THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Analytical Results on the Performance of Shared Resource Allocation Systems

IOSIF SALEM

Division of Networks and Systems
Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Göteborg, Sweden 2015

Analytical Results on the Performance of Shared Resource Allocation Systems
Iosif Salem

Technical report 138L
ISSN 1652-876X
Department of Computer Science and Engineering
Distributed Computing and Systems Group

Division of Networks and Systems
Chalmers University of Technology
SE-412 96 GÖTEBORG, Sweden
Phone: +46 (0)31-772 10 00

Author e-mail: `iosif@chalmers.se`

Printed by Chalmers Reproservice
Göteborg, Sweden 2015

Analytical Results on the Performance of Shared Resource Allocation Systems

Iosif Salem

Division of Networks and Systems, Chalmers University of Technology

ABSTRACT

Resource allocation is a prevalent problem in a wide range of domains of computer science. Analytical tools that evaluate the performance of resource allocation systems allow us to compare with experimental ones, and utilize the design of such systems.

We consider shared-object systems that require their threads to fulfill the system jobs by first acquiring sequentially the objects needed for the jobs and then holding on to them until the job completion. Such systems are in the core of a variety of shared-resource allocation and synchronization systems. We provide methods for estimating the performance of such systems in terms of expected task throughput and delay for completion. To the best of our knowledge, this is a new perspective that can provide better analytical tools for the problem, in order to estimate performance measures similar to ones that can be acquired through experimentation on working systems and simulations.

We also study the problem of maximizing the energy utilization in the Smart Grid, where the energy supply becomes available in an online fashion (due to unpredictable energy sources) and the energy demand can have some flexibility (energy dispatch problem). Utilizing a proposed modeling of the energy dispatch problem as an online scheduling problem, we model supply-following demand in terms of the Adwords problem, in order to provide algorithmic solutions of measurable quality. In systems where demands are small compared to the individual supply, we prove a $\left(1 - \frac{1}{e}\right)$-competitive ratio. For cases where this does not hold, we extend the Adwords problem to utilize dynamic budgets, and present an algorithm with a $\frac{1}{2}$-competitive ratio.

Keywords: Resource allocation, analytical performance evaluation, shared-object systems, energy dispatch problem.

Preface

Parts of the contributions presented in this thesis have previously appeared in the following manuscripts.

- ▷ **Iosif Salem**, Elad M. Schiller, Marina Papatriantafilou, Philippas Tsigas, "Shared-object System Equilibria: Delay and Throughput Analysis," to appear in the *17th International Conference on Distributed Computing and Networking (ICDCN 2016)*.
  An earlier version of this paper appeared as *Technical Report, CoRR, abs/1508.01660*, `http://arxiv.org/abs/1508.01660`, 2015
- ▷ Giorgos Georgiadis, **Iosif Salem**, Marina Papatriantafilou, "Tailor your curves after your costume: Supply-following demand in Smart Grids through the Adwords problem," appeared as *Technical Report 2015:01*, Department of Computer Science and Engineering, Chalmers University of Technology, Sweden, 2015

Acknowledgments

This thesis would not have been possible without the support of many people. I would like to thank my supervisor Assoc. Prof. Elad Michael Schiller, my co-advisor Assoc. Prof. Marina Papatriantafilou and Prof. Philippas Tsigas for their support and guidance during my studies. I am grateful to Prof. Peter Damaschke for being my examiner. I am also grateful to Asst. Prof. Jukka Suomela from the University of Aalto for being my licentiate presentation's discussion leader, but also for taking the time to go through my work.

I thank Giorgos Georgiadis for being in the intersection of collaborators and friends, and to Valentin Tudor and Thomas Petig for making our common office a nice place to be and work. I must also thank Yiannis Nikolakopoulos for being a constant source of help and a good friend.

I am grateful to my current and former colleagues Aras Atalar, Aljoscha Lautenbach, Anders Gidenstam, Bapi Chatterjee, Beshr Al Nahas, Daniel Cederman, Paul Renaud Goud, Ivan Walulya, Magnus Almgren, Katerina Mitrokotsa, Mohamed Mustafa, Farnaz Moradi, Olaf Landsiedel, Oscar Morales, Nhan Nguyen, Tomas Olovsson, Vincenzo Gulisano, Zhang Fu and rest of DCS, NS and CSE colleagues for the nice working environment. Moreover, I thank Eva Axelsson, Tiina Rankanen, Peter Lundin and in general the administrative personnel for their excellent support. Special thanks to Marianne Pleen-Schreiber and her family for being always so welcoming.

I am grateful to my friends Manos Androulakis, Stavros Tzilis, Angelos Arelakis, Christina Karazisi, Vasiliki Kalogeropoulou, Yiannis Sourdis, Christina Sakalaki, and to the newcomers Vassilis Papaefstathiou, Chloe Alverti, Petros

Voudouris and Vagelis Vasilakis, for all the good time that we are having in Sweden. Furthermore, I thank all the current and former organizers of Gothenburg's PhD Pub for adding interesting events to my calendar and for working together towards that.

I especially thank my family, Sarra and Louiza, for supporting and encouraging me through all these years. I am also grateful for the support of all my friends.

Finally, I consider myself lucky to have met Maria Angelopoulou. Thank you for making my life pretty!

Iosif Salem
Göteborg, September 2015

Contents

List of Figures

Part I

INTRODUCTION

1

Introduction

Resource allocation is an often challenging problem appearing in a diverse range of areas in computer science, but also in other domains such as economics and project management. In computing systems, any running process needs to use resources such as shared memory, I/O, CPU time etc., and their allocation affects not only the process' execution but also the execution of other processes and the whole system's performance [1, 2]. Focusing on the process interaction and to improve resource utilization, we need these shared resources to be used (alt. accessed) concurrently by many processes. Nevertheless, each process' access to resources should not compromise the result of other processes' executions and the communication among processes for sharing these resources should be done in a correct and efficient manner [3].

The current and future electric power grid is another example of resource allocation systems. In such a system, that is often referred to as the Smart Grid [4], the energy supply originates from the utility company, but also from renewable energy sources (e.g. photovoltaic arrays, wind generator farms). Furthermore, energy loads can be distributed or may have less regular patterns compared to households (e.g. electric car fleets, data centers), and on the other hand, the demand for energy can have some flexibility. In this context, scheduling the allocation of energy supply to match the demand determines not only the performance on a system level but also the utilization of the system's available resources [5]. Focusing on these domains, this thesis includes two papers (Paper I and Paper II of Part II) that contribute in the area of analytically evaluating the performance of shared resource allocation systems.

The remainder of this section is organized as follows: in Sections 1.1 and 1.2 we explain the area and the challenges for Paper I, and respectively, Paper II. In Section 1.3 we highlight the contributions of these papers and in Section 1.4 we conclude and discuss future work.

1.1 Shared Resource Allocation: Models and Performance Analysis

Modeling a shared resource allocation system (SRAS) allows us to draw conclusions on its performance but it can also utilize the design of such systems. The performance of a SRAS is commonly evaluated either through experimentation, or simulation, or analytically. We focus on a class of SRAS in which computing entities that we refer to as threads carry out tasks (jobs) that require mutual exclusion to a number of shared resources (objects). These jobs arrive at the system in a given rate and are then assigned to the threads. Our work in Paper I of this thesis contributes to the analytical performance evaluation of such systems.

We study an SRAS model that includes N threads, M objects and J jobs. A job is a tuple including a vector of objects and an operation to be executed on those objects. Jobs are continuously arriving in the system and are then assigned to threads in a FIFO order. A thread carries out a job by (1) acquiring and holding on to each object in the job's object vector and (2) executing the job operation after having acquired all of the job's objects. Note that a thread acquires the job's objects in the order of the job's object vector and we require the objects in these vectors to follow an ascending order for the absense of deadlocks, i.e., to ensure that the delay of a thread to complete a job is bounded. Moreover, the competing threads acquire each object in a FIFO manner. Therefore, the delay for a thread to complete a job in such a system depends on that of other threads and so on. Furthermore, the contention caused due to these dependencies varies with the rate in wich jobs arrive in the system (job arrival rate). We say that such a system is in *equilibrium*, when the rate of arriving jobs matches the rate in which these jobs are completed (completion rate). Paper I contributes to the identification of these dependencies and gives tools for analytically evaluating the performance of such systems in equilibrium.

The existing practice considers job delay and completion rate as the performance measures of working systems. Empirical experiments often study shared resource systems at their saturation point in which the system is at its peak utilization. Let us describe peak utilization scenarios using two vectors; one for job arrival rates and another for their completion rates. A saturation point is the case in which: (1) the system is in equilibrium, i.e., the arrival rate of any particular job matches the completion rate of this job, as well as (2) the system is at the stage at which a higher arrival rate of any job to the system cannot increase the completion rates. Our study considers the entire range of these equilibria rather than just peak utilization scenarios. We then propose a procedure for finding such equilibria in an approximated fashion, if such exist in the given system. Once we approximate an equilibrium, we can estimate its performance measures, i.e., job delay, completion rate and blocking time.

This problem has well-known results studying the worst-case job delays, which may even be exponential on metrics, such as the chromatic number of the resource graph [6, 7]. In this graph, the vertices (objects) are connected if there is at least one thread that may request them both at any point in time. In the context of working systems, the expected time is rather different than the worst case and therefore computer experiments are the common way for evaluating the system performance. A generalization of the dining philosophers problem, as in [7, 8], in which every job includes a fixed set of objects that it may need is also related to this problem. We provide a new perspective that enables an analysis of the evaluation metrics by considering measures both at the system level and at the level of each resource. In particular, we consider performance measures that are associated with each resource, such as the delay, completion rate and blocking time. On the system level, we consider the job arrival and completion rates, as well as the total number of threads, N, and objects, M.

1.2 Shared Resource Allocation: The Smart Grid Case

The supply and demand in the electric power grid is a real-life paradigm of a resource allocation system. Until recently the utility company was the only supplier of energy to the electric power grid and energy should be always available on demand. Nowadays, the energy mix includes also renewable energy sources (e.g., photovoltaic arrays, wind generator farms) and possibly storage, i.e., battery arrays. Moreover, the needs from the consumers' side are also changing, since energy loads may have less regular patterns compared to households (e.g., electric car fleets, data centers). Thus, the supply and demand in the electric power grid (often referred to as *Smart Grid* [4]) has began shifting towards a market-oriented paradigm, where generated energy comes from many different sources and is brokered to consumers through utilities or electricity vendors.

In Paper II of this thesis, we address the resource allocation problem of utilizing the available energy supply through flexibility in demand (or *supply-following demand* [9]). We provide two online algorithms that can solve this resource utilization problem in a variety of supply-related assumptions (e.g. availability of storage) along with an analysis of their competitive ratios [10]. In this direction, we connect our analysis of the energy utilization problem with the Adwords problem [11]. Paper II also provides an experimental study that evaluates the algorithmic solution against real consumption data from a pilot housing project.

A number of approaches in the literature exist regarding the problem of supply-following demand. In [12], Kok et al. focus on both the supply and demand, and use a hierarchical mechanism and a market structure to match consumers with producers, with the ultimate goal of reducing peaks in consumption. On the other hand, Barker et al. [13] focus on background consumption loads (i.e. loads that the consumer does not interact with), and by applying scheduling techniques, such as a variation of the Earliest Deadline First algorithm, they shift demand during the day in order to reduce peaks in consumption. The works of Lu et al. [14] and Tu et al. [15] are closer to the context of Paper II of this thesis, since both present online algorithms with proven competitive ratios. However, they do so for special cases of interest: in [14], Lu et al. focus on fast-responding generators (e.g. gas or diesel turbines) and present an algorithm that operates for any combination of demand, supply and price, and in [15], Tu et al. focus on data centers and on a cost minimization problem where price is a parameter. In addition, Georgiadis et al. [16] present a novel modeling and an online algorithm that can schedule flexible demand in order to reduce peaks in consumption in scenarios where forecasts are unreliable or not available (e.g. renewable energy sources, energy storage). Nevertheless, one common element of all approaches above is that the optimization goal is the reduction of peak demand in the considered time period, while the criterion in this work is utilization of all available supply[1].

1.3 Contributions

The contributions of this thesis lie in the area of providing analytical results for the performance of shared resource allocation systems. In Paper I of this thesis, we provide methods for analytically estimating the expected task throughput and delay for completion in a model of shared resource allocation systems

[1] I.e. a peak might be desirable in a specific time where supply is too high, e.g. due to increased wind generation.

where computing entities require mutual exclusion to a number of shared resources due to continuously arriving tasks that they carry out. We also provide a rigorous analysis of the proposed algorithm using tools from Queuing Theory and especially Queueing Networks [17]. In Paper II of this thesis, we provide two algorithms for approximating the optimal utilization of the energy supply with respect to the flexibility of demand for the energy mix in the Smart Grid [4]. Paper II also includes two online algorithms along with proofs of their competitive ratios.

1.4 Conclusions and Future Work

This thesis provides analytical tools for evaluating the performance of resource allocation systems in two directions. In Paper I we study a system model where computing entities require mutually exclusive access to a number of resources, due to continuously arriving tasks that they carry out. The results of Paper I can be used as a pilar for modeling the performance of more complex systems, where different mechanisms are used for allocating resources, rather than the sequential acquisition that is studied here. In Paper II we provide two online algorithms for the problem maximizing the energy utilization in the Smart Grid. A possible continuation of this work could study the interaction of various pricing mechanisms with the proposed algorithms and the combined impact on maximizing resource utilization. The discussed papers follow in Part II of this thesis.

Bibliography

[1] Gadi Taubenfeld, *Synchronization algorithms and concurrent programming*, Pearson Education, 2006.

[2] George F Coulouris, Jean Dollimore, and Tim Kindberg, *Distributed systems: concepts and design*, pearson education, 2005.

[3] Alexandra Fedorova, Sergey Blagodurov, and Sergey Zhuravlev, "Managing contention for shared resources on multicore processors," *Communications of the ACM*, vol. 53, no. 2, pp. 49–57, 2010.

[4] Sarvapali D Ramchurn, Perukrishnen Vytelingum, Alex Rogers, and Nicholas R Jennings, "Putting the'smarts' into the smart grid: a grand challenge for artificial intelligence," *Communications of the ACM*, vol. 55, no. 4, pp. 86–97, 2012.

[5] Matthias D Galus and Göran Andersson, "Demand management of grid connected plug-in hybrid electric vehicles (phev)," in *Energy 2030 Conference, 2008. ENERGY 2008. IEEE*. IEEE, 2008, pp. 1–8.

[6] Nancy A. Lynch, "Upper bounds for static resource allocation in a distributed system," *J. Comput. Syst. Sci.*, vol. 23, no. 2, pp. 254–278, 1981.

[7] Nancy A. Lynch, *Distributed Algorithms*, Morgan Kaufmann, 1996.

[8] Marina Papatriantafilou and Philippas Tsigas, "On distributed resource handling: Dining, drinking and mobile philosophers," in *On Principles Of Distributed Systems, Proceedings of the 1997 International Conference, Chantilly, France, December 10-12, 1997*, 1997, pp. 293–308.

[9] R.H. Katz, D.E. Culler, S. Sanders, S. Alspaugh, Y. Chen, S. Dawson-Haggerty, P. Dutta, M. He, X. Jiang, L. Keys, A. Krioukov, K. Lutz, J. Ortiz, P. Mohan, E. Reutzel, J. Taneja, J. Hsu, and S. Shankar, "An information-centric energy infrastructure: The berkeley view," *Sustainable Computing: Informatics and Systems*, vol. 1, no. 1, pp. 7–22, 2011.

[10] Allan Borodin and Ran El-Yaniv, *Online computation and competitive analysis*, cambridge university press, 2005.

[11] Aranyak Mehta, "Online matching and ad allocation," *Foundations and Trends in Theoretical Computer Science*, vol. 8, no. 4, pp. 265–368, 2013.

[12] 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, AAMAS '05, pp. 75–82, ACM.

[13] S. Barker, A. Mishra, D. Irwin, P. Shenoy, and J. Albrecht, "SmartCap: Flattening peak electricity demand in smart homes," in *2012 IEEE International Conference on Pervasive Computing and Communications (PerCom)*, Mar. 2012, pp. 67 –75.

[14] Lian Lu, Jinlong Tu, Chi-Kin Chau, Minghua Chen, and Xiaojun Lin, "Online energy generation scheduling for microgrids with intermittent energy sources and co-generation," in *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems*, New York, NY, USA, 2013, SIGMETRICS '13, pp. 53–66, ACM.

[15] Jinlong Tu, Lian Lu, Minghua Chen, and Ramesh K. Sitaraman, "Dynamic provisioning in next-generation data centers with on-site power production," in *Proceedings of the Fourth International Conference on Future Energy Systems*, New York, NY, USA, 2013, e-Energy '13, pp. 137–148, ACM.

[16] Giorgos Georgiadis and Marina Papatriantafilou, "Dealing with storage without forecasts in smart grids: Problem transformation and online scheduling algorithm," in *Proceedings of the 29th Annual ACM Symposium on Applied Computing*, New York, NY, USA, 2014, SAC '14, pp. 518–524, ACM.

[17] Gunter Bolch, Stefan Greiner, Hermann de Meer, and Kishor S Trivedi, *Queueing networks and Markov chains: modeling and performance evaluation with computer science applications*, John Wiley & Sons, 2006.