A Reinforcement Learning Approach to Setting Multi-Objective Goals for Energy Demand Management

Ying Guo, CSIRO ICT Centre North, Australia

Astrid Zeman, CSIRO ICT Centre North, Australia

Rongxin Li, CSIRO ICT Centre North, Australia

ABSTRACT

In order to cope with the unpredictability of the energy market and provide rapid response when supply is strained by demand, an emerging technology, called energy demand management, enables appliances to manage and defer their electricity consumption when price soars. Initial experiments with our multi-agent, power load management simulator, showed a marked reduction in energy consumption when price-based constraints were imposed on the system. However, these results also revealed an unforeseen, negative effect: that reducing consumption for a bounded time interval decreases system stability. The reason is that price-driven control synchronizes the energy consumption of individual agents. Hence price, alone, is an insufficient measure to define global goals in a power load management system. In this article we explore the effectiveness of a multi-objective, system-level goal which combines both price and system stability. We apply the commonly known reinforcement learning framework, enabling the energy distribution system to be both cost saving and stable. [Article copies are available for purchase from InfoSci-on-Demand.com]*

Keywords: Demand Management; Energy; Multi-Agent Systems; Reinforcement Learning; Stability

INTRODUCTION

As technology advances, consumers become increasingly more power hungry. This causes many countries, including Australia, to suffer from an increasing gap between electricity supply and demand (An Introduction to Australia's National Electricity Market, 2005; Hou, I., 2007; US Department of Energy, 2002. The traditional way of tackling such a problem is to increase supply by investing heavily in infrastructure and building more generators. Alternatively, power load management can reduce levels of power consumption on the demand side, and hence reduce the level of energy required to run appliances, saving money and reducing the risk of inadequate supply (Sutton, R. S., &

Barto, A. G., 1998; Wilson, B., Rassenti, S. J., & Smith, V. L., 2003).

Electricity distribution is a complex system, consisting of loads, generators, and transmission and distribution networks (Guo, Y., Li, J., & James G., 2005). To control this physical system, generators and retailers bid into a market that balances supply and demand while ensuring safe network operation. Demand and price fluctuate quickly and loads that are responsive in real time can have high value to retailers and networks. This scenario is ideal for the adoption of multi-agent technology (Dimeas, L., & Hatziargyriou, N. D., 2005;. McArthur, S. D. J., & Davidson, E. M., 2005). A network of autonomous agents can be overlaid on the physical distribution network, controlling customer loads and, where available, local generators (Borenstein, S., Jaske, M., & Rosenfeld, A., 2002; Hägg, S., and Ygge, F., 1995; Li, R., Li, J., Poulton, G., & James, G., 2008; Platt, G., 2007; Ygge, F., 1998.

Australia's Commonwealth Scientific and Industrial Research Organisation (CSIRO) is developing an energy management and control system that consists of an agent network to be installed at multiple levels of the electricity distribution network (James, G., Cohen, D., Dodier, R., Platt, G., & Palmer, D., 2006; Li, J., Poulton G., & James, G., 2007; Li, J., Poulton, G., James, G., Zeman, A., Wang, P., Chadwick, M., & Piraveenan M., 2007; Zeman, A., Prokopenko, M., Guo, Y., and Li, R.., 2008. This system has a three-level architecture consisting of the following types of agents: 1) the top-level *broker agent*; 2) the middle-level *group agent* and 2) the bottom-level *appliance agent*.

Appliance agents are responsible for the low-level management of consumption for each end-use device. At the bottom level, appliances can be intelligently switched on or off based on customer preferences (Guo, Y., et al., 2005). A cluster of such agents is then managed by a group agent on the middle level (Li, R., Poulton, G., 2007; Ogston, E., Zeman, A., Prokopenko, M., & James, G., 2007). The group agents also receive an energy quota from an upper-level broker agent representing the needs of electricity market participants and network operators. This energy quota is a limit on the total energy consumption for a group of appliances.

In this article, we focus on the intelligent management of the broker agent -- that is, how to choose the optimal strategy to provide the real-time quota to the group agent. As the top level in the system, the broker agent needs to incorporate information from the market (such as the current local price of energy) as well as input from group agents (Guo, Y., 2007). The broker agent is required to manage the risk of exposure to volatile wholesale pool prices and reduce strain on the network.

Energy price can change dramatically when the demand is very high such as a hot summer afternoon. Because market energy prices and weather conditions are dynamic, we chose to use the reinforcement learning (RL) framework as an online learning approach. The two usual approaches to reinforcement learning are model-free and model-based. Model-free algorithms perform well for simple problems. Because the energy market is a dynamic system which exhibits unpredictable properties such as the price of energy, we cannot easily generate a model beforehand. Hence we use the model free RL algorithm – Q-learning. To cope with the dynamic nature of the energy market, we enable the reward matrix to update towards the optimum (which is different from conventional RL) while the agent learns the environment. We present this approach and its application for setting the system-level goal. The experimental results show that by using this method, the broker agent is not only adaptive to multiple systems, but is also able to find an optimal balance between both system stability and energy cost.

The remainder of this article is organised as follows. Section 2 outlines our approach and Section 3 describes the RL learning process. In Section 4, we present experimental results in a simulated environment. Finally, conclusions based on these experiments are given in Section 5.

Related Content

A Framework for Dynamic Agent Organizations
Shaheen Fatimaand Michael Wooldridge (2009). *Handbook of Research on Multi-Agent Systems: Semantics and Dynamics of Organizational Models (pp. 446-459).*
www.irma-international.org/chapter/framework-dynamic-agent-organizations/21110

Introducing AI and IA into a Non Computer Science Graduate Programme
Ioanna Stamatopoulou, Maria Fasliand Petros Kefalas (2011). *Multi-Agent Systems for Education and Interactive Entertainment: Design, Use and Experience (pp. 89-100).*
www.irma-international.org/chapter/introducing-into-non-computer-science/50396

Simulating Tax Evasion with Utilitarian Agents and Social Feedback
Attila Szabó, László Gulyásand István János Tóth (2010). *International Journal of Agent Technologies and Systems (pp. 16-30).*
www.irma-international.org/article/simulating-tax-evasion-utilitarian-agents/39030

Coalition Formation Among Agents in Complex Problems Based on a Combinational Auction Perspective
Hiromitsu Hattori (2007). *Architectural Design of Multi-Agent Systems: Technologies and Techniques (pp. 48-60).*
www.irma-international.org/chapter/coalition-formation-among-agents-complex/5172

Inductive Logic Programming and Embodied Agents: Possibilities and Limitations
Andrea Kulakov, Joona Laukkanen, Blerim Mustafaand Georgi Stojanov (2011). *Developments in Intelligent Agent Technologies and Multi-Agent Systems: Concepts and Applications (pp. 116-131).*
www.irma-international.org/chapter/inductive-logic-programming-embodied-agents/49359