Sequential Decision Making for Intelligent Agents
Papers from the AAAI 2015 Fall Symposium

Autonomous Electricity Trading Using Time-of-Use Tariffs in a Competitive Market

Daniel Urieli and **Peter Stone**
Dept. of Computer Science
The University of Texas at Austin
Austin, TX 78712 USA
{urieli,pstone}@cs.utexas.edu

Abstract

This research studies the impact of Time-Of-Use (TOU) tariffs in a competitive electricity market place. Specifically, it focuses on the question of how should an autonomous broker agent optimize TOU tariffs in a competitive retail market, and what is the impact of such tariffs on the economy. We formalize the problem of TOU tariff optimization and propose an algorithm for approximating its solution. We extensively experiment with our algorithm in a large-scale, detailed electricity retail markets simulation of the Power Trading Agent Competition (Power TAC) and: 1) find that our algorithm results in 15% peak-demand reduction, 2) find that its peak-flattening results in greater profits and/or profit-share for the broker and allows it to win in head-to-head competition against the 1st and 2nd place brokers from the Power TAC 2014 finals, and 3) analyze several economic implications of using TOU tariffs in competitive retail markets.

1 Introduction

To address the need for secure, clean and sustainable energy, governments around the world are re-engineering the traditional electricity grid into a *smart-grid* (U.S 2003; Eur 2011). One of the milestones in the smart-grid vision is "customer participation in power markets through demand-side-management". Demand-side management refers to adapting customer demand to supply conditions, and may be implemented using new power market structures. Due to the high cost of failure in the real world (Borenstein 2002), it is important to test new market structures in simulation before deploying them. This is the focus of the Power Trading Agent Competition (Power TAC) (Ketter, Peters, and Collins 2013), and of this research. In Power TAC, multiple autonomous broker agents compete with each other for making profits in a state-of-the-art simulation of future power markets. Such brokers serve as financial intermediaries representing large numbers of customers, and thus minimize risk-adjusted costs and earn profit while reducing energy costs for their customers (Ketter, Peters, and Collins 2013).

One of the primary goals of demand-side management (DSM) is *peak-flattening* (distributing consumption more evenly throughout the day). One of the main methods proposed for implementing DSM is *Time-Of-Use (TOU) tariffs*, which are energy selling contracts specifying different prices for different times of day. TOU tariffs incentivise customers to adapt their consumption to prices. By doing so, customers can lower their energy costs, while possibly increasing their inconvenience, due to changing their consumption pattern. This research focusses on two questions. First, how should an autonomous broker optimize a TOU tariff that is both 1) attractive to customers in a competitive retail market with fixed-rate tariffs and 2) more profitable for the broker than the best fixed-rate tariff? Second, what is the economic impact of TOU tariffs in a competitive market? Our contributions are as follows:

- We formalize the problem of selecting an optimal TOU tariff in competitive markets, show that it is intractable, and propose an efficient optimization algorithm that approximates its solution. Our algorithm is fully implemented in our broker agent.
- Our algorithm results in 15% peak-demand reduction in a complex, large-scale simulation of competitive power markets (Power TAC). To the best of our knowledge, our work is the first to show that TOU can achieve the primary goal of peak-flattening in competitive markets in such a complex, detailed, realistic simulation.
- Our agent's peak-flattening results in greater profits and/or profit-share and allows it to beat the 1st and 2nd place agents from the 2014 Power TAC finals, while reducing customers energy costs.
- Using extensive experimentation, we analyze several economic implications of using TOU in competitive retail markets.

2 Background

Our testbed domain is the Power Trading Agent Competition (Power TAC) simulation environment (Ketter, Peters, and Collins 2013). Power TAC is an annual competition in which the competitors are autonomous brokers programmed by participants from around the world. Figure 1 shows the structure of the Power TAC simulation environment. At a high level, autonomous broker agents compete with each other by acting in three markets: (1) a *wholesale market*, in which brokers procure energy from traditional generation companies through a sequential double-auction bidding pro-

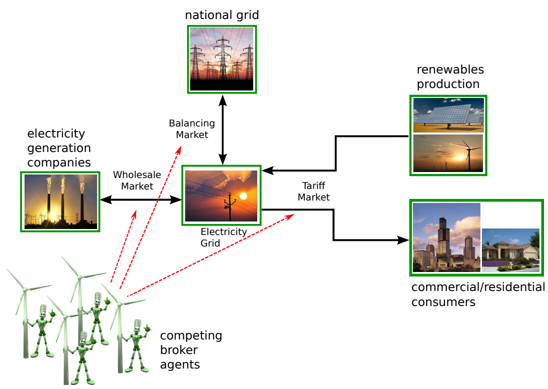

Figure 1: Structure of the Power TAC simulation

cess, (2) a *tariff market*, which is a retail market in which brokers offer *tariffs* (energy-selling contracts) to which autonomous, utility-optimizing customers may subscribe, and (3) a *balancing market*, which serves to ensure that supply and demand are balanced at all times and to financially penalize brokers for any imbalance in their portfolios.

Typically, a peaked daily customer demand is much more expensive to satisfy than demand that is flat throughout the day, since energy-generation costs are sharply increasing with the amount of generation. To reduce its own costs, a broker may incentivise its customers to flatten demand by publishing TOU tariffs in the tariff market. Here we define a TOU tariff to be a tuple $P := \langle p_0, p_1, \cdots, p_{23}\rangle$, where p_t is the *hourly rate* in cents/kWh during hour-of-day t. In general, simulated customers (represented by autonomous agents) are willing to subscribe to TOU tariffs and shift their consumption to save costs if cost-reduction (compared with competing fixed-price tariffs) is large enough to offset the increased inconvenience due to consumption-shifting.

3 TOU Tariff Optimization

Given the internal states of the simulator and competing brokers, the broker's energy trading problem is a complex MDP with continuous high-dimensional states and actions. However, since competitors' states and parts of the simulator state are unobservable, the trading problem is actually a POMDP. Nevertheless, we approximate the trading problem as an MDP. To combat the curse-of-dimensionality we use a lookahead policy to optimize trading quantities and prices in the tariff and wholesale markets. In the lookahead policy, tariff and wholesale market actions must be optimized in conjunction, to maintain low imbalance. Part of this lookahead policy is optimizing TOU tariffs.

We frame the TOU optimization problem in terms of demand and cost curves. Let the *demand curve* of tariff τ be the function $D : \mathbb{R} \to \mathbb{R}$ that maps energy-selling prices to the resulting energy demand from customers that will subscribe to τ. Let the *unit-cost curve* be the function $C : \mathbb{R} \to \mathbb{R}$ that maps an energy amount to the unit-price for which the broker is able to procure it in the wholesale market. We define the TOU tariff optimization problem as:

$$P^* := \arg\max_P U_H(P) \tag{1}$$

where P is a vector of 24 hourly rates, and where $U_H(P)$ is the expected-utility (profits) over some future horizon H, predicted with a lookahead simulation using a transition function model composed of future demand and cost-curves. Equation 1 is an optimization over actions at a given state, and should be executed continually as a part of the broker's lookahead policy, to adapt the broker's actions in the tariff market to the continually changing game-state.

Equation 1 provides a theoretical objective for solving an autonomous broker's TOU optimization problem. In practice, there are at least two main obstacles to solving this equation. First, future demand and cost-curves are typically unknown. Second, even if they were known, optimization requires search in a non-convex, discontinuous space, which is generally intractable. To address the first problem of estimating demand and cost-curves we use online learning from game data, and use known statistical information about customers' behavior patterns. To address the second problem, we resort to finding a local optimum, designing an empirical gradient-based optimization algorithm that uses the curves-based lookahead as an objective evaluation. Our algorithm outperforms well-known optimization algorithms.

We extensively evaluated our TOU algorithm in hundreds of games. We compared its performance with two baselines, one using fixed-priced tariffs, and one using a naive TOU optimization. We showed that TOU tariffs can compete successfully with fixed-rate tariffs: using TOU tariffs, our agent performs better than the top 2 agents of the Power TAC 2014 finals. It reduced peak-demand by 15% compared with using only fixed-rate tariffs, increased its profits and/or profit-share, and saved costs for all customers (including competitors'). Interestingly, while past work warned against TOU tariffs inducing *customer-herding*, our TOU optimization method did not exhibit such a phenomenon, due to a combination of (1) our broker's profit-optimization using predictive demand models leading to coordinated flattening, and (2) a smooth customer inconvenience metric. This underlines a potential benefit of employing autonomous TOU brokers in competitive power markets. In addition, we have seen that a TOU broker's customer share is an important factor in its ability to flatten demand: to be able to counter-balance peaked consumption of fixed-rate brokers' customers, it needs to gain large customer-share by creating attractive TOU tariffs that are still profitable.

References

Borenstein, S. 2002. The trouble with electricity markets: Understanding california's restructuring disaster. *Journal of Economic Perspectives* 16(1):191–211.

European Commission. 2011. *Communication: Energy Roadmap 2050.*

Ketter, W.; Peters, M.; and Collins, J. 2013. Autonomous agents in future energy markets: The 2012 Power Trading Agent Competition. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*. AAAI.

U.S. Department of Energy. 2003. *"Grid 2030" A National Vision For Electricity's Second 100 Years.*