Mechanism Design for Task Procurement with Flexible Quality of Service

(Extended Abstract)

E. H. Gerding[1], A. Rogers[1], K. Larson[2] and N. R. Jennings[1]
[1]University of Southampton, Southampton SO17 1BJ, United Kingdom
{eg,acr,nrj}@ecs.soton.ac.uk
[2]University of Waterloo, Waterloo, ON N2L 3G1, Canada
klarson@cs.uwaterloo.ca

ABSTRACT

In this paper, we consider the problem faced by an agent contracting multiple self-interested service providers, that are able to flexibly manipulate their quality of service in order to maximise their own utility, to complete a single computational task. We extend an existing model of such service providers, and derive optimal task procurement mechanisms in the setting where the agent has full knowledge of the cost functions of these service providers (considering both simultaneously and sequentially procurement). We then extend these results to the incomplete information setting where the agent must elicit cost information from the service providers, and we characterise a family of incentive-compatible and individually-rational mechanisms. Sequential procurement always generates greater utility for an agent than simultaneous procurement, and contracting multiple providers is preferable to contracting just one.

1. INTRODUCTION

Service-oriented computing, in which computational resources are seamlessly and dynamically procured from third party suppliers as they are required, has generated significant recent activity within the research community. Examples of such initiatives include Grid and utility computing, and these technologies are increasingly being proposed for a wide range of scientific and business workflows. However, to reach the full potential of this vision, such systems require that both the suppliers and consumers of these computational resources are able to engage in autonomous negotiation and contracting (given their own individual goals and requirements). To this end, agent-based approaches that make use of computational mechanism design have been advocated [1].

Matsubara provides a mechanism based on a payment rule that incentivises each potential contractor to truthfully reveal the resources that they will commit to the task. The approach used is similar to that more recently applied in an information setting where *strictly proper scoring rules* are used to incentivise information providers to truthfully reveal a probabilistic estimate whose generation requires the investment of costly resources [2, 3]. However, these mechanisms are restricted to the case that a single service provider is contracted to complete each task. In reality, when faced with the uncertain execution of tasks, it is common to introduce redundancy by contracting multiple providers for the same task; either simultaneously or sequentially (i.e. the agent awaits the failure of one provider before approaching another) [4].

It is this shortcoming that we address in this paper, and in more detail, in the full paper we first derive optimal task procurement mechanisms when the agent has full knowledge of the cost functions of the service providers. We consider procurement from multiple providers, and consider settings where these providers are contracted simultaneously and sequentially (i.e. the agent awaits the failure of one provider before approaching another). We then extend these results to the incomplete information setting where the agent must elicit the cost information from the providers. We characterise a family of mechanisms and prove that they are incentive compatible (i.e. the providers have a dominant strategy to truthfully reveal their cost functions to the agent), and individually rational (i.e. the expected utility of providers that participate in the mechanism is greater or equal to zero). Based on the insights obtained from the optimal, full information case, we present three mechanisms from this family: a uniform and discriminatory pricing mechanism for the simultaneous procurement case, and a mechanism for the sequential procurement case.

2. PROBLEM DESCRIPTION

We consider that contracting agent, A, has a task, T, that it would like to have completed. If the task is completed successfully then the agent receives value V, and otherwise it receives zero. We assume that there are n service providers capable of performing task, T. The probability of any provider successfully completing the task (the quality of service offered) depends on the amount of some costly resources that it decides to allocate to the task.

Formally, we assume that each service provider, i, has a potentially unlimited supply of resources[1], and denote $r_i \geq 0$ as the amount of resources that i will devote to executing the task. We assume that as a provider allocates more resources to the task, the probability that the task will be successfully executed increases. That is, there is a *quality of service* (QoS) function $P : \mathbb{R}^+ \rightarrow [0, 1]$ such that $P(r_i)$ is the probability that i successfully completes the task if it devotes r_i resources to the problem. We assume that $P(\cdot)$ is common to all providers, that if a provider devotes *no* resources to the task then it will fail (that is, $P(0) = 0$), and that the more resources are devoted to the task, the more likely it is to successfully complete the task. Finally, we assume that the probability of

[1]However, the *costs* of these resources (explained below) can become arbitrarily large.

Cite as: Mechanism Design for Task Procurement with Flexible Quality of Service, (Extended Abstract), E. H. Gerding, A. Rogers, K. Larson, N. R. Jennings, *Proc. of 8th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2009),* Decker, Sichman, Sierra and Castelfranchi (eds.), May, 10–15, 2009, Budapest, Hungary, pp. 1259–1260

success of any service provider depends only on its own resource allocation, and not on the success or failure of any other provider.

We model the costly resources of provider i with a cost function, $c_i : \mathbb{R}^+ \to \mathbb{R}$. We assume that $c_i(\cdot)$ is continuous, increasing and convex, and that $c_i(0) = 0$. In addition, we assume that for any two service providers i and j, if they have different cost functions, $c_i(r)$ and $c_j(r)$, these functions are non-crossing for $r > 0$, i.e., either $c_i(r) = c_j(r), \forall r$ or $c_i(r) \neq c_j(r), r > 0$.

In order to create incentives for service providers to invest a certain amount of resources, the contracting agent uses a payment scheme whereby the payment depends on whether or not the task was successfully completed. In the case that service provider i is contracted, it is automatically paid β_i, and then, if the task is successfully completed, it receives an additional bonus α_i. We assume that α_i is always non-negative, but place no restrictions on β_i.

3. ELICITING COST FUNCTIONS

In this section, we introduce a family of *task procurement mechanisms* TPM(α, β), such that providers are willing to truthfully reveal their cost functions to the agent, that then uses this information in order to set α_i and β_i appropriately. Specifically, the agent first decides on a maximum number of providers, m ($1 \leq m \leq n$), it wants to procure services from. The agent then executes TPM(α, β), which proceeds as follows:

1. *Cost elicitation:* All providers $i \in \{1, \ldots, n\}$ report their cost functions $\hat{c}_i(\cdot)$ to the agent. We do not assume that providers reveal their true cost functions.
2. *Service provider selection and payment specification:* The agent selects the m providers with the lowest reported cost functions. We denote the set of chosen providers by M. Since we assume that the cost functions are non-crossing, there is no ambiguity in this selection. The agent then calculates α_i and β_i for each provider $i \in M$, and reports these parameters to the providers.
3. *Task execution:* If the agent uses a simultaneous procurement strategy, then all providers in M are asked to perform the task. If the agent uses a sequential procurement strategy then one provider, $i \in M$, is chosen at random to perform the task. If provider i fails, then another provider $j \in M \setminus \{i\}$ is chosen at random. This process continues until either a provider successfully completes the task, or all providers in M have attempted the task once and failed.
4. *Payment:* Any provider $i \in M$ that was contracted by the agent and successfully completed the task is paid $\alpha_i + \beta_i$. If provider i failed at the task then it receives β_i. All providers in M that were not asked to attempt the task, and those not in M initially, receive zero.

3.1 Simultaneous Procurement

We start by noting that the two parameters α and β allow us to define a *family* of incentive compatible mechanisms. Then, Theorem 1 characterises the family of such mechanisms for the simultaneous procurement strategy. We let $\hat{c}_{m+1}$ denote the $(m+1)^{th}$ lowest reported cost function, and $r^*(\alpha, c)$ is the optimal investment decision of a provider with a cost function $c(\cdot)$ and when the agent announces parameter α.

Theorem 1. *If, for all $i \in M$, α_i is independent of $\hat{c}_i(\cdot)$, and β_i is given by:*

$$\beta_i = \hat{c}_{m+1}(r^*(\alpha_i, \hat{c}_{m+1})) - VP(r^*(\alpha_i, \hat{c}_{m+1})) \quad (1)$$

then TPM(α, β) with simultaneous procurement is individually rational and incentive compatible.

We introduce two mechanisms that satisfy the above requirement for incentive compatility: (1) *uniform pricing* and (2) *discriminatory pricing*. From Theorem 1, since β is given by Equation 1, we only need to worry about setting α_i for all $i \in M$. Now, in the uniform pricing mechanism, $\alpha = \alpha_1 = \ldots = \alpha_m$, and we calculate α as follows:

$$\alpha = V\left[1 - P(r^*(\alpha, \hat{c}_{m+1}))\right]^{m-1} \quad (2)$$

Corollary 1. *Uniform Pricing Mechanism. TPM(α, β) with α and β satisfying Equations 2 and 1 respectively is individually rational and incentive compatible.*

Although uniform pricing is a natural extension of the single-provider case, there are better alternatives. In particular, although from Theorem 1, α_i needs to be independent of $\hat{c}_i$, we can use the cost functions of other providers to calculate α_i. In more detail, we calculate α_i for a specific provider $i \in M$ by solving the following system of equations:

$$\alpha_i = V \prod_{j \in M \setminus \{i\}} [1 - P(r^*(\alpha_j, \hat{c}_j))] \quad (3)$$

$$\alpha_j = V[1 - P(r^*(\alpha_i, \hat{c}_{m+1}))] \prod_{k \in M \setminus \{i,j\}} [1 - P(r^*(\alpha_k, \hat{c}_k))] \quad (4)$$

Corollary 2. *Discriminatory Pricing Mechanism. If, for each $i \in M$, α_i is given by independently solving equations 3 and 4, and if β_i is given by Equation 1, TPM(α, β) is incentive compatible and individually rational.*

3.2 Sequential Procurement

We introduce a sequential procurement strategy in which providers are randomly selected in the third stage of the mechanism, and we refer to $i \in M$ as the i^{th} provider in the random sequence, but $\hat{c}_{m+1}$ is the $(m+1)^{th}$ lowest reported cost as before. Given these considerations, Theorem 2 reformulates the requirements for incentive compatibility in terms of the sequential procurement setting.

Theorem 2. *If, for all $i \in M$, α_i is independent of any $\hat{c}_j(\cdot)$, $m \geq j \geq i$, and β_i is given by Equation 1, then TPM(α, β) with sequential procurement is individually rational and incentive compatible.*

4. REFERENCES

[1] I. Foster, N. R. Jennings, and C. Kesselman. Brain Meets Brawn: Why Grid and Agents Need Each Other. In *Proc. of the 3rd Int. Conference on Autonomous Agents and Multiagent Systems*, volume 1, pages 8–15, 2004.

[2] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359, 2005.

[3] A. Papakonstantinou, A. Rogers, E. H. Gerding, and N. R. Jennings. A truthful two-stage mechanism for eliciting probabilistic estimates with unknown costs. In *Proc. of the 18th European Conference on Artificial Intelligence, Patras, Greece*, pages 448–452, 2008.

[4] S. Stein, N. R. Jennings, and T. Payne. Flexible service provisioning with advance agreements. In *Proc. of the 7th Int. Conference on Autonomous Agents and Multi-Agent Systems*, pages 249–256, 2008.