Engineering Applications of Artificial Intelligence 45 (2015) 325–338

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

Using reputation and adaptive coalitions to support collaboration in competitive environments

Ana Peleteiro[a,*], Juan C. Burguillo[a], Michael Luck[b], Josep Ll. Arcos[c], Juan A. Rodríguez-Aguilar[c]

[a] Department of Telematics Engineering, University of Vigo, Spain
[b] Department of Informatics, King's College, London, United Kingdom
[c] IIIA-CSIC, Artificial Intelligence Research Institute, Barcelona, Spain

ARTICLE INFO

Article history:
Received 8 April 2014
Received in revised form
24 May 2015
Accepted 2 July 2015
Available online 1 August 2015

Keywords:
Coalitions
Collaboration
Reputation
Crowdsourcing
Competitive environments

ABSTRACT

Internet-based scenarios, like co-working, e-freelancing, or crowdsourcing, usually need supporting collaboration among several actors that compete to service tasks. Moreover, the distribution of service requests, i.e., the arrival rate, varies over time, as well as the service workload required by each customer. In these scenarios, coalitions can be used to help agents to manage tasks they cannot tackle individually. In this paper we present a model to build and adapt coalitions with the goal of improving the quality and the quantity of tasks completed. The key contribution is a decision making mechanism that uses reputation and adaptation to allow agents in a competitive environment to autonomously enact and sustain coalitions, not only its composition, but also its number, i.e., how many coalitions are necessary. We provide empirical evidence showing that when agents employ our mechanism it is possible for them to maintain high levels of customer satisfaction. First, we show that coalitions keep a high percentage of tasks serviced on time despite a high percentage of unreliable workers. Second, coalitions and agents demonstrate that they successfully adapt to a varying distribution of customers' incoming tasks. This occurs because our decision making mechanism facilitates coalitions to disband when they become non-competitive, and individual agents detect opportunities to start new coalitions in scenarios with high task demand.

1. Introduction

In real world domains, individuals usually face the problem of solving tasks, composed of subtasks, that cannot be achieved by them individually; to address this, they need to group together in order to be able to accomplish such tasks with guarantees. This may be the case when supporting collaboration in new Internet-based scenarios, like co-working (The Economist, 2011), or crowdsourcing (Slivkins and Vaughan, 2013), which are becoming increasingly important. In these scenarios, customers submit tasks to be serviced, with several actors competing to service them. To make this more complex, however, the number and rate of service requests changes over time, as does the service workload (the number of subtasks per tasks) required by each customer.

Over the past decade, crowdsourcing has emerged as a cheap and efficient method of obtaining solutions to simple tasks that are difficult for computers to solve but possible for humans. Crowdsourcing markets bring together requesters, who have tasks they need to perform, and workers, who are willing to perform these tasks in a timely manner in exchange for payment (Slivkins and Vaughan, 2013). There are several examples of crowdsourcing platforms, such as Amazon's Mechanical Turk (Ipeirotis, 2010) or oDesk (https://www.odesk.com/), and the popularity of crowdsourcing markets has led to empirical and theoretical research on the design of algorithms to optimize various aspects of these markets, such as the assignment of tasks (Karger et al., 2011a,b). Thus crowdsourcing has appeared as a new application domain to model and analyze the problem of online decision making, as well as design algorithms to tackle it. Online decision algorithms have a rich literature in operations research, economics, and several areas of computer science including machine learning, theory of algorithms, artificial intelligence, and algorithmic mechanism design (Slivkins and Vaughan, 2013). However, in the case of crowdsourcing, as tasks are usually not too complex, workers are typically recruited individually, without considering the possibility of recruiting groups of people to jointly perform more complex tasks. Such complex tasks include those demanded in several domains, such as in international commerce, bidding for government contracts or continuous auctions.

For example, producing a magazine, an academic paper, or a new movie may involve many individuals working in structured

* Corresponding author.
E-mail addresses: apeleteiro@gti.uvigo.es (A. Peleteiro), J.C.Burguillo@gti.uvigo.es (J.C. Burguillo), michael.luck@kcl.ac.uk (M. Luck), arcos@iiia.csic.es (J.Ll. Arcos), jar@iiia.csic.es (J.A. Rodríguez-Aguilar).

http://dx.doi.org/10.1016/j.engappai.2015.07.009
0952-1976/

teams, each with different skills and roles, collaborating on a common goal. In this way, a fixed pool of workers may need to be allocated to perform tasks that arrive dynamically and that must be completed by some deadline. Crowdsourcing therefore poses unique challenges for both workers and requesters ranging from job satisfaction to direction-setting, coordination, and quality control (Kittur et al., 2013). However, currently crowdsourcing faces many challenges that must be addressed in order to achieve all of its potential. Kittur et al. (2013) present a framework that envisions a future of crowd work enabling more complex, creative, and highly valued sustainable collaborations and co-working. In fact, they present several research challenges in crowdsourcing areas; in response, in this paper we explore a new crowdsourcing model based on their approach.

Most crowdsourcing platforms share the common feature of repeated interaction. In this respect, Afsarmanesh et al. (2009) confirm that long-lived groups (those that last in duration beyond the servicing of a single job) are successfully used in real world scenarios, such as in manufacturing or ICT, among others. According to them (Afsarmanesh et al., 2009), when groups are long-term creations, as in the case of our coalitions, successful repeated collaborations help these groups to enhance their service performance over time, since these repeated interactions improve agents collaboration. Therefore, to fully benefit from coalition-based collaborations, we must learn how to form coalitions as well as how to sustain them. However, sustaining a coalition poses two main challenges: (i) how to cope with agents within the coalition that do not honor their commitments; and (ii) how to compete with other coalitions that offer the same services. To tackle these problems requires that a coalition, as a whole, continuously adapts to remain competitive, i.e., in order to have high probabilities of being assigned tasks. Indeed, in an open environment, several competing coalitions may be formed with the aim of performing the very same service. Thus, on the coalition side, this requires the capability of (i) composing the most appropriate set of agents to perform a service; and (ii) deciding when to disband the coalition because it is no longer beneficial. Moreover, agents immersed in such competitive environments must also individually adapt by deciding (i) whether to remain in a coalition or join another existing one; and (ii) whether to remain part of a coalition or to leave it in order to start up a new one. Therefore, both coalitions and agents require decision-making mechanisms that allow them to adapt and to remain competitive over time.

Slivkins and Vaughan (2013) propose specific directions to tackle the design of a crowdsourcing model: adaptive task assignment, dynamic procurement, repeated principal-agent problem, reputation systems, and the exploration–exploitation tradeoff. In this paper, we mainly focus on the first of these, also using reputation as a way to assess the risk of cooperating with others. However, while we propose to use coalitions of agents to perform complex tasks, most previous work on task allocation with coalitions does not consider how coalitions can be maintained over time in the face of a change once they are formed. For this reason, Klusch and Gerber (2002) develop a dynamic coalition formation scheme (DCF-S) that helps agents react to changes in their set of goals and in the agent society. Similarly, Soh and Li (2003) present a dynamic coalition formation mechanism in which learning mechanisms are used at several levels to improve the quality of the coalition formation process in a dynamic, noisy, and time-constrained domain. Nonetheless, such approaches suffer from several shortcomings. First, they mainly focus on supporting the formation of a single coalition for a single task. Thus, they do not consider the bigger picture (and more realistic situation), where there are several coalitions competing to provide the very same service. In fact, most previous work has commonly assumed that a coalition disbands when the current task is finished. Hence, a coalition disappears after the coalition fulfills its goal. Mérida-Campos and Willmott (2004) explore this in the context of iterative games, where several coalitions compete to be assigned tasks in several rounds. They present a dynamic coalition formation mechanism where coalitions must adapt at each time step in order to be assigned more tasks. However, with their mechanism, agents use a pre-established strategy for joining or abandoning partners and, while there is adaptability regarding coalition composition, the adaptation of the number of coalition is not specifically addressed, i.e., how many coalitions are necessary when they compete with each other to obtain tasks.

In response to these omissions, in this paper we present a model to build and adapt coalitions so that they can be assigned complex tasks with the goal of improving the quality and quantity of the tasks completed. Thus the key contribution in this paper is a decision mechanism that allows agents in a competitive environment to autonomously enact and sustain coalitions, not only in their composition, but also in the number of coalitions that are necessary depending on the incoming tasks. Two key components in such a mechanism are the reputation, both of coalitions as a whole, and the reputation of individual agents; and the strength of collaboration *synergies* (resulting from successful repeated collaborations) within coalitions. Reputation has been shown to be effective to assess the risk of cooperating with other individuals. The notion of synergy captures the insight that working together repeatedly improves cooperation among humans. In our model, when agents employ our decision mechanism, we show that it is possible for them to maintain high levels of customer satisfaction (in terms of percentage of services finished on time). Note that we focus on customer satisfaction because quality of service is a current and major problem in crowdsourcing solutions, since there are no guarantees that a service will be good enough. In more detail, we provide the following contributions:

- First, we provide a decision making mechanism for coalitions to help them continuously adapt to remain competitive, i.e., to have high probabilities of being assigned tasks. On the one hand, our mechanism allows a coalition to assemble the most reliable team of agents to service a certain task based on the reputation of agents. On the other hand, the mechanism also helps a coalition decide whether the coalition should be sustained or otherwise disbanded because it is no longer beneficial.
- Second, we provide a decision making mechanism that allows agents to remain competitive, i.e., to have high probabilities of being assigned subtasks. On the one hand, our mechanism allows an agent to decide whether to continue to remain as part of a coalition, or instead to join another coalition. Such a decision is based on (i) the strength of the successful repeated collaborations of an agent within its coalition; and (ii) the overall reputation of the coalition. On the other hand, our mechanism allows an agent to decide when to start a new coalition.
- Finally, we provide an empirical analysis showing that the usage of our mechanisms by agents makes it possible to maintain high levels of customer satisfaction (percentage of tasks serviced on time). Here, we show that coalitions demonstrate high resilience: even when the percentage of reliable agents is low ($\sim 40\%$), the percentage of serviced tasks on time is beyond 80%. In addition, coalitions and agents demonstrate that they adapt to a varying distribution, i.e., the arrival rate, and characteristics, of customers' incoming tasks. Thus, we obtain $\sim 95\%$ of tasks serviced on time despite significant variations in the incoming arrival rate and characteristics of tasks.