Towards Adaptive Energy Plan Aggregation over a Peer-to-Peer Tree Overlay

Evangelos Pournaras   Martijn Warnier   Frances Brazier
Intelligent Interactive Distributed Systems
VU University Amsterdam
The Netherlands
{E.Pournaras,M.Warnier,FMT.Brazier}@few.vu.nl

Elth Ogston
Department of Computer Science
University of Warwick
United Kingdom
elth@dcs.warwick.ac.uk

Abstract

Global stabilization of energy networks without centralised control is a challenge. This paper explores the potential of coordination of energy usage of intelligent thermostatically controlled appliances (TCA) using a p2p network. A p2p tree overlay provides the basic structure for distributed plan aggregation and distribution. The approach is presented and discussed in the light of the risks and benefits identified.

1. Introduction

Efficient energy production, distribution and utilisation is a challenge. Current approaches to address this challenge aim to decrease fluctuations, i.e., oscillations in demand [1].

Acquiring stability in energy usage of *thermostatically controlled appliances* (TCA [7]) would make a difference, as TCAs consume 25% of the energy used in the U.S.A. [9]. Australia's Commonwealth Scientific and Industrial Research Organization (CSIRO) has recognised this problem and is currently developing industrial platforms for agent-controlled appliances [3] designed for this precise purpose. The Pacific Northwestern National Laboratory (PNNL) is currently working on GridWise$^{\text{TM}}$, a smart energy network that aims to alleviate system oscillations [2].

In this paper, TCAs –appliances such as refrigerators, air conditioners, and water heaters– are nodes in a p2p network. TCAs are assumed to be equipped with intelligent circuit controllers [2], i.e., software agents [15]. They can turn their function 'on' or 'off' or, more sophisticately, of importance in the context of this research, TCAs can select energy usage plans that satisfy user consumption requirements, minimizing oscillations in use.

Global energy stabilization is achieved by coordinating these energy plans. Our p2p network distinguishes (1) nodes representing individual energy consuming devices equipped with controllers that produce energy consumption plans, i.e., TCAs and (2) nodes that are capable of aggregating information/plans provided by other nodes to propose a best joint plan.

The p2p approach is based on a tree overlay. The main question is: *Is a peer-to-peer tree overlay a beneficial scheme for integrating fully distributed energy plan aggregation?*

Performing aggregation by means of a p2p tree overlay has a high potential to bridge the communications gap which currently prevents distributed energy resources from avoiding collective load oscillations.

This paper is outlined as follows: Section 2 illustrates the multidimensionality of energy plan aggregation by focusing on the distributed view of uncertainties in TCAs and the benefits of such a perspective. Section 3 argues for the choice of a tree to perform adaptive plan selection and energy stabilization. Finally, Section 4 concludes by discussing the benefits the proposed approach would have for solving the problem of the energy plan aggregation.

2. The Multidimensionality of Energy Plan Aggregation

Previous work has mainly addressed the problem of energy plan aggregation from a local point of view, optimizing individual electronic devices, see e.g. [7]. In contrast, this work follows a global approach to the problem. The aggregation procedure requires the right plan coordination and communication between TCA controlling agents. One complexity in this discussion is that the behavior of individual devices varies by unpredictable interdependent overlapping factors. A global energy plan aggregation should consider, and ideally exploit, this behavior.

2.1. Distributed Uncertainties

The state queueing model [7] describes the behavior of water heaters over time. According to this model, there are some factors that vary the parameters of TCAs between ranges rather than precise values. Calculation of the exact values follows different probability density functions. This can lead to major global instabilities between different appliances. The model reveals that such influences fluctuates in the range of hours in the cycle times of devices. Considering the cumulative outcome, the network can experience additional oscillations if there are no mechanisms to coordinate the uncertainties in a distributed manner. Another characteristic of TCAs is the temperature setpoint. It is the temperature value on which the thermostat changes from the 'on' state to the 'off' state. This setpoint value varies and depends on the environment of the devices and the users' behavior. The distributed view of the local oscillations of the queueing model are discussed below:

- *The standby heat loss coefficient*: This coefficient can vary over time even on the same type of machines. This means that the variation of this coefficient is represented in the individual plans for controlling agents. If the energy system has this knowledge, it can coordinate the plans in such a way that the effect of the varied coefficients is eliminated. This can be translated as a prediction of the next 'on' states of the appliances in order to minimize the introduced oscillations.

- *The ambient temperature:* In the state queuing model the ambient temperature is a function of time. From a local perspective, temporal variations can be modeled with respect to diurnally or seasonally intervals. On the other hand, considering the distributed energy scenario, time by itself is not a deciding factor as the ambient temperature has locality-aware characteristics. For example, large countries are heterogeneous in ambient temperature due to time differences and climate variations. In this case, the locality of nodes in planning is critical.

- *The random customer behavior:* Various models classify the behavior of customers based on empirical investigations [12]. The end-effect in TCAs is the modification of the temperature setpoints that denotes the transition from the 'on' to the 'off' state. This will change the cycle times of the appliances and will influence the distributed coordination issues for the standby heat loss coefficient.

2.2. Benefits in Aggregation

Assuming that the devices can generate a set of supported plans, the aggregation function is a complex operation that should *gather*, *combine*, *coordinate*, *match* and finally *choose* the globally optimum plans. These different plans are the outcome of the uncertainties discussed in Section 2.1. A scenario of clustering agents that perform aggregation of energy plans generated by refrigerators is illustrated in [10].

Given a single central aggregator, an exhaustive brute force search can be used to find globally optimal plans. In a large distributed environment, however, multiple aggregators are required to keep search costs manageable. If agents are grouped randomly under these aggregators, and locally optimal aggregates are calculated separately in each group, global results are likely to be far from optimal [10]. Grouping agents based on known complementary characteristics leads to better solutions. As characteristics of TCAs change over time, such groupings require dynamic adaptations.

Knowledge about uncertainties in TCAs can support the prediction of the future global energy needs. Individual nodes should interact in order to coordinate the cycle times, the random load behaviors and the different temperature setpoints of nodes. This imposes an *adaptive informed aggregation* based on individual node profiles and characteristics. Performance can be improved if the aggregators have knowledge of the locality of nodes and the influence of ambient temperatures. Matching in respect to one or multiple oscillations factors can give a dynamic to the system and a self-awareness towards its global sustainability.

3. Integration in a Peer-to-Peer Tree Overlay

Distributed energy plan aggregation requires an appropriate networking infrastructure. One possible way of handling this complex aggregation operation successfully is by using a p2p overlay. This section focuses on this issue and the beneficial characteristics and attributes that a tree overlay can provide to the distributed energy resources scenario.

3.1. The nature of nodes

An energy system consists of two levels [14]: The *transmission level* and the *distribution level*. Although in a p2p network nodes are allowed to function in both levels, this work focuses on the functionality of the distribution level.

Three entities are distinguished on this level. The distribution feeders, the household meters and the appliance software agents. In the remainder of this paper a p2p network with nodes (TCAs) that can both act as *plan aggregators* and *plan dispatchers* is assumed.

3.2. The Tree Characteristics

P2p aggregation can be done in either a structured or unstructured manner. Gossiping protocols can perform a range of unstructured aggregation functions [4, 6]. However, aggregating over a tree structure has an advantage for complex aggregation functions in that it benefits from a minimum communication overhead. The aggregation requests visit each node only once. Building a tree can be feasible either through autonomous actions of nodes or by taking advantage of higher level entities (feeders, substations). In the first case, an aggregation tree can be created on-the-fly by a top-down reverse-query procedure [8]. The nodes send queries recursively to their neighbors, which record them as parents. The success of the formed tree is related to the connectivity of the graph that can be satisfied by various overlay maintenance protocols such as the peer-sampling service [5]. In the latter case, the higher level entities can act as bootstrap nodes that provide potential joining nodes with the tree.

Tree overlays suffer from the effect of failures. Each failure of a node disconnects the branch 'under' the failing node. Various algorithms have been proposed that reconnect branches or, ideally, build robust trees [13]. In the latter, the procedure is reliability-driven according to metrics such as bandwidth and lifetime. Nodes with high 'reliability' are moved up in the tree. Besides such solutions that can support the concept of a p2p energy system overlay, there is another fundamental benefit. TCAs are more reliable and robust than computers and their usage is standardized. Thus failures will occur less often than in typical p2p networks used for data sharing on PCs.

3.3. Adaptive Plan Aggregation and Energy Stabilization

The hierarchy of the tree can be exploited to support some adaptive actions between appliances. Notice that this goes against the standard p2p philosophy. However, certain actions and properties can be outlined that cannot only frame a solution but also enhance the performance and achieve optimal plan coordination and global energy stabilization. The properties of the adaptive energy plan aggregation over a tree overlay are as follows:

- *Dual peer role:* In an aggregation cycle –the exchange of messages starting from the leaves of the tree and moving up to the root– nodes have two roles. The first role, the plan dispatcher, operates when child nodes send messages to their parent nodes. These messages contain the set of supported plans together with other properties, such as QoS properties or information concerning uncertainties of a node. There can also be information about the global plan stabilization as is illustrated below. The plan aggregator fulfills the second role, that is carried out by the parent node. It collects all the information from all its children, it preprocesses (plan combinations and transformations) the information to apply some local-to-global adaptations and finally it matches, compares and selects the best plan for each child.

- *Plan-driven overlay formation:* Performance can be enhanced by clustering the parent-child relationships. Uncertainties can be utilized to provide another degree of knowledge between parent and child to coordinate the different plans successfully.

- *Overlay memory:* The aggregation procedure, viewed from a temporal perspective, can be quantified based on the aggregation cycles. Performing aggregation by considering previous knowledge from former cycles or former nodes in each cycle can make the energy system more robust to potential oscillations. This knowledge concerns selected or rejected plans, the summation of selected plans as has been calculated in former cycles or any other parameters that are able to *improve or retain the energy stabilization*. The tree can provide this functionality as results can be retained and broadcasted back to all nodes at the end of every cycle.

- *Local-to-global stabilization:* Plan aggregators only know about the plans of their children. A more advanced procedure could impose plan dispatchers to send the information that has been gathered in the bottom branches upwards. This, together with the overlay memory, provides a more informative decision by plan aggregators. Summations can be reconfigured to perform matchmaking in a more global view. The result is a stabilization with a higher global effect.

- *Node-to-branch-to-tree convergence:* The aggregation starts always from the leaves with minimum knowledge, which can include some overlay memory and local knowledge. As the aggregation evolves, the previous decisions affect the future ones because information about the energy consumption is gathered in the branches. Finally, the procedure converges from the leaves/nodes and branches to the tree and the result is expected to express the optimum adaptive summation plan.

4. Conclusions

Distributed energy systems impose *a problem transition from the distributed energy resources to the distributed information resources*. Individual intelligence in electrical appliances is not enough to integrate energy networks on a global level. The problem of load-balancing and information overhead requires sophisticated network infrastructures with nodes that coordinate and communicate to satisfy the global goals of the system.

This paper argues that establishing a p2p overlay could facilitate the required integration of distributed energy systems. A tree is an interesting solution because its hierarchical characteristics can support the application of adaptive dynamic mechanisms that minimize the oscillations in energy systems. Properties such as *dual peer role, plan-driven overlay formation, overlay memory, local-to-global stabilization and node-to-branch-to-tree convergence* can be realized in a tree in order to perform adaptive plan aggregation and energy stabilization.

The next step is to investigate how these ideas can be implemented in a real distributed system. At this moment, an implementation of a p2p tree overlay in AgentScape [11] is in progress. AgentScape is an advanced software agent platform on which nodes can be modeled similarly to the agents of TCAs. Experiments on different plans will reveal more insights in this approach.

Acknowledgments

This project is supported by the NLnet Foundation `http://www.nlnet.nl` and by the UK Engineering and Physical Science Research Council grant number EP/F000936/1.

References

[1] P. R. Bijwe, S. M. Kelapure, D. P. Kothari, and K. K. Saxena. Oscillatory stability limit enhancement by adaptive control rescheduling. *International Journal of Electrical & Power Systems*, 21(7), 1999.

[2] D. Hammerstrom. Part ii. grid friendly™appliance project. PNNL 17079, Pacific Northwestern National Laboratory, October 2002.

[3] G. James, D. Cohen, R. Dodier, G. Platt, and D. Palmer. A deployed multi-agent framework for distributed energy applications. In *5th International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2006),Hakodate, Japan*, May 2006.

[4] M. Jelasity, A. Montresor, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Trans. Comput. Syst.*, 23(3):219–252, 2005.

[5] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based peer sampling. *ACM Trans. Comput. Syst.*, 25(3):8, 2007.

[6] D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *FOCS '03: Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, page 482, Washington, DC, USA, 2003. IEEE Computer Society.

[7] N. Lu, D. P. Chassin, and S. E. Widergren. Modeling uncertainties in aggregated thermostatically controlled loads using a state queueing model. *IEEE Transactions on Power Systems*, 20(2), 2005.

[8] S. Madden, R. Szewczyk, M. J. Franklin, and D. Culler. Supporting aggregate queries over ad-hoc wireless sensor networks. In *WMCSA '02: Proceedings of the Fourth IEEE Workshop on Mobile Computing Systems and Applications*, page 49, Washington, DC, USA, 2002. IEEE Computer Society.

[9] P. Mazza. The smart energy network: Electrical power for the 21st century. *Climate Solutions*, 2002.

[10] E. Ogston, A. Zeman, M. Prokopenko, and G. James. Clustering distributed energy resources for large-scale demand management. In *SASO*, pages 97–108, 2007.

[11] B. J. Overeinder and F. M. T. Brazier. Scalable middleware environment for agent-based internet applications. In *Proceedings of the Workshop on State-of-the-Art in Scientific Computing (PARA'04)*, volume 3732 of *Lecture Notes in Computer Science*, pages 675–679, Copenhagen, Denmark, June 2004. Springer.

[12] F. J. Peterson, E. Miller, R. A. Gillman, W. M. Warwick, and W. F. Sandusky. End-use load and consumer assessment program. *Energy and Buildings*, 19(3), 1993.

[13] G. Tan, S. A. Jarvis, and D. P. Spooner. Improving the fault resilience of overlay multicast for media streaming. In *DSN '06: Proceedings of the International Conference on Dependable Systems and Networks*, pages 558–567, Washington, DC, USA, 2006. IEEE Computer Society.

[14] S. E. Widergren, J. M. Roop, R. T. Guttromson, and Z. Huang. Simulating the dynamic coupling of market and physical system operations. In *Power Engineering Society General Meeting, 2004. IEEE*, pages 748–753 Vol.1, 2004.

[15] M. Wooldridge and N. Jennings. Intelligent Agents: Theory and Practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.