudorandomly selected number to a path number. This process is described in Figure 3.12.

```
1  r = uniform randomly distributed value between (0,1)
2  e = 1
3  eFound = FALSE.
4  While (not eFound)
5      If ( r ≤ cdf(e))
6          eFound=TRUE
7          break out of while loop
8      EndIf
9      e = e + 1
10 EndWhile
11 %the value for 'e' is now known
```

Figure 3.12. Reverse CDF Pseudocode

The virtual ant (evaluated at line 6) travels to node $\boldsymbol{j}$ along edge $\boldsymbol{e_{ij}}$. Edge and node path histories are maintained on two stacks. If the value selected $\boldsymbol{j}$ is the same as the most recent entry on the stack, the ant is backtracking. If the ant is backtracking, the most recent entries are popped from the stacks. If the ant is venturing further from the nest, the node and edge traversals are pushed onto the path-home stack P^k.

In Line 7 of Figure 3.11, it is the nature of the Steiner problem for the entire prize to be claimed in order to maximize the profit. The cost along the edges are incurred for their use and independent of the amount the edge carries.

In Line 8 of Figure 3.11, repulsive pheromone is laid along an edge by use of (18). The EMA factor, ρ, must be a value between 0 and 1. A large value for ρ will introduce small changes in the environment. (ρ_d is the Exponential Moving Average factor for deposition of pheromone). The algorithm uses a large value initially to encourage exploration, and a smaller value as the iteration count grows to encourage exploitation.

$$\tau_{r_{ij}} = \begin{cases} \tau_{r_{ij}} \times \rho_d + (1 - \rho_d), if\ head\ connects\ to\ node \\ \tau_{r_{ij}} \times \rho_d - (1 - \rho_d), if\ tail\ connects\ to\ node \end{cases} \tag{18}$$

The resultant τ_r is a value that becomes nonzero as the repulsive pheromone takes on significance. It approaches ± 1 as the number of iterations approaches infinity[5].

In Line 11 of Figure 3.11, when a find occurs for an unsolved tree, attractive pheromone is laid along each edge homeward by use of (19).

$$\tau_{a_{ij}} = \begin{cases} \tau_{a_{ij}} \times \rho_d - (1 - \rho_d), if\ head\ connects\ to\ node \\ \tau_{a_{ij}} \times \rho_d + (1 - \rho_d), if\ tail\ connects\ to\ node \end{cases} \tag{19}$$

The resultant $\tau_{a_{ij}}$ is a value that approaches ± 1 as the number of iterations approaches infinity.

In Line 15 of Figure 3.11, global pheromone evaporation occurs by use of (3).

3.4.4. Expected Results. A PCST graph can have any topology. It will contain nodes with one or more edges. A node with multiple parents and no children could serve as a termination point in some problems. However, this is not the case with the Steiner Tree Problem. Any node with multiple edges could serve as a waypoint to connect to another point in the solution. For the PCST problem, it seems that only a node with a degree of one (a leaf node) can serve as a starting point for deprecation. Repulsive pheromone is first placed on the pendant edge and spreads from there to other parts of the graph by the ant as she forages. The performance of the TabuACO depends entirely on

---

[5] (If evaporation occurs the pheromone concentration will settle on a lesser value.)

the Steiner Tree topology. If no leaf nodes are found, we would expect the TabuACO to behave no better than a conventional ACO. If a leaf node is found, and the rules of geometry and the placement of prizes allow large areas of the search space to be deprecated, we would expect the TabuACO to significantly outperform a conventional ACO.

3.4.5. Experiments. A library of Steiner Tree problems was published by Resende

[42]. Improvements were made to these problems by Ljubic [41]. Three of the K100 series from Ljubic were used for the testing. None of these puzzles contained leaf nodes, so in order to create a greater variety of graphs for testing, the K1C100.3 and K14C100.3 puzzles were created.

The K1C100.3 was made by taking the K100.3 puzzle and attaching a child node to every node in the original puzzle. Every prize in the parent node was copied to its child. A cost of 1000 points was assigned to each pendant edge.

The K14C100.3 puzzle was made by taking each node of the K100.3 puzzle and attaching two children to it, two children to those children, and two children to each of those children.

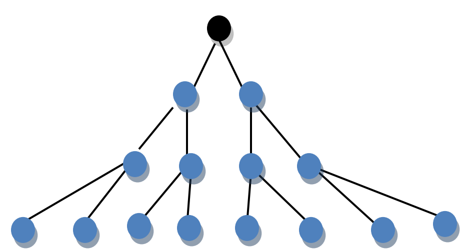

Figure 3.13. K14c Construction.

The construction process is depicted in Figure with the original parent node in black and the children in blue. The result was that each original (parent) node was given 14 children. The children are comprised of 8 leaf nodes and 6 inner nodes

The K196c100.3 puzzle was made by taking each node of the K14c100.3 puzzle and attaching two children to it, two children to those children, and two children to each of those children. The result is that each original node of the K100.3 puzzle was given 196 children connected by a tree structure. Each structure has 64 leaf nodes and 132 inner nodes.

The test problems are summarized in Table 3.8.

The ACO and TabuACO will have an objective function which takes node #1 in the graph as the root and determines the optimal path (minimum spanning tree) to each of the prizes in the forest. The score for each tree is computed as the value of the prize minus the cost of traveling over the minimum spanning tree to obtain that prize. The score for the forest is computed as the sum of the scores for each tree.

Table 3.8. Benchmark Steiner Forest Problems

Name	Nodes	Leaf Nodes	Prizes	Edges	Average Degree	Tree Content
K100.1	42	0	9	185	4.4	0
K100.2	24	0	8	83	3.5	0
K100.3	26	0	8	123	4.7	0
K1c100.1	84	42	18	227	2.7	50%
K1c100.2	42	24	16	107	2.5	50%
K1c100.3	52	26	16	149	2.9	50%
K14c100.2	360	208	16	419	1.2	93%
K196c100.2	5400	1664	16	5459	1.0	99%

The cost of each tree is evaluated as the value of the prize minus the cost of the path from the root to the prize bearing node. The algorithm will run for 5,000 sorties or until the forest is solved. Optimal solutions were not available at the libraries and had to be determined through experimentation.

To perform the test, the performance of a standard ACO with positive (only) pheromones will be compared to the performance of the TabuACO with positive and repulsive pheromones.

3.4.6. Data. Each solver was allowed to run a fixed number of sorties, and the solution compared to the best-known solution. The results are summarized in Table 3.9.

Table 3.9. Benchmark Steiner Forest Results for K100, K1c100, and K14c100 series

Puzzle Name	Best Score	Sorties	Percentage above best-known score	
			ACO	TabuACO
K100.1	-558838	5000	3	3
K100.2	-339296	5000	1	1
K100.3	-763735	5000	0.6	0.7
K1c100.1	-1137638	5000	46	45
K1c100.2	-687461	5000	2.5	2.1
K1c100.3	-1543180	1000	3	4
K1c100.3	-1543180	5000	1.4	1.4
K14c100.2	-729491	1000	14	13
K14c100.2	-729491	5000	10	11

Results were averaged over 10 trials. It was observed that no repulsive pheromone was deposited during the K100 trials. Noticeable amounts of repulsive pheromone were deposited on leaf pendant edges during the K14c100 trials. Analysis of the underlying data finds no statistically significant difference between the performance of the

TabuACO compared to the ACO for the K100, K1c100, and K14c series. Analysis of the K196c100.2 puzzle found that the 1000 sortie test was insufficient time to allow either solver to complete, but the TabuACO did outperform the reference ACO. Rather than total the scores and somehow accommodate a penalty for missing trees, it seemed best to change the metric and simply count the number of trees solved in the allotted time. Using this metric, all of the previous tests performed equally well for both solvers.

Table 3.10. Steiner Forest Comparison for K196c100.2

Puzzle Name	Prizes in Forest	Sorties Permitted	Average Number of Prizes Located	
			ACO	TabuACO
K196c100.2	15	1000	10	14
K196c100.2	15	5000	13	15

Results in Table 3.10 were averaged over 10 trials. Analysis of the data[6] finds a statistically significant difference. The TabuACO outperformed the reference ACO for the K196c100.2 puzzle.

3.4.7. Discussion of Results. A graph of a Steiner Forest can take on any shape. In the experiment above, a series of graphs resembling a meshed network were tested (the K100 series). The topology of these graphs prevented repulsive pheromone from being deposited, and thus no difference was observed between the performance of the two solvers.

[6] Using the Student's T-test, values greater than 4 were found for t, while a value of 1.8 or greater was needed to conclude the result is significant with 95% certainty.

With the belief that leaf edges were necessary to jump start repulsion, the K1c100 series was created, and prizes copied to the children nodes. This puzzle however was trivial for either solver to complete and no difference in performance was measurable.

Continuing with the belief that a Steiner Graph containing a significant amount of tree topology would perform differently than graphs with mesh topology, the K14c100.2 puzzle was created. The data showed that significant amounts of repulsive pheromone could be observed in the environment. However, no significant difference in performance was measured. The puzzle, with prizes placed downstream of existing prizes, was apparently fairly trivial for the ACOs to solve.

In the final series of tests, each original node of the K100.2 was given 196 children in a tree structure, and 7 prizes were randomly placed throughout the tree portion of the graph. The original 8 prizes in the mesh portion of the graph were retained. Ants had to navigate through the mesh to reach prizes in the trees. This puzzle proved to be challenging enough, and of a suitable topology, for a difference in performance to be noticed. It is believed that ants could easily get lost foraging in the graph. Without repulsive pheromone, they could waste a lot of time researching areas that have already been searched. In comparing puzzles, Table 3.8 reveals that when the average degree in the graph is approximately 1.0, or when the percentage of tree content relative to total graph content is approximately 99%, we see a significant difference in performance of the TabuACO over the reference ACO.

3.4.7.1. Conditions within the model. There are a number of conditions which can affect the performance of the solvers. Once a problem has been rendered as a model, the TabuACO attempts to take advantage of the model topology to reduce the search space the solver must examine. A number of topological conditions can influence the performance of the TabuACO. These are summarized in Table 3.11.

Table 3.11. Summary of Topological Analysis

Topological Condition	TabuACO offers advantage over ACO	Discussion
Puzzle lacks any leaf nodes	No	The lack of leaf nodes prevents deprecation from occurring
Puzzle contains a small number of leaf nodes	No discernable advantage	The amount of search space that can be trimmed is small in proportion to the total search space.
The puzzle is NP hard and could take years of computing time to solve.	No discernable advantage	The hardness of the puzzle prevents significant penetration into the search space and prevents any significant buildup of pheromone from occurring.
Puzzle contains a mixture of tree and mesh topologies	Possible	The presence of leaf nodes will allow repulsive pheromone to be deposited. The amount of territory which can be trimmed is generally a function of the number of nodes with degree = 1. If these are strategically located, or a large portion of the node population, a performance difference can be expected.
Puzzle is modeled as a tree structure	Very Likely	Leaf nodes allow deprecation to occur and allow the ant to note that she has travelled a given dead-end path before.

3.4.7.2. Conditions within the solver. When the graph appears to contain the conditions necessary for success, improper tuning of the solver can still prevent proper results. Usually, best performance is obtained when repulsive pheromone deposition rates are high, and repulsive evaporation rates are low. In fact, if the “food” isn’t allowed to move around, and if a pathway identified as an empty leaf node will continue to remain empty until the puzzle is solved, then there is no need to evaporate repulsive pheromone at all.

Repulsive pheromone can accumulate to a high level and remain high until the solver converges on a solution.

3.5. RESEARCH FINDINGS

Results from the Steiner Tree Problem and the Quadratic Assignment Problem show that the TabuACO (having both attractive and repulsive pheromones) can outperform the reference ACO (having only attractive pheromones). The challenge is finding a problem of the right size with the right topology. A performance advantage is found on problems which are modeled as graphs with a tree topology. Likewise graphs which are 99% tree shaped demonstrate a clear advantage. As the number of tree pathways diminish, and the number of leaf nodes decline, the advantage to the TabuACO declines as well. The TabuACO offers no discernable advantage for graphs which have very little tree topology. A graph which contains 100% mesh topology and 0% tree topology offers no advantage. Furthermore, no advantage is found for problems (such as the TSP) in which every node in the model must also appear in the solution. An advantage is only found in graphs where nodes in the model can be temporarily removed from the search space because they do not contribute to the solution.

4. THE APPLICATION OF THE TABU ANT COLONY OPTIMIZER TO TRANSACTIVE ENERGY MARKETS

4.1. THE "SMART GRID"

Many systems are driven to become "ultra quality" systems. They are challenged to improve the quality of what they provide while managing cost. Furthermore, many successful systems, after they are well established, find that they must add new functionality, or in some cases, be repurposed to perform tasks that they were not originally designed for. We find all of these paradigms to be true of the Smart Grid effort. In it, the power grid is tasked to work in ways in which it was not originally designed, to improve the quality of the service provided, to provide new features not currently available, and at the same time to manage costs.

The incorporation of new sources of power, and the increased flow of information are an important part of this transformation. There are many aspects of the grid that warrant improving. In the US, the Energy Independence and Security Act (EISA) of 2007 characterized the "Smart Grid" in part as:

1.) Increased use of digital information and controls technology to improve reliability, security, and efficiency of the electric grid.
2.) Dynamic optimization of grid operations and resources, with full cyber-security.
3.) Deployment and integration of distributed resources and generation, including renewable resources.
4.) Development and incorporation of demand response, demand-side resources, and energy-efficiency resources.
5.) Deployment of ''smart'' technologies (real-time, automated, interactive technologies that optimize the physical operation of appliances and consumer

devices) for metering, communications concerning grid operations and status, and distribution automation.

6.) Integration of ‘‘smart’’ appliances and consumer devices.

7.) Deployment and integration of advanced electricity storage and peak-shaving technologies, including plug-in electric and hybrid electric vehicles, and thermal-storage air conditioning.

8.) Provision to consumers of timely information and control options.

9.) Development of standards for communication and interoperability of appliances and equipment connected to the electric grid, including the infrastructure serving the grid.

10.) Identification and lowering of unreasonable or unnecessary barriers to adoption of smart grid technologies, practices, and services [43].

Item #3 on the list underscores the importance of renewable energy. Indeed, since this act of Congress, many states and regions around the world have decided to vigorously pursue the adoption of renewable energy. This new direction in the industry is causing a paradigm shift which cannot be supported very well by the current system, and is the subject of the application section of this dissertation.

4.2. SMART GRID DOMAINS

The classic teaching by Adam Smith [44] is that when a free market is allowed to operate, an “invisible hand” can effectively guide a very large system (the economy) to an equitable outcome. This is an interesting choice of words today because it says that the economy operates as a self organizing system.

Electricity, when viewed as a commodity, can be scarce or abundant. It can be easy or difficult to deliver. Local prices can change customer behavior to affect the local demand for electricity. The energy market is used every day to solve problems and match generation to forecasted load. But many markets also employ a flat rate. This can have the effect of hiding problems with energy availability. This lack of market functionality has the effect of preventing economics from having its desired effect.

Analysis by National Institute of Standards and Technology (NIST, a division of the US Dept. of Commerce) defines seven domains in the Smart Grid Conceptual Model [45]. These are depicted in Figure 4.1.

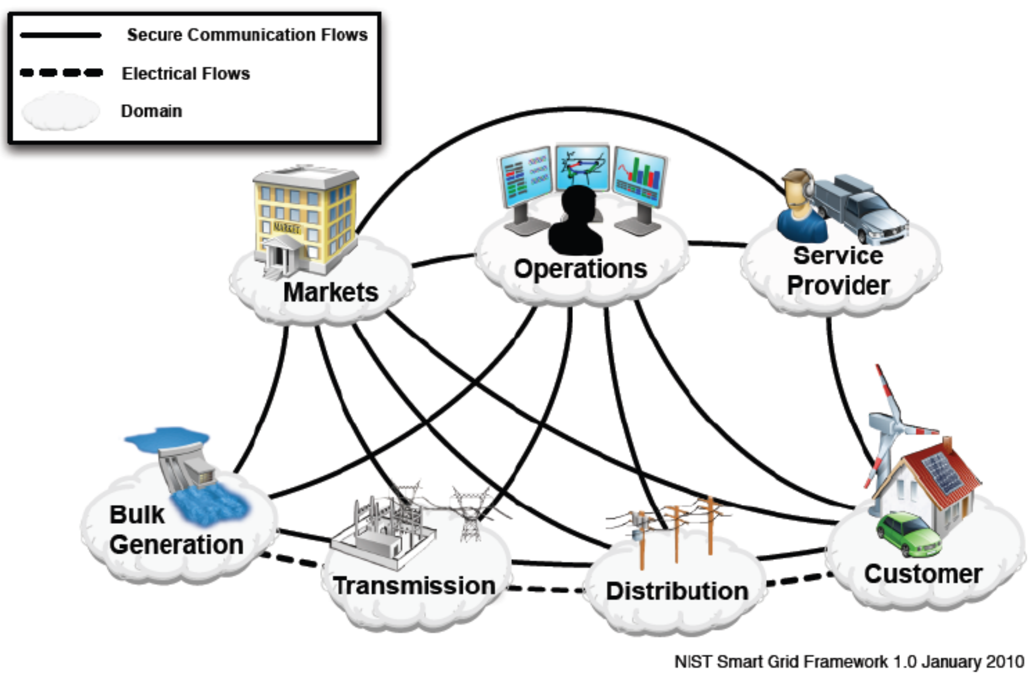

Figure 4.1. The Seven NIST Smart Grid Domains.

- Energy Markets
- Network Operations
- Service Provider
- Bulk Generation
- Transmission
- Distribution
- Customer

4.2.1. The Market Control Loop. Energy is bought and sold through contracts within regional organizations for the energy to be transmitted throughout the United States. This centralized control mechanism has worked well with a base load electrical power paradigm, but as more distributed forms of electric power generation are introduced, the power supplied to the grid becomes more variable.

The current wholesale electrical energy market for the electrical power grid is operated much like a control system. The energy market allows, enables, or contractually requires generators to contribute power at specific times and locations on the grid. This is done through the formation of contracts between energy buyers (who purchase large blocks of power on behalf of consumers), energy producers (who operate utility-scale generation plants), and owners of transmission systems (who are hired to move the power from one location to another over long distances). This market process makes it possible for network operators to control the grid in near real-time.

Network operations can be viewed as a fast acting “inner control loop” which manages the grid in near real-time, while the energy market is a slow acting “outer control loop” that establishes the flow of power (Figure 4.2). Network operators manage the flow of power by continually monitoring the draw of loads, the output of generators, and by making adjustments to match the level of generation to meet the load.

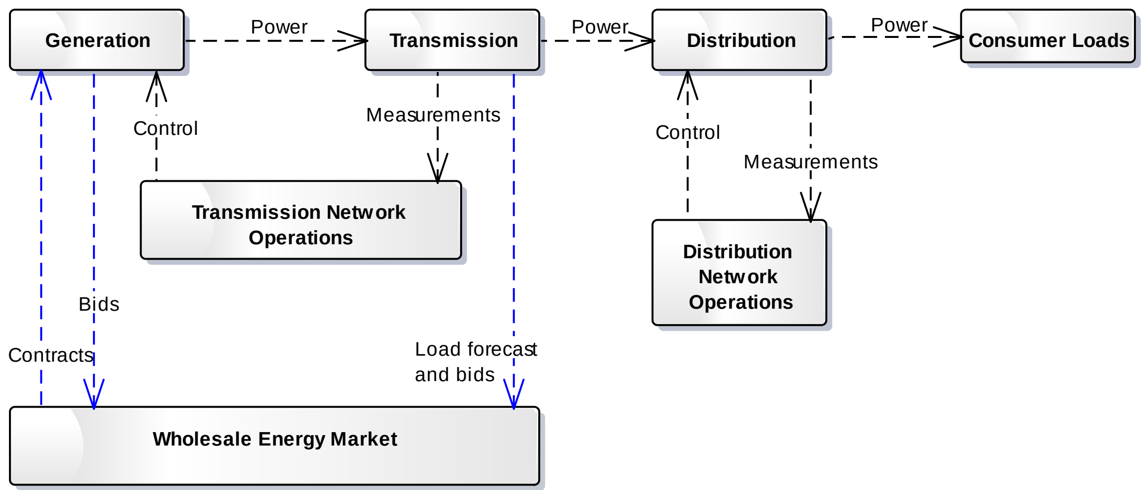

Figure 4.2. The Energy Market and Energy Management as Outer and Inner Control Loops

The wholesale energy market is focused on the generation and transmission of electrical power and does not account for distribution and consumer load. While accounted for in the transmission circuit model, consumer loads and distribution circuits are commonly aggregated and represented as a single load. This has served well in the traditional grid comprised of centralized generation and single owner transmission and distribution assets [46] [47] [48] [49] [50], but new technologies are moving the electric grid away from single party ownership of all generation, transmission, and distribution assets.

Many regions have goals of transitioning significant amounts of power from fossil-fuel powered sources to distributed energy resources (DERs) [51] [52] [53]. The introduction of these new power sources challenges both the operation of the grid as well as the traditional energy marketplace [54] [55]. The shift is most often being done by adding significant amounts of wind and photovoltaic (PV) power into the distribution grid at both utility and consumer levels [56]. As the use of these DERs expands, it becomes increasingly difficult to regulate the voltage and frequency of power placed on the grid [57]. In addition to causing an overall voltage rise [58], increased use of PV sources increase the possibility of rapid voltages changes on the grid due to changes in weather than cannot be compensated for [59]. The market itself complicates the ability of network operators to manage the grid because some tariffs in some regions require that operators “must take” renewable power when it is available [60] and many DER providers are too small to participate in the existing bulk power markets but have an impact on the grid [61] [62] [63] [64]. There are a number of proposals for modernizing the grid as well as the energy market [64] [65] [66] [67] [68]. While these are fine ideas, [64] [65] [66] [67] at the present we have not demonstrated an ability to have grid assets that self-organize in such a way that consumers of all sizes participate in determining the market price for every interval of time. The proposal found in [68] is interesting in that it is one of the few to mention self-organization. The proposal [68] could apply to large or small net zero market participants, but its focus is the optimization of Combined Heat and Power (CHP). A building outfitted with optimized CHP would accept prices from the market and decide when to draw or contribute power to the grid. Unfortunately, this proposal doesn’t explain how this process might scale to involve all market participants. The introduction of a transactive energy market provides a means of decentralizing control and allowing for self-organization of the grid assets.

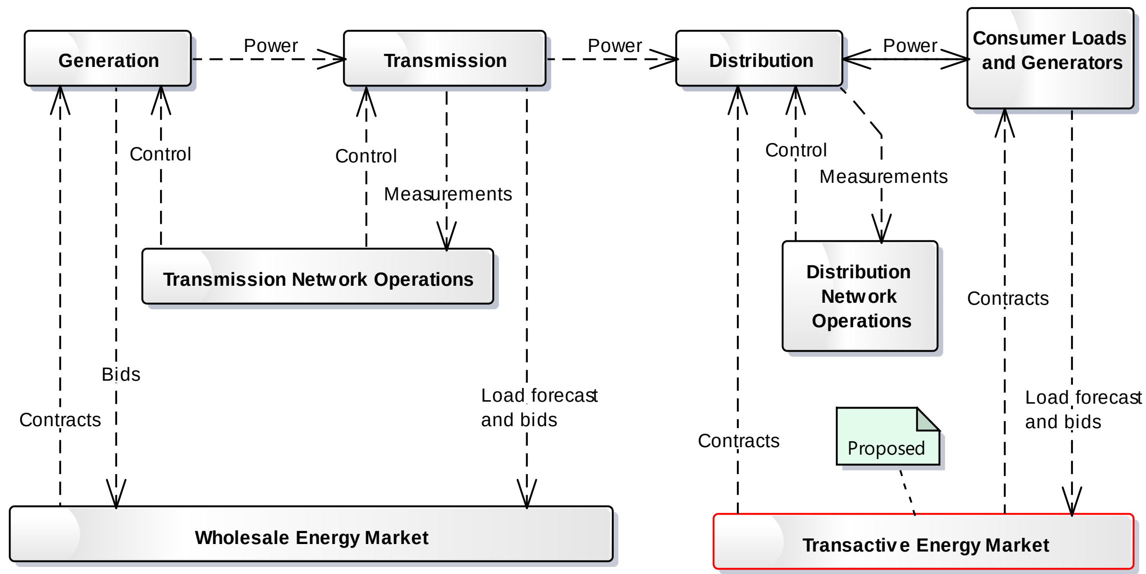

Figure 4.3. A Transactive Energy Market Added to Perform Distribution Aggregation

Transactive energy markets can be described as "an internet-enabled free market, where customer devices and grid systems can barter over the proper way to solve their mutual problems, and settle on the proper price for their services, in close to real time" [69]. In this paper we propose a method of implementing a self-organizing transactive energy market which will operate at the distribution circuit level. Each distribution circuit is represented as a single component in a system which models a distribution network, but each distribution circuit typically serves thousands of service locations. Scaling the current wholesale market to include every conceivable participant has been deemed by many to be impractical. Instead, some have called for a transactive energy paradigm to either supplement or potentially replace the legacy system [70] [71] [63]. Self-organizing systems have been shown to alleviate communication problems through methods such as limited stigmergy [72]. An ant colony optimization method is implemented to create a decentralized control scheme.

4.2.1.1. Timescales. Figure 4.3 describes various “control loops” which operate at different timescales. There is of course the fast-acting control loop that everyone thinks of – the real time SCADA loop used by network operators to monitors loads and adjust generation. There are however many additional “control loops” outside this loop. Before electricity is generated, it must be planned (via the energy market function). Before a generator can sell electricity on the market, it must come online.

Table 4.1. Electric Utility Function Times Relative to the Flow of Power

Time relative to actual flow of power	Function
Years in advance	Planning and construction - Transmission adequacy - Generation adequacy
Months in advance	Maintenance
Hours in advance	Energy Market - Load forecast - Unit commitment - Congestion - T&D efficiency
Real time	Operations - Energy Management -- SCADA -- Load following -- Demand Response - Voltage regulation - System stability - Outage detection, crew dispatch, and FLISR
Moments afterward	Billing and account management

Before a generator can come online, it must be built. Before it can be built it has to be permitted. These functions operate at different timescales as shown in Table 4.1.

4.2.1.2. Control challenges imposed by the market-selected fuel mix. The Energy Market will almost always optimize for "economical dispatch." In this process, the least expensive fuel (e.g. nuclear) is selected first, and more expensive forms of generation last. A typical, centrally dispatched fuel mix can be found in Figure 4.4.

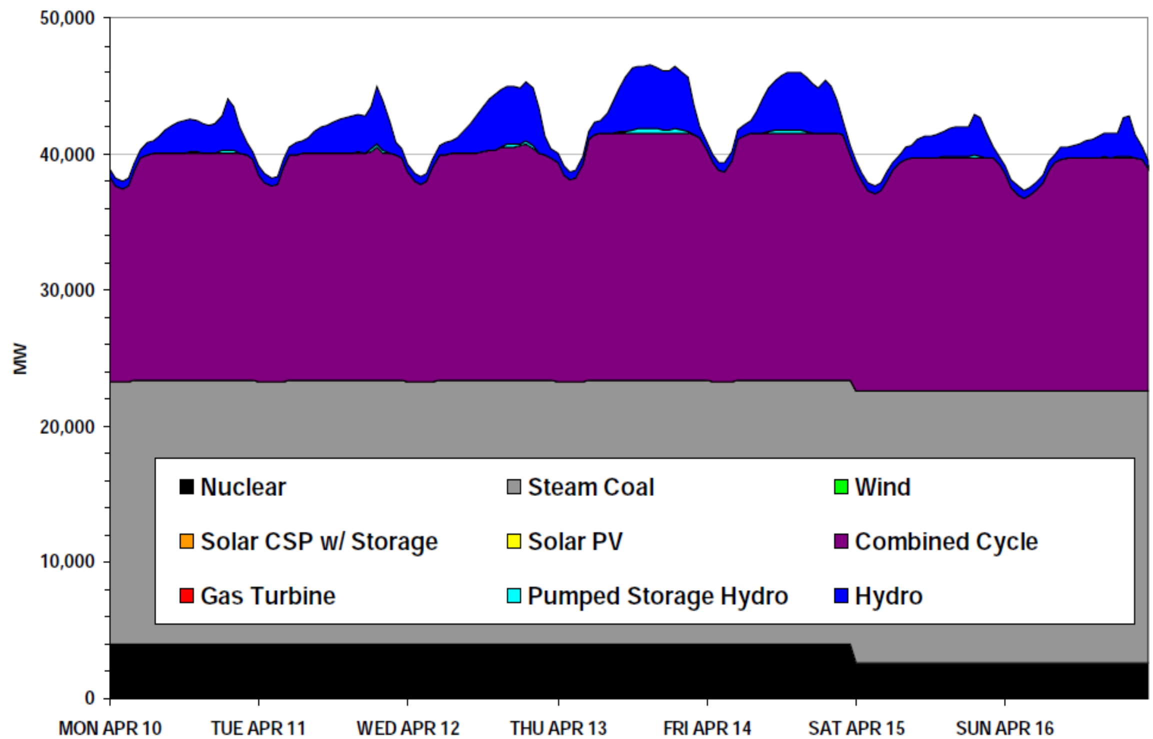

Figure 4.4. Fuel Mix in Study Area With No Wind Conditions (Showing Low Cost Power on Bottom (Nuclear in Black) and Higher Priced Energy Above It. Legend Located in Gray Band.) [73]

In this analysis by the National Renewable Energy Lab [73], we see in Figure 4.4 a system that draws approximately 38,000 MW at night and approximately 46,000 MW during the peak of the day. An entire week is shown starting on a Monday at Midnight. The load is met rather nicely by dispatching sufficient amounts of power from various nuclear, coal, combined cycle, and hydro sources.

In Figure 4.5, a small amount of windpower is introduced. The smooth peaks and valleys that we saw in Figure 4.4 are gone. They are replaced by more jagged edges which only approximate the load on the system. In order to accommodate the variable contribution from the wind, the network operator must make adjustments to the other forms of generation which he can control.

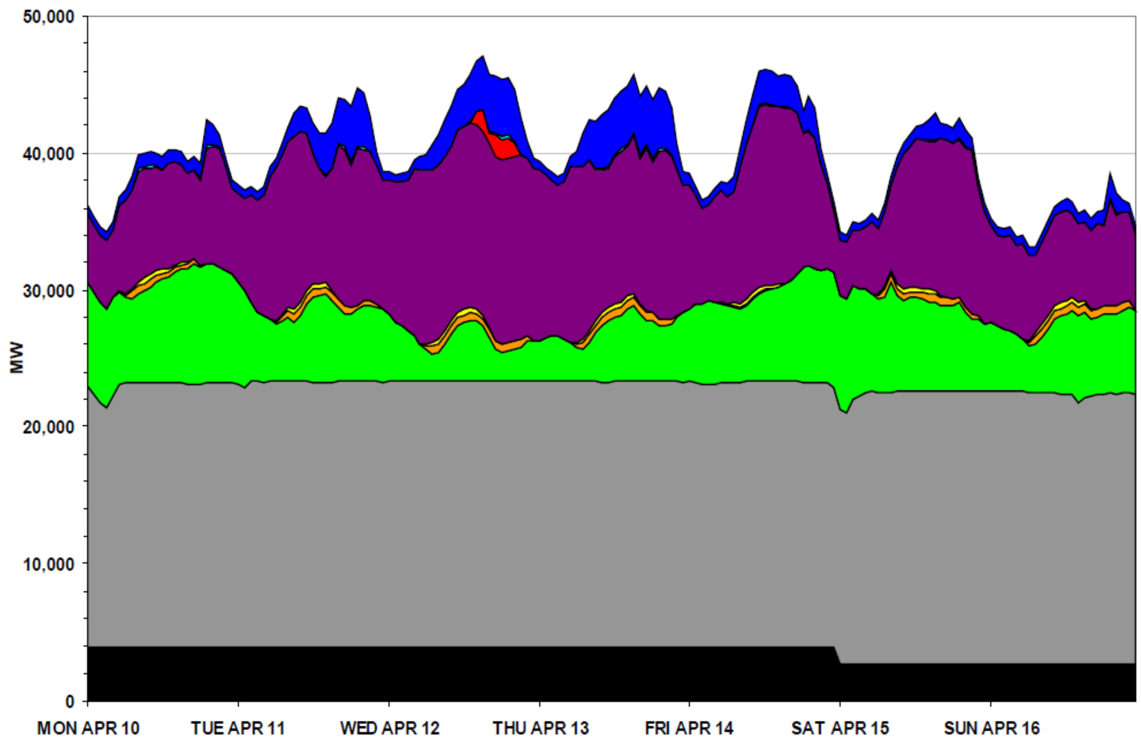

Figure 4.5. Fuel Mix in Study Area with 10% Wind and 1% Solar Contribution [73]

The network operator would endeavor to shed the most expensive forms of power first, and the least expensive forms of power last – all while balancing the amount of generation against the amount of load. In Figure 4.6 we see even more wind power added to the mix at the expense of combined cycle and coal. In Figure 4.7 it reaches the level of a 30% contribution by wind. The smooth, periodic valleys and troughs across the top of Figure 4.4 are all but gone by Figure 4.7.

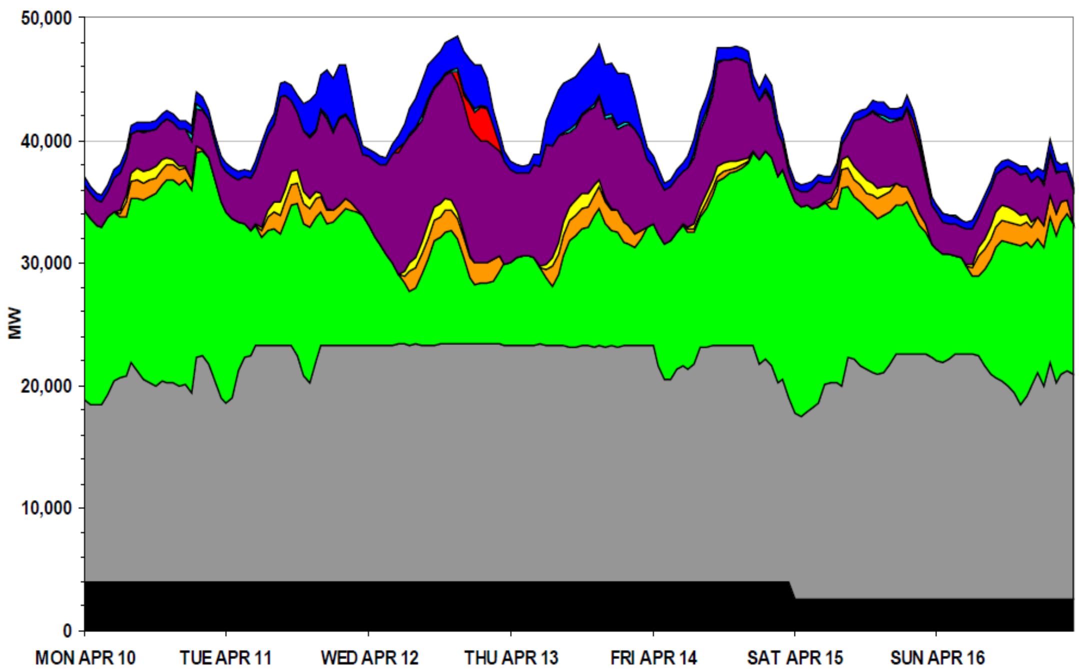

Figure 4.6. Fuel Mix with 20% Wind Penetration and 3% Solar [73]

Additional contribution by wind and solar displaces dispatchable generation (such as coal), and ultimately causes the network operator to lose control of the grid. There are points on the graph in Figure 4.7 (such as Monday at Midnight) where the remaining dispatchable generation (gray and black) is much smaller than the wind contribution (in green).

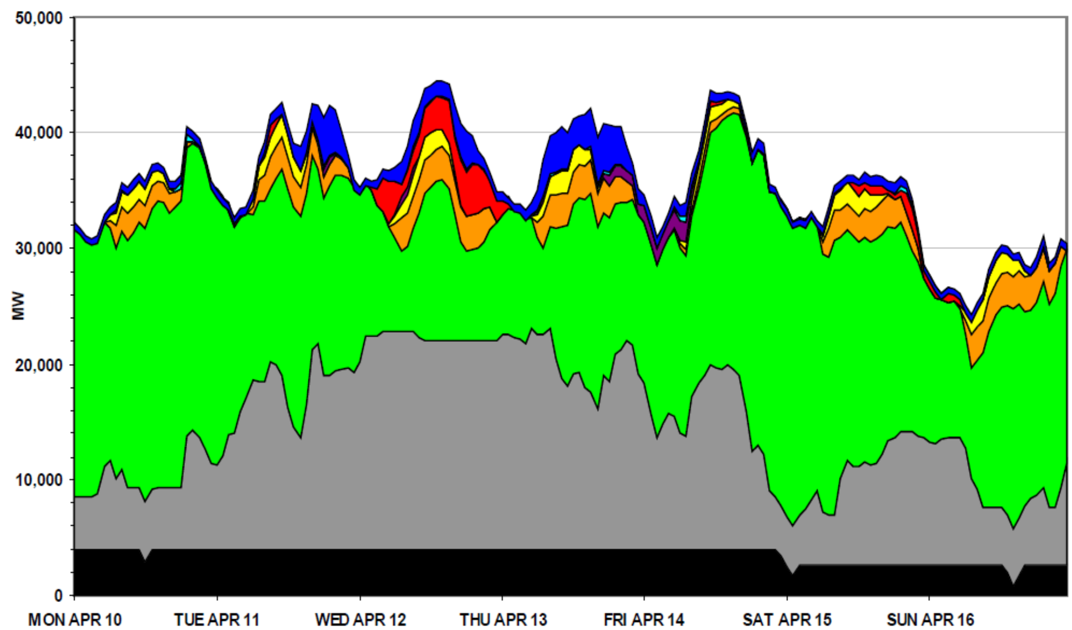

Figure 4.7. Fuel Mix with 30% Wind Penetration and 5% Solar [73]

Given these conditions, and the loss of control margins within the system, a complete and sudden drop in windspeed could make it very challenging to maintain voltage on the grid.

4.2.2. Operations Control Loop. The discussion above described that there are multiple loops which are necessary to bring power online. Different loops control the construction and destruction of power plants; the financial agreements to fuel and spinup the generation, and finally the real-time operation of the power plants and transmission network.

Network operators continually monitor the health of the electrical network. A number of network conditions can occur which could serve as a sign of poor system health:

- Overvoltage
- Undervoltage
- Phase imbalance
- Insufficient real power
- Insufficient reactive power
- Under frequency
- Over frequency

Network operators ordinarily adjust generation levels in an effort to balance the generation against the load. However, some network operators have the option of adjusting the load. Adjustment to the load – usually through some market-based reward for the consumer – is called "demand response."

Federal Energy Regulatory Commission (FERC) identifies 15 or more different activities[7] which fall under the umbrella term "Demand Response." These include:

1.) Direct Load Control
2.) Interruptible Load
3.) Critical Peak Pricing with Control

7 It should be noted that the list of what constitutes "demand response" varies somewhat by organization. The FERC definition includes Ancillary Services under the umbrella term. It should also be noted that the FERC definition of DR, and FERC Order 745 which allows utilities to cut back generation in favor of DR was challenged in court and went all of the way to the US Supreme Court [74].

4.) Load as Capacity Resource

5.) Spinning Reserves (Previously included in Ancillary Services classification)

6.) Non-Spinning Reserves (Previously included in Ancillary Services classification)

7.) Emergency Demand Response

8.) Regulation Service (Previously included in Ancillary Services classification)

9.) Demand Bidding and Buyback

10.) Time-of-Use Pricing

11.) Critical Peak Pricing

12.) Real-Time Pricing

13.) Peak Time Rebate

14.) System Peak Response Transmission Tariff

15.) Other [75]

There are over a dozen different techniques used for DR, but DR calls are generally provided a day in advance – when the energy market determines that there will be a generation shortfall. Sometimes the spot market can identify an hour in advance that there will be a price spike and/or a generation shortfall. Customers who subscribe to a Real Time Price (RTP) program might be informed in "real time" of the price change (though in some experimental programs customers are informed the day after the event [76]).

DR is not a widespread capability at the time of this writing. One form of DR is "Direct Load Control." The largest Direct Load Control program in the world is at FPL (formerly "Florida Power and Light"). With 800,000 devices deployed on customer's water heaters, air conditioners, etc. FPL is able to control over 1,000 MW of power [77]. No other utility has a direct load control program of this magnitude.

This level of control is operated at the wholesale market level. It addresses generation shortfalls, or transmission line issues. The utility does not have real-time visibility to distribution assets, and cannot alleviate the strain on assets even if it knew an issue has developed. For example, if a service transformer were overloaded, contractual limitations, AMI network limitations, customer participation levels, and other constraints would make it so that a solution could not be guaranteed in every case.

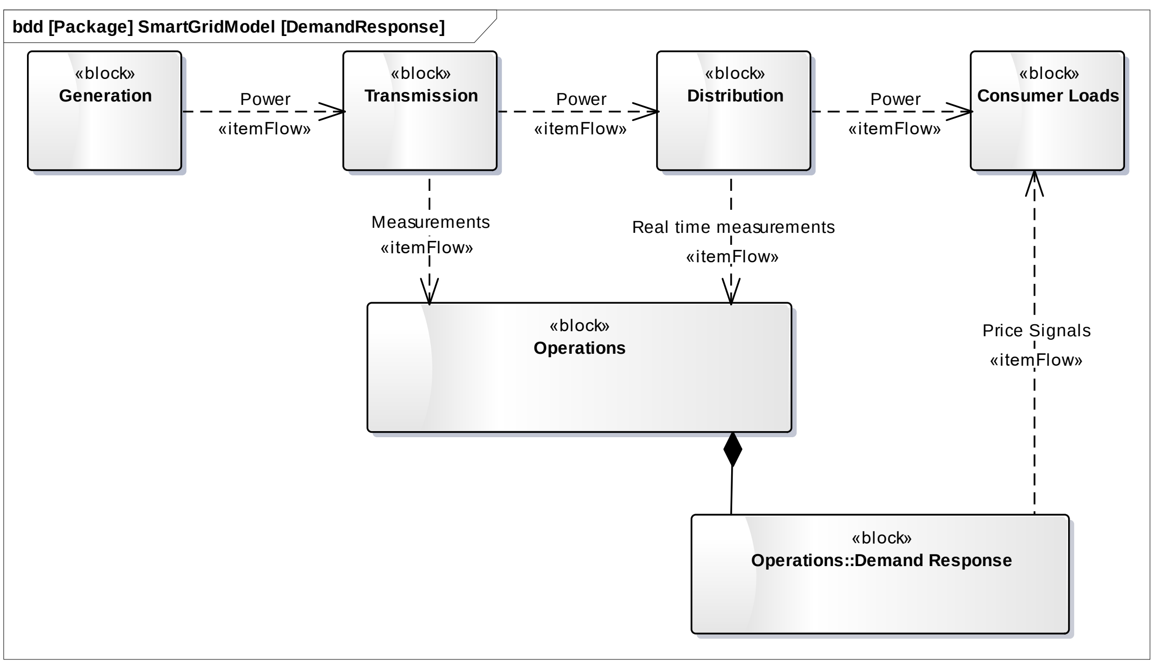

Figure 4.8. Demand Response as an Extension of Operations

Demand Response is called upon by network operators. Although the Energy Market prepares the contracts with the consumer to enable this activity, it should be considered an extension of network operations as described in Figure 4.8.

4.2.3. Power Markets. This section endeavors to explain the various markets at work today, and ultimately show the need for an additional market to cover an uncovered area.

The history of the regulated utility can be traced back to the origins of electrification. When Edison, Westinghouse, and others developed the many technical innovations necessary to operate the electrical grid, Samuel Insull developed rate innovations and promoted monopolies under government regulation. Today we find the grid undergoing considerable change as new pressures push it to be reinvented. The original design assumed that electricity would always flow unidirectionally from a small number of large generators to numerous small consumers. This design is challenged today to deliver power in the face of numerous distributed energy resources, investment in alternative power sources and underinvestment in aging generation and infrastructure. The electric utility industry is very capital intensive. Investments in recent years are 58% in generation, 31% distribution, and 11% into transmission [78].

It is well established in the literature that a properly functioning free market works to the mutual benefit of both buyers and sellers of a product, and that this is a very desirable trait in the energy market [79]. Customers find the best products to meet their objectives, while sellers find the broadest possible market for their goods.

In order to promote competition and free enterprise in the industry, many regions have moved toward deregulated markets. The wholesale power markets in all of these regions allow producers to compete for the available business. Wholesale power markets are a fairly recent innovation. Deregulation actually started in other countries and gained wide appeal in the US around the ‘90s and 2000s. Congress formally created the Federal Energy Regulatory Commission (FERC) in 1978 and gave the FERC the authority over electricity transmission. Over the years, FERC issued a series of rules which provided for gradual deregulation of the power industry.

In 1992, FERC order 636 allowed for “customer choice.” In many regions, the program allows retail consumers to choose their provider [80]. The provider serves as a retail aggregator, and may offer multiple plans to their subscribers. The retail aggregators then in turn participate in the wholesale market on behalf of consumers. The plans offered by the retail aggregators may offer a variety of banking services to the consumer. They may be able to offer a guaranteed rate, and a fixed amount on the bill for a period of one year. While we don’t see it in practice, the technology exists to expose retail customers to real-time prices driven by wholesale market prices. When done properly, and with sufficient granularity, the market can be effective at attracting generation to relieve network congestion [81] [82] [83].

Industry experts observe that the grid is moving from a “fly by wire” type of control for dispatchable generation to a stochastic form of generation. They note that new methods are needed to control such a grid [84]. The rise of “prosumers” (those who sometimes produce and sometimes consume energy) is changing the nature of the grid.

4.2.3.1. The free market. The idea of commerce dates back prior to recorded history. Many countries endeavor to facilitate free market economies. It is well known that a healthy marketplace competition promotes the ability of consumers to find the attributes they desire in a product, and conversely, the ability of producers to find additional markets for their products. Two hundred years ago, Adam Smith in his book The Wealth of Nations described the naturally occurring self-regulation which occurs in a free market economy. Smith described it as if an “invisible hand” guides participants to “promote an end which was no part of (their) intention.” [44] The economy itself is commonly viewed as a self-organizing system [85]. Naturally occurring commerce in a free market allows prices to rise and fall in response to supply and demand. Surplus leads to low prices, which in turn spurs consumption, which in turn eliminates the surplus and brings prices higher. Conversely, shortage leads to high prices, which spurs conservation, which in turn eases the shortage and allows prices to fall.[8]

Despite the effectiveness of a free market economy, the power industry has found that some degree of monopolization is practical. In 1892 Samuel Insull became president of the Chicago Edison Company. He argued quite successfully that given the high cost of generating equipment, and that these power plants were centrally located, utilities should be allowed to operate as a monopoly, free of competition, and be regulated by the state [46]. He reasoned that because the high capital investments required by the utility, competition would not result in lower costs or better service, but in duplication of service

[8] It is interesting to note that many insects, including ants, work individually toward their own ends, but collectively find the behavior of the colony to appear to have directed organization (guided by an “invisible hand”). Ants individually forage for food, just as humans may individually hunt or gather. Ants and humans perform many other supporting roles to sustain their lives. Humans have worked together to form organized societies which obtain or produce the required goods. Smith argues that humans may do so with their own individual ambitions in mind, but collectively, through the marketplace may appear to have an invisible hand which guides the outcome. Ultimately, goods are produced and delivered for a price which is mutually agreed upon as fair.

and higher costs. He warned that there could be too much service in areas of greatest competition, and little or no service in areas where infrastructure has not been built [46]. This wisdom has prevailed since the inception of the grid but is being questioned as generators become more numerous and diverse. The original underlying belief was that large, centralized generation is always more cost effective than smaller distributed generation. This original assumption is no longer always true.

There are those today who promote the decentralization of the electric grid [54]. Some argue that there are numerous thermodynamic inefficiencies in large power systems and argue that there are cost saving benefits of decentralized generation [86].

The electricity market is different from most markets because it is a commodity that cannot be stored in bulk in its raw form. This prevents "market makers"[9] from participating in the market – buying and selling the commodity for a profit. Instead, a number of other approaches have been developed. Utilities may deal directly with consumers, or in some states (such as Texas) a retail layer is introduced between the generation and transmission companies and the consumer. This layer provides customer service to the consumers and aggregates retail purchases to buy power at the wholesale level.

---

[9] Despite the absence of "market makers" in the electricity marketplace, similar opportunities exist if one has the right equipment. With energy storage equipment, it would be possible to buy energy when the price is low, store it, and use it when the price is high. There are a number of technologies which could someday be utilized in this role. Centralized storage can be found today in pumped hydro and compressed air. Smaller scale storage might be found in a fleet of electric vehicles, banks of batteries, or small scale air liquefaction.

4.2.3.2. History. Economic theory starts with the basic premise from Adam Smith that free market exchange is the best mechanism to balance supply and demand at a fair price [44]. The rules which define a fair market mechanism are arrived at by consensus between government and private entities. Every region is allowed to arrive at rules that they feel are appropriate.

The original economies of vertical integration argued originally by Insull [46] have been studied over the years. Economists Coase and Williamson find that vertically integrated utilities are easier to operate, [47] [48] and others find that they cost less to operate [49] [50].

The new challenge for utilities is for them to accept power from a diversity of clean and renewable sources when and where such power is available. In many cases this requires a vertically integrated utility to break up, and separate generation from transmission and distribution. Such utilities must now rely heavily on energy markets to orchestrate the energy needed to supply the projected demand.

As the Smart Grid is developed, the view in the US is that a wholesale market can be operated effectively which represents the needs of the stakeholders, and that DER can be viewed as a load reduction within the customer domain[10]. Small generators are not allowed to participate in the energy market [87]. The view in Europe of the Smart Grid is quite different. DER is considered a separate (eighth) domain. In Europe, owners of small generation can participate in the Energy Market just as large generators [88]. DER is also controllable by Operations just as large generation. These fundamentally different views

[10] Referring now back to Figure 4.1. The Seven NIST Smart Grid Domains.

are important in that significant levels of generation are largely ignored by the US energy market and recognized by the European (as depicted in Figure 4.9).

The European view is to take the four NIST domains that have physical actuality and make them columnar domains in the 3D SGAM view. A fifth domain, “DER” is added to give place to the ability of non-utility owned generation to participate in the market place. The NIST domains “Markets” and “Operations” finds their place as “Zones” in the SGAM design. Europe (in particular Germany) calls for the creation of a “Smart Grid,” but also of “Smart Markets” to allow for the interaction of DER with the Energy Market [89].

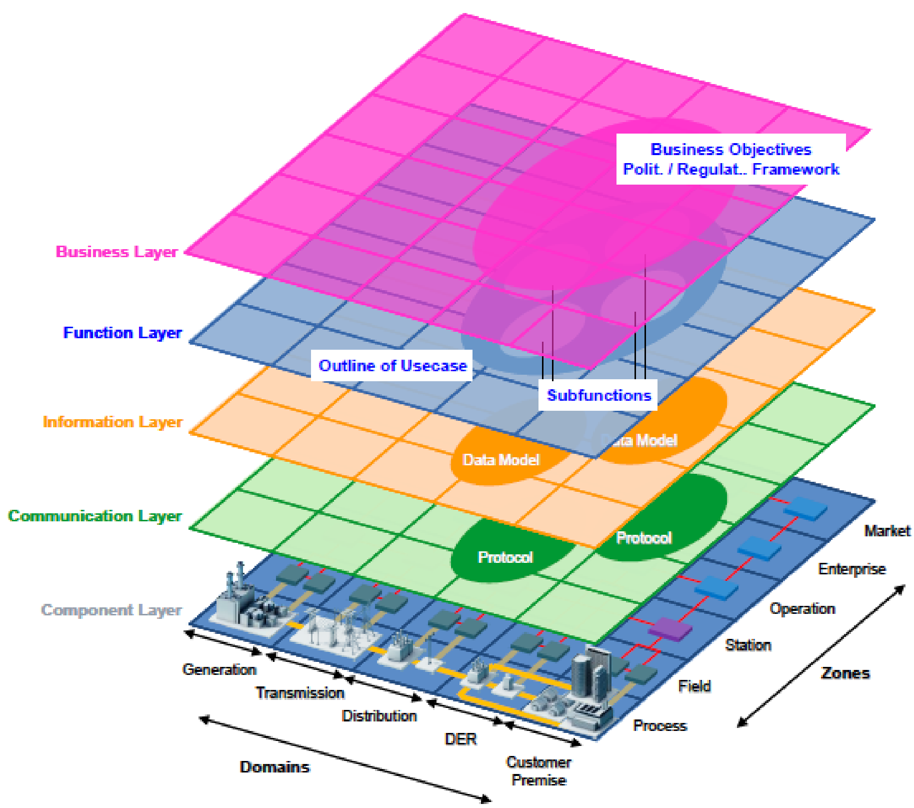

Figure 4.9. (European) Smart Grid Architecture Model (SGAM) [88]

4.2.3.3. The argument for better markets. Different theories exist as to how a market should be operated for the maximum benefit of society [90] [91]. Some believe in shielding consumers from price spikes, while others believe in exposing consumers to such spikes.

In most tariffs and most markets operating today the real cost of power is hidden from consumers. The consumer is instead presented a uniform price regardless of their location and regardless of the spot market shortage or surplus that may exist at the moment. As a result, consumers lack any incentive to respond to market conditions when there is a shortage of electricity or a constraint on the powerline. This lack of information for consumers can create devastating financial strains on utilities as they attempt to mask all of the challenges in creating and delivering power [92].

Some experts have concluded that the biggest problem with the current grid is not technical but regulatory [93] [94]. When customers pay a flat rate regardless of their location it causes those who live near a source of power to subsidize those who live far away. When customers pay a flat rate regardless of their time of use, it causes those who use power off-peak (when wholesale prices are low) to subsidize users who tend to use more power on-peak (when wholesale prices are high). Furthermore, with the deployment of supplemental power sources such as rooftop solar, we find that “the wealthy” are able to reduce their utility bills. This ultimately leaves the less affluent to pay for the infrastructure that serves everyone. While the matter of equity is one for state regulators to debate, many believe that “as a matter of principle, ethical pricing should be cost-based and not create subsidies between consumers” [67].

The industry needs policies that incentivize investment in the appropriate areas [93]. Some believe the way customers are insulated from the real cost of power is a core part of the problem. "Price-responsive demand is essential to realize the benefit of the Smart Grid" [95]. "The average cost-based pricing formula precludes economically accurate price signals from guiding consumption decisions" [64]. When markets are allowed to function, the market itself offers an emergent behavior which guides the required behavior. High fuel prices are passed along to the customer and motivate conservation. Low fuel prices can also be passed along in a deregulated market and allow consumers to enjoy savings. When the real cost of power is scaled out from a zonal market to a nodal market, it promotes market based congestion management [79]. "Market-based determinations of capacity needs (have been found to be) superior to even the best guesses of regulators" [96].

4.2.3.4. US wholesale markets. North America is divided into separate regions for reliability regions called "interconnects." In the lower 48 states, the interconnects[11] join and trade energy in three sections: The Western interconnect, the Texas interconnect, and the Eastern interconnect [97].

These regions are overlaid with reliability councils, and also subdivided into independent systems (Independent System Operators or Regional Transmission Operators). The ISOs/RTOs are depicted in Figure 4.10. These organizations help move power and call upon spinning reserve to balance load.

[11] The term "interconnect" is used to say that all generators and loads within this area are interconnected.

Power producers and very large consumers of energy (such as entire utilities) participate in market operations within power pools located in these operating regions. It is possible for buyers and sellers to form contracts between each other within their power pool. A power producer may participate in a distant market through the use of agreements with owners of transmission lines.

Figure 4.10. US Independent System Operators (RTOs and ISOs) [98]

Transmission capacity is a very real constraint. Even though the “Eastern Interconnect” stretches from Kansas to the East coast, it does not mean that unlimited power may freely flow between all points as needed within the interconnect. Shortages can occur within an interconnect and even within an ISO. Shortages result in higher prices.

Figure 4.11 shows that different parts of the country can pay more for energy than other regions (even in a fluid spot market).

Figure 4.12 captures a moment of time within the Midwest ISO in which a price fluctuation occurred within the ISO. There is a spot market for energy, and corresponding prices. But there is also a day-ahead energy market. When proper planning can occur, regions tend to settle on day-ahead prices which are less (per MW) than the spot market. (Such regional pricing is depicted in Figure 4.12).

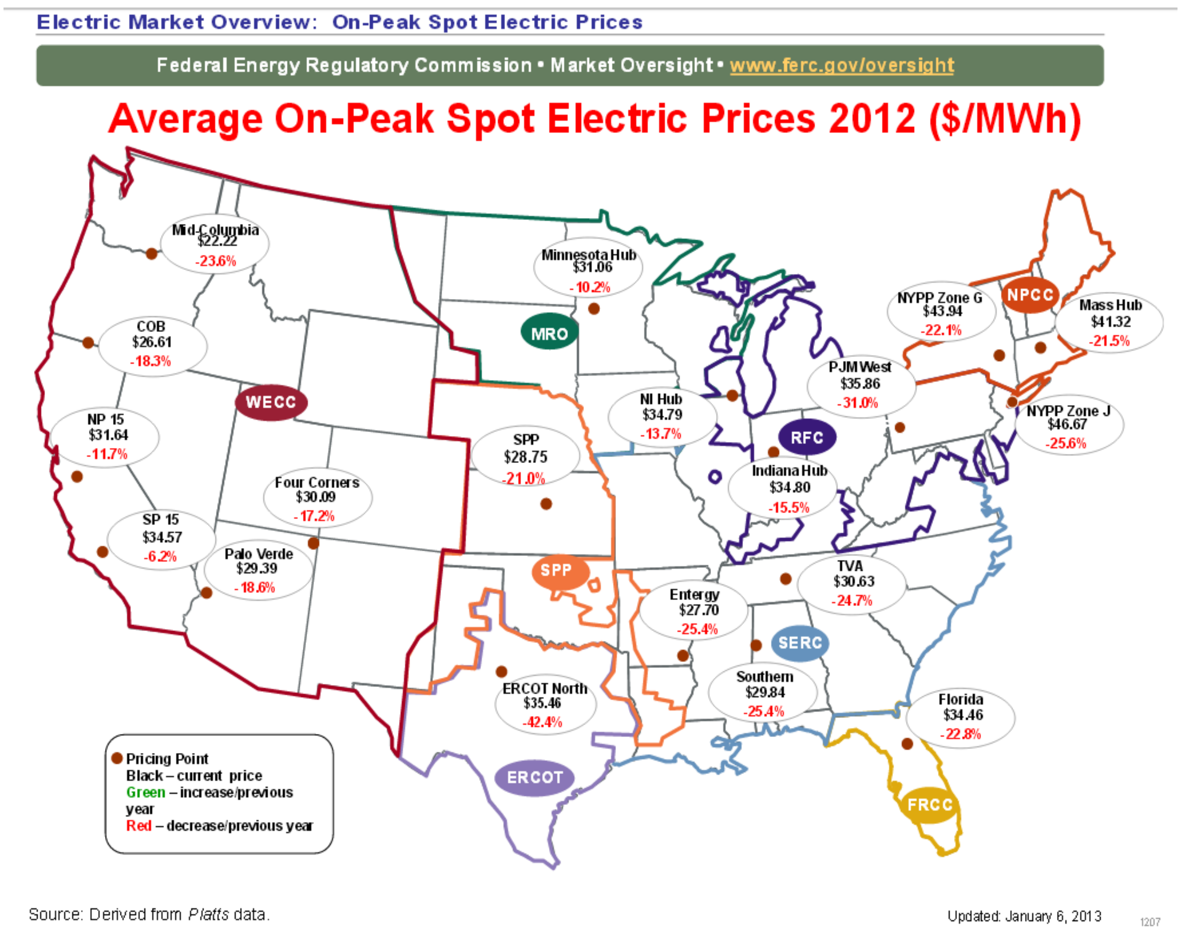

Figure 4.11. Example Spot Pricing Developed Among and Within Organizations in the USA [99]

All of the wholesale markets in the US have transitioned from a "zonal market" to a "nodal market" [101] [102]. This creates much finer granularity in the analysis and improvements in the ability to address localized shortages and surpluses of power. For example, transitioning to a nodal system in Texas is estimated to cost between $530-660 million, yet will save customers an estimated $5.6 billion in the first 10 years of operation [103].

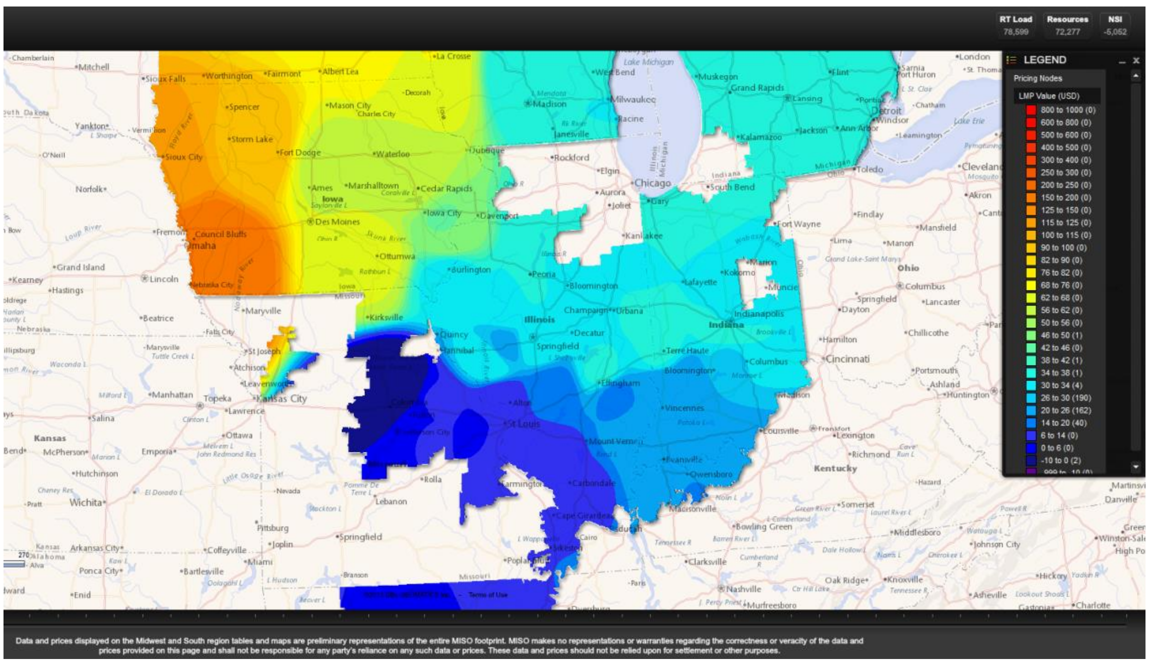

Figure 4.12. Example Price Map at MISO Generated for Dec 20, 2017 at 23:40 Showing Parts of Missouri, Iowa, and Illinois [100]

The wholesale market drives the movement of power from large ("centralized") generators to distribution networks (as shown in Figure 4.1). SCADA instruments in the distribution substation allow network operators to see how much power flows through the substation, but once it is in the distribution network visibility to real-time power flow is generally not available. Temporary shortages and surpluses within a given distribution network are generally hard to find, and there is no “retail energy market” market mechanism to correct these differences. All of the costs are “internal” and owned by a single entity.

Distribution networks are designed to be worry free. Power engineers who design the distribution network will size the substation transformers based on the power needed downstream. They will ensure that the conductors that carry power from the substation to the service transformer are adequately sized. (Customers select a service capacity and thus directly influence the size of the service transformer.) Everything is carefully engineered to deliver the power the customers say they need when they say they need it. However, distribution networks are big and slow to change. A growth spurt of new construction in one area can outstrip a substation transformer’s ability to deliver power to that area. Infrastructure can suffer weathering or damage and lose capacity. New and unexpected types of loads (such as electric vehicles) can catch a utility off guard and unable to serve the accumulated demand within a service territory. Customer-sided generation (such as wind and solar) is outside the utility’s control. Customers can contribute power to the grid (or not contribute) in ways the utility energy purchaser cannot rely on or control. All of these can create localized shortages of energy, yet these shortages are not reflected in the cost of power. Residential consumers tend to pay a flat rate, regardless of their time of use, and regardless of the distance energy must travel to reach their load.

4.2.3.5. Day ahead and hour ahead markets. All dispatchable generation is planned based on a load forecast, which in turn, is based on the weather forecast. In order to accurately gauge the need for power, these forecasts are refined over time. The "next day" weather forecast tends to be far more accurate than the forecast for the next week. The forecast for the "next hour" tends to be more accurate than the forecast for the next day. These short-term markets tend to reflect the real cost of generating power.

For the most part, all residential consumers pay the same price per kWh at every time of day regardless of the availability and price of power. By insulating the customer from the "pain" that the utility feels, the customer is prevented from participating in the solution.

4.2.3.6. Large generation. The markets operate by first determining the amount of load that is expected to be required in each region. The load forecast is largely based on the weather forecast. In the wholesale energy market, buyers meet sellers armed with this information. The load requirements generally outstrip the ability of any single power plant to provide. The market thus identifies the lowest cost provider and accepts that bid. More power is usually needed, so the next lowest cost bid is accepted, and so on until sufficient generation is contracted to meet the forecasted load and ancillary service requirements. Large generation is usually able to offer power at a lower price than small generation. In part, large generators can use less expensive fuels than small generators (such as U235 or coal). But these plants usually also require that a significant portion of the plant capacity be utilized, and that it is not operated in a dynamic fashion that requires numerous ramp ups and ramp downs in reaction to a wildly varying load.

4.2.3.7. Retail markets. The cost of electricity consists of energy costs and network costs [104]. Ordinarily a utility will offer a variety of tariffs to retail customers. All of the tariffs are different yet all of them are considered fair. The most commonly used tariff employs a flat rate. The flat rate charges consumers the same price for energy regardless of the time or amount used. Residential customers have very predictable usage patterns. Flat rates are easily metered by a simple watthour meter.

4.2.3.8. Real time pricing. Opinions as to the ideal and most equitable pricing mechanism continue to be a source of vigorous debate. Ahmad Faruqui writes:

> Flat rate pricing, which has been in place for the past century creates an enormous subsidy between customers with varying load shapes. It is unethical and needs to be replaced by dynamic pricing. Not only will this be more ethical, it will also improve the economics of the power system and lower costs for all customers [67].

Efforts to develop time-based rates (in particular Time Of Use based rates) date back to the original rollout of the grid. Improvements in technology have allowed the consideration of other rates which are more reflective of wholesale market activity. A landmark paper proposing “Dynamic Pricing” was written by Vickrey in 1971 [66].

Since then, there have been pilot programs to expose retail customers to real-time prices [105]. These have been well received in concept, but slow to find implementation. A rare example of a retail RTP program can be found at Elevate Energy, Chicago Illinois USA [76]. Residential customers participate and use either the price from the wholesale market from the G&T that supplies the electricity. One G&T posts the day-ahead price from the wholesale market. Another G&T uses a dynamic spot market price which is adjusted throughout the day. Customers can learn about the real time price by:

- Visiting a website
- Text message alerts
- Email alerts

They can call the utility on the phone and talk to a person and ask them about the price forecast.

Customers settle up after the fact with the price they paid during the market interval during which they consumed the power.

The barriers to implementation appear to be political rather than technical [67] [106].

In terms of raw power, Georgia Power (Southern Company) operates the largest RTP market in the US. It was felt that opening up the pool of all available power to all market participants was the fairest mechanism [107]. However, to participate in this RTP market, a customer has to be 1 MW or larger. This limits the number of participants to 80.

However, one could question the fairness of real time prices which were based on a criterion the customer didn’t want. A better approach would be to develop a “transactive energy market” which optimizes along the criterion (or criteria) specified by the consumer.

4.2.3.9. Transactive energy. “Transactive energy” is the name given to the notion that numerous small customers can effectively participate in the wholesale market to buy and sell energy by contracting to activate customer owned DER. Efforts (such as the OASIS Energy Market Information Exchange) have been made to develop protocols for transactive energy [108]. In theory, this would allow a customer to buy DER power from a neighbor. Unfortunately, the OASIS eMIX standard doesn't identify a means to communicate these messages in real time to the participants [109].

The GridWise Architecture Council has developed ideas similar to those proposed here [110]. The tenets proposed by the GWAC include:

- "The information age will revolutionize the way in which energy systems work today.
- Intelligence (better information supports better decision making) will invade all levels of the energy systems from generation, to bulk transmission, local distribution, and residential, commercial, and industrial consumption.
- Value is best judged in a fair market environment open to participation and review by regulatory authority. This should be exposed at all levels of the system.
- Transparency of value allows market participants to develop and deploy economical solutions that cross traditional enterprise and regulatory boundaries.
- The system will evolve from its present form of operation, through a series of tractable changes over time. Changes include organizational boundaries (ownership and operational responsibility), technology deployment, and forms of collaboration between system components.
- Collaboration based upon autonomous decision making enhances the resilience of complex systems to system-wide failure and accommodates evolutionary changes" [65].

Technology minded customers would know that each of the alternative energy sources have environmental issues of their own.

4.2.3.10. New challenges. It is not always possible to upgrade everything at any time. Not all utilities have the funds to undertake significant construction projects. Yet, a significant upgrade may be required if there is a rapid adoption of Electric Vehicles (EVs). A study by Bloomberg forecasts that the cost of EVs will drop to match the cost of comparable gasoline powered vehicles in the year 2020 [111]. When this happens, it could prove to be a tipping point in the market and a rapid shift to the adoption of EVs. A rapid charging EV can draw more power than the power drawn by the house it is parked next to. However, if properly managed, EVs could increase electric sales overall by 25% [112].

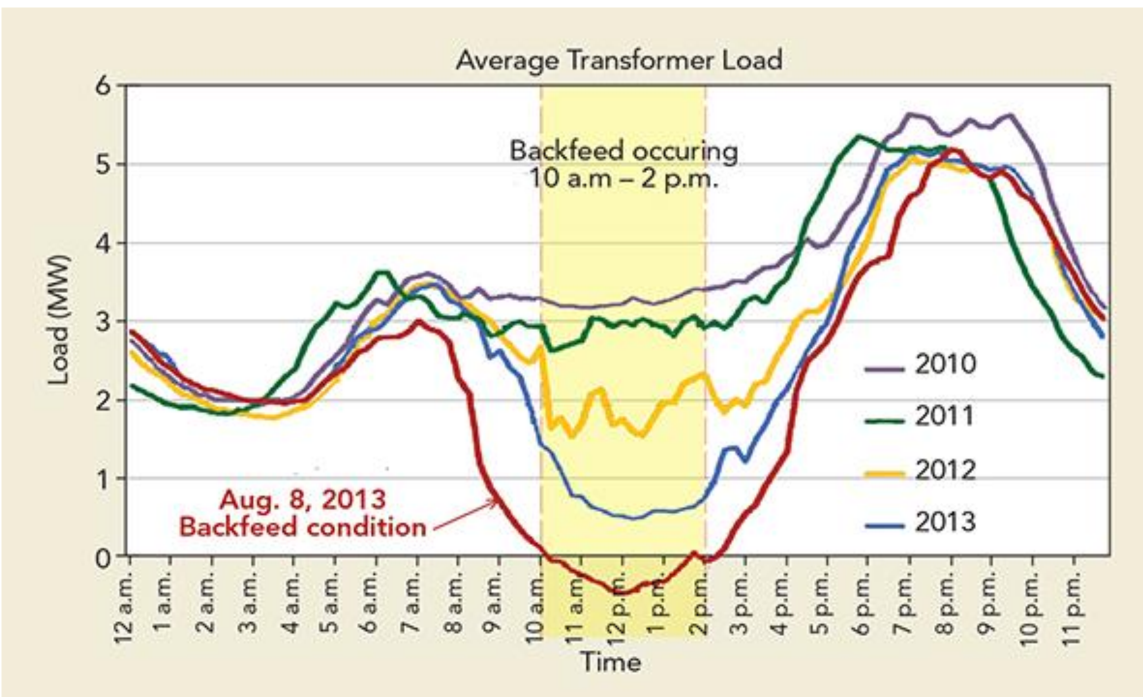

Figure 4.13. Results of Increased Penetration of PV at a Hawaii Electric Circuit. [113]

This type of loading will complicate both the energy market and operations. Add to this the adoption of DER. It is possible for energy to appear where and when it isn't needed. This too can cause difficulties in operating the grid. (Note the red curve in Figure 4.13 that dips below zero during the 10 a.m. to 2 p.m. window).

There is research that suggests that given a centralized knowledge of the electric vehicles, their locations, and their charging requirements, it would be possible to incentivize vehicle owners (as Demand Response participants) to stagger their vehicle charging in a way that levels the load on the available generation [114]. This in turn would allow lower cost generation to serve the load and minimize the need for expensive peaking generation.

In a similar fashion, given a centralized knowledge of the real-time consumption of electricity, and the strain on every distribution asset, it would be possible (in theory) to issue some type of Demand Response program to alleviate the strained asset. But current DR participation ranges from zero to 100% at different service transformers. DR participation ranges from 0.1% of peak load to 10.2% at different ISOs/RTOs, and an ability to shed as much as 6.2% of peak load for all ISOs/RTOs combined [115]. The only mechanism that does reach every consumer is the marketplace.

The marketplace must continue to support large, centralized generation, but also accommodate DER. If DR is built out in an area, a local marketplace needs to support all sizes of DR entrants to the market. The energy market needs a mechanism to support an "all of the above" energy policy. Large and small generators, large and small consumers, large and small DR participants. Remote participants would require the support of transmission operators, while local participants require the support of distribution operators.

There are two ways this can be approached:

1.) Build out the centralized energy market to have improved granularity in the markets and improved visibility to distribution issues, or

2.) Develop a type of distributed energy market which offers improved granularity in the markets and improved visibility to the distribution issues.

Either approach could be made to work. Building out a centralized system allows for a simpler algorithm, but it creates new challenges. A centralized approach could introduce a single point of failure in the architecture. It also requires enormous amounts of data to traverse vast distances bidirectionally in real time. If the current data is not available to the centralized system, it cannot compute an optimal result. On the other hand, a distributed system is more difficult to destroy. A computing node is responsible for fewer participants, so the loss of a computing node causes less damage to market participation.

The transactive energy (TE) market endeavors to be this mechanism. TE allows numerous, small, local markets to be created to form a local price, and through the free market, solve local problems locally. With generation and ancillary services put in place (as DR programs) the market can then hand off planned exchanges to network operators to execute.

TE can create customized price for electricity which is as individualized as every service location. This can serve as an alternative to the flat rate tariff, and an answer to the problems the flat rate tends to perpetuate. As long as the majority of residential customers operate on a flat rate tariff, it will be difficult for the wholesale energy market to scale out and engage small users in a way that will protect the interests of all of the market participants. The flat rate tariff actually creates subsidies between tariff classes [67]. It hides the true cost of delivering power to each service location. It perpetuates many of the problems the market currently has in promoting competition and responding to fluctuating fuel costs.

4.3. STATEMENT OF NEED

A market mechanism is needed that will accept participants of any size, apply a consistent set of rules, and scale to include all participants.

4.3.1. Distributed Energy Resources. The introduction of significant amounts of wind and photovoltaic (PV) power into the distribution grid as distributed energy resources (DER) can create voltage and frequency regulation challenges [57]. Not only can a significant penetration of PV power cause the overall voltage to rise [58], even moderate amounts of PV penetration along with windblown clouds, can create a situation in which rapid voltages changes are introduced onto the grid, and substation tap changers are not able to keep up [59]. The introduction of these new power sources challenges both the operation of the grid as well as the traditional energy marketplace [54]. “The present electric power delivery infrastructure was not designed to meet the needs of a restructured electricity marketplace, the increasing demands of a digital society, or the increased use of renewable power production”[55]. New approaches are needed which consider the needs of consumers, producers, and the grid itself.

The traditional grid has a long history of centralized generation accompanied by a single owner of transmission and distribution assets [46] [47] [48] [49] [50]. However, the trend in deregulation is to split up the ownership of generation, transmission, and distribution assets. This paves the way for small independent power producers to participate in the grid, but small participants are generally excluded from the bulk power market [61] [62]. With many areas now legislating renewable portfolio standards ranging from 20 to 100% penetration in sourcing power to the grid, the opportunity for small power producers to participate should not be ignored.

4.3.2. The US Wholesale Power Market. Transmission and distribution circuits are very different from each other. The thinking supporting the design of these circuits is very different. They are constructed using different rules and operate at different voltages. In the wholesale energy market, an entire distribution circuit (consisting of a substation transformer and everything electrically downstream) might be aggregated and represented as a single node in a transmission model. Transmission level “nodes” or “regions”[101] might participate in a wholesale energy market, but distribution-level stakeholders typically do not. By aggregating all of the details of the distribution circuit into a single node (and a single “load” in the model), the issues within the distribution circuit become obscured. The wholesale market participants are not informed of activity within the distribution circuit and, as a result, cannot do anything to address local issues that may arise.

Independent System Operators (ISOs) and Regional Transmissions Organizations (RTOs) cover regions that serve approximately two-thirds of the nation’s electricity load (depicted in Figure 4.10). The ISOs & RTOs foster competition for electricity generation by, in some cases, forming or, in other cases, supporting a wholesale energy marketplace. Bids are made in the marketplace to optimize for cost (called “economic dispatch”) [116].

The energy buyers are often utilities who need to acquire power to serve to their customers. The energy sellers are the owners of large generators. (The generator might be owned by a utility or it might be independently owned.) The owners of transmission participate in, and facilitate the sale. The contracts formed cover a variety of possibilities including sometimes the option to buy between zero and all of the power available at a power plant. Frequency and voltage regulation are often handled in a contract for “ancillary services” in which one or more buyers hold the option to buy power as needed in real time.

4.4. INTRODUCING TRANSACTIVE ENERGY

Distributed generation is the new reality. Nevertheless, to offer a fair price to participants of all sizes, a way needs to be developed to engage the smallest of consumers and producers in the energy market (including those connected to distribution voltages) [44] [63] [64]. Such a change can engage large numbers of small producers and consumers to alter their behavior in response to the corporate needs of the community. Higher (or lower) energy prices can stimulate changes in consumer behavior in response to prices, as well as smarten investment. Only when the granularity of the market is extended to include all consumers and producers will we see the smallest assets (such as individual service transformers) protected by market economies.

A decentralized energy market known as “transactive energy” (TE) has been proposed to supplement and eventually replace [70] the legacy, centrally controlled system. TE supports both market [71] and control functionality. In layman’s terms, one might describe transactive energy as “an internet-enabled free market, where customer devices and grid systems can barter over the proper way to solve their mutual problems, and settle on the proper price for their services, in close to real time" [69].

The marketplace needs to continue to support centralized generation but must also make room for DER. DER might be able to relieve a strained asset, but so might Demand Response where it has been sufficiently built out downstream of the asset, and the network operator has visibility to the situation. The market needs a means to support an "all of the above" approach to energy and consumption. TE aims to become that mechanism through individualized markets with a large number of participants.

4.4.1. "Fairness" Under the Transactive Energy Paradigm. There are many different tariffs in place at each electric utility. These tariffs also vary from utility to utility. Each of these distribution tariffs are considered "fair" yet they are different. The transactive energy market departs from these traditional practices and endeavors to apply the rules considered "fair" at the wholesale level to the retail level. The notion of what is "fair" in the energy marketplace has been argued for the past 100 years, and the debate will not be settled here. Creating a retail energy market which is similar to the wholesale market implies that consumers would have to pay individually for power from the sources they select, and they would have to pay for the transportation of that power as well. One consequence of a retail energy market is that producers must compete against each other at a retail level; and conversely, consumers compete against each other as well. When competition occurs in the marketplace usually it is considered beneficial. In this competition we expect:

- Consumers willing to pay more should be able to extend their reach and obtain power from distant locations (provided no network assets constrain their purchase).
- Consumers willing to pay more should be able to obtain expensive local power when inexpensive distant power is unavailable due to network constraints.
- Generators should end up with contracts from a mix of consumers – some local, some distant.
- Any given asset should end up with contracts from a mix of consumers which are electrically downstream from the generators.
- Consumers may end up with a of mix generation sources (of various costs) even when they specifically asked the solver to provide the lowest cost solution.

The TabuACO achieved all of these outcomes in the contracts that were formed. The process of incrementally obtaining power (using an ant authorized to purchase small amounts of power at a time) allowed all of the consumers and producers to compete in the marketplace simultaneously and arrive at an equitable sharing of the opportunities. Another approach to achieving this particular outcome would be to formally control the sale of power so that only small increments of power could be sold under any contract, and thus the solver (whatever it is) must labor to form many thousands of contracts to obtain the power any given consumer requires. This gradual release of energy to the marketplace allows for some degree of competition between consumers and an opportunity for some consumers to purchase low cost power some of the time.

In the experiments, consumers seeking the lowest cost would sometimes get contracts with generators that were low cost but not the lowest cost. There are several things to be said about this outcome:

1.) The stochastic search process which authorizes higher prices when low cost solutions cannot be found will occasionally form a contract with a distant generator that is low cost and bypass a local generator that was lower cost. The Tabu ACO algorithm used in the TE application is a "greedy" algorithm in the sense that it acquires any suitable power it can obtain. It takes the view that it is better to acquire the power required in the marketplace at a reasonable price than to run the risk of falling short before the market window closes.

2.) Depending on a consumer's strategy, this outcome is not necessarily suboptimal. It is usually a good practice to diversify one's supply portfolio in order to improve product availability and maintain lines to multiple suppliers. Portfolio diversification is practiced by many purchasing agents around the world. We can say that the selection of multiple sources for generation is actually a desirable outcome.

3.) This "proof of concept" experiment did not employ the "evaporation" commonly used by ACOs to converge to a solution. Instead it allowed the ant to form a contract with any generator that met the purchase criteria. This is deferred to be a future work in section 6.2.5. One approach to attaining the lowest cost would be to implement a form of selective "evaporation" with the contracts. Once all of the required generation has been found, an application-specific evaporation algorithm can be used. The algorithm can identify cost paid for the least expensive contract, and the energy obtained by the most expensive contract. A new (cost saving) sortie can be issued in which the ant seeks a replacement for the most expensive contract. If a replacement is found during the bidding period, it can be approved, and the expensive contract can be cancelled or possibly demoted to an agreement to provide spinning reserve instead of power. This process can continue until the lowest available outcome is acquired within the time allowed. Of course, nearly everybody wants the lowest cost generation, but not everyone will be able to acquire it. It can be said that it is "fairer" that people pay an average price. This is part of the debate. The other side would argue that not all consumers are "average" and that consumers should pay in proportion to the actual cost of goods and services.

4.) Alternatively, the existing algorithm can be used to seek the lowest cost generation. A user who really values hitting a certain cost target, can name their price, and if the ant doesn't find generation available at that price, adjustment their strategy before the market window closes. In this way, the existing algorithm is perfectly satisfactory. Of course, this strategy risks experiencing an energy shortfall at the close of the market window.

Furthermore, on the subject of marketplace competition, a real commercial application would likely want tighter guarantees of fair play. In the proof-of-concept experiment, equity was assured by running all nests for equal iterations on the same computer. If processing is distributed so that every nest has its own processor, similar controls are needed to ensure that some nests do not gain an unfair advantage due to differences in processing power and network performance.

4.4.1.1. Overages and shortfalls. There are several scenarios that may occur in which a consumer uses more or less power than contracted, or a generator contributes more or less power than agreed. Who pays for these differences? In the wholesale market the participants pool their power. If a generator goes offline, it loses out on revenue, and additional power is (often) purchased on the spot market to compensate for the shortfall. Sometimes network operators can simply raise the contribution from the spinning reserve and frequency reserve allocation to address the shortfall. A similar thing can occur here. Consumers can use the transactive energy market to agree to purchase the minimum amount of power they think they will require, agree to purchase spinning reserve to cover any additional amount of power they think they may require, and allow network operators to control the spinning reserve contribution in order to stabilize the grid. What if a consumer draws more power than their energy and spinning reserve contracts allow? As described in section 6.2.3, this is a subject for future work. However, we would expect uncontracted power to be supplied by a power pool, and settled at a price higher than average. The energy market provides a first pass approximation in matching the level of generation to the load, but network operators have the means to tune the output of certain generators until the level of generation matches the load.

A consumer with onsite generation may elect to operate their system off-grid, but if they go on-grid they have access to additional generation beyond what they own. This simple connection usually results in higher reliability and higher availability. But the ability to purchase the power necessary to meet the local need is something that may cause consumers to recalculate their investment in local generation. For those who are off-grid, local power is “free,” but only after a substantial initial capital investment.

The normal behavior for the stakeholders is described in Figure 4.14. Consumers are expected to form a contract for low cost power which is the minimum they believe they will need to use, and another contract with a generator for spinning reserve. The spinning reserve costs something because the equipment is “rented” and standing by to serve the consumer. If called upon to contribute, fuel will be spent, and the delivered power which was held in reserve utilized. This source will likely be more expensive than the other contracts the consumer holds. It should also be mentioned that the “network” that transports the power is not expected to have zero resistance. There will be electrical losses along this network. If a consumer elects to buy power from a generator that is quite remote from their location, they do so with the understanding that some of the power will be lost along the way.

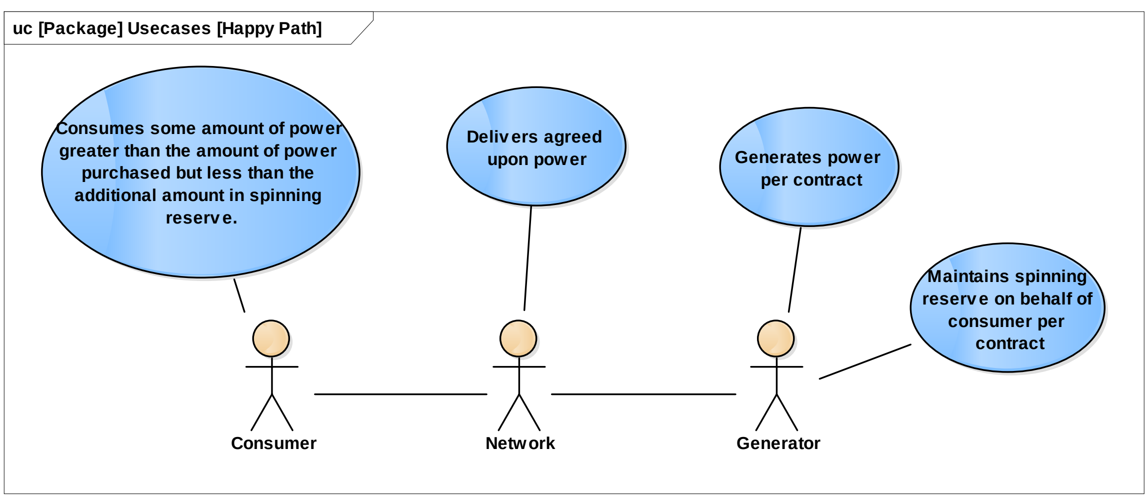

Figure 4.14. Happy Path Usecase Showing Contractual Roles

The corner case analysis is always more difficult than the happy path. What should happen to our contractual agreements if a network outage occurs? It may be that the loss of a transmission line allows distant generators to continue to operate but due to the severed line, cause their power to be delivered to different consumers instead. What is fair in this case? On one hand, it could be argued that electrons are fungible and that electrons added at one point on the grid don't necessarily make it to their contractual destination anyway. Let consumers pay the rate they have used. On the other hand, it could be said that high cost electricity can make or break certain industrial applications, and cost matters a great deal to the economics of the business process. Should consumers be informed of a transmission line outage and be asked to curtail? What if it is a local outage and they become informed as their lights go out? This is described in Figure 4.15.

If the communications network were still operational, and the solver able to perform a rapid analysis for the energy market, it may be possible for the consumer's equipment to immediately negotiate new contracts with local generation. Suddenly, the power which would have left the area for sale elsewhere is no longer able to leave, just as the power generated remotely is unable to enter the area. Can the old contracts be temporarily set aside, and new contracts negotiated under the new (temporary) grid topology? The distribution network (and/or network operator) could inform the consumer's equipment of the grid failure, and of power shortfall.

If the communication network is a high-performance network, it may be possible for the network to isolate the fault, come to understand its new topology,[12] and allow new contracts to be renegotiated and operate until the old topology is restored.

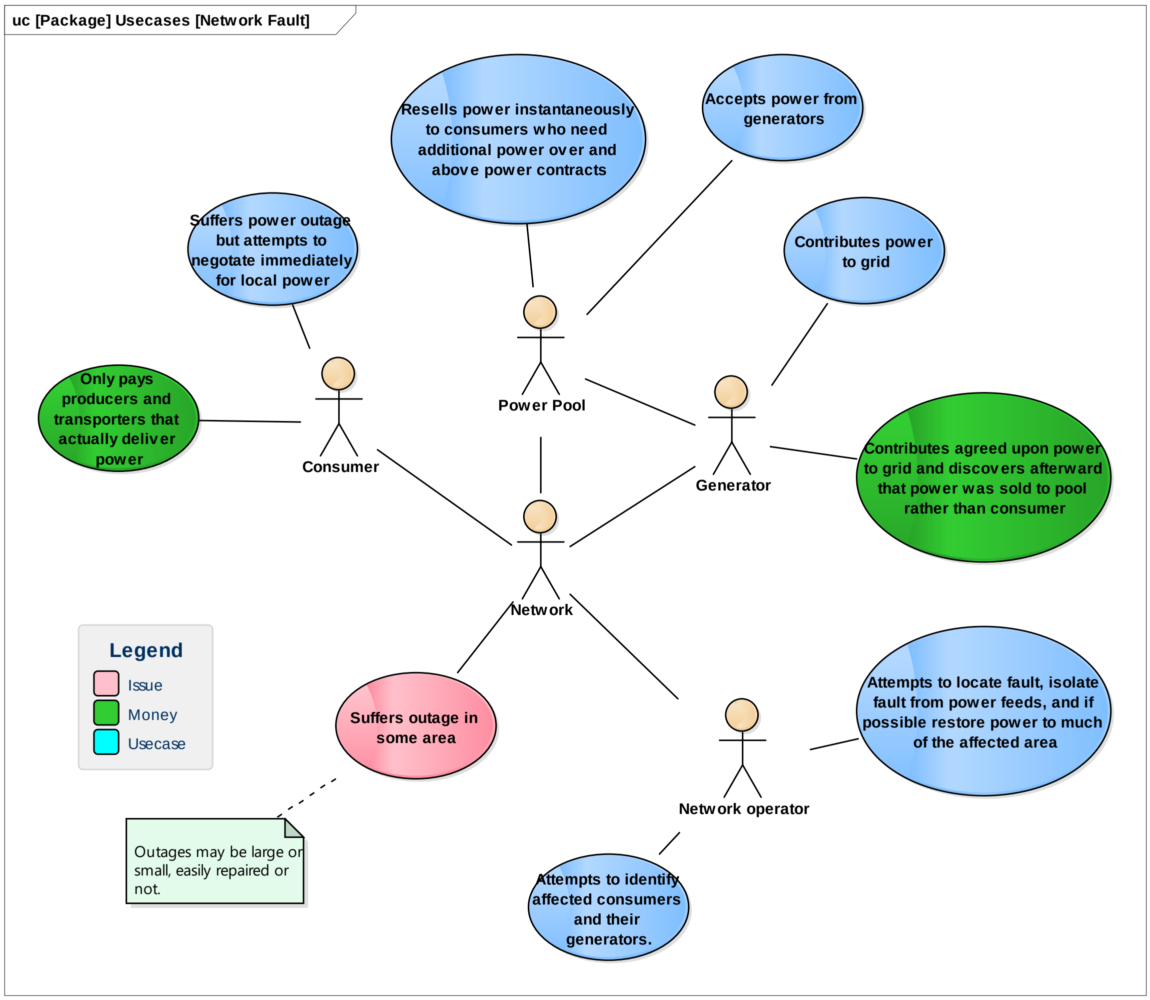

Figure 4.15. Power Outage Usecase

[12] Products to monitor transmission asset health are common. Products to monitor distribution asset health are available today, and costs are falling.

The discussion of what is “possible,” “affordable,” and “fair” is left for implementers to negotiate. Much of what is considered “fair” or “reasonable” is often based upon the cost of the technology available at the time.

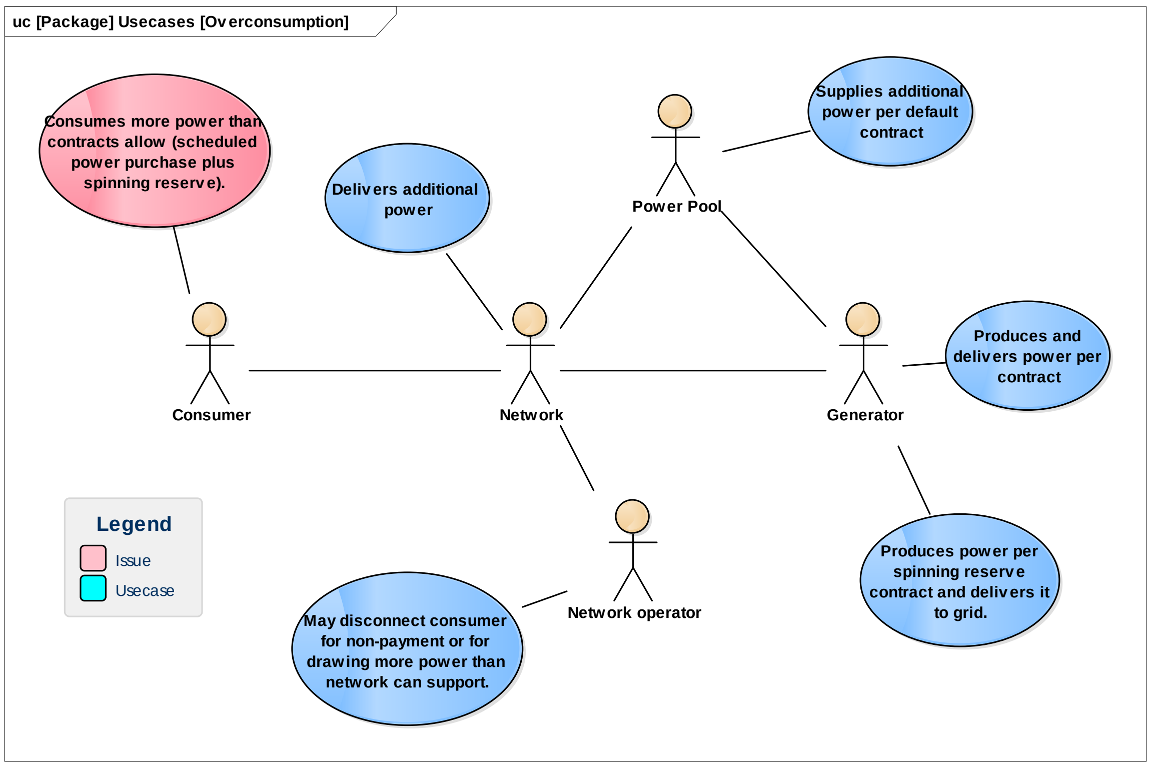

Figure 4.16. Overconsumption Usecase

It should also be noted that this usecase introduces the “power pool” actor. This is a bit of an artificial construct. The power pool is a collection of various participants who have excess power for sale. The list includes:

- Generators that contribute power to the grid as part of their startup or shutdown process that is not specifically under contract.

- Generators that have unsold spinning reserve that could readily contribute power at the request of the network operator.
- Consumers that do not consume all of the power they have bought.

Another scenario is the case of the consumer using more power than they have agreed to purchase. Figure 4.16 depicts the usecase scenario. What is reasonable in this case? Should the consumer be disconnected the moment they exhaust the contract? Most of today's "smart meters" can easily support "prepay" metering and automatically disconnect when the account balance reaches zero. They can also support a grace period. Should the proposed system be invoked to acquire a purchase off of the spot market? Should the consumer instead be assigned a purchase from the power pool at the current market price? Some consideration has to be given to the ability of the communication network. If it is highly capable, it may be possible to negotiate a new contract in less than a second. If the network is instead optimized for cost, it may be that power is allowed to flow to the extent the network can tolerate it, and payment is reconciled after the fact. Integration with other systems is cited as a future work.

Consumers are expected to sign up for spinning reserve so that they may draw power at their discretion. Spinning reserve is expected to be inexpensive to reserve, but above average in cost to utilize. As a cost saving measure, consumers may designate a fixed amount of power to be applied before any spinning reserve is utilized. Generators under spinning reserve are controlled by the network operator. Generators assigned a fixed contract may contribute their block of power at the appointed time without further instructions from the network operator.

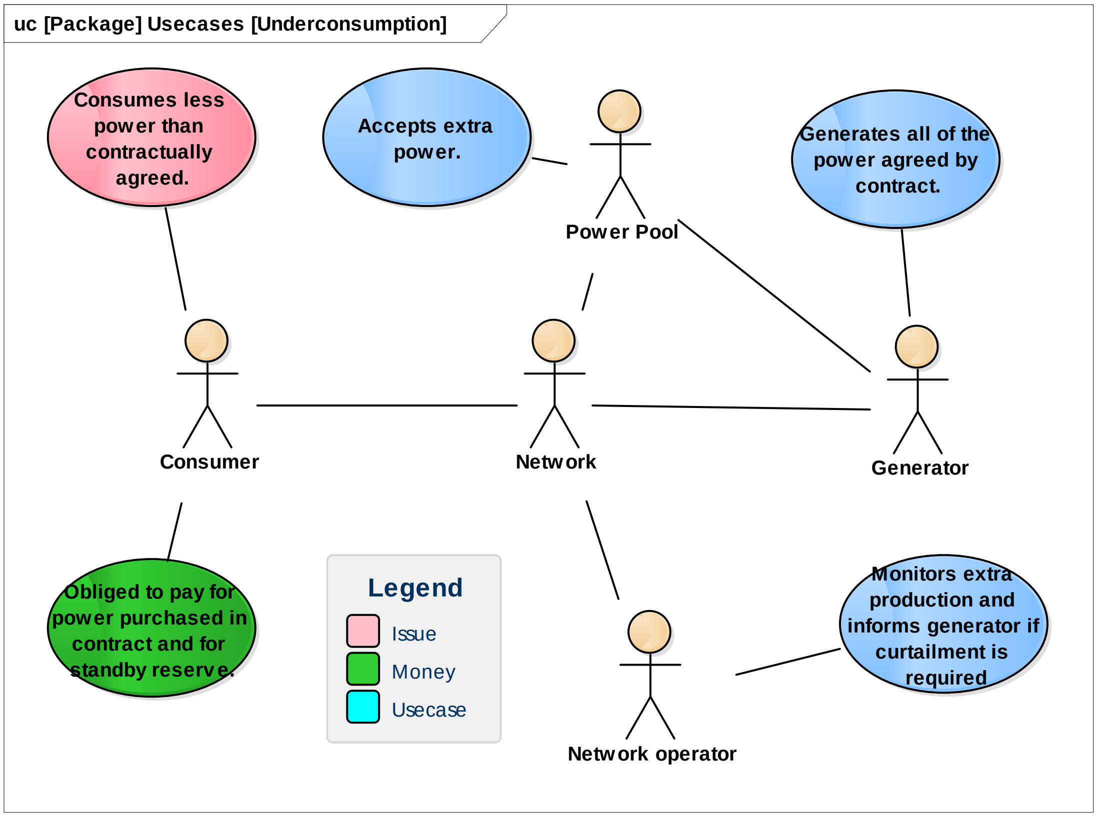

Figure 4.17. Underconsumption Usecase

What happens if the consumer doesn't use the amount of fixed power they contracted for? (Such a situation is depicted in Figure 4.17). Case law would probably find that since the contractor(s) delivered on their end of the agreement, the consumer is obliged to pay them for their services. The excess power would flow to the "power pool" entity for use at the discretion of the network operator.

Sometimes generators have problems and must go offline for unscheduled maintenance or repair. (The usecase scenario is considered in Figure 4.18)

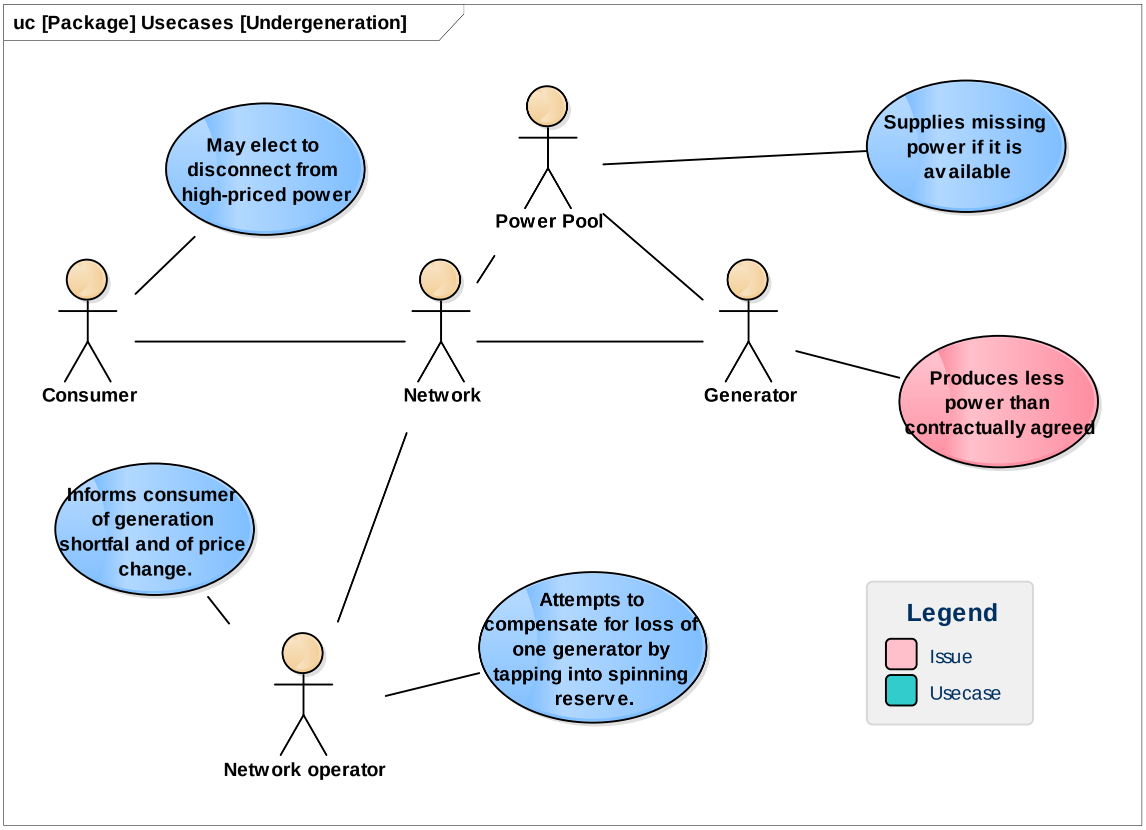

Figure 4.18. Generation Shortfall

There are many options to be considered.

- The generator might use the solver to resell its contracts to other generators
- The network operator might inform the consumer of the generation shortfall and inform the consumer that they will be purchasing power from the power pool at the spot market price until the consumer finds a replacement generator.
- The consumer, once informed of the failure, might elect to conserve power and work within the constraints of the available power from other generators also under contract.

Large generators often must use or contribute a bit of power in order to get synchronized to the grid. This is usually considered a negligible amount of power by the generator, but what is to be done with it? What if an altruistic owner of solar generation simply adds their excess power to the grid? What is to be done with excess "free" power? (The usecase scenario is considered in Figure 4.19). The network operator will be expected to monitor power production conditions, detect excess contributions, and determine what to do about it.

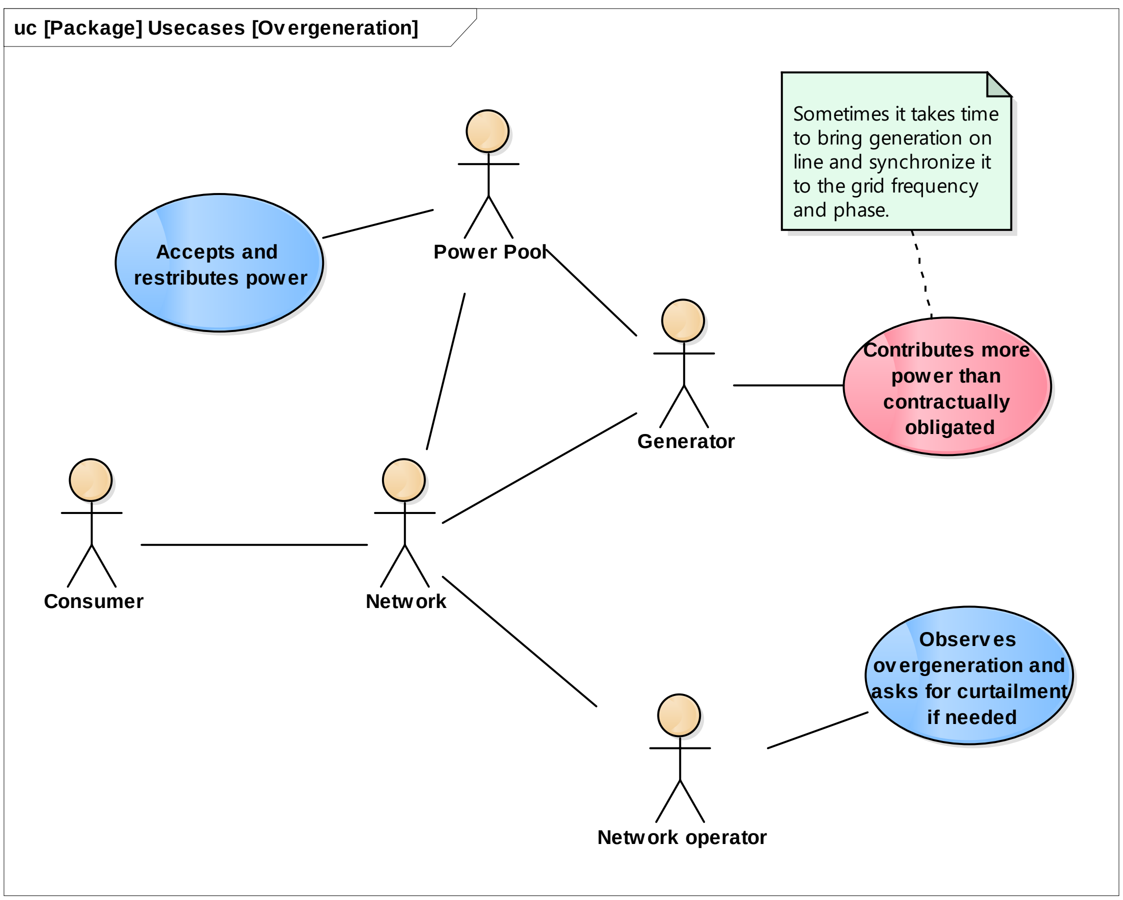

Figure 4.19. Generation Excess

4.4.1.2. Ancillary Services. The legacy energy market sells ancillary services as well as energy. Ancillary services are used to handle the corner cases described above which lead to generation/load mismatch. It is likely that anciallary services will still have a place in the marketplace in addition to energy sales.

The conclusion of the analysis is that different regions will likely continue to have different needs and subscribe to different notions of fairness. Any commercialized system will need flexibility to support the range of options that real users would require. The notion of fairness is best settled regionally. A real time spot market will likely be required. A high-performance network would allow for immediate renegotiation of contracts in the event of a network fault. If the network and solver perform sufficiently well, it may be possible for the network operator to leverage "immediate renegotiation" as a new tool to use in the event of an unplanned network event.

4.4.2. Disaster Recovery. There are many things that can damage the grid. A massive solar flare could destroy much of the grid [117]. Large storms have been known to destroy the grid [118]. Damage can occur to not only the transmission of power but also to the monitoring and control of power. However, the introduction of significant amounts of distributed generation along with distributed controls can change the paradigm. Power will not necessarily be lost when a transmission line is lost. For example, if a transmission line between a major generator and a substation is severed, any distributed generators downstream of the substation will continue to serve the load. The disconnected network may continue to operate in an islanded fashion.[13] By the same token, a neighborhood with significant amounts of distributed generation should be able

[13] Usually with poor voltage and frequency regulation, but continuing to operate nevertheless.

to operate independently from the grid (as a “microgrid”) if disconnected due to a disaster.[14]

Distributed systems are well known in the literature for their advantages over centralized systems when it comes to robustness. Distributed generation, distributed control, along with a distributed energy market could allow an entire system to bootstrap from an outage after a storm. Smart (communicating) distribution assets could determine which lines and transformers were still in service. An automated fault location, isolation, and service restoration (FLISR) capability could bring up portions of the grid that survive the storm, and allow islanded systems to operate for months before being reconnected to the regional interconnect. This type of capability is possible with the transactive energy paradigm and not possible with today’s centralized control paradigm.

Table 4.2. GridWise Architecture Council Transactive Energy Framework Nodes [63]

Node Name	Scope
Regional	Represents an entire region such as an ISO or RTO
Control Area	Represents a legacy control area and its automatic generation control
Distribution	Represents one or more distribution networks
Market Participation	Represents a wholesale market participant. (A market participant may be active in more than one region).
Supply	Represents a market participant which is not part of the legacy system, and any size.
Building	Represents all of the load at a premises or otherwise connected to the grid

[14] A disaster in this case would include a hurricane, thunderstorm, earthquake, as well as a man made disaster (e.g. terrorism).

4.4.3. The GWAC Transactive Energy Framework. The GridWise Architecture Council envisions a system architecture for transactive energy which overlays the structure of the current power control and energy market [63]. The architecture is summarized in Table 4.2. It is thought that participants can change roles at any time. A supply node could become a building node (load) as conditions change.

4.5. STATEMENT OF SCOPE

This dissertation proposes a method of implementing transactive energy (TE) using the TabuACO. It does not attempt to utilize TE as a control mechanism. This instead is left to other researcher endeavors. A literature review finds that other research is underway to leverage TE as part of a grid control system [119]. Linear programming, artificial neural networks, and other approaches are being considered.

4.6. PROPOSED TRANSACTIVE ENERGY MODEL

Implementation of a transactive market with ant colony optimization requires a new architecture to describe the control relationship (ref. Figure 4.20). The transmission and distribution networks will still operate in a similar manner, with the main difference being the creation of a new market. This new market is supportive of the design constraints of the distribution circuit. Much like the process in the wholesale market, the transactive retail market will form contracts between three parties:

1.) The owner of the generation

2.) The owner(s) of the grid assets

3.) The consumer

In this representation, the consumers are the nest for the ACO, and the energy at the generation sites represent the food the ant agents are seeking. The grid itself is modeled as a directed graph, with each grid asset represented as either a vertex or an edge (Table 4.3).

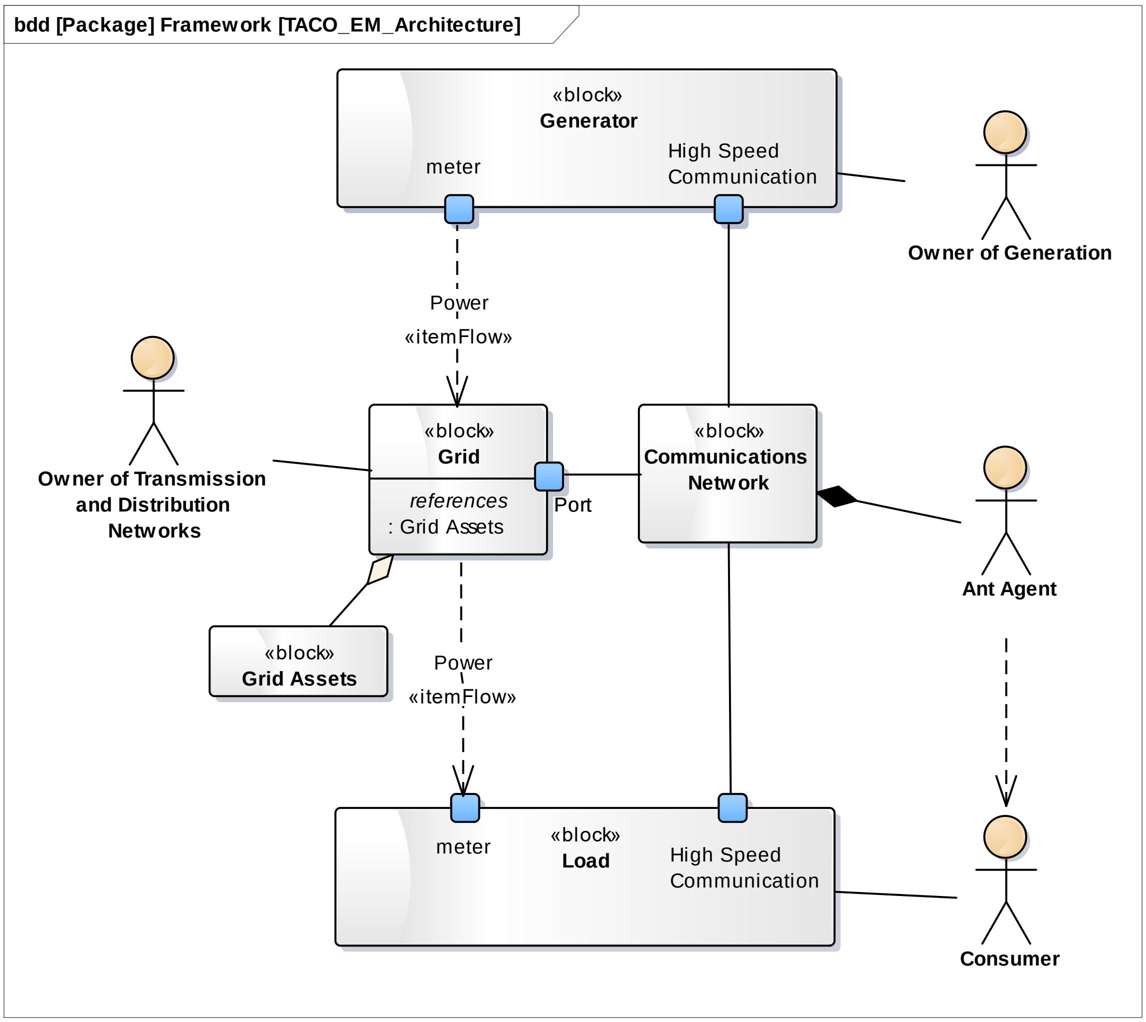

Figure 4.20. Proposed Transactive Energy Framework (SysML Block Definition Diagram)

When an ant finds a food source (generation), it forms a contract with that generation asset and the grid assets along the path it traversed to deliver the energy to the consumer. This architecture operates as an object-oriented model (per Figure 4.20) with a set of rules and assumptions:

- A consumer may change roles and become a producer (and *vice versa*). Each actor is represented as playing only one role during a market interval – the role which represents their net requirements.
- A market interval is typically a 15-minute period of time (but could be operated as being more or less).
- The grid consists of assets (wires, transformers, switches and other gear) used to move electricity from one location to another. Each grid asset has capacity constraints and can be contractually reserved for use by contract through negotiation.
- Each asset may be reserved for up to its rated capacity, but the owner of the asset may have reasons of their own to limit the usage to a level below the maximum rating of the asset.
- Forward and reverse flows are accounted for separately – so an asset which is completely reserved in the forward direction may still accept reservations in the reverse direction until both directions reach capacity.
- Consumers shop for both generation, and transportation via the power grid through contractual reservations which are communicated via the communications network.
- Software agents (ants) act on behalf of the consumer to make the reservations.
- The consumer is not required to reserve 100% of the capacity of a given generator or grid asset.

- Only real energy will be sought in the marketplace, although reactive energy and other services exist.
- Negotiation is dynamic and occurs anew every market interval.
- Negotiation for a given market interval on the calendar may occur a day ahead, an hour ahead, or minutes ahead of the actual need.

Negotiation for a given market interval must complete within the time of the market interval, so that it can be closed out, and negotiation for the next interval begin.

4.7. APPLICATION OF THE TABU ANT COLONY OPTIMIZER

A mapping from the market elements to ACO elements are described in Table 4.3 along with a few simplifications.

Table 4.3. ACO Correspondence

ACO element	Market element
Ant	Agent working on behalf of consumer to find energy sources
Food (or prize)	Energy or ancillary services. For sake of this exercise, the ant will only forage for real energy. As a simplification, reactive energy and ancillary services are considered acquired when real energy is required. These other commodities can be purchased by additional foraging in a future work.
Food retrieval	A contract is formed to secure energy from the producer, and a contract is formed with each asset along the route to deliver power to the service location.
Nest	Consumer’s revenue meter (point of common coupling between utility wiring and premises wiring).
Pheromone	Hints left by agents (in a public data store located in field assets) as to what can be found further down the pathway. (Note: these agents may be agents of any consumer.)
Node	A grid asset which is not a conductor
Edge	A grid asset which is a conductor

4.7.1. Application of a Reference Direction. A reference direction is overlaid on the wires model to turn it into a directed graph. In the distribution network, wires which run toward the substation are considered "upstream." Wires which run away from the substation are considered "downstream." When a distribution network is fed by multiple substations or a single substation with multiple busses, one source is arbitrarily modeled as "upstream" and all of the others are labeled as "downstream."

4.7.1.1. Leaf and root nodes. A leaf node is any node with a degree of one. A tree network will ultimately have leaf nodes along its outmost edges. However, due to the setup of the problem and the assignment of reference directions, certain information can be derived from the reference directions to indicate a stopping criterion for the foraging ant. A substation will likely have multiple feeders connecting to a bus which radiate outward. This substation bus can be viewed as a significant source of power. Unless some capacity constraint is reached, an ant will not visit a substation bus looking for power and walk away. In the model, the root of the tree which models the distribution network will be depicted as a node with a degree greater than or equal to one, and for these edges, the reference direction for all edges face the same direction. Example topologies for the leaf and root nodes are depicted in Figure 4.21.

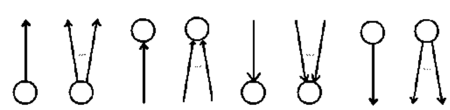

Figure 4.21. Example Leaf and Root Nodes

Leaf nodes are interesting to the ant because they may contain a food source, a nest, or be empty.

4.7.1.2. Intermediate nodes serving as waypoints. Nodes which are not leaf nodes are intermediate nodes. Example topologies are depicted in Figure 4.22. All such nodes have a degree greater than one and reference directions into and out of the node. Such nodes may be service transformers, tie points, protective devices, poles, or any other distribution asset which is not a leaf node.

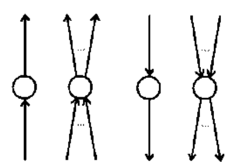

Figure 4.22. Example Intermediate Nodes

4.7.2. Pheromone Accounting. Two pheromones are used for each commodity being sought: one pheromone is used to attract, and another pheromone is used to repel the ant to or from traveling along a given direction on the edge.

4.7.2.1. Attractive pheromone. An attractive pheromone takes on a positive value when it recommends to the ant that she be drawn in the direction the arrowhead is pointing, and a negative value when she is to be attracted to the arrow tail (attracted to the opposite direction the arrow is pointing).

4.7.2.2. Repulsive pheromone. The repulsive pheromone takes on a positive value along an edge when it recommends to the ant that she be repelled away from the arrow tail. A positive repulsive value will therefore cause the ant to follow the direction the arrowhead is pointing. A negative repulsive value repels the ant away from the arrowhead. A negative repulsive value therefore drives the ant towards the arrow tail.

4.7.3. Deposition. The TabuACO has the ant deposit pheromone under two different sets of rules. One rule is applied when the ant encounters a find. Another rule is applied when the ant doesn't encounter a find.

4.7.3.1. Attractive pheromone. When an ant forages she maintains a path history of the path she must take in order to return home. Since she is foraging in a graph, and multiple paths may exist to get home even among a given set of nodes, the path history accounts for both nodes and edges.

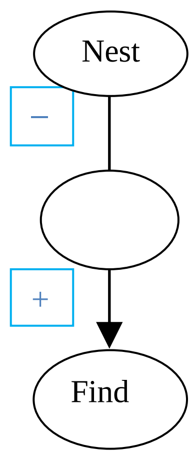

Figure 4.23. Attractive Pheromone Signs

Upon finding a food source and striking an agreement to purchase power, she returns home along the path on record. As she travels, she deposits attractive pheromone to lead ants back to the find, and she firms up agreements with each asset to rent a portion of the asset's capabilities for the interval period. The ant remembers the path home as she forages. Once a find is identified, and the ant travels homeward, edges pointing away from the find are painted with negative attractive pheromone.

Edges pointing toward the find are painted with positive attractive pheromone. The sign used depends on the reference direction as depicted in Figure 4.23. The formula for deposition of attractive pheromone is described in (19).

Each find by the ant results in a new trail of attractive pheromone being laid (in addition to any pheromone already present) between the food source and the ant’s nest.

4.7.3.2. Repulsive pheromone. Repulsive pheromone is laid while an ant forages. If she finds a leaf node which does not have a food source, she deprecates the node and all edges connected to it. Pheromone is a signed value. If the edge to the leaf node is upstream (pointing away from the node), she applies positive repulsive pheromone. If the edge to the leaf node is downstream (pointing into the node) she applies negative repulsive pheromone. The sign used depends on the reference direction as depicted in Figure 4.24.

Figure 4.24. Repulsive Pheromone Signs

If the node has a food source, and the ant could not afford it, she does not deprecate the node. She continues foraging.

If the ant encounters an intermediate node that does not have a food source, she examines all of the edges connected to it. If all of the edges except one have been deprecated with repulsive pheromone, she decides to deprecate this node. This is done by deprecating the remaining edge. If it is an upstream node positive repulsive pheromone is laid. If it is a downstream node, negative repulsive pheromone is laid.

The formula for deposition of repulsive pheromone is described in (18).

4.7.4. Evaporation. The TabuACO contains two types of pheromone. These pheromones evaporate at different rates.

4.7.4.1. Attractive pheromone. It is possible for a food source to become exhausted. Pheromones which led ants to a particular node become obsolete when the food source is no longer present. The attractive pheromones along every edge will be partially evaporated with every iteration. This is done by multiplying the pheromone values by a number between zero and one.

4.7.4.2. Repulsive pheromone. Repulsive pheromones are not evaporated. Areas of the network which do not have food will not have food at some later point during the market interval. They remain deprecated for the duration of the evaluation.

4.8. EXPERIMENT TO SHOW THAT THE TABU ACO CAN SERVE A RETAIL MARKET AND ABIDE BY DISTRIBUTION NETWORK DESIGN LIMITS

Section 4.2.3 described how the existing wholesale energy market performs a function to arrange the scheduled movement of power through the transmission network between generators and loads. The arrangement observes all of the design constraints of the assets involved. It also arrives at a mutually agreed price for the power and use of the transmission assets to transport it. We would expect a retail energy market to do the same at the distribution level.

4.8.1. Experiment. A test was performed on the “IEEE 34 Node Test Feeder” [120] model. Its purpose was to confirm the ability of an ACO to recognize a new issue in the network and endeavor to work around it as it sought to form contracts between buyers, sellers, and transporters of power. It also served as a small-scale proof of concept test for the algorithm.

It should be noted that despite the “O” in “ACO”, the TabuACO solver is used in this experiment to analyze and manage the contribution of power to the grid – not optimize the purchase of power to ensure a particular consumer received the lowest price.

IEEE offers Figure 4.25 when describing their 34 Node Distribution Feeder.

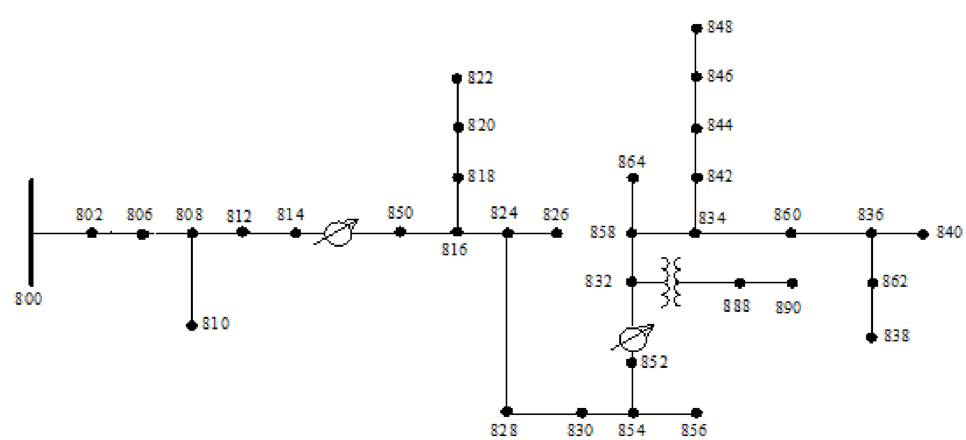

Figure 4.25. Original IEEE 34 Node Distribution Feeder (not to scale) [126]

The IEEE model contains information regarding a feeder circuit, complete with node ratings, node locations, conductor ratings, and loads at various locations. There are six “spot loads” identified with nodes 830, 840, 844, 848, 860, and 890. There are nineteen “distributed loads” identified in the IEEE model that occur along a conductor between nodes.

The IEEE model represents aggregate loads and other pieces of equipment at various “nodes” along the feeder. Additional detail is needed in order to conduct this experiment. Small generation must be added at various points throughout the feeder, and prices must be set as bids for power from these locations. Furthermore, additional detail needs to be supplied regarding the various loads.

Figure 4.26 shows the additional detail drawn (not to scale) of the generators and loads under the modeled nodes. The color indicates the price bid for the generation. Substation power (at node 800) is bid at a very low price. The other blue ovals are somewhat higher in price (but still low), yellows are higher, and reddish purple is the highest price. Pricing for this particular experiment is arranged so that ants will prefer substation power. Price data are found in Appendix C.

A study of the IEEE data finds that node 890 requires 503 kVA, but transformer XFM-1 is rated for 500 kVA. This invariably creates a situation where an IEEE-defined transformer will be overloaded with an IEEE-defined load. This will serve as the condition for a test.

In this test, two groups will be compared. The control group will simulate power delivery to a legacy wholesale energy nodal market, where the only node visible to the market is node 800. A day-ahead load forecast will be available for the feeder as a whole as presented to node 800. The test group will model a transitive energy market at Node 8900. The two market methods will be compared.

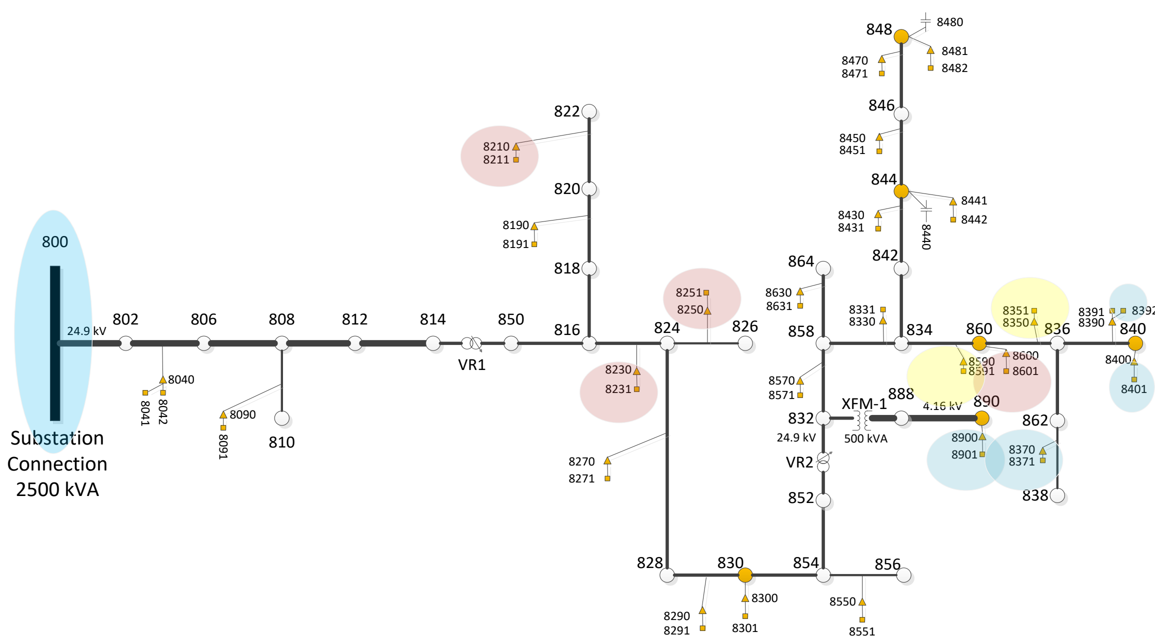

Figure 4.26. Modified Distribution Feeder Model Depicting Source of Power and Prices Bid (not to scale).

The successful method will be the one that alleviates the strain on transformer XFM-1. The load at node 890 will participate in the transactive energy market, as well as at least one other node upstream of XFM-1. Nests at these nodes will represent consumers and seek to form contracts with generators and power network providers as described in section 4.9.1.1. Furthermore, in this multinest environment, the objective of the solver will be to secure the first available generation that meets the cost criterion. The highest priority of the solver is to secure the power required by the nest. A second order priority is to keep the costs low. The consumer defines a “buy immediately” action that is to be taken if a certain cost target is met. The upper cost limit is the rate that the ant is authorized to spend if the lower priced power cannot be easily found.

4.8.1.1. Proposed methodology and detailed modifications to the IEEE-34 feeder model. To solve this problem models are created to represent the feeder connectivity, the assets deployed along the feeder, the loads, the distributed generation, and other grid assets. Each service requesting power is represented as a nest, and ants leave the nest to forage for energy. The details of the model are posted in Appendix 2.

4.8.1.2. Theory of operation. A transactive energy market as described in Figure 4.3 will be operated as the experiment. The TabuACO solver will match individual consumer loads with all available generators. The solver will allow multiple nests to compete simultaneously for the available resources. The solver will allow each nest to form contracts within the solver's model for the resources and energy the nest needs. The ant performs resource locking of the assets as she forages, then at the end of the sortie, either releases or commits to the portion reserved. Ants incrementally reserve resources as funds and network constraints allow, until all desired energy is acquired.

4.8.1.3. Class design. This experiment is a small-scale test to prove out the concepts suggested. A node has properties which identify its type (pole, transformer, etc.), its kVA rating, its coordinates, polyphase connection, as well as its contracted kVA in each direction. An edge represents a conductor. Every conductor has exactly two ends, and a node at each end. It too has a kVA rating (based on its ampacity and insulator voltage rating). Edges in the model have a direction. All edges point upstream, and since the IEEE model is a distribution network, all edges point to Node 800 (the substation transformer). The Edge model records how much kVA has been contracted in each direction. A contract for some amount of power through the conductor in the upstream direction is treated independently from a contract for power flowing in the downstream direction. While Kirchoff's current laws say that these flows would actually cancel each other out, they cannot be relied upon to be present. So, a current flow up to the rated amount is allowed to occur contractually in each direction.

4.8.1.4. Ant agency. An ant from each nest is given authorization by the nest to spend money and form contracts on the nest's behalf. As the ant travels, she has a certain cost criterion in mind. Any generation that meets the cost criterion is immediately purchased. However, if an ant cannot locate affordable generation very readily, she is authorized to spend more at various hop count setpoints as determined by the nest. The ant will not go above the maximum setting authorized by the nest.

The ant keeps records of the path home, but also of the contracts she has signed along the way. The ant examines the constraints of every asset she travels. If the asset is fully committed, she will not attempt to transit it. If some margin remains, but it is less than the full amount she is looking for, she will reduce her contractual requirements to the amount the asset can carry, then transit the asset. The ant will limit the contract with the generator to be no greater than the constraints allowed. Should the ant doubleback in her foraging, any constraints imposed by that asset will be lifted. In the path from the nest to the generator, the ant has "locked" the portion of the asset necessary to deliver the intended payload.

Once the ant forms a contract with a generator, the deliverable payload is finally known. The generator may have less power available than the ant was authorized to purchase. As the ant travels home, she finalizes the contracts with each asset. Contracts with assets will be adjusted downward (as necessary) to become only the amount needed for the generator. Unneeded portions will be freed for use by other ants.

The ant also considers the losses of carrying energy across the asset. The series resistive losses of the conductors are considered and used to derate the payload amount delivered.

The cost per kVAh is modelled, and the ant computes the financial cost to move the payload through the series of assets along the path. This travel cost will also serve to limit the range of the ant to what is affordable and keep her near her nest.

4.8.2. Expected Results. The data contains numerous loads and sources of generation under each node in the IEEE model. Many of these added sources of generation are priced below the price of power from the main substation transformer. This should cause the TabuACO to find and exhaust all of the low-cost sources of power. This should result in different nests to arrive at different average costs. A “price map” should be generated by the transactive energy market that resembles the price map generated by the wholesale market.

Node 890 requires 503 kVA, yet transformer XFM-1 is rated for 500 kVA. Power at node 8901 is priced higher than many sources upstream of XFM-1. It is expected that the TabuACO will get 500 kVA from low cost sources upstream of XFM-1, then when the limit is reached, obtain the balance of the required power from the higher priced source downstream of XFM-1.

4.8.3. Results. In hundreds of trials of the TabuACO, the ant always found the constraint at transformer XFM-1. Contracts totaling 500 kVA were formed with low cost supplies upstream of XFM-1, while a contract for higher priced power at node 8900 was formed for 3 kVA.

The costs incurred by the various individual loads are plotted in Figure 4.27.

Y-coordinate

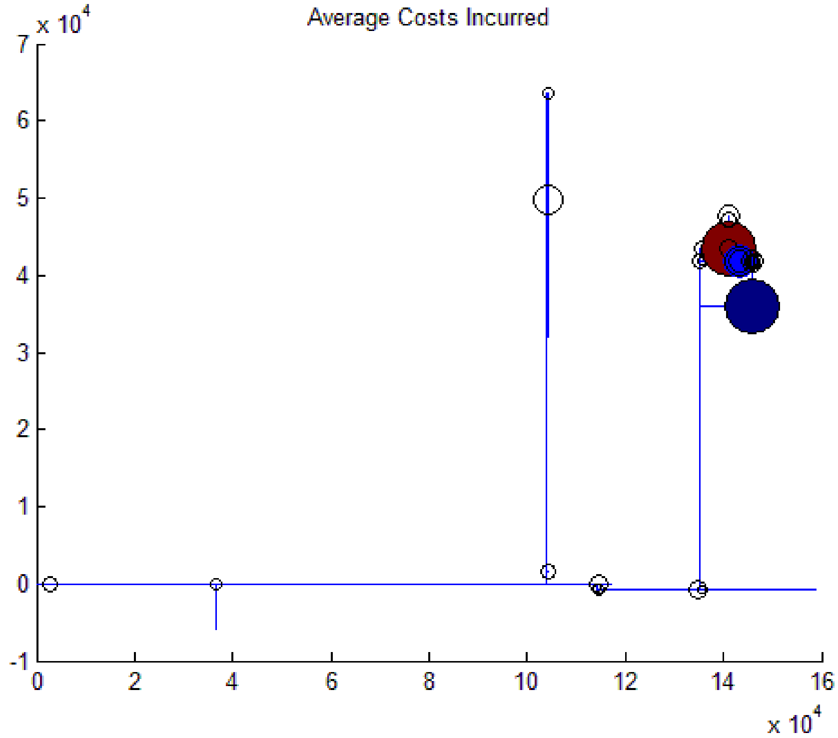

X-coordinate

Figure 4.27. Location Cost Map for Experiment After Run (Plotted to a Proportionate Scale) for Modified IEEE 34 Node Feeder

Figure 4.27 looks different than Figure 4.25 because it is plotted to scale. The feeders serve a mixture of sparse and crowded service locations. The size of the load is depicted by the size of the circle. Larger circles represent larger loads. Filled circles represent loads that participated in the market. The color of the fill represents the average cost paid by the customer.

Figure 4.28 offers a closeup of the upper right area of Figure 4.27. This area contains components that have constraints that are exercised during the run.

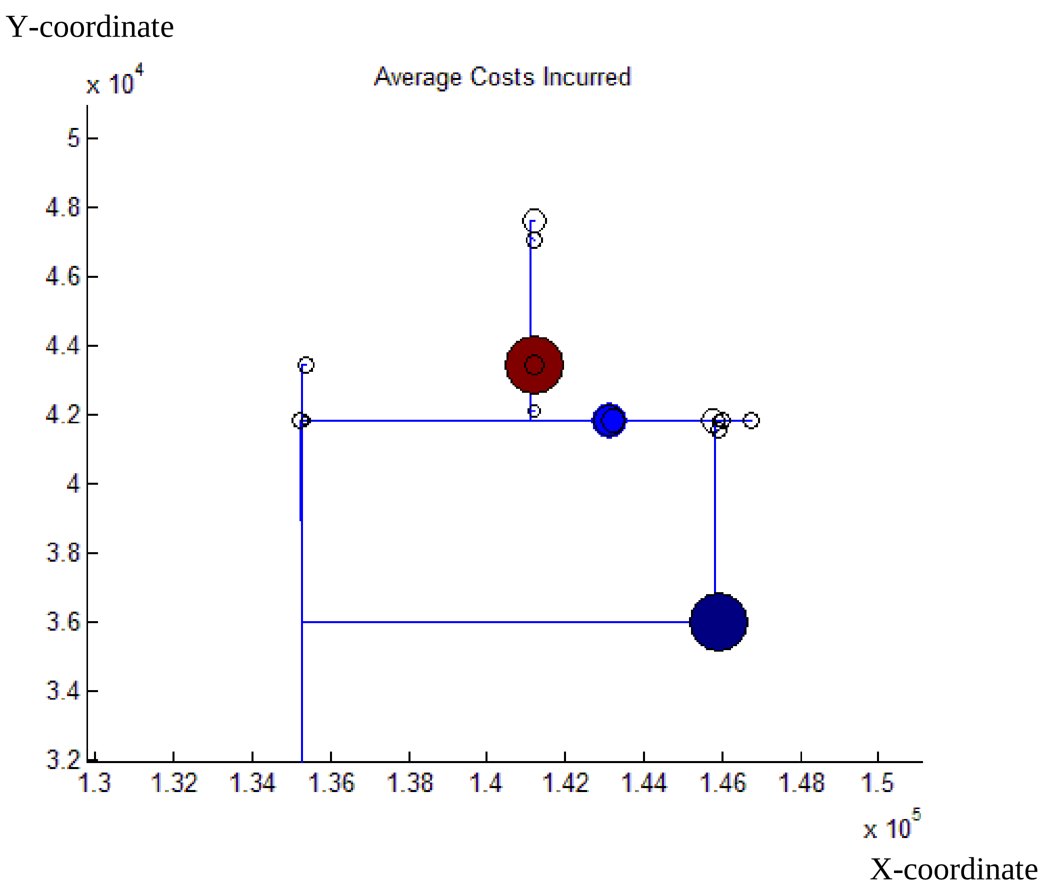

Figure 4.28. Closeup View of Nest 890 (red) and Other Nests that Competed (blue) in the Proof of Concept Test of the Modified IEEE 34 Node Feeder.

4.8.4. Discussion of Results. Ants were given a small portion of the load and tasked to acquire a contract for generation. As ants incrementally built contracts between node 890 and the substation at node 800, they eventually found that the node representing transformer XFM-1 refused to grant passage of any more contracts for power. Ants were forced to look elsewhere for power. Eventually, after walking far enough, they were authorized to spend more. They could eventually afford the power available locally at node 8900 and finally fulfilled the needs of the nest.

At no time did the ants fail to see the constraint and attempt to work around it. Occasionally they ignored the pheromone recommendations and wandered into unproductive areas. A histogram of the ant hop count is pictured below in Figure 4.29. We see it is a bell shape which is skewed to the left. This is typical of the TabuACO. Often the ant is "lucky" and finds results quickly – thus the skew to the left; however sometimes the ant is quite unlucky, and she is required to walk the entire graph.

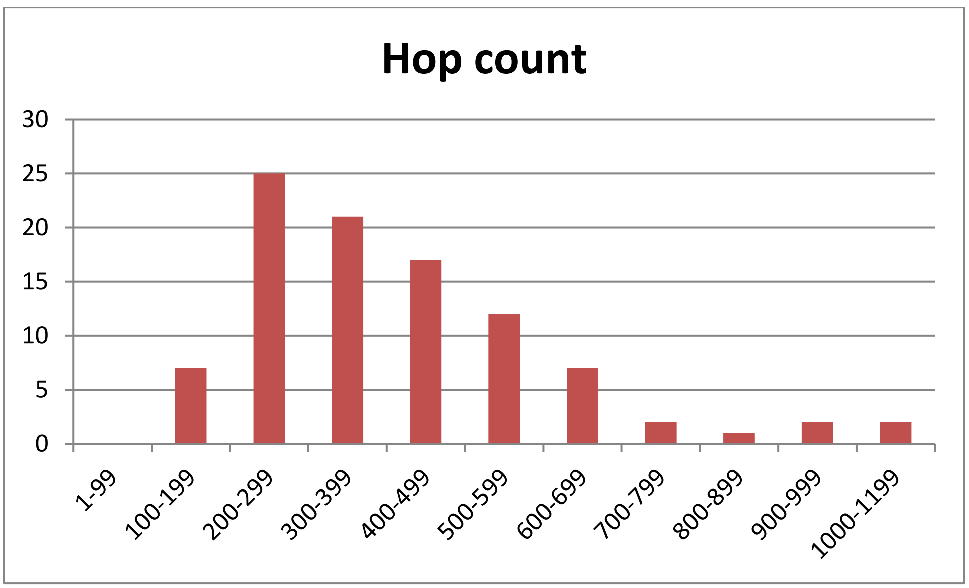

Figure 4.29. Hops Until Solution, Histogram Showing Distribution of 100 Trials

4.9. EXPERIMENT TO SHOW THAT THE TABU ACO CAN SERVE A RETAIL MARKET AND ABIDE BY DISTRIBUTION NETWORK OPERATIONAL LIMITS

Section 4.8 described how the TabuACO could implement a retail energy market and abide by network design constraints. However, there are situations where the static design limits are not sufficient to describe the dynamic capability of the grid design. The distribution network was originally designed for the unidirectional flow of energy, from centralized generation to distributed loads. The addition of distributed generation complicates matters. The original design assumptions are no longer valid, and the original constraints no longer sufficient to regulate the operation of the substation. A dependency exists on the load present. A test will be conducted to determine if the algorithm can leverage an external engineering analysis tool to determine the limits of the distribution network in order to accept as much PV power as possible. Again, by way of reminder, there are multiple loops at work as described by Figure 4.3. The energy market serves as an outer, planning control loop, while other systems operate inner loops to control substation operation as it runs. Only the energy market planning loop will be tested here. Implied in the discussion is that spinning reserve is still required for network operations, and that acquiring a contract for spinning reserve can be formed using the same technique as a contract for power. Customers will be expected to form contracts for spinning reserve, as a percentage of the power they buy, to the level specified by the regional RTO/ISO.

4.9.1. Experiment. A large-scale test on real utility data was conducted to test the scalability of the concept[15]. The scalability test will use real substation data supplied by the Electric Power Research Institute (EPRI) [59]. This data set (known as "Feeder J1") represents a distribution circuit located in the northeastern US with a 12-kV feeder circuit that serves approximately 1300 residential, commercial, and light industrial customers via

[15] The scalability test also tests the refinements in the code after numerous improvements were made as a result of benchmark testing. The Proof Of Concept testing used simple code which was not necessarily tuned. The POC code had a number of limitations including an inability to handle a graph where loops may be present in the network topology.

58 total miles of primary line. The peak load on the system is approximately 6 MW and the circuit contains 1.7 MW of customer-owned PV systems (Figure 4.35). The J1 Feeder data set includes a substation, several PV arrays (red circles), and a large number of electrical loads (indicated with the blue circles). The Feeder J1 data set is of interest because customers have complained of overvoltage conditions after several sets of PV arrays were added to the feeder. This was noted even though the feeder J1 has four voltage regulators throughout the circuit, and five capacitor banks (three of which are voltage controlled) to help control voltage [59].

This experiment analyzes the contribution of power from multiple sources of power, into the grid, drawn by multiple loads, with the contribution being limited to safe levels. It does not optimize the cost paid by the participants. This effort is deferred as a future work in section 6.2.5.1.

This data set was represented in the ACO architecture as a tree graph by identifying the substation's connection to the transmission network as the root of the graph, and all other nodes and edges as being downstream in the distribution network, resulting in 4840 nodes (1147 of which are nests), 4852 edges, and 1147 ants.[16] The EPRI Feeder J1 data set was analyzed with EPRI's OpenDSS software [121], which allowed for the visualization of the network (Figure 4.35) to be generated.

The scalability test consists of two experiments: a baseline experiment which does not use the ACO for correction, and a test which uses the ACO to limit the operation of the system to safe levels.

---

[16] Not all loads became nests because some of the loads had an issue which prevented them from being viably served (for example, some were present in the model, but not connected to the circuit).

The baseline experiment was an evaluation of the current market (Figure 4.2). This represents the use of the wholesale energy market to serve the needs of the Feeder J1 distribution circuit modeled as a single node. In this baseline experiment, all loads were served by the substation. No visibility is provided to the wholesale market related to any issues which arise on the distribution circuit. OpenDSS was also used to simulate the engineering results of using this data set without any correction by the TabuACO algorithm. This experiment provides the baseline information on the current methods being employed and their ability to mitigate voltage variations.

The ACO experiment was an evaluation of the market with PV assets contributing power to the grid which are interrupted and then restored. A comparison is made using the current methods as described in section 3 to the transactive energy market module added to operate a "retail level" market on the Feeder J1 circuit (Figure 4.3). In this experiment, the TabuACO algorithm was used to match loads with any available source of power. For the TabuACO experiment, ants were able to purchase from the substation or the PV arrays. Different power sources were assigned different prices with substation power being more expensive than PV power. Ants were given spending authority sufficient to buy PV power first, and as foraging progressed ants would eventually gain the authority to purchase substation power. In both experiments the OpenDSS software was used to validate each proposed purchase scenario. The OpenDSS software was used to validate each proposed resource reservation from the TabuACO algorithm and either accept or reject each matchup based on a test for overvoltages. The proposed power sources and loads are sent to OpenDSS which performs an engineering analysis and produces a data file. The ACO reads the voltage file produced by OpenDSS, and compares every voltage to the allowed upper threshold. If all of the voltages are within the acceptable voltage range, the proposed sale (represented by power and loads in the OpenDSS model) is accepted. If any voltage exceeds the allowable range, the proposal is rejected, and the ant is forced to forage further to find another source. To encourage the use of PV power, PV power was priced at 1¢ while substation power was priced at 7¢. The ant was initially authorized to purchase power at a cost below both. If she is unsuccessful in purchasing PV power, she will eventually become authorized to buy the higher priced power at the substation transformer. As the ACO operated, ants searched for 1% of each nest’s energy requirements with each sortie. All of the ants competed concurrently for the same resources.

The EPRI Feeder J circuit is different from the IEEE 34 Node circuit, and they have different issues. A study of Feeder J doesn’t find any constrained assets in the design. Instead, with the EPRI Feeder J circuit, voltage can be difficult to regulate.

The Engineering Analysis software used in this research is EPRI’s OpenDSS. OpenDSS is able to simulate a variety of timescales. In the Feeder J data, a voltage source is modeled at the substation bus. A voltage source is able to source tremendous amounts of current, up to the MVA rating of the transformer, and holds a steady voltage at that point in the circuit. The substation also models a tap changer which can automatically regulate the voltage. Due to the timing of the operation of the tap changer, a baseline analysis is conducted with the tap changer functioning. A voltage increase is studied by temporarily disabling voltage regulation. Details on the setup of OpenDSS, calls to OpenDSS, and other instructions on how to operate the system is described in the OpenDSS material [121].

The EPRI Feeder J data is defined in terms of OpenDSS defined elements. The mapping alluded to in Table 4.3 is described in more detail in Table 4.4.

Table 4.4. Engineering model to TabuACO model mapping

EPRI OpenDSS Class	TabuACO Class
buscoords	Node
generator	Node
load	Node
line	Edge
transformer	Edge
Vsource (voltage source at a substation)	Node
PVsystem (photo-voltaic power source)	Node

When an ant finds a food source (a power source), the reservation must be validated both financially and electrically. There are numerous concerns that must be considered such as overloading of individual assets, as well as the operation of the system outside the acceptable voltage range. The Engineering Analysis software used in this research is EPRI's OpenDSS [121], which can be used to simulate a variety of electrical loads and sources at a variety of timescales. The EPRI Feeder J1 data is used by the ACO architecture to create a model of the feeder and propose solutions. The solutions are evaluated by the OpenDSS software to determine whether demand was met and the voltage conditions on the feeder.

4.9.1.1. Proposed methodology. To solve this problem models are created to represent the feeder connectivity, the assets deployed along the feeder, the loads, the and the distributed generation. Each service requesting power is represented as a nest, and ants leave the nest to forage for energy. Two types of files are exchanged between the

TabuACO solver and OpenDSS. One file contains a model representing connectivity. This file is expressed as OpenDSS commands, and is imported into both OpenDSS and the solver. The solver edits this model to update the loads at each location and passes the updated model to OpenDSS for evaluation. The solver asks OpenDSS to render voltages. OpenDSS then publishes a file representing voltages. The voltages are examined by the solver for compliance with the required bounds. If any voltage at any point exceeds the bounds the proposal is rejected. If all voltages are within bounds the proposal is accepted.

4.9.1.2. Theory of operation. A transactive energy market as described in Figure 4.3 will be operated as the experiment. The TabuACO solver will match individual consumer loads with all available generators. OpenDSS will be used to perform an engineering analysis of the PV and substation contributions. Both the solver and OpenDSS will use a model of the EPRI Feeder J1 circuit. The solver will use the model to perform resource locking as multiple nests compete simultaneously for limited resources. Also, since the solver is an ACO, it will accrue virtual pheromones as a solution emerges.

4.9.1.3. Class design. An object-oriented design is used, and the objects are described in the following Ant, Edge, and Node subsections.

- **Ant**.

In the Ant class, attributes are used to record the ant's current node, a history of the nodes and edges which serve as the path home, and the nest node which the ant calls home. An attribute describes the ant's mode: if the ant is foraging, resting at home, or travelling home. A series of attributes keep track of the power sought, and the power acquired. An array which corresponds to the path home identifies the kVA constraints encountered along the path. In order to determine the ant's level of effort, the ant maintains a count of its movements (hops). An array of authorized price points defines what the ant is authorized to pay. An attribute also maintains the cost of traveling home. The authorized funds minus the cost home leaves the amount of funds available to purchase power. Functionally, the ant is able to deprecate a path, forage for food, or return to the nest. Functionally, the ant is able to deprecate a path, forage for food, or return to the nest. A UML class diagram of the Ant class is presented in Figure 4.30.

Ant
authorizedPricePoints, an array that describes the prices the ant is allowed to pay as a function of hopCount. **contractNode**, the ID of the node with which the ant signed a contract. **costHome**, the running total of costs incurred to lease assets along the path home. **kvarAcquired**, the amount of reactive power acquired. **kvarSought**, the amount of reactive energy sought. **kwAcquired**, the amount of active acquired. **kwSought**, the amount of real power sought during the current market interval. **locationID**, the nodeID the ant is currently at. **mode**, ant mode: foraging/home/homewardBound **nestID**, the nodeID of the ant's origin. **pathConstraintsHome**, list of kVA constraints on the path home. **pathEdgesHome**, the list of edges to the nets. **pathEfficienciesHome**, a list of efficiencies on the path home. **pathNodesHome**, the list of nodes to the nest. **pathReservationsHome**, a list of power amounts reserved at each segment along the path home. **previousLocation**, the node the ant was previously at. **ttl**, time to live.
deprecatePath(·), lay repulsive pheromone **forage**(·), look for food **returnToNest**(·), lay attractive pheromone if warranted

Figure 4.30. Class Design of Ant

- **Edge.**

An Edge class is used to model conductors and transformers. Edge properties are typically imported from a wires model. Each edge has exactly two ends. Due to the directionality of the graph, one end is identified as these are called the “up” node (arrowhead) and the other the “down” node (arrow tail).

Edge
attPheromoneKvar, attractive pheromone for kVar **attPheromoneKw**, attractive pheromone for kW **delivCommitKw**, downstream commitments **dnNode**, the downstream node (arrow tail connection) **ID**, the edge ID within the network **kvaRating**, the conductor’s rated capacity in kVA **length**, conductor length in feet. **linecode**, the class of conductor described by the dataset. **name**, the name assigned by the dataset to this edge **recdCommitKw**, commitment to allow current flow upstream **repPheromoneKvar**, repulsive pheromone for kVar **repPheromoneKw**, repulsive pheromone for kW **travelCostPerKva**, cost in dollars for each kVA to travel through the edge. **travelEfficiency**, the efficiency with which energy traverses the node. **type**, conductor/transformer **upNode**, the upstream node (having an arrowhead)
evaluateQuery(·)

Figure 4.31. Class Design of Edge (UML Class diagram)

Each edge has an ID, a name imported from the wires model, a kVA rating, a length, and a travel cost which is expressed as a travel efficiency. Pheromones are deposited on edges. One attribute is used to maintain the attractive pheromone level, and another to record the repulsive pheromone level. Edges are grid assets and as with any resource have to be reserved. An attribute is used to describe the commitment of power flowing upstream, and other for the commitment of power to flow downstream. A UML class diagram of the Edge class is presented in Figure 4.31.

- **Node.**

A Node class is used to represent service locations, generators, poles, and every other asset that is not an Edge.

Node
ID, the node ID within the network **name**, the name assigned by the dataset **type**, asset type: transformer/pole/substation/PV **kvaRating**, the node's kVA rating **xloc**, geographic x location **yloc**, geographic y location **upstreamEdges**, a list of edges which lead upstream to the substation **downstreamEdges**, a list of edges which lead downstream **travelCostPerKva**, cost in dollars for each kVA to travel through the node **travelEfficiency**, the efficiency with which energy traverses the node **delivCommitKw**, the kW commitment sold in the interval **delivCommitKvar**, the kVAr commitment sold by node **recdCommitKw**, the kW purchased by node **recdCommitKvar**, the kVAR commitment purchased by node **desiredkW**, amount of real power this node needs **desiredKvar**, amound of reactive power this node needs **fundsAvail**, the unspent funds currently available within the current market interval. **antAgents**, a list of ants which serve this nest **pricePerKwh**, price per kWh for energy from this source **availKw**, uncommitted real power available **availKvar**, uncommitted reactive power available **delCntrctFromNode**, an array of entries describing a contract for delivered power **delCntrctKw**[·], a list of contractual kW **delCntrctFunds**[·], funds owed to the node **delCntrctAnt**[·], ants which have formed contracts
evaluateQuery(·)

Figure 4.32. Class Design of Node

The Node's properties are typically imported from a wires model. Each node has a geographic x-y location, an ID, a name, a rating, a cost for using the asset, an efficiency in using the asset. Since the network is a graph it can have “zero to many” upstream edges, and “zero to many” downstream edges. Class attributes store reservations for power to flow downstream and upstream through the node. If the node has load characteristics, an attribute tracks the amount of power the node is seeking (as a Nest). The funds available to spend are tracked, as are the ant IDs which belong to the nest. If the node has generator characteristics, the selling price is kept as an attribute, as is the available uncommitted power. Class attributes also track the funds owed to this node, or owed by this node to other grid assets. A UML class diagram of the Node class is presented in Figure 4.32.

4.9.1.4. Ant agency. The ACO foraging process imbeds an evaluation of the conductor (edge) and transformer (node) constraints. Figure 3.1 and Figure 3.2 show the basic function of the TabuACO. Figure 4.33 now shows these concepts applied to the energy market. The pathways originally described in Figure 1.1 now model conductors. Nodes of Figure 1.1 model transformers, poles, and other transformational assets. The consumer’s system represents the “nest,” and authorizes an ant to look for “food.” This purchase involves traveling a viable path, and not only finding energy for sale, but finding it at an affordable price. Every step that the ant takes causes a pedometer to increment. As she forages a great distance, an authorization function allows her to spend more. She may forage in a way that brings her back to a food source she passed by previously because she could not (initially) afford the purchase.

Upon finding a source with an acceptable price, the ant forms a contract with the producer and lays attractive pheromone back to the nest. She thus communicates with other ants using clues left in the environment.

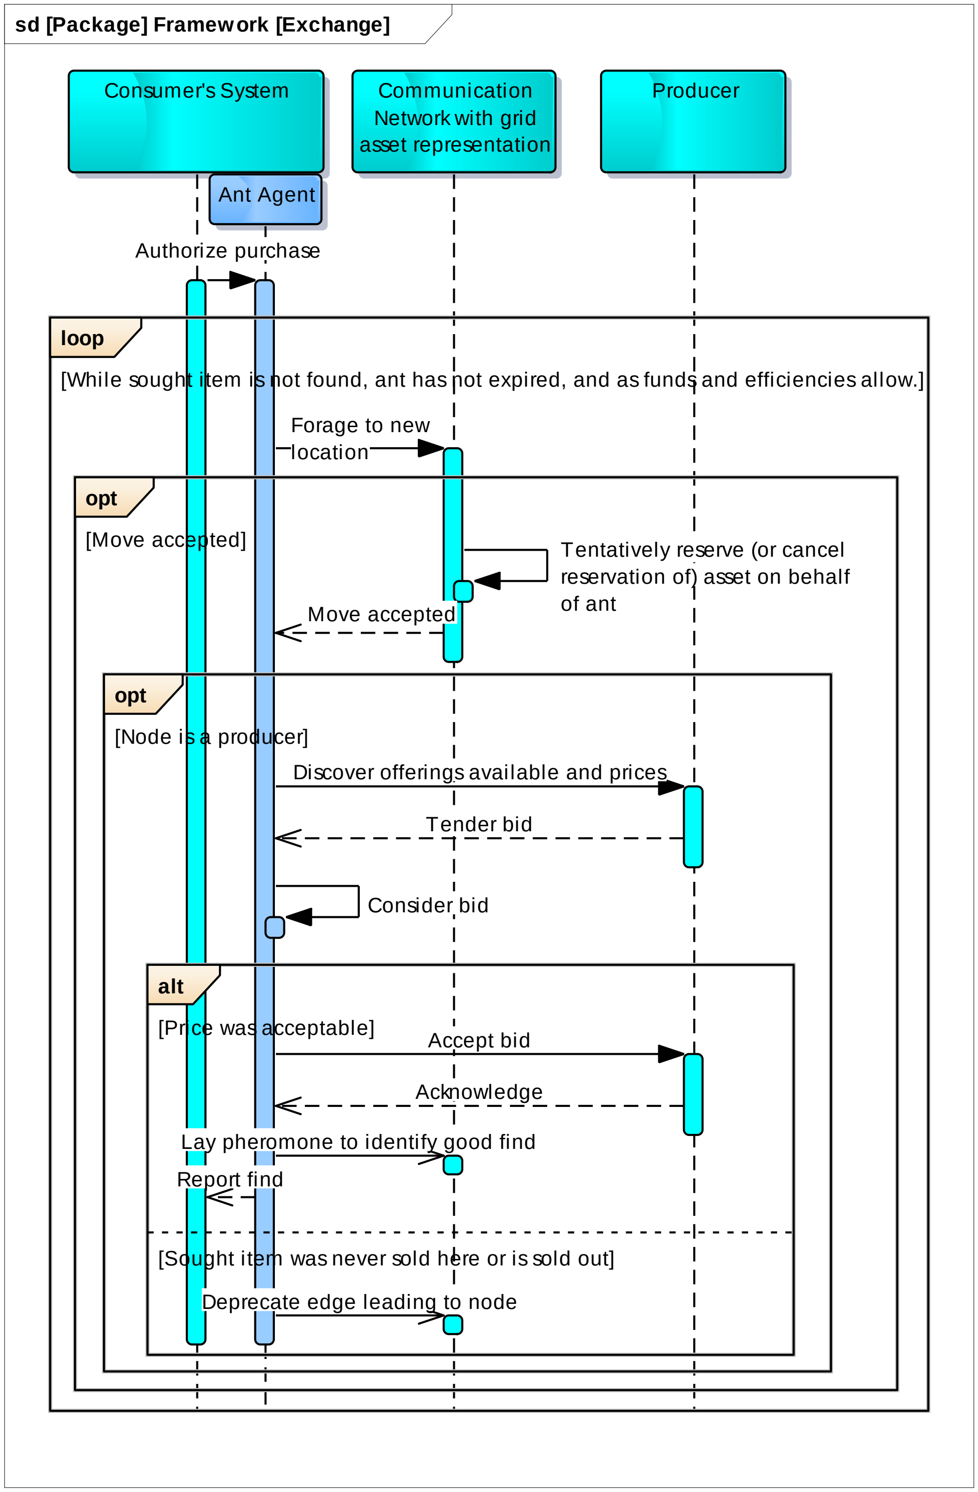

Figure 4.33. Messaging (SysML Sequence Diagram)

4.9.1.5. The TabuACO applied to transactive energy. In the context of this application, each service location will be represented as a "nest." Each ant is tasked with fulfilling a small percentage of the energy required for that service location using a given amount of money (and pay a certain rate). As the ant forages, she makes tentative contracts with the network elements along the path traversed to reserve the resources she needs. If the edge or node will not permit the full amount required, the amount available is reserved. If a network element that cannot be reserved is encountered, the ant turns back and forages elsewhere.

In addition to the availability of energy resources, the ant also considers the cost of travel through the network assets. If she has insufficient funds to travel further, she turns back toward the nest. She may also turn back because a leaf node is encountered, or simply because the random foraging process results in such a choice. When the ant travels back along the same path toward the nest, she unlocks the resources she previously reserved.

When the ant finds an acceptable and available source of power, she will form a contract, accept the offer from the generator, and lock in a portion of that resource. On the return to the nest, she will reserve every conductor, transformer, transmission and distribution asset required to move the power from the generator to the nest. Certain contracts may need to be adjusted downward if the original energy amount cannot be reached due to asset or generator constraints. Pseudocode for this process is depicted in Figure 4.34.

```
Initialize network and ants
While one or more nests are unsatisfied
    Set one or more ants to have an antMode of “foraging.” Each ant is given a time limit to search for food, a specific quantity of a specific commodity to find, and a formula which authorizes increasing price payments as a function of ant hop count. The formula identifies an initial bid and a maximum cost the buyer is willing to pay.
    While ants are on sorties
        Case antMode of
            Foraging:
            Select an edge to traverse. As the ant moves, she must maintain a path history that shows how to return to the nest. If she moves further away from the nest, she must form a tentative contract with the edge owner to carry the item she is looking for homeward. The asset must consider the sum of all contracts formed to see if capacity remains to accept another. If she moves towards the nest (retracing her footsteps) she must cancel the contract previously formed along that edge (and with the node she just left behind).
            Upon arrival at a node, the ant attempts to form a contract with it. If it is a grid asset (e.g. transformer), the contract is to carry the prize across it. If the node is a generator, the ant examines the node to see if it contains the sought prize for an acceptable price. If so, the ant claims the prize and sets antMode to be “homeward bound.” Otherwise, she continues foraging.
            If the new node is a leaf node, and the prize was not found, the ant lays down a repulsive pheromone (particular to the prize type) on all adjoining edges.
            If the new node is not a leaf node, does not contain a prize, all adjoining edges except the homeward edge are considered. If none of the edges (save the homeward edge) contain attractive pheromone, and all (save the homeward edge) contain some repulsive pheromone, then all edges (including the homeward edge) are deprecated with repulsive pheromone.
            If an ant forages unsuccessfully for an extended period of time so that its time limit expires, then antMode will be set to “homeward bound.”
            Homeward Bound:
            As an ant returns home, and a prize was found, the ant lays down attractive pheromone along each node and edge along the way. Upon reaching the nest, the antMode is set to “home.”
```

Figure 4.34. TabuACO Pseudocode for Solving the Transactive Energy Market

```
            Home:
            Ant reports findings to nest, all contracts formed and the amount spent. The
            ant sortie is over.
        EndCase
        Decrement the ant's time-to-live
        Increase the ant's authorization to pay more for the prize.
    Endwhile
    Evaporate pheromones (positive only)
Endwhile
Report optimal paths to each food source.
```

Figure 4.34. TabuACO Pseudocode for Solving the Transactive Energy Market (con't)

4.9.1.6. Engineering validation of proposed contract. When the ant finds a food source, the potential sale must be validated with engineering analysis even though the sale is known to be valid financially. There are numerous concerns that could be considered by the engineering analysis such as overloading of individual assets, as well as operation of the system outside the acceptable voltage range. Feeder J, being a real circuit, has many conditions which could affect its performance. However, EPRI identifies one condition in the circuit which is particularly bothersome to customers [59]. This particular condition will be the focus of the simulation and the event the ants work to guard against.

The "Feeder J" circuit is fed primarily from a substation in the SE portion. In Figure 4.35, solar arrays are in red. Loads are small blue circles. Ordinarily voltages are regulated fairly well. In a feeder with PV arrays installed, troubling overvoltages can happen on a partly cloudy day. When the wind drives the clouds along, it can move a shadow over the location of the PV array. This will cause the contribution from the array to drop off precipitously. If the PV array was supplying a significant percentage of the load, this loss will cause the local voltage to drop as well. The voltage regulation equipment in the substation will adjust one step upward (+2.5%) to correct the voltage back up into the proper range. When the shadow moves off the array the PV will return to contributing full power. If this occurs quickly, it can cause the voltage to rise before the voltage regulators have time to react. A voltage swell will occur, and potentially damage done to consumer's devices before the substation tapchanger can move a step downward (-2.5%.)

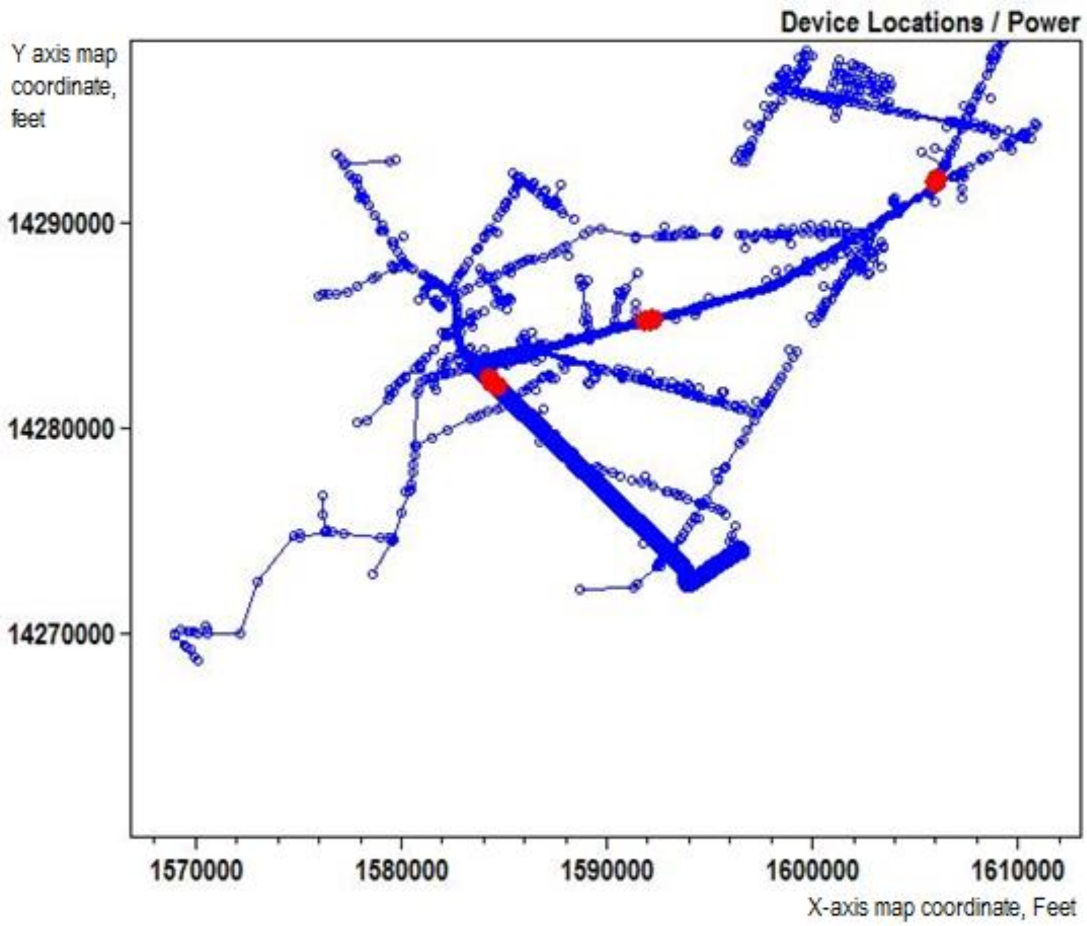

Figure 4.35. OpenDSS Rendering of Feeder J Loads

During the night, and on overcast days, the Feeder J circuit relies entirely on power from the substation. The substation is located at the lower-right end of the wide blue trail depicted in Figure 4.35. The result is a well-behaved distribution voltage which stays well away from the upper and lower limits. Figure 4.36 shows the allowable high and low voltage tolerancces as horizontal red lines at 1.05 and 0.95 pu respectively. Nominal voltage is of course provided with a voltage of 1.000.

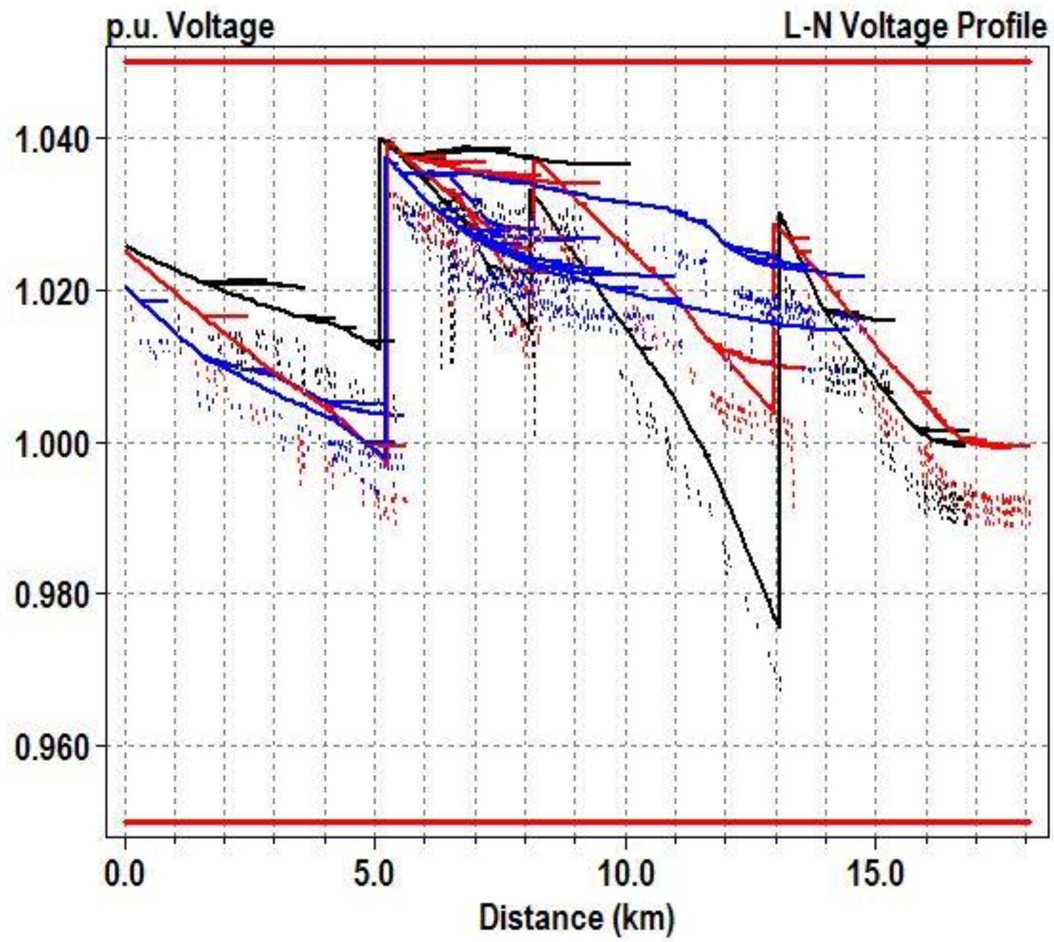

Figure 4.36. OpenDSS Rendering of Feeder J Voltage as a Function of Distance from the Substation

When the power from the PV arrays joins with the contribution from the substation, it causes voltages at some locations to jump above the tolerable limits. A baseline evaluation of the feeder J1 was performed to find the voltage versus distance from the substation bus. Different customers will have different experiences depending on their distance from the substation, and their proximity to the PV arrays. The red, blue, and black solid lines in the middle of the Figures 4.36 through 4.38 are the three phase voltages. The solid lines for each phase are the medium voltage lines and the dashed lines below are the low voltage lines. A nominal voltage would be at 1 pu right through the middle of the graph. In Figure 4.36, the voltage at the substation starts at 1.02 pu (102% of nominal). As power travels down the feeder, voltages drop due to resistive losses in the conductors. One can see the medium voltages (represented as solid lines) steadily decline out to a distance of 5 km from the substation. The low voltage lines (represented as dashed lines) lose voltage even more quickly over short distances. These are seen below the solid lines forming “the beard” in the diagram. In Figure 4.36, the voltages jump at a distance of 5 km up to 1.04 pu due to the action of voltage regulators and capacitor banks. This pattern can be seen to repeat as the distance from the substation increases as a means of maintaining an acceptable line voltage. Voltages drop from 5 km to 8 km and are boosted at 8 km due to voltage regulation, as can be seen in the voltage drops from 8 km to 13 km which are boosted at 13 km. Voltages continue to drop from 13 km to the end of the feeder, where we see that medium voltages are at nominal (1.00 pu), and low voltages are lower but acceptable, running at 0.99 pu. When voltages are properly regulated, they stay well within the red lines running horizontally along Figure

4.36 at 1.05 and 0.95 pu. However, when distributed generation is present, the ability of the voltage regulators to maintain the voltage within this range may be compromised.

The voltage increases we saw in Figure 4.36 occurring at 5, 8, and 13 km are due to static volt/var corrections in the circuit. One can see that if these corrections were not made in the circuit, voltages would become unacceptably low. However, these same locations now become vulnerable to excessive voltages when the entire circuit runs high The overvoltage scenario is depicted in Figure 4.37.

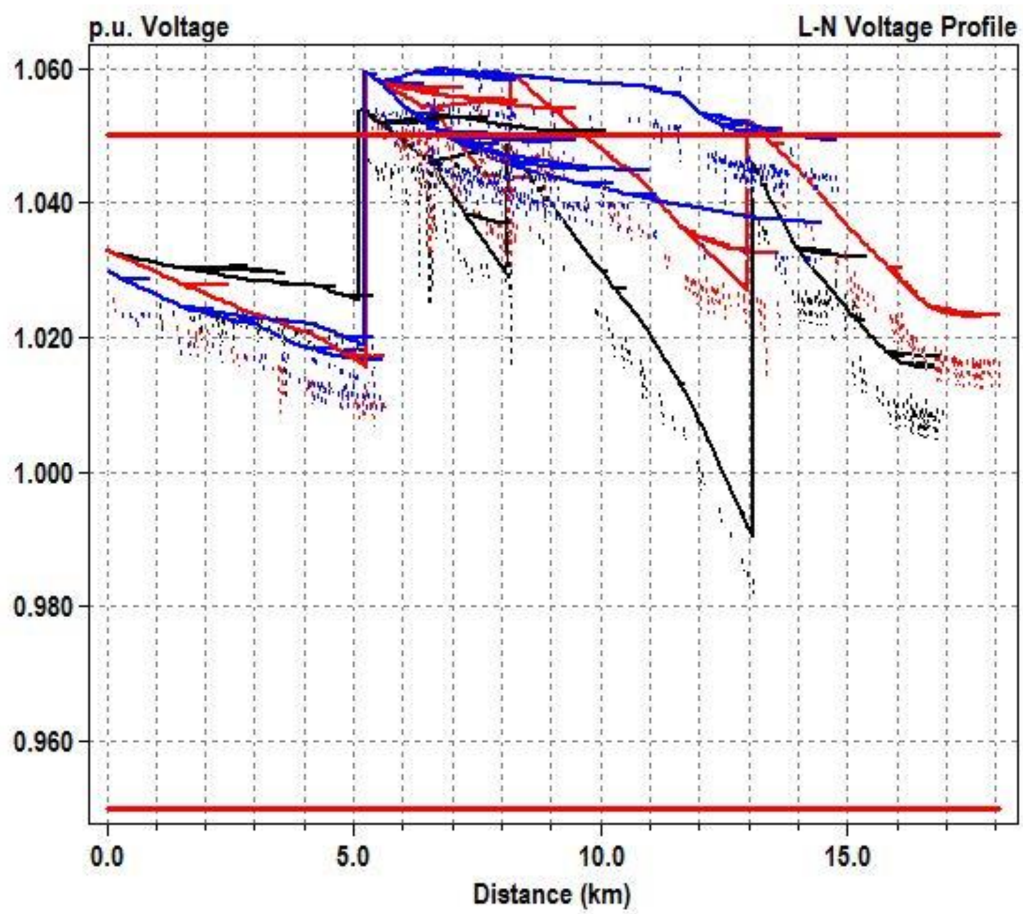

Figure 4.37. OpenDSS Analysis of Feeder J Voltage After a Cloud Moves off PV Array Allowing Full PV Contribution

4.9.2. Expected Results. As the ACO operates, ants will take a small (1%) portion of each funded nest and search for energy sources. All of the ants will be authorized to operate concurrently, but due to computational constraints will take turns. Ants will trim leaf nodes and begin to follow useful trails to find power sources. As time goes by, the ant agents authorization to spend more money will increase, and she will eventually buy the higher priced local power (provided that the consumer has authorized such a rate).

4.9.3. Results. In the baseline experiment substation power was supplied to the loads, and no overvoltages were detected. The (simulated) voltage regulators were able to compensate for any load changes and maintain nominal voltage (as described above and in Figure 4.38.)

In the ACO experiment the addition of unreliable PV power was considered. When PV power freely supplements substation power and there is an interruption it results in overvoltages as depicted in Figure 4.37. This is the result when the proposed algorithm is not applied. When the TabuACO was implemented, ants foraged for power from a variety of sources. Ants (from nests representing loads) formed contracts with low cost PV sources first until an overvoltage occurred. The overvoltage predicted was attributable to the PV source. The ant then worked instead to obtain power from other sources until arrangements were made that proved to be viable. The TabuACO algorithm was able to limit voltages to safe levels in the ACO experiment (Figure 4.38.) High voltages are found on the medium voltage (solid) lines at a distance of 5 to 11 km from the substation. However, the algorithm tested the low voltage (dashed) lines to ensure that no loads would suffer overvoltages.

The TabuACO is a stochastic optimization process and so different runs sometimes produce different outcomes. A typical outcome is found in Table 4.5. and in Figure 4.38.

In this example, a total of 1526 kW was purchased from the PV sources before a 3% overvoltage would have been encountered. The balance of the projected load was served by the substation. Repeating the experiment for 80 trials resulted in a mean of 1752 kW purchased from the PV sources with a standard deviation of 266 kW (with the balance of the required load coming from the substation). Every run resulted in a successful outcome. Other runs of the solver resulted in different PV allocations, but all met with the same upper limit although the PV allocation differences caused differences in the shape of the voltage curves within the voltage bounds.

Table 4.5. TabuACO Findings Regarding Limits to PV Contributions on Feeder J

Source Name	Source Type	kVA rating	kW purchased
C_Existing13	PV	16	16
B_Existing12	PV	8	8
A_Existing9	PV	11	11
C_Existing2	PV	12	12
C_Existing11	PV	11	11
B_Existing7	PV	11	11
3P_ExistingSite3	PV	836	572
3P_ExistingSite2	PV	209	209
3P_ExistingSite4	PV	523	464
B_Existing3	PV	11	11
C_Existing5	PV	22	22
C_Existing10	PV	11	11
3P_ExistingSite1	PV	314	168
Substation	Market mix	16,000	9969

It may also be noted that as the TabuACO ran, ants would initially forage from each nest everywhere, in all directions. However, as ants discovered leaf nodes, deposited repulsive pheromone, and discovered unproductive trail nodes which led to leaf nodes, soon entire trails were deprecated. This, along with any positive pheromones which may have been deposited by other ants serving nearby nests, led the remaining ants to quickly find viable power sources and converge to a solution. It was observed (but not quantified) that as time passed, the algorithm appeared to run faster. Ants spent less time aimlessly foraging, and travelled directly to nearby power sources.

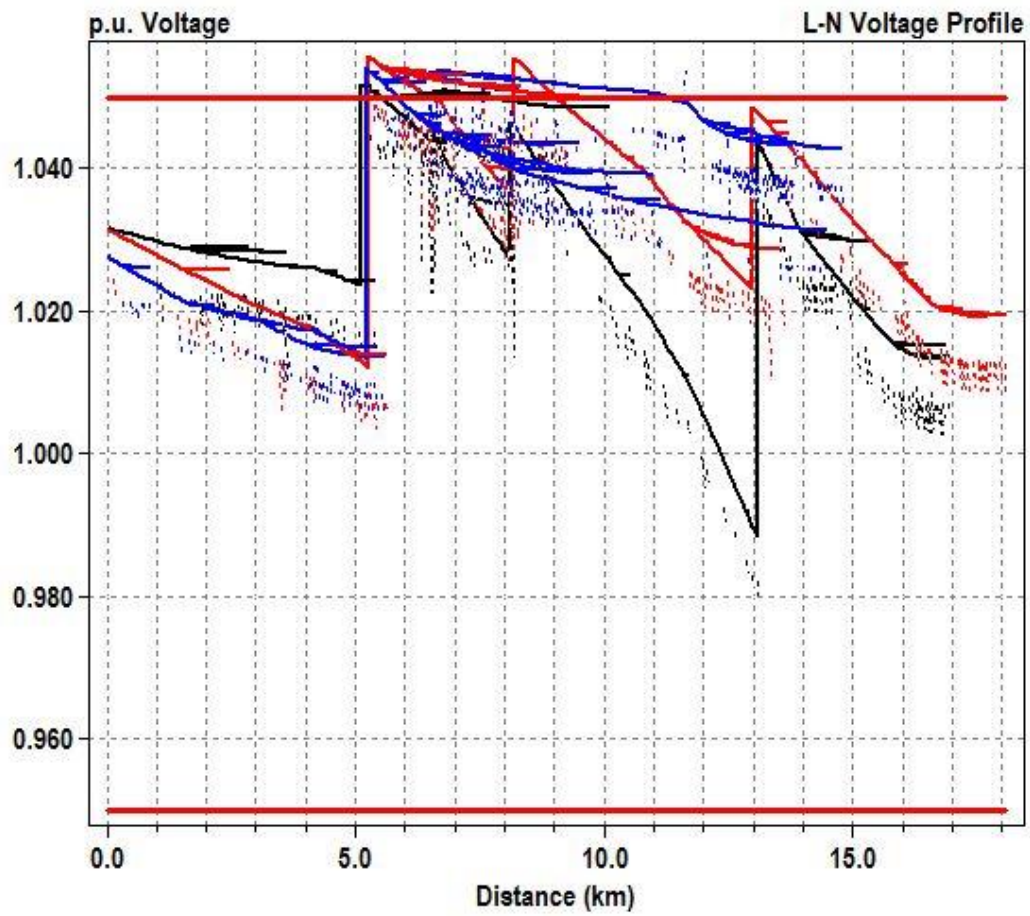

Figure 4.38. OpenDSS Rendering of Voltages Along Feeder J with Limits Applied.

4.9.4. Discussion of Results. An EPRI model of a real distribution circuit was studied. Using an ant colony approach, a retail energy market with location marginal pricing was simulated within a distribution circuit. Figure 4.37 shows the behavior of the circuit when the algorithm limits are not applied. In this baseline experiment, it shows the voltage escalating to 6% above nominal. Overvoltages can damage equipment. Compare this to Figure 4.38 which shows the behavior of the circuit when the algorithm limits are applied. The ACO experiment showed that the voltage could be successfully limited to a target value (3%) above nominal.

The market endeavors to plan a suitable fuel mix prior to the real-time operation of the grid. The transactive energy market modeled in this approach resolves market contracts in a bottom-up fashion – whereas the conventional (wholesale) energy market creates arrangements in top-down fashion. Furthermore, ACO's are known for their emergent properties. Solutions arise out of the data in a self-organizing fashion. The ACO solver operated each nest independent of every other nest. Competing ants actually collaborated to deposit pheromones along trails and share information (this is known as "stigmergy"). The system self-organized without any centralized coordination. This is an important point because a distributed energy market, accompanied by distributed control of the grid could fit well with distributed generation.

It should be noted that there may be an increased cost to install and operate communication equipment which can capture, communicate, and resolve the market and operational issues. To some degree the cost is the cost of improving the granularity of the analysis – regardless if the analysis is performed in a distributed fashion or a centralized fashion. A distributed architecture also requires a different approach to implementing or enforcing global goals for the system. A systemwide maximum or minimum may be difficult to achieve without a systemwide measurement. On the other hand, that doesn't prohibit such measurements from being made, disseminated and acted upon. Some measurements (such as line frequency) are disseminated systemwide automatically. Other measurements however shift the burden from moving data from the field to a central location, to instead move measurements from a central location to the field. The scale of the experiment is also notable. The results show that the solver operated an energy market for an entire distribution circuit. The 14 sources identified in Table 4.5. are used to feed 1147 loads. If an entire distribution circuit can operate as an energy market, then it should be possible to integrate this market with a transmission level energy market. An interface could be created at the substation to tie these markets together. (This research shows the possibility of operating at least all of the medium voltage distribution circuits as independent markets.) But this is not to say the entire system must convert to one methodology or another. When tied together through a well-defined interface, it would be possible for a mix of methodologies to be employed. All of the markets could be operated bottom-up, some of the markets operated bottom-up and others top-down, or the entire market in a top-down centralized fashion. This research added equipment to provide granularity, to expose distribution circuit issues and, to allow small-stakeholder participation. The TabuACO solver employed a methodology which lends itself to deployment in a distributed manner, but in truth, it ran on a "centralized computer" which merely modeled the possibility of distributed operation. By using a technique that incrementally converges to a solution, and validating each step towards a solution with an engineering analysis package such as OpenDSS, it was possible to arrive at a solution that was both financially and technically viable.

4.10. CONCLUSIONS REGARDING THE APPLICATION OF THE TABU ACO TO THE ENERGY MARKET

The ACO implemented a transactive energy market for one market window for an entire feeder. A “Proof Of Concept” test (section 4.8) and a “Test To Scale” (section 4.9) both found that the transactive energy market was able to communicate the information needed to limit contributions from certain sources and promote the contribution from other sources in order to avoid electrical catastrophes on the grid.

The TabuACO was able to identify the appropriate fuel mix ratio which maximized the PV contribution while avoiding the conditions that caused the PV contribution to create voltage spikes that would overpower or outpace the substation voltage regulation. This analysis mix kept service voltages within the allowable limits which would be violated using the uncontrolled mix. The tabu property of the algorithm was able to quickly eliminate unproductive leaf nodes and unproductive pathways. The ACO property of the algorithm was able to guide the ant agents to low cost PV power sources, verify the proposal, and resort to more expensive substation power when the PV mix reached its viable limits.

The TabuACO performed well on this problem because it was able to trim away ineffective paths. Ant agents were able to easily find power sources, and generally avoid wasted effort. While an ant colony optimizer is not the only method that can arrive at a solution, the difficulty in most other methods is that they require data to be collected in a central location. This can work well for hundreds or possibly thousands of participants, but scaling to millions of participants can be difficult. In a real-world grid application, each interconnect may contain hundreds of millions of meters. Large networks require more computing time than small networks. Distributed methods offer advantages in keeping the analysis smaller, moving less data, and eliminating single points of failure from the architecture. If a distributed or federated architecture is used, the designer would still need to determine how it would be sized and how the architecture itself would be optimized. The TabuACO manages this by allowing the consumer to set the funds the ant is authorized to spend, and the network identifies how much it costs to move power from one location to another. This effectively puts the ants “on a leash” so that they cannot wander very far from their nest. The size of the ants roaming territory is determined by consumer preferences, not by architectural constraints.

The legacy wholesale market provides a top-down market solution and limits itself to only consider transmission level loads. This research introduces a bottom-up market solution that considers every load and every constraint – no matter how small. By moving the energy market to a local level, local issues can be identified and addressed. The fact that the solver operated an energy market in a distributed, bottom-up fashion is significant. Distributed architectures are well known for their robustness. By enabling local problems to be solved locally, it:

- Eliminates the challenges in moving vast amounts of data long distances;
- Eliminates the difficulties in keeping records in a central location current;
- Eliminates the opportunity for catastrophic failure due to a weather disaster taking out the centrally-located control center;

- Creates an opportunity for local markets to operate with the assets that remain standing after a natural disaster destroys grid assets that ordinarily tie a locale to the grid.

The goal of the paper was to show that calculations for decentralized market operation are possible; and that swarm intelligence such as Ant Colony methods can offer viable alternatives to traditional methods. This is accomplished through the self-organization of energy buyers and sellers within a model of a wires network.

While cloud cover is mentioned in the data set, this approach was not reliant on a weather forecast. The speed of the regulator recovery and the tolerance of the grid to accept both load and PV generation is used by the ACO to determine the limits. If instead the weather is overcast it becomes a moot point – the PV will not contribute very much to the grid regardless of what anyone does. If the weather is sunny with no clouds, the limit serves as a safeguard if an unexpected cloud appears.

5. CONCLUSIONS

The TabuACO has been found to perform well, and even outperform conventional ACOs in several benchmarking problems. The TabuACO performs better than a conventional ACO because it preserves information unlike a conventional ACO. When an unproductive path is found, it is not merely ignored, it is noted. This difference in the algorithm is important. The repulsive information deposited in the environment was found to be sufficient to solve network problems just as attractive information (in a conventional ACO) can solve a network problem. The TabuACO was tried with repulsive data only and it performed similarly to its performance with attractive data only. The TabuACO with both repulsive and attractive data significantly outperformed either pheromone by itself.

Testing of the TabuACO against a variety of benchmark problems led to the discovery that the rules in use are very effective for trimming unproductive leaves from a tree. The rules allow deprecation to spread so that entire branches may be deprecated from a tree. However, it was discovered that as a graph becomes less tree-like and more mesh-like, the effectiveness of the TabuACO is reduced. The increased interconnectedness of the graph prevents deprecation from spreading very far. A fully interconnected graph may prevent deprecation from being applied altogether. The development of computationally-safe rules which could be used to trim highly-interconnected nodes is a topic for future research.

The testing also yielded other finds. Testing against the QAP problem demonstrated that the algorithm could scale to tackle large problems. Furthermore, the code did so without storing a model of the entire problem. The solver stored data for only the most promising paths. In the benchmark problems it retained a model of the 1000 most interesting nodes yet the model was able to effectively represent a puzzle containing many hundreds of millions of nodes. Such storage efficiency could be useful for computationally constrained applications. This characteristic can be particularly useful in distributed applications which either cannot store, or do not need to store the entire model.

The TabuACO was applied to the Smart Grid Energy Market. One important remaining challenge in the Energy Market is to find a way to extend the existing market mechanisms, refine their granularity, and reach all of the individual stakeholders. The market hasn't been operated to such a scale. This research modeled a distribution circuit and matched individual loads to power sources while observing the constraints of the circuitry. It found that it was able to operate each service location as a nest, and with thousands of nests operating concurrently, the TabuACO generated contractual agreements between power producers, network asset owners, and consumers. The improved granularity allowed market mechanisms to protect distribution assets much like the mechanism the wholesale market currently uses to protect transmission assets.

A test of the TabuACO using the IEEE 34 Node feeder model was able to identify a constrained transformer and protect it with market mechanisms. (Higher priced local generation was brought on line to supply power so the constrained asset did not have to be overloaded.) In a similar fashion EPRI Feeder data along with OpenDSS were used to discover the limits to which local PV arrays could safely contribute to the grid.

By operating a complete distribution network as a single power pool, it can double as the node which currently represents the aggregate load at the same location in the existing wholesale market. This means that an entire distribution network can operate as a microgrid, independent of other connections, yet buy and sell power through the transmission grid to other participants. It means that the existing nodal wholesale markets can operate much as they currently do, but the concept of locational margin pricing can occur within the system which is represented as a single node at a higher level.

6. FUTURE WORK

There are many exciting opportunities for additional research related to this dissertation.

6.1. ACO RESEARCH

The fact that clear improvements in performance can be made in certain situations with the simple rules provided in the TabuACO imply that more sophisticated rules can yield additional improvements.

6.1.1. Advanced Subnet Deprecation Rules. The rules presented above are merely a starting point for additional research. Additional rules could be developed in which nodes with a degree greater than one could be deprecated.[17] The ant maintains a history of the path home. This history could be expanded, and the ant analysis improved so that portions of a subnet which do not contribute to the solution could be identified and deprecated.[18]

With the current deprecation process, foraging ants which happen across a useless node will trim it. With more sophisticated deprecation rules, it may be that an ant must transition to a special form of foraging so that she temporarily places a higher priority on closing out the investigation of the subnet than on foraging for food. A third form of pheromone may be beneficial in this case. When an investigation finds that a subnet has been investigated, and found to potentially contribute to a solution, the environment

---

[17] At one point a mistake in the Steiner tree research allowed multi-edged nodes at the bottom of a subnet to be deprecated. This yielded faster convergence, however, the rules were not computationally safe.

[18] It may be that the ant would have to be aware of the objective function (such as travel cost) in identifying edges to deprecate.

should be marked with a neutral pheromone to indicate this fact. This neutral pheromone should prevent the ant from investigating the territory again.

The new rules may also introduce another dimension of tuning which is required for this third type of pheromone. Tuning might control the point at which an ant places a higher priority on subnet investigation than on foraging for food.

6.1.2. Refinement of the Path Selection Formula. The research above experimented with several forms of computing the probability of path selection as a PDF with a variety of formulae[19]. The literature review shows quite a bit of experimentation has occurred into various approaches for depositing attractive pheromone and leveraging the data found in the environment to select a path. Quite often the attractive pheromone value is raised to an exponential power to amplify trace amounts of pheromone and allow the pheromone to dominate the selection. Some experimentation with different approaches to repulsive pheromone deposition and utilization can be expected to occur with subsequent research.

6.1.3. Parallel Processing. The Energy Market Application was written in a way that thousands of nests each sent out a single ant at a time on a sortie, but with multiple processors, it would be possible to take advantage of the hardware resources to deploy multiple ants simultaneously. Research could be conducted with multiprocessor platforms to explore practical interfaces for model sharing among a multitude of processors.

---

[19] Compare the simple formula used in section 3.2.3 and 3.4.3 to the more sophisticated approach in section 3.3.1.3.

6.2. ENERGY MARKETS

Transactive Energy, if implemented to scale, could fundamentally change the operation of the grid. However, additional research is needed to drive it to commercialization.

6.2.1. A Complete Transactive Energy Market. The exercise described in section 4.7 was limited to one commodity. However, in a real system, multiple commodities must be balanced. In a practical system, real and reactive energy must both be balanced, and ancillary services must also be provided. In addition to frequency regulation and spinning reserve, distributed generation (itself) brings about the need for additional services. Distribution networks all have three phase power near the substation. However, such networks tend to be radially fed with single phase circuits covering large remote areas. Power contributed on one phase will not propagate to the other phases. Furthermore, phases are separated in time. Energy markets need to address each phase. It is entirely possible that a phase imbalance could develop. Abundant energy could appear on one phase with a deficit on another. Contributions of power to one phase could also result in congestion that does not appear on other phases. Contributions can cause a voltage imbalance between phases, as well as movement of the phase angle. A three-phase participant in the affected region would have the option of addressing a market phase imbalance by buying energy inexpensively on one phase and reselling it on another.

An ACO might allow multiple ants (from each nest) to forage for each commodity concurrently.

6.2.2. Secure Transactions. All contracts between machines should be protected by authentication. Establishing security can be challenging. A "chain of trust" must be established. New developments in blockchain technology offer an interesting research opportunity in its application to transactive energy markets [122].

6.2.3. Expansion to Multiple Markets. This dissertation studied a circuit which spans the distribution circuit. An interface can be formed at the substation transformer (and substation meter) so that the entire downstream circuit is aggregated into one market and represented by a node in the wholesale energy market. This is profound because it would finally allow all markets to be joined. A hierarchy of suppliers and demands can be built to join all related markets [123].

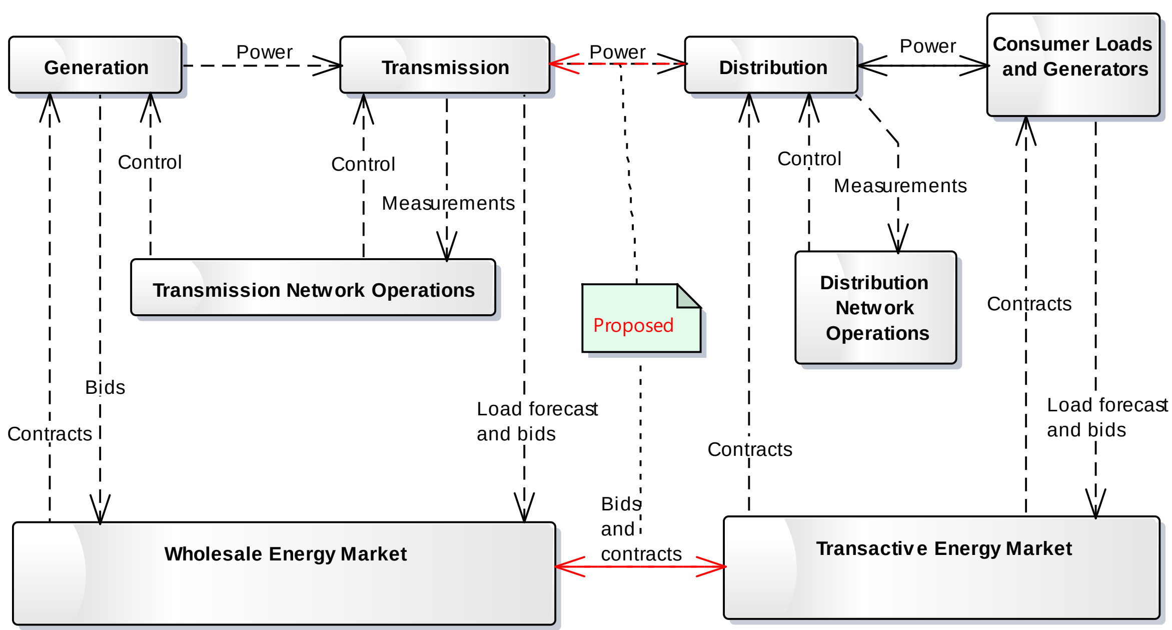

Figure 6.1. Joining Retail (TE) and Wholesale Markets

The needs and/or capabilities of the aggregated distribution circuit can interact at the wholesale level to complete the scaling of the market. The support described in Figure 4.3 can be extended as described in Figure 6.1.

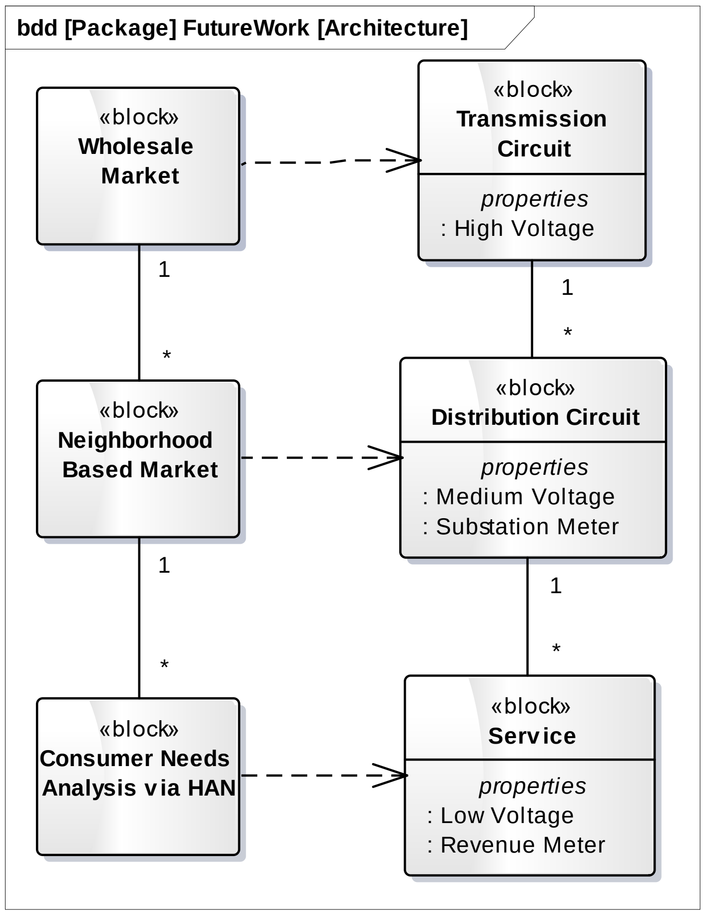

Figure 6.2. Joining Wholesale, Retail, and Submeter Markets

An interface can be formed at the substation transformer so that all downstream loads and generators are aggregated into one market and represented by a node in the wholesale energy market. Bids and contracts could flow across this interface, while power flows across the existing electrical interface. Today's Wholesale Energy Market resolves bids every 15 minutes. The Transactive Energy Market would be expected to do the same.

The needs of the consumer should also be identified in an automated fashion. This aggregation technique can be extended downward so that a Home Area Network (HAN) strives to identify consumer requirements and load sensitivity to price variation. Each market resolution process correlates to a domain as depicted in Figure 6.2.

Thus, a “smart grid” could be formed which has the ability to provide location marginal pricing with full granularity, and provide a market mechanism to respond to issues large and small. Such an exchange would seemingly benefit all stakeholders. Large generators benefit by moving away from flat tariffs toward a cost-based pricing system. Small generators would benefit by participating in additional types of markets only available today to large generators. Consumers would benefit by having a say in which power is purchased on their behalf, by being able to potentially store and resell power, and by having local (microgrid) market alternatives to centralized power. Transmission network operators would benefit by having an additional (aggregate) entity which they can call upon in time of need.

By having consumers purchase ancillary services as they buy energy, and by joining federated markets, it could become possible to operate the grid entirely in a bottom-up instead of a top-down fashion. This could be explored and tested for scalability and robustness.

6.2.4. Expansion to a 24-Hour Market. The dissertation research took the view that a single market interval (typically 15 minutes) would be evaluated before it was due. The approach could be extended to cover a 24-hour planning period of 96 fifteen-minute intervals. A fairly simple modification to the code could occur to allow a vector array of 96 intervals to be carried and tracked by each object instead of a single value.

During this 24-hour period, a consumer may elect to defer certain loads to a less expensive period of time. A “shiftable load” could be defined as “a load that the consumer can move to occur earlier or later in time.” (Examples of shiftable load include electric vehicle charging, clothes washing machines, dishwashing machines, and hot water heating. An assumption can be made that these loads can be programmed to start automatically under algorithm control.) A “non-shiftable load” could be defined as “a load that the consumer wants to draw at a time of day that is, from their perspective, non-negotiable.” In wholesale markets prices rise and fall throughout the day in relation to the load (and the cost of fuel). The research can be changed to have ant agents shop for power for consumers over a 24-hour span. In each interval of time, the ant would shop for the nonshiftable load required during that interval, and if the price is favorable, consider also purchasing power for the shiftable portion of the load. (The shiftable load would be defined by time boundaries and any other constraints that it may have.) Resource locking could also occur among the grid assets using a 96-element vector. Ants would return home with a 96-element vector representing contracts with generators and network providers.

By having all of the nests compete concurrently in the free market, they obtain power fairly, on a first come, first served basis.

The Energy Market is commonly operated to optimize for cost while maintaining reliability. Other optimizations could be considered. In some parts of the USA “Customer Choice” programs are in effect which allow consumers to choose retail suppliers. By enabling individual consumers to choose individual power providers, it would allow them to choose preferred fuel types, or other optimizations the customer wishes to pursue. The use of alternative energy has consequences [124] [125] [126] [127] [128]. Consumers should be allowed to weigh these consequences as they choose their power sources.

6.2.5. Optimizations. The research presented here demonstrated the viability of the Transactive Energy, but a number of optimizations should be considered on the path to commercialization.

6.2.5.1. Performance. The code used to drive the Feeder J research could be improved to run faster. The largest gains can be made by improving the Engineering Analysis software runtime.

6.2.5.2. Consumer objectives. It would be possible to modify the algorithm used in this research so that a consumer could drive their purchase in a given direction (such as the lowest possible cost). This can be done by overbooking a given commodity and then cancelling the contracts which are the least favorable (e.g. most expensive). This process can continue until the market window closes or until it is found impractical.

6.2.5.3. Communication. It is common for communication protocols to be optimized to operate more efficiently [129]. Quite often context is leveraged to gain operational efficiencies. Ants communicate in a common environment by laying pheromones. Ants leverage this context to efficiently arrive at a solution. It may be possible to develop a new protocol which formally communicates pheromone information between smart grid applications.

6.2.6. Integration to Real-time Control. The introduction of significant amounts of distributed generation (DG) will no doubt affect many aspects of the grid. Figure 4.2

shows an outer planning loop performed by the energy market, and an inner control loop performed by network operations. The market domain (depicted in Figure 4.1) interacts with generators and energy distributors differently when a transactive energy market is utilized. Consumers also interact with wholesale providers of energy. Transactive energy involves a new level of monitoring and control that needs to accompany the deployment of a new market paradigm. Real-time monitoring of smart meters needs to occur at the distribution level, much like the real-time SCADA monitoring that occurs presently at the transmission level. Furthermore, for a transactive market to coexist with a wholesale market (as depicted in Figure 6.1,) a transactive power pool entity would have to be managed and funded much like a wholesale power pool entity. Such a system would anticipate the ability to monitor and report net energy flow every few seconds, remotely connect and disconnect customers, and issue commands to distributed generation. The grid tied DG will need to accept up/down commands and power factor adjustments. A new caliber of AMI network would need to be developed to support these new requirements. The current wholesale energy market updates every 15 minutes. A transactive energy market would need update at a similar rate so that network operations could act to control the circuit in a similar way.

APPENDIX A

GLOSSARY

Ancillary Services: wholesale market services (other than the sale of bulk power) which are necessary to sustain the grid. These services include spinning reserve (which can be called upon to supplement the power already being supplied to the grid), and frequency regulation.

Ant: a software component which serves as an agent on behalf of its client.

Ant Colony Optimizer: a nature-inspired optimization technique that solves problems by using (virtual) pheromone trails deposited in (a modeled) environment.

Customer Choice: the ability of a utility customer to freely choose the sources of power and payment plan.

Demand Response (DR): A generic term used to describe an approach to incentivize consumers to reduce consumption when energy is scarce. Some definitions of DR include Ancillary Services.

Distributed Energy Resource (DER): small generation or energy storage which is not located near large generators and often customer owned.

Graph: a mathematical structure which contains Nodes and Edges. A pair of Nodes exists for each Edge. Graphs may have directional or unidirectional edges.

Leaf: a node in a tree which has one or more upward edges but no downward edges.

Locational Marginal Pricing: the localized price of wholesale electric energy when influenced by patterns of load, generation, and the physical limits of the transmission system.

Microgrid: an electrical distribution system containing loads and distributed energy resources (such as distributed generators, storage devices or controllable loads) that can be operated in a controlled, coordinated way either while connected to the main power network or while islanded. [130]

Near Real time: a term used in various contexts to mean the performance of a task during a period of time just before the actual moment (that it is required for a planning system) or just after the actual moment (that something has occurred for a reporting system) without intentional delay.

Nest: a node which an ant identifies as its home.

Nodal Market: a market operated by an ISO in which energy is traded at a nodal level (where a node represents a bus or aggregation point).

Path: a chain of nodes connected by edges.

Pheromone: in the model, a signed value between -1 and +1 associated with graph edges, which represents the colorless, odorless chemical placed in the environment by ants to communicate the success of a find.

Proof Of Concept (POC): a small test performed to confirm the viability of an idea.

Real time: to perform a task "live" within a period of time with sufficient speed to serve as a control

Root: a node in a tree (or directed graph) which has downward running edges and no upward edges.

Spanning Tree: a subtree which does not include all of the nodes in the original tree.

Smart Grid Architecture Model (SGAM): a European effort to define the Smart Grid.

Subgraph: a portion of a graph which does not include all of the nodes in the original graph.

TabuACO: an ACO (described herein) that uses both attractive and repulsive pheromones to converge on a solution. The attractive pheromone is applied to draw the ant into productive areas of the model. The repulsive pheromones are applied to discourage travel to areas that have been previously explored and found to be unproductive.

Transactive energy: a system of economic and control mechanisms that allows the dynamic balance of supply and demand across the entire electrical infrastructure using value as a key operational parameter. [63]

Traversal: the problem of finding a path from one node to another node in a graph.

Tree: a specialized form of a graph that contains only a single edge to link two nodes.

Utility Death Spiral: a theoretical process in which the adoption of DER could allow consumers to go "off grid" and raise the cost of energy for those who remain on the grid. This in turn would further incentivize the remaining customers to invest in DER and leave the grid.

Zonal Market: a market operated by an ISO in which energy is traded at a zonal level (where zones may span an entire state).

APPENDIX B

SOFTWARE LISTING

INTRODUCTION

A copy of the code used to implement the research as well as the data can be found in the accompanying CD.

CONTENTS

APPENDIX C

MODIFICATIONS TO IEEE 34-NODE FEEDER

The IEEE 34-Node model needed additional detail so that generators could be present as well as loads. A drawing of the detailed model is available in Figure 4.26. Table C.1 below provides additional underlying detail regarding the model used by the research. The detail added to the model adds poles, service transformers, service locations, and a few small sources of generation to the model.

Table C.1. Modifications to IEEE 34 Node Types, Locations, and Loads

Derived ID	X Loc	Y Loc	Extended node type	IEEE Defined Load					
				Ph 1		Ph 2		Ph 3	
				kW	kVAr	kW	kVAr	kW	kVAr
800	0	0	Substation						
802	2580	0	Pole						
8040	2585	-5	transformer						
8041	2685	-15	Service			30	15		
8042	2785	-15	Service					25	14
806	4310	0	Pole						
808	36540	0	Pole						
8090	36545	-5	transformer						
8091	36645	-15	Service			16	8		
810	36540	-5804	Pole						
812	74040	0	Pole						
814	103770	0	VoltageRegulator						
850	103780	0	Pole						
816	104090	0	Pole						
818	104090	1710	Pole						
8190	104090	1705	transformer						
8191	104190	1690	Service	34	17				
820	104090	49860	Pole						
8210	104095	49855	transformer						
8211	104195	49840	Service	135	70				
822	104090	63600	Pole						
8230	104095	63595	transformer						
8231	104195	63580	Service			5	2		
824	114300	0	Pole						
8250	114305	-5	transformer						
8251	114405	-20	Service			40	20		
826	117330	0	Pole						
8270	114295	-425	transformer						
8271	114290	-420	service					4	2
828	114300	-840	Pole						
8290	114305	-845	transformer						
8291	114405	-860	Service	7	3				
830	134740	-840	Pole						
8300	134745	-845	transformer						
8301	134845	-860	Service	10	5	10	5	25	10
854	135260	-840	Pole						
8550	135265	-845	transformer						
8551	135365	-860	Service			4	2		
856	158590	-840	Pole						

Table C.1. Modifications to IEEE 34 Node Types, Locations, and Loads (con't)

852	135260	35990	VoltageRegulator						
832	135260	36000	Pole						
888	135260	36000	transformer						
890	145820	36000	Pole						
8900	145825	35995	transformer						
8901	145925	35980	Service	150	75	150	75	150	75
8570	135255	41825	transformer						
8571	135250	41820	service	4	2	15	8	13	7
858	135260	41830	Pole						
8630	135265	41825	transformer						
8631	135365	41810	Service	2	1				
864	135260	43450	Pole						
8330	135265	43445	transformer						
8331	135365	43430	Service	4	2	15	8	13	7
834	141090	41830	Pole						
842	141090	42110	Pole						
8430	141095	42105	transformer						
8431	141195	42090	Service	9	5				
844	141090	43460	Pole						
8440	141090	43460	Cap Bank						
8441	141095	43455	transformer						
8442	141195	43440	Service	135	105	135	105	135	105
8450	141095	43455	transformer						
8451	141195	43440	Service	0	0	25	12	20	11
846	141090	47100	Pole						
8470	141095	47095	transformer						
8471	141195	47080	Service			23	11		
848	141090	47630	Pole						
8480	141090	47630	Cap Bank						
8481	141095	47625	transformer						
8482	141195	47610	Service	20	16	20	16	20	16
8590	143105	41825	transformer						
8591	143107	41820	service	16	8	20	10	110	55
860	143110	41830	Pole						
8600	143115	41825	transformer						
8601	143215	41810	Service	20	16	20	16	20	16
8350	145780	41825	transformer						
8351	145785	41820	service	30	15	10	6	42	22
836	145790	41830	Pole						
8390	145795	41825	transformer						
8391	145895	41810	Service	18	9				
8392	145995	41810	Service			22	11		
840	146650	41830	Pole						
8400	146655	41825	transformer						
8401	146755	41810	Service	9	7	9	7	9	7
862	145790	41550	Pole						
8370	145795	41545	transformer						
8371	145895	41530	Service			28	14		
838	145790	36690	Pole						

Table C.2 below describes the ratings of each node, a randomly assigned generation capability, along with a randomly assigned asking price for the power.

Table C.2. Modified IEEE 34 Node Ratings and DER Offer Prices

Derived ID	Type	IEEE Load	KVA Rating	DER Generation Ph A (kVA)	Ph B (kVA)	Ph C (kVA)	Gen Subtotal by service (kVA)	price in cents / kWh
800	Substation		2500.0	833.3	833.3	833.3	2500.0	12.4
802	Pole		2500.0					0.0
8040	transformer		112.5				0.0	0.0
8041	service	33.5	48.0	0.0	3.8	0.0	3.8	86.8
8042	service	28.7	48.0	0.0	0.0	0.0	0.0	0.0
806	Pole		2500.0					0.0
808	Pole		2500.0					0.0
8090	transformer		25.0				0.0	0.0
8091	service	17.9	24.0	0.0	1.9	0.0	1.9	74.4
810	Pole		2500.0					0.0
812	Pole		2500.0					0.0
814	Voltage Regulator		2500.0					
850	Pole		2500.0					0.0
816	Pole		2500.0					0.0
818	Pole		2500.0					0.0
8190	transformer		50.0				0.0	0.0
8191	service	38.0	48.0	0.0	0.0	0.0	0.0	0.0
820	Pole		2500.0					0.0
8210	transformer		225.0				0.0	0.0
8211	service	152.1	192.0	38.4	0.0	0.0	38.4	111.6
822	Pole		2500.0					0.0
8230	transformer		10.0				0.0	0.0
8231	service	5.4	9.0	0.0	0.4	0.0	0.4	111.6
824	Pole		2500.0					0.0
8250	transformer		50.0				0.0	0.0
8251	service	44.7	48.0	0.0	1.0	0.0	1.0	111.6
826	Pole		2500.0					0.0
8270	transformer		10.0				0.0	0.0
8271	service	4.5	9.0	0.0	0.0	0.0	0.0	0.0
828	Pole		2500.0					0.0
8290	transformer		10.0				0.0	0.0
8291	service	7.6	9.0	0.0	0.0	0.0	0.0	0.0
830	Pole		2500.0					0.0
8300	transformer		112.5				0.0	0.0
8301	service	49.3	96.0	0.0	0.0	0.0	0.0	0.0
854	Pole		2500.0					0.0
8550	transformer		10.0				0.0	0.0
8551	service	4.5	9.0	0.0	0.5	0.0	0.5	74.4
856	Pole		2500.0					0.0
852	Voltage Regulator		2500.0					
832	Pole		2500.0					0.0
888	transformer		500.0				0.0	0.0
890	Pole		2500.0					0.0
8900	transformer		750.0				0.0	0.0
8901	service	503.1	720.0	24.0	24.0	24.0	72.0	24.8
8570	transformer		45.0				0.0	0.0
8571	service	36.2	45.0	1.8	1.8	1.8	5.4	62.0
858	Pole		2500.0					0.0
8630	transformer		10.0				0.0	0.0
8631	service	2.2	9.0	1.1	0.0	0.0	1.1	74.4
864	Pole		2500.0					0.0
8330	transformer		45.0				0.0	0.0
8331	service	36.2	43.0	0.0	0.0	0.0	0.0	0.0
834	Pole		2500.0					0.0
842	Pole		2500.0					0.0

Table C.2. Modified IEEE 34 Node Ratings and DER Offer Prices (con't)

8430	transformer		25.0				0.0	0.0
8431	service	10.3	24.0	3.4	0.0	0.0	3.4	74.4
844	Pole		2500.0					0.0
8440	Cap Bank		100.0					
8441	transformer		750.0				0.0	0.0
8442	service	513.1	720.0	43.2	43.2	43.2	129.6	111.6
8450	transformer		112.5				0.0	0.0
8451	service	50.6	96.0	0.0	0.0	0.0	0.0	0.0
846	Pole		2500.0					0.0
8470	transformer		37.5				0.0	0.0
8471	service	25.5	36.0	0.0	5.8	0.0	5.8	74.4
848	Pole		2500.0					0.0
8480	Cap Bank		150.0					
8481	transformer		112.5				0.0	0.0
8482	service	76.8	96.0	0.0	0.0	0.0	0.0	0.0
8590	transformer		225.0				0.0	0.0
8591	service	163.2	216.0	14.4	14.4	14.4	43.2	37.2
860	Pole		2500.0					0.0
8600	transformer		112.5				0.0	0.0
8601	service	76.8	96.0	0.6	0.6	0.6	1.9	99.2
8350	transformer		112.5				0.0	0.0
8351	service	92.6	96.0	4.5	4.5	4.5	13.4	37.2
836	Pole		2500.0					0.0
8390	transformer		112.5				0.0	0.0
8391	service	20.1	48.0	0.0	0.0	0.0	0.0	0.0
8392	service	24.6	48.0	0.0	0.5	0.0	0.5	24.8
840	Pole		2500.0					0.0
8400	transformer		45.0				0.0	0.0
8401	service	34.2	43.0	0.3	0.3	0.3	0.9	24.8
862	Pole		2500.0					0.0
8370	transformer		50.0				0.0	0.0
8371	service	31.3	48.0	0.0	2.9	0.0	2.9	24.8
838	Pole		2500.0					0.0

There are two transformers identified in the IEEE 34 Node Feeder Model. The ratings for these transformers are described in Table C.3 below.

Table C.3. IEEE 34 Node Feeder Model Transformer Ratings

Transformer	Rating (kVA)	High Side		Low Side	
		kV	Configuration	kV	Configuration
Substation	2500	69	Delta	24.9	Wye
XFM-1	500	24.9	Wye	4.16	Wye

Each node is connected by a conductor. Conductors (network edges) are defined in Table C.4. Each line has two endpoints. The network has an upstream/downstream sense to it. The upstream direction is always towards the primary source of power (in this case the substation bus.)

Table C.4. Modified IEEE 34 Node Feeder Model Line Sections

ID	Upstream Node	Downstream Node	Distance	IEEE Conductor Type	kVA Rating
1	800	802	2580.0	300	7470
2	802	806	1730.0	300	7470
3	802	8040	7.1	301	5602.5
4	806	808	32230.0	300	7470
5	808	810	5804.0	303	1369.5
6	808	812	37500.0	300	7470
7	808	8090	7.1	303	1369.5
8	812	814	29730.0	300	7470
9	814	850	10.0	301	5602.5
10	816	818	1710.0	302	1369.5
11	816	824	10210.0	301	5602.5
12	816	8230	63595.0	302	1369.5
13	818	820	48150.0	302	1369.5
14	818	8190	5.0	302	1369.5
15	820	822	13740.0	302	1369.5
16	820	8210	7.1	302	1369.5
17	824	826	3030.0	303	1369.5
18	824	828	840.0	301	5602.5
19	824	8250	7.1	303	1369.5
20	824	8270	425.0	301	5602.5
21	828	830	20440.0	301	5602.5
22	828	8290	7.1	301	5602.5
23	830	854	520.0	301	5602.5
24	830	8300	7.1	301	5602.5
25	832	858	5830.0	301	5602.5
26	832	888	0.0	301	5602.5
27	832	8570	5825.0	301	5602.5
28	834	842	280.0	301	5602.5
29	834	860	2020.0	301	5602.5
30	834	8590	5840.0	301	5602.5
31	836	840	860.0	301	5602.5
32	836	862	280.0	301	5602.5
33	836	8390	7.1	301	5602.5
34	840	8400	7.1	301	5602.5
35	842	844	1350.0	301	5602.5
36	844	846	3640.0	301	5602.5
37	844	8430	1355.0	301	5602.5
38	844	8440	0.0	301	5602.5
39	844	8441	7.1	301	5602.5
40	844	8450	7.1	301	5602.5
41	846	848	530.0	301	5602.5
42	846	8470	7.1	301	5602.5
43	848	8480	0.0	301	5602.5
44	848	8481	7.1	301	5602.5
45	850	816	310.0	301	5602.5
46	852	832	10.0	301	5602.5

Table C.4. Modified IEEE 34 Node Feeder Model Line Sections (con't)

47	854	852	36830.0	301	5602.5
48	854	856	23330.0	303	1369.5
49	854	8550	7.1	301	5602.5
50	858	834	5830.0	301	5602.5
51	858	864	1620.0	302	1369.5
52	858	8330	1615.0	301	5602.5
53	860	836	2680.0	301	5602.5
54	860	8350	2670.0	301	5602.5
55	860	8600	7860.0	301	5602.5
56	862	838	4860.0	304	1867.5
57	862	8370	7.1	301	5602.5
58	864	8630	1630.0	302	1369.5
59	888	890	10560.0	300	1248
60	890	8900	12066.4	301	936
61	8040	8041	100.5	301	5602.5
62	8040	8042	200.2	301	5602.5
63	8090	8091	100.5	303	1369.5
64	8190	8191	101.1	302	1369.5
65	8210	8211	101.1	302	1369.5
66	8230	8231	101.1	301	5602.5
67	8250	8251	101.1	303	1369.5
68	8270	8271	7.1	301	5602.5
69	8290	8291	101.1	302	1369.5
70	8300	8301	101.1	301	5602.5
71	8330	8331	101.1	301	5602.5
72	8350	8351	7.1	303	1369.5
73	8370	8371	101.1	303	1369.5
74	8390	8391	101.1	301	5602.5
75	8390	8392	200.6	301	5602.5
76	8400	8401	101.1	301	5602.5
77	8430	8431	101.1	301	5602.5
78	8441	8442	101.1	301	5602.5
79	8450	8451	101.1	301	5602.5
80	8470	8471	101.1	301	5602.5
81	8481	8482	101.1	301	5602.5
82	8550	8551	101.1	303	1369.5
83	8570	8571	7.1	301	5602.5
84	8590	8591	0.0	301	5602.5
85	8600	8601	0.0	301	5602.5
86	8630	8631	0.0	302	1369.5
87	8900	8901	0.0	301	936

The kVA rating of a given line segment is determined by multiplying the distribution voltage of the given line segment (24.9 kV or 4.16 kV) times the ampacity of an aluminum conductor of the size specified in the IEEE configuration table. Each conductor has a "type" defined by Table C.5.

Table C.5. IEEE 34 Node Model Conductor Types and Gauges

Type	Conductor Gauge	Ampacity
300	0/1	100
301	#2	75
302	#4	55
303	#4	55
304	#2	75

BIBLIOGRAPHY

[1] M. Dorigo and T. Stutzle, *Ant Colony Optimization*. MIT Press, 2004.

[2] E. J. H. Robinson, D. E. Jackson, M. Holcombe, and F. L. W. Ratnieks, "Insect communication: 'No entry' signal in ant foraging," *Nature*, vol. 438, no. 7067, pp. 442–442, Nov. 2005.

[3] G. I. Pirk, J. Lopez de Casenave, and L. Marone, "Evaluation of Three Techniques for the Study of Harvester Ant (Pogonomyrmex spp.) Diet," *Environ. Entomol.*, vol. 36, no. 5, pp. 1092–1099, Oct. 2007.

[4] "Study links success of invasive Argentine ants to diet shifts," *Health & Medicine Week*, University of California, San Diego, p. 184, 14-Jan-2008.

[5] J. J. Howard, "Leafcutting Ant Diet Selection: The Role of Nutrients, Water, and Secondary Chemistry," *Ecology*, vol. 68, no. 3, pp. 503–515, Jun. 1987.

[6] J. S. Pedersen, M. J. B. Krieger, V. Vogel, T. Giraud, and L. Keller, "Native Supercolonies of Unrelated Individuals in the Invasive Argentine Ant," *Evolution*, vol. 60, no. 4, pp. 782–791, 2006.

[7] M. Dorigo, "Optimization, Learning, and Natural Algorithms," PhD, Politecnico di Milano, Italy, 1992.

[8] A. Ezzat, A. Ashraf, and D. Wunsch, "A Bare-Bones Ant Colony Optimization Algorithm That Performs Competitively on the Sequential Ordering Problem," *TBD*, p. 20, 2014.

[9] M. Pedemonte, S. Nesmachnow, and H. Cancela, "A survey on parallel ant colony optimization," *Appl. Soft Comput.*, vol. 11, no. 8, pp. 5181–5197, Dec. 2011.

[10] R. Schoonderwoerd, O. E. Holland, J. L. Bruten, and L. J. M. Rothkrantz, "Ant-Based Load balancing in Telecommunications Networks," *Adapt. Behav.*, vol. 2, pp. 169–207, 1996.

[11] O. Roux, C. Fonlupt, D. Robilliard, and E. Talbi, “ANTabu - Enhanced version,” LIL-98–04, 1999.

[12] J. Montgomery and M. Randall, “Anti-pheromone as a Tool for Better Exploration of Search Space,” in *Lecture Notes in Computer Science*, vol. 2463, Springer-Verlag Berlin, 2002, pp. 100–110.

[13] D. Rodrigues, J. Louçã, and V. Ramos, “From Standard to Second Order Swarm Intelligence Phase-Space Maps,” in *Eighth European Conference on Complex Systems*, Vienna Austria, 2011, p. Poster.

[14] V. Ramos, D. M. S. Rodrigues, and J. Louçã, “Second Order Swarm Intelligence,” *Eighth Int. Conf. Hybrid Artif. Intell. Syst.*, vol. 8073, pp. 411–420, Jun. 2013.

[15] A. Ezzat and A. Abdelbar, “A Less Exploitative Variation of the Enchance Ant Colony System Applied to SOP,” *TBD*, p. 8, 2014.

[16] D. D. Haynes and S. M. Corns, “Algorithm for a Tabu Ant Colony Optimizer,” presented at the Evolutionary Computation (CEC), 2015 IEEE Congress on, Sendai, Japan, 2015, pp. 529–535.

[17] L. M. Gambardella, É. Taillard, and M. Dorigo, “Ant Colonies for the Quadratic Assignment Problem,” *J. Oper. Res. Soc.*, vol. 50, no. 2, pp. 167–176, Feb. 1999.

[18] É. Taillard, “Robust taboo search for the quadratic assignment problem,” *Parallel Comput.*, vol. 17, pp. 443–455, 1991.

[19] E.-G. Talbi, O. Roux, C. Fonlupt, and D. Robillard, “Parallel Ant Colonies for the quadratic assignment problem,” *Future Gener. Comput. Syst.*, vol. 17, no. 4, pp. 441–449, Jan. 2001.

[20] W. Wiesemann and T. Stützle, “Iterated Ants: An Experimental Study for the Quadratic Assignment Problem,” in *Ant Colony Optimization and Swarm Intelligence*, M. Dorigo, L. M. Gambardella, M. Birattari, A. Martinoli, R. Poli, and T. Stützle, Eds. Springer Berlin Heidelberg, 2006, pp. 179–190.

[21] C. Qi, “A Modified Ant Algorithm for Solving the Quadratic Assignment Problem,” in *The 2007 International Conference on Intelligent Pervasive Computing, 2007. IPC*, 2007, pp. 452–455.

[22] A. Sameh, A. Ayman, and N. Hasan, “Parallel Ant Colony Optimization,” *Int. J. Res. Rev. Comput. Sci.*, vol. 1, no. 2, pp. 77–82, Jun. 2010.

[23] T. Stutzle, “Parallelization Strategies for Ant Colony Optimization.” Springer, 03-Jun-2006.

[24] M. Lopez-Ibanez, “Multi-Objective Ant Colony Optimization,” PhD, Technische Universitat Darmstadt, Germany, 2004.

[25] D. Pinto, B. Baran, and R. Fabregat, “Multi-Objective Multicast Routing based on Ant Colony Optimization.” 2005.

[26] I. Alaya, C. Solnon, and K. Ghedira, “Ant Colony Optimization for Multi-Objective Optimization Problems,” in *19th IEEE International Conference on Tools with Artificial Intelligence, 2007. ICTAI 2007*, 2007, vol. 1, pp. 450–457.

[27] D. Haynes and S. Corns, “Applying the Tabu Ant Colony Optimization Algorithm to Operate a Transactive Energy Market,” Submitted to the *Engineering Management Journal*.

[28] S. A. Mulder and D. C. Wunsch, “Million city traveling salesman problem solution by divide and conquer clustering with adaptive resonance neural networks,” *Neural Netw.*, vol. 16, no. 5–6, pp. 827–832, Jun. 2003.

[29] D. C. Wunsch and S. Mulder, “Evolutionary algorithms, Markov decision processes, adaptive critic designs, and clustering: commonalities, hybridization, and performance,” in *International Conference on Intelligent Sensing and Information Processing, 2004. Proceedings of*, Chennai, India, 2004, pp. 477–482.

[30] D. Haynes and S. Corns, “Algorithm for a Tabu Ant Colony Optimizer with Deferred Pheromone Calculation,” *IEEE Trans. Evol. Comput.*, 2016.

[31] T. C. Koopmans and M. Beckmann, “Assignment Problems and the Location of Economic Activities,” *Econometrica*, vol. 25, no. 1, pp. 53–76, Jan. 1957.

[32] F. S. Hillier and M. M. Connors, “Quadratic Assignment Problem Algorithms and the Location of Indivisible Facilities,” *Manag. Sci.*, vol. 13, no. 1, pp. 42–57, Sep. 1966.

[33] M. Fischetti, M. Monaci, and D. Salvagnin, “Three Ideas for the Quadratic Assignment Problem,” *Oper. Res.*, vol. 60, no. 4, pp. 954–964, Aug. 2012.

[34] R. E. Burkard, E. Cela, S. E. Karisch, and F. Rendl, “QAPLIB - A Quadratic Assignment Problem Library,” *Journal of Global Optimization 10:391-403, 1997*, Aug-2011. [Online]. Available: http://anjos.mgi.polymtl.ca/qaplib/. [Accessed: 21-Feb-2015].

[35] N. Christofides and E. Benavent, “An Exact Algorithm for the Quadratic Assignment Problem on a Tree,” *Oper. Res.*, vol. 37, no. 5, pp. 760–768, 1989.

[36] NIST, “Critical Values of the Student’s t Distribution,” National Institute of Standards and Technology, Gaithersburg, MD, Jan. 2015.

[37] C. Fleurent and J. Ferland, “Genetic hybrids for the quadratic assignment problems,” *DIMACS*, vol. 16, pp. 190–206, 1994.

[38] R. Battiti and G. Tecchiolli, “The Reactive Tabu Search,” *Oper. Res. Soc. Am.*, vol. 6, no. 2, pp. 128–140, Spring 1994.

[39] D. Connolly, “An improved annealing scheme for the QAP,” *Eur. J. Oper. Res.*, vol. 46, pp. 93–100, 1990.

[40] M. Gendreau, J.-F. Larochelle, and B. Sansò, “A tabu search heuristic for the Steiner Tree Problem,” *Networks*, vol. 34, no. 2, pp. 162–172, Sep. 1999.

[41] I. Ljubic, “Exact and Memetic Algorithms for Two Network Design Problems,” PhD, Technische Universitat Wien, Vienna Austria, 2004.

[42] M. G. C. Resende, “Test problem distribution,” *Test problem distribution*, 10-Dec-2011. [Online]. Available: http://www2.research.att.com/~mgcr/data/index.html. [Accessed: 06-Sep-2014].

[43] 110th US Congress, *Energy Independence and Security Act of 2007*. 2007, p. 310.

[44] A. Smith, *The Wealth of Nations: Of Systems of political economy*, vol. 4, 5 vols. London: W. Strahan and T. Cadell, 1776.

[45] NIST, “NIST Framework and Roadmap for Smart Grid Interoperability Standards,” US Department of Commerce, NIST Special Publication 1108, Jan. 2010.

[46] F. McDonald, “Samuel Insull and the movement for state utility regulatory commissions,” *Bus. Hist. Rev. Pre-1986*, vol. 32, no. 000003, p. 241, Autumn 1958.

[47] R. Coase, *The Firm, The Market, and the Law*. Chicago: University of Chicago Press, 1988.

[48] O. Williamson, *Transaction Cost Economics*. Northhampton, MA: Edward Elgar Publishing, 1975.

[49] J. Kwoka, *Power Structure: Ownership, Integration, and Competition in the U.S. Electricity Industry*. Boston: Kluwer Academic, 1996.

[50] R. Michaels, “Vertical Integration and the Restructuring of the U.S. Electricity Industry,” *Policy Anal.*, vol. 572, Jul. 2006.

[51] C. Froome, “Renewable Energy in Australia: 20% by 2020. Can this be Achieved?,” *Manag. Environ. Qual.*, vol. 21, no. 2, p. 177, 2010.

[52] K. Ek, “Public and private attitudes towards ‘green’ electricity: the case of Swedish wind power,” *Energy Policy*, vol. 33, no. 13, pp. 1677–1689, Sep. 2005.

[53] T&D World, “Scotland Sets Goal of 100% Electricity from Renewable Energy by 2020,” *Transmission & Distribution World*, 08-Jul-2011.

[54] P. Fox-Penner, *Smart Power: Climate Change, the Smart Grid, and the Future of Electric Utilities*, 1st ed. Washington: Island Press, 2010.

[55] G. Clark, "Smart Grid: Integrating Renewable, Distributed, & Efficient Energy," in *Smart Grid: Integrating Renewable, Distribted, & Efficient Energy*, S. Fereidoon, Ed. Elsevier, 2012, p. vii.

[56] U. S. Energy Information Administration, "EIA - Electricity Data 2016," *Electric Power Monthly*, 25-Oct-2016. [Online]. Available: https://www.eia.gov/electricity/monthly/epm_table_grapher.cfm?t=epmt_1_1. [Accessed: 23-Nov-2016].

[57] M. Aachiq, J. G. D. S. F. Junior, T. Oozeki, and Y. Iwafune, "Evaluation of Residential PV-EV System for Supply and Demand Balance of Power System," *Electr. Eng. Jpn.*, vol. 195, no. 2, pp. 11–20, Apr. 2016.

[58] G. Yang *et al.*, "Voltage rise mitigation for solar PV integration at LV grids: Studies from PVNET. dk," *J. Mod. Power Syst. Clean Energy*, vol. 3, no. 3, pp. 411–421, Jul. 2015.

[59] EPRI, "EPRI | Distributed PV Monitoring and Feeder Analysis," 2016. [Online]. Available: http://dpv.epri.com/feeder_j.html. [Accessed: 22-Jul-2016].

[60] P. Hanser, K. Madjarov, W. Katzenstein, and J. Chang, "Riding the Wave: Using Demand Response for Integrating Intermittent Resources," in *Smart Grid: Integrating Renewable, Distributed, & Efficient Energy*, F. Sioshansi, Ed. Elsevier, 2012, pp. 235–256.

[61] T. Penn, "RTP Discussion," 19-Mar-2014.

[62] "spp effective 2016 planning criteria 1.pdf." [Online]. Available: https://www.spp.org/documents/33003/spp%20effective%202016%20planning%20criteria%201.pdf. [Accessed: 23-Oct-2016].

[63] "GridWise Transactive Energy Framework," GridWise Architecture Council, PNNL-22946 Ver 1.0, Jan. 2015.

[64] H. Harvey and S. Aggarwal, “America’s Power Plan: Overview, Rethinking Policy to Deliver a Clean Energy Future,” 2013.

[65] GridWise Architecture Council, “GridWise Architecture Tenets and Illustrations,” PNNL, Richland, WA, USA, PNNL-SA-39480, Oct. 2003.

[66] W. Vickrey, “Responsive Pricing of Public Utility Services,” *Bell J. Econ. Manag. Sci.*, vol. 2, no. 1, pp. 337–346, Apr. 1971.

[67] A. Faruqui, “The Ethics of Dynamic Pricing,” in *Smart Grid: Integrating Renewable, Distributed, & Efficient Energy*, F. Sioshansi, Ed. Waltham, MA: Academic Press, 2012, pp. 61–84.

[68] E. Negeri and N. Baken, “Distributed Storage Management Using Dynamic Pricing in a Self-Organized Energy Community,” in *Self-Organizing Systems*, 2012, pp. 1–12.

[69] J. St. John, “A how-To Guide for Transactive Energy,” *Greentechmedia.com*. 20-Nov-2013.

[70] F. A. Rahimi and A. Ipakchi, “Transactive Energy Techniques: Closing the Gap between Wholesale and Retail Markets,” *Electr. J.*, vol. 25, no. 8, pp. 29–35, Oct. 2012.

[71] F. Sioshansi, “Complex Transactions Made Simpler, Thanks To Transactive Energy,” *Electr. J.*, vol. 27, no. 9, pp. 1–3, Nov. 2014.

[72] M. Dorigo, E. Bonabeau, and G. Theraulaz, “Ant algorithms and stigmergy,” *Future Gener. Comput. Syst.*, vol. 16, no. 8, pp. 851–871, Jun. 2000.

[73] M. Lew *et al.*, “Large-Scale Wind Integration studies in the United States: Preliminary results.” National Renewable Energy Laboratory, Sep-2009.

[74] R. Walton, “What the Supreme Court decision on FERC Order 745 means for demand response and DERs | Utility Dive,” *What the Supreme Court decision on FERC Order 745 means for demand response and DERs: Legal cerntanty will allow grid operators to focus on integrating DR into their market structures*, 03-Feb-2016. [Online]. Available: http://www.utilitydive.com/news/what-the-supreme-court-decision-on-ferc-order-745-means-for-demand-response/413092/. [Accessed: 18-Sep-2016].

[75] Federal Energy Regulatory Commission, “2010 Assessment of Demand Response & Advanced Metering,” Department of Energy, Washington DC, Feb. 2011.

[76] Kotewa, Larry, “RTP in Chicago,” 29-Mar-2014.

[77] M. Andreolas, “Mega Load Management System Pays Dividends,” *Transmission & Distribution World*, 01-Feb-2004. [Online]. Available: http://tdworld.com/distribution-management-systems/mega-load-management-system-pays-dividends. [Accessed: 11-Sep-2016].

[78] “Transmission Projects: At A Glance.” Edison Electric Institute, Mar-2015.

[79] B. McDonald, M. Bardee, J. Seesel, D. Meyer, and K. Larsen, “Report to Congress on Competition in Wholesale and Retail Markets for Electric Energy: Pursuant to Section 1815 of the Energy Policy Act of 2005,” DOJ, FERC, DOE, FTC, DOAg, Washington DC, Report to Congress, Apr. 2007.

[80] D. Cooke, “Empowering Customer Choice in Electricity Markets.” International Energy Agency, Oct-2011.

[81] A. Jokić, M. Lazar, and P. P. J. van den Bosch, “Real-time control of power systems using nodal prices,” *Int. J. Electr. Power Energy Syst.*, vol. 31, no. 9, pp. 522–530, Oct. 2009.

[82] M. Cadwalader, S. Harvey, W. Hogan, and S. Pope, “Coordinating Congestion Relief Across Multiple Regions,” Harvard University, Cambridge MA, Oct 7, 1999.

[83] S. S. Oren and A. M. Ross, "Economic congestion relief across multiple regions requires tradable physical flow-gate rights," *IEEE Trans. Power Syst.*, vol. 17, no. 1, pp. 159–165, 2002.

[84] C. Greer, "Grid 3.0 Roundtable 2014, Part 1," presented at the NIST Electricity Sector Issues Roundtable, Gaithersburg, MD, 13-Nov-2014.

[85] M. Prokopenko, "Design Versus Self-Organization," in *Advances in Applied Self-Organizing Systems*, M. Prokopenko, Ed. Springer London, 2013, pp. 3–21.

[86] A. Lovins, *Small is Profitble: The Hidden Economic Benefits of Making Electrical Resources the Right Size*. Rocky Mountain Institute, 2002.

[87] SGiP, "NIST Smart Grid Collaboration Wiki, Smart Grid Interoperability Panel (SGIP) Site," *NIST Smart Grid Collaboration Wiki*. 19-2012.

[88] CEN-CENELEC-ETSI, "SG-CF/M490/E - Part E: Smart Grid Use Case Management Process," CEN-CENELEC-ETSI Smart Grid Coordination Group, Brussels, SG-CG/M490/E, Nov. 2012.

[89] Bundesnetzagentur, "'Smart Grid' und 'Smart Market', Eckpunktepapier der Beundesnetzagentur zu den Aspekten des sich verandernden Energieversorgungssystems," Bonn, Dec. 2011.

[90] W. W. Hogan, "Multiple Market-Clearing Prices, Electricity Market Design and Price Manipulation," *Electr. J.*, vol. 25, no. 4, pp. 18–32, May 2012.

[91] F. Wolak, "Measuring the Benefits of Greater Spatial Granularity in Short-Term Pricing in Wholesale Electricity Markets," *Am. Econ. Rev. Pap. Proc.*, vol. 101, no. 3, pp. 247–252, 2011.

[92] R. R. Mabey and P. S. Malone, "Chapter 11 Reorganization of Utility Companies," *Energy Law J.*, vol. 22, p. 277, 2001.

[93] D. Malkin and P. Centolella, "Results-Based Regulation: A Modern Approach to Modernize the Grid," General Electric, Oct. 2013.

[94] A. Brown and R. Salter, “Can Smart Grid Technology Fix the Disconnect Between Wholesale and Retail Pricing?,” *Electr. J.*, vol. 24, no. 1, pp. 7–13, Jan. 2011.

[95] H. Chao, “Price-Responsive Demand Management for a Smart Grid World,” *Electr. J.*, vol. 23, no. 1, pp. 7–20, Jan. 2010.

[96] M. J. King, K. King, and M. B. Rosenzweig, “Customer Sovereignty: Why Customer Choice Trumps Administrative Capacity Mechanisms,” *Electr. J.*, vol. 20, no. 1, pp. 38–52, Jan. 2007.

[97] “Learn More About Interconnections.” [Online]. Available: http://energy.gov/oe/information-center/recovery-act/recovery-act-interconnection-transmission-planning/learn-more. [Accessed: 16-Feb-2014].

[98] FERC, “Electric Power Markets: National Overview,” *Federal Energy Regulatory Commission*, 20-Mar-2010. [Online]. Available: http://www.ferc.gov/market-oversight/mkt-electric/overview.asp. [Accessed: 16-May-2012].

[99] FERC, “Electric Market Overview: On-Peak Spot Electric Prices,” *Federal Energy regulatory Commission*, 06-Jan-2013. [Online]. Available: www.ferc.gov/market-oversight/mkt-electric/overview/elec-ovr-on-pk-elec-pr.pdf. [Accessed: 17-Nov-2013].

[100] “Midwest ISO LMP Contour Map,” 22-Dec-2017. [Online]. Available: https://www.misoenergy.org/LMPContourMap/MISO_All.html. [Accessed: 22-Dec-2017].

[101] “US power markets switch to nodal trading,” *www.risk.net*. [Online]. Available: http://www.risk.net/energy-risk/feature/1720263/us-power-markets-switch-nodal-trading. [Accessed: 25-Jul-2013].

[102] “What a Nodal Market Means for Texas and ERCOT,” *Energy, Technology, & Policy*. [Online]. Available: http://webberenergyblog.wordpress.com/2010/03/14/what-a-nodal-market-means-for-texas-and-ercot/. [Accessed: 25-Jul-2013].

[103] OPUC, "The Texas Office of Public Utility Counsel," *Electric Rliability Council of Texas*. [Online]. Available: http://www.opuc.texas.gov/ERCOT.html#nodal. [Accessed: 16-Feb-2014].

[104] C. Brandstatt, G. Brunekreeft, and N. Friedrichsen, "Smart Pricing to Reduce Network investment in Smart Distribution Grids -- Experience in Germany," in *Smart Grid: Integrating Renewable, Distributed, & Efficient Energy*, F. Sioshansi, Ed. Waltham, MA: Academic Press, 2012, pp. 317–342.

[105] Global Energy Partners, "AmerenUE Demand Side Management (DSM) Market Potential Study Volume 4: Program Analysis," Global Energy Partners, Walnut Creek, CA, 1287–4, Jan. 2010.

[106] E. Hirst, "Price-Responsive Demand in Wholesale Markets: Why Is So Little Happening?," *Electr. J.*, vol. 14, no. 4, pp. 25–37, May 2001.

[107] T. Penn, "Interview with Terry Penn, Staff Analyst, Georgia Power," 20-Jun-2011.

[108] E. Cazalet and W. Cox, "OASIS Energy Market Information Exchange (eMIX) TC," *OASIS*, 2014. [Online]. Available: https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=emix. [Accessed: 06-Apr-2014].

[109] "Energy Market Information Exchange (EMIX)," OASIS, Standard Version 1.0, Committee Specification Draft 01, Nov. 2010.

[110] GridWise Architecture Council, "GridWise Architecture Council," Jul-2014. [Online]. Available: http://www.gridwiseac.org/. [Accessed: 30-Jul-2014].

[111] T. Randall, "Here's How Electric Cars Will Cause the Next Oil Crisis," *Bloomberg*, 25-Feb-2016. [Online]. Available: http://www.bloomberg.com/features/2016-ev-oil-crisis/. [Accessed: 20-Sep-2016].

[112] H. Trabish, “The new EV playbook: How utilities can gain from the coming boom in electric vehicles | Utility Dive,” *The new EV playbook: How utilities can gain from the coming boom in electric vehicles.*, 27-Jun-2016. [Online]. Available: http://www.utilitydive.com/news/the-new-ev-playbook-how-utilities-can-gain-from-the-coming-boom-in-electri/421559/. [Accessed: 20-Sep-2016].

[113] D. Nakafuji, “Visibility Enables Photovoltaic Integration,” *Transmission & Distribution World*, 18-Aug-2016. [Online]. Available: http://tdworld.com/generation-renewables/visibility-enables-pv-integration. [Accessed: 21-Sep-2016].

[114] M. L. Crow, “Economic Scheduling of Residential Plug-In (Hybrid) Electric Vehicle (PHEV) Charging,” *Energies*, vol. 7, no. 4, pp. 1876–1898, 2014.

[115] Federal Energy Regulatory Commission, “2015 Assessment of Demand Response & Advanced Metering,” Department of Energy, Washington DC, Dec. 2015.

[116] “FERC: Electric Power Markets - National Overview.” [Online]. Available: http://www.ferc.gov/market-oversight/mkt-electric/overview.asp. [Accessed: 24-Nov-2015].

[117] R. Lovett, “What if the Biggest Solar Storm on Record Happened Today?,” *National Geographic News*, 04-Mar-2011.

[118] N. Allen and L. Watson, “Hurricane Maria: Puerto Rico left entirely without power,” *The Telegraph*, 21-Sep-2017.

[119] S. A. Chandler, J. H. Rinaldi, R. B. Bass, and L. Beckett, “Smart grid dispatch optimization control techniques for transactive energy systems,” in *2014 IEEE Conference on Technologies for Sustainability (SusTech)*, 2014, pp. 51–54.

[120] P. E. S. Distribution System Analysis Subcommittee, “IEEE 34 Node Test Feeder.” Institute of Electrical and Electronics Engineers, 17-Sep-2010.

[121] EPRI, “EPRI | Smart Grid Resource Center > Simulation Tool – OpenDSS,” *EPRI Smart Grid Resource Center*, 2011. [Online]. Available: http://smartgrid.epri.com/SimulationTool.aspx. [Accessed: 22-Jul-2016].

[122] N. Aitzhan and D. Svetinovic, “Security and Privacy in Decentralized Energy Trading through Multi-signatures, Blockchain and Anonymous Messaging Streams,” *IEEE Trans. Dependable Secure Comput.*, vol. PP, no. 99, p. 1, Oct. 2016.

[123] J. K. Kok, C. J. Warmer, and I. G. Kamphuis, “PowerMatcher: Multiagent Control in the Electricity Infrastructure,” in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems*, New York, NY, USA, 2005, pp. 75–82.

[124] J. Sathaye *et al.*, “Renewable Energy in the Context of Sustainable Development,” in *IPCC Special Report on Renewable Energy Sources and Climate change Mitigation [O. Edenhofer, R. Pichs-Madruga, Y. Sokona, K. Seyboth, P. Matschoss, S. Kadner, T. Zwickel, P. Eickemeier, G. Hansen, S. Schlömer, C.von Stechow (eds)]*, Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2011.

[125] A. L. Drewitt and R. H. W. Langston, “Assessing the impacts of wind farms on birds,” *Ibis*, vol. 148, pp. 29–42, 2006.

[126] “Environmental Impacts of Solar Power | UCSUSA,” *Union of Concerned Scientists*. [Online]. Available: http://www.ucsusa.org/clean_energy/our-energy-choices/renewable-energy/environmental-impacts-solar-power.html. [Accessed: 04-Aug-2013].

[127] “Comparison of Lifecycle Greenhouse Gas Emissions of Various Electricity Generation Sources.” World Nuclear Association.

[128] W. L. Ellsworth, “Injection-Induced Earthquakes,” *Science*, vol. 341, no. 6142, Jul. 2013.

[129] D. D. Haynes, “A Case for Optimized Protocols in the Creation of a Smarter Grid,” *IEEE Trans. Power Deliv.*, vol. 25, no. 3, pp. 1476–1482, Jul. 2010.

[130] N. Hatziargyriou, “Network of the Future: electricity supply systems of the future,” presented at the CIGRE AG, Athens, 26-Sep-2013.

VITA

David Donald Haynes received a Bachelor of Science degree in Computer Science and a Bachelor of Science degree in Electrical Engineering from Washington University in St. Louis, Missouri in May, 1983. He received a Graduate Certificate in Systems Engineering and a Master of Science degree in Electrical Engineering from the Missouri University of Science & Technology in May, 2007. He received a Ph.D. in Systems Engineering from the Missouri University of Science & Technology in May, 2019.

David has participated in the development of numerous standards within the IEC, ANSI, IEEE, and the MultiSpeak® software initiative. He has worked as an engineer and has been named as the primary inventor on more than 6 US patents. The patents include: *A Method for generating Cryptographic "One-Time Pads" and Keys for Secure Network Communications (*9,762,560), *Power Plant Using Compressed or Liquified Air for Energy Storage* (8,453,444), *Metrology Timekeeping Systems and Methods* (9,677,908), and others.