

Scalable Mechanism Design for the Procurement of Services with Uncertain Durations

Enrico Gerding
& Sebastian Stein
School of ECS
University of Southampton
Southampton, SO17 1BJ, UK
{eg,ss2}@ecs.soton.ac.uk

Kate Larson
Cheriton School of CS
University of Waterloo
Waterloo, ON, Canada
klarson@cs.uwaterloo.ca

Alex Rogers
& Nicholas R. Jennings
School of ECS
University of Southampton
Southampton, SO17 1BJ, UK
{acr,nrj}@ecs.soton.ac.uk

ABSTRACT

In this paper, we study a service procurement problem with uncertainty as to whether service providers are capable of completing a given task within a specified deadline. This type of setting is often encountered in large and dynamic multi-agent systems, such as computational Grids or clouds. To effectively deal with this uncertainty, the consumer may dynamically and redundantly procure multiple services over time, in order to increase the probability of success, while at the same time balancing this with the additional procurement costs. However, in order to do this optimally, the consumer requires information about the providers' costs and their success probabilities over time. This information is typically held privately by the providers and they may have incentives to misreport this, so as to increase their own profits. To address this problem, we introduce a novel mechanism that incentivises self-interested providers to reveal their true costs and capabilities, and we show that this mechanism is ex-post incentive compatible, efficient and individually rational. However, for these properties to hold, it generally needs to compute the optimal solution, which can be intractable in large settings. Therefore, we show how we can generate approximate solutions while maintaining the economic properties of the mechanism. This approximation admits a polynomial-time solution that can be computed in seconds even for hundreds of providers, and we demonstrate empirically that it performs as well as the optimal in typical scenarios. In particularly challenging settings, we show that it still achieves 97% or more of the optimal.

Categories and Subject Descriptors

I.2.11 [AI]: Distributed AI—*multiagent systems*

General Terms

Economics, Reliability

Keywords

Mechanism design, redundancy, service procurement

1. INTRODUCTION

Increasingly, participants in large distributed systems are able to discover and automatically procure the services of others. This allows service consumers to complete complex computational tasks on demand, but without the need to invest in and maintain expensive hardware. Already, such a service-oriented approach is gaining

Cite as: Scalable Mechanism Design for the Procurement of Services with Uncertain Durations, E. Gerding, S. Stein, K. Larson, A. Rogers and N. R. Jennings, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. XXX-XXX. Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

popularity in a large range of application areas, including Grids, peer-to-peer systems, and cloud and utility computing [2, 11, 4]. Despite its benefits, flexible service procurement poses new challenges that have not been addressed satisfactorily by current research. In particular, as services are offered by external providers that are beyond the consumer's direct control, their execution time can be highly uncertain, due to concurrent access by other consumers, hardware or network problems and the provider's scheduling policies. This uncertainty becomes a critically important issue if the task needs to be completed by a certain deadline.

As a result, a consumer needs to make appropriate decisions about which services to procure, balancing the probability of success with the overall cost. In particular, instead of only procuring the service of a single provider for a particular task, the consumer may benefit by redundantly procuring services from multiple providers (either simultaneously or sequentially). Furthermore, because service providers are inherently self-interested agents, they may choose to mis-represent their capabilities if this promises to increase their profits. For instance, a provider may exaggerate its speed in order to entice potential customers to procure its service, or it may inflate its costs to elicit higher payments. In these cases, consumers may end up procuring unsuitable services from providers who are unable to complete the task in a timely or effective manner. To address these challenges, in this paper we consider a generic procurement scenario with service execution uncertainty, in which multiple services can be obtained dynamically and redundantly. Furthermore, we apply mechanism design to incentivise the providers to truthfully reveal information about their costs, as well as their quality of service.

A number of related papers apply mechanism design to service procurement with execution uncertainty. In particular, Porter et al. suggest a mechanism that incentivises providers to report a truthful estimate of their success probability for a given task [9]. Ramchurn et al. extend this by considering scenarios where providers also report on their perceived reliability of other providers [10]. While these do not consider redundancy to increase the consumer's success probability, this extension is explicitly examined in [3]. However, all of these papers only consider the success probability, and not *time-critical* tasks. In contrast, we explicitly model the time component by considering uncertain service *durations*. This model is much richer, as it allows additional services to be procured dynamically over time. Uncertain durations are investigated in [12], but that work assumes that the duration distributions are known and proposes mechanisms for eliciting information about the costs of providers only. Furthermore, these mechanisms are not efficient, unlike those presented here. However, since calculating the optimal outcome is often intractable, here we additionally con-

sider, for the first time, how to find approximate solutions while maintaining the economic properties of the mechanism.

In more detail, this paper extends the state of the art in the following ways. First, we show that the well-known Vickrey-Clarke-Groves (VCG) mechanism can be applied to our procurement setting to elicit costs when the duration uncertainty of service providers is known, but that this mechanism breaks down if this is not the case. We then introduce our novel *Execution-Contingent* mechanism, where the payments depend on the actual task completion time, and show that this mechanism is incentive compatible in ex post implementation w.r.t. reporting the costs as well as duration uncertainty, and is individually rational. However, for these economic properties to hold, the mechanism requires the solution to be calculated optimally and breaks down in the case of heuristic search algorithms. Now, since finding the optimal is a computationally hard problem that is intractable in large settings, we show how we can approximate the solution while maintaining the properties of the mechanism. In particular, we provide a technique for finding the solution in polynomial time, which typically takes only seconds, even in environments with hundreds of service providers. Although we also prove that this approximation can be arbitrarily far from the optimal in theory, we evaluate it empirically and find that it performs as well as the optimal in typical procurement scenarios. Furthermore, even in particularly challenging settings, the approximation achieves 97% or more of the optimal.

Our paper is structured as follows. In Section 2, we introduce the procurement problem and show how to find an optimal, dynamic procurement plan using redundancy. In Section 3, we then consider the mechanism design problem, as well as a polynomial-time solution for approximating the solution. In Section 4, we empirically evaluate our approaches and, finally, Section 5 concludes.

2. PROBLEM DESCRIPTION

In this section, we first formalise the problem and introduce all relevant terminology (Section 2.1). Then, we show how to find an optimal procurement strategy with redundancy (Section 2.2).

2.1 Model

We consider a single service consumer A , who would like to complete a task T . The consumer derives a value $V \in \mathbb{R}^+$ if the task is successfully completed within a given deadline $D \in \mathbb{R}^+$, and 0 otherwise. Furthermore, we assume that there are m service providers, given by the set $M = \{1, \dots, m\}$, which can complete the task on the consumer's behalf. The consumer can invoke a provider $i \in M$ at any time in the interval $[0, D]$. In particular, the consumer may have multiple services running concurrently for the task and the value V is obtained if at least one of the services completes within the required time. We assume that, once invoked, the provider remains committed to the task until it is completed (possibly beyond the deadline). Thus, a service cannot be interrupted.

2.1.1 Procurement Strategy

Given the above setting, we are interested in finding a procurement strategy ρ , which specifies a plan that determines which providers should be invoked and when. We denote \mathcal{P} to be the set of all valid strategies, and compactly represent each strategy by a vector $\rho = \langle (s_1, t_1), \dots, (s_n, t_n) \rangle \in \mathcal{P}$ with $n \leq m$, where each element represents the invocation time $t_i \in [0, D]$ of a provider $s_i \in M$. Importantly, a provider s_i is only invoked at time t_i (and incurs cost c_{s_i}) if no provider has so far completed the task. Without loss of generality, we assume that $t_i \leq t_{i+1}$ (i.e., elements of the vector are ordered by their invocation time), and $s_i \neq s_j$ if $i \neq j$. We use $\rho = \emptyset$ to denote the case where no provider is invoked.

2.1.2 Service Providers

As service completion times are generally uncertain, we let X_i be a random variable describing the execution time of provider i , where we assume that $\text{Prob}(X_i \leq 0) = 0$. This is the time from invocation to completion and includes any time needed for pre- and post-processing, queueing and data transfers. The random variables X_i for $i \in M$ are distributed according to the cumulative distribution functions $F_i(t)$, where $F_i(t) = \text{Prob}(X_i \leq t)$ is the probability that the task is successfully completed at most t time units after invocation. In the following, we also refer to F_i as agent i 's *duration function*. Given the duration functions, and a strategy ρ , we can calculate the probability that task T is completed by a certain time. To do so, we let $X_\rho = \min_{i \in \{1, \dots, n\}} (t_i + X_{s_i})$ denote the random variable describing the *overall completion time* of the task. Then, the probability that the task T is completed by a certain time t is:

$$\text{Prob}(X_\rho \leq t) = \text{Prob}\left(\bigcup_{i=1}^n X_{s_i} \leq t - t_i\right) \quad (1)$$

In our empirical analysis we will assume that the distribution functions of different providers are independently distributed, in which case note that:

$$\text{Prob}(X_\rho \leq t) = 1 - \prod_{i=1}^n (1 - F_{s_i}(t - t_i)). \quad (2)$$

However, we emphasise that the theoretical results are more general, and also hold, for example, if the distributions are correlated.

On execution, provider i incurs a cost c_i , and to compensate a provider for this cost, each provider $i \in M$ receives a payment, which is given by transfer functions τ_i . These transfers are determined by a *procurement mechanism* (discussed in detail in Section 3) and could depend on the actual outcome (i.e., on the set of providers actually invoked and whether or not the task succeeded), as well as the procurement strategy. Finally, due to the inherent uncertainty in the executions outcome, we assume that all participants aim to maximise their expected utility *prior to* executing strategy ρ , which is given by:

$$E[u_i(\rho)] = E[\tau_i|\rho] - c_i \cdot (1 - \text{Prob}(X_\rho \leq t_\rho(i))), \quad (3)$$

where $E[\tau_i|\rho]$ is the *expected transfer* to provider i and $t_\rho(i)$ is its invocation time ($t_\rho(i) = \infty$ if i is not in ρ).

2.1.3 Consumer's Utility and Social Welfare

We define the consumer's utility as the difference between the value it derives from the task and all transfers it pays out, resulting in an *expected utility* of:

$$E[u_A(\rho)] = V \cdot \text{Prob}(X_\rho \leq D) - \sum_{i \in M} E[\tau_i|\rho] \quad (4)$$

As is common in the mechanism design literature, we are interested in choosing a strategy which maximises the *social welfare*, which is the sum of all utilities that agents derive in the system. However, since the actual completion time is unknown until execution, we need to consider the *expected* social welfare when selecting a strategy ρ , which is given by:

$$\begin{aligned} E[w(\rho)] &= E[u_A(\rho)] + \sum_{i \in M} E[u_i(\rho)] \\ &= V \cdot \text{Prob}(X_\rho \leq D) - \sum_{i=1}^n c_{s_i} \cdot (1 - \text{Prob}(X_\rho \leq t_i)) \end{aligned} \quad (5)$$

Note that the transfers do not appear in this equation since these simply redistribute the wealth between the agents.

2.2 Optimal Service Procurement

As noted above, we are interested in mechanisms that choose the optimal procurement strategy ρ^* , which maximises the expected social welfare, i.e., $\rho^* = \operatorname{argmax}_{\rho \in \mathcal{P}} E[w(\rho)]$. To solve this, we assume, *for now*, that we have full information about the providers' costs and duration functions, as these are required for calculating the expected social welfare, as shown in Equation 5. In Section 3, we will then consider settings with private information.

Now, as described in Section 1, a key feature of the optimal strategy is that it often includes multiple providers that are invoked *redundantly* at different times to complete the task. This redundancy allows the consumer to mitigate the uncertain behaviour of single providers and thereby derive a high probability of success. Unfortunately, including redundancy leads to a difficult optimisation problem, because it involves selecting a suitable subset of providers as well as establishing the optimal invocation times for each of these. Although solving this problem is not our main focus here, we briefly outline how this has been addressed in [12], as we will use that algorithm in our evaluation (Section 4).

In particular, the problem can be simplified by first identifying the optimal invocation times of a *given* sequence of service providers. Here, it is possible to derive a simple analytical solution when making certain assumptions about the duration functions of providers. More specifically, if it is assumed that service durations are exponentially distributed¹, then the optimal invocation times can be quickly computed using backwards induction (noting that it is always optimal to invoke the first provider at time $t = 0$).

As this computation can be done efficiently, it is then feasible to use a branch-and-bound algorithm to find the optimal sequence of providers. This algorithm starts by searching all possible sequences of providers (which grows faster than exponentially with m), but then quickly discards large parts of the search space that are known to be sub-optimal. Doing this significantly reduces the search effort, compared to an exhaustive search, and it can solve medium-sized problem with around 12 providers in less than a second and 20 providers in minutes. However, when m becomes larger than around 25 providers, then it quickly becomes infeasible to use the branch-and-bound approach, because its performance can still be exponential or worse. For such settings, a greedy heuristic has been proposed that is not optimal in general, but that has been shown to achieve 99.88% of the optimal in experiments [12].

In the rest of this paper, we will make no further assumptions about how ρ^* is found, except that we have some algorithm to compute it, given the performance characteristics of the providers. As this information is likely to be private in realistic settings, we now describe how providers can be incentivised to reveal this.

3. MECHANISM DESIGN

So far we have asserted that the consumer has all the information available about the providers' costs and duration functions in order to compute the optimal procurement strategy. Here, however, we consider the situation where this information needs to be elicited, and we would like to design transfers such that each provider maximises its expected utility by truthfully reporting this information. That is, we would like the mechanism to be *incentive compatible*. In addition, since participation is voluntary, the mechanism should award the providers with a positive utility, at least in expectation. That is, the mechanism should be *individually rational*. To this end, we start by considering the case where we know the duration

¹This is a common assumption in such settings. More specifically, it means that the providers' duration functions are given by $F_i(t) = 1 - e^{-\lambda_i t}$, where λ_i is a rate parameter.

distributions but not the costs², and then proceed to a setting where we need to elicit both duration distributions and costs. Then, in section 3.4 we show how to use approximate solutions such that computing the payments becomes computationally tractable, while at the same time retaining the desired economic properties.

3.1 Preliminaries

In the following, we denote the reported costs and duration functions revealed to the mechanism by \hat{c}_i and \hat{F}_i . We let $\hat{c} = \langle \hat{c}_1, \dots, \hat{c}_m \rangle$ be the reported costs of all service providers, and define $\hat{F} = \langle \hat{F}_1, \dots, \hat{F}_m \rangle$ analogously. As is standard, we use the notation $\hat{c}_{-i} = \langle \hat{c}_1, \dots, \hat{c}_{i-1}, \hat{c}_{i+1}, \dots, \hat{c}_m \rangle$ to denote all cost reports except from provider i (and \hat{F}_{-i} is, again, defined in a similar manner). Thus, we sometimes write $\hat{c} = \langle \hat{c}_i, \hat{c}_{-i} \rangle$ and $\hat{F} = \langle \hat{F}_i, \hat{F}_{-i} \rangle$.

Given the information announced by the service providers, it is possible to evaluate different procurement strategies. For example, we let $E[w(\rho)|\hat{c}, \hat{F}]$ denote the expected social welfare of procurement strategy ρ given the reports of the providers, and $E[w(\rho)|c, F]$ is the true expected social welfare. The optimal procurement strategy, given \hat{c} and \hat{F} , is $\rho^*(\hat{c}, \hat{F}) = \operatorname{argmax}_{\rho \in \mathcal{P}} E[w(\rho)|\hat{c}, \hat{F}]$. If the expected welfare and optimal procurement strategy are computed based on the same information, we will typically abbreviate this as $E[w(\rho^*(\hat{c}, \hat{F}))] = E[w(\rho^*(\hat{c}, \hat{F}))|\hat{c}, \hat{F}]$. Finally, we will use $\rho^*(\hat{c}_{-i}, \hat{F}_{-i}) = \operatorname{argmax}_{\rho \in \mathcal{P}_{-i}} E[w(\rho)|\hat{c}_{-i}, \hat{F}_{-i}]$ to refer to the optimal procurement strategy if provider i had never existed (where \mathcal{P}_{-i} is the set of all strategies that do not contain i).

3.2 Unknown Costs, Known Distributions

We first show that, when the duration probability functions, F_i , are publicly known, we can apply the well-known Vickrey-Clarke-Groves (VCG) mechanism [7] to our procurement setting. This mechanism proceeds as follows. Using the reported costs and the known duration functions, the consumer finds the optimal procurement strategy, $\rho^*(\hat{c}, F)$. Then, before executing $\rho^*(\hat{c}, F)$, the consumer computes and pays each service provider $i \in M$ a transfer:³

$$\tau_i = E[w_{-i}(\rho^*(\hat{c}, F))] - E[w(\rho^*(\hat{c}_{-i}, F_{-i}))]. \quad (6)$$

The second term of the transfer is the expected social welfare of the optimal procurement strategy if provider i did not exist. The first is the expected social welfare obtained by the optimal procurement strategy $\rho^*(\hat{c}, F)$, *excluding the reported cost of provider i* :

$$E[w_{-i}(\rho^*(\hat{c}, F))] = V \cdot \operatorname{Prob}(X_{\rho^*(\hat{c}, F)} \leq D) - \sum_{j=1, j \neq i}^n \hat{c}_{s_j} \cdot (1 - \operatorname{Prob}(X_{\rho^*(\hat{c}, F)} \leq t_j)) \quad (7)$$

We emphasise that, when computing $E[w_{-i}(\rho^*(\hat{c}, F))]$, only provider i 's cost is ignored, but the provider is not removed completely from the social welfare. In particular, provider i 's existence in the procurement strategy may affect the probability of success and therefore the consumer's utility, as well as that of other providers, since it may influence whether or not they are invoked.

By defining the transfers for each service provider i as was done in Equation 6, it is straightforward to show that service provider i maximises its expected utility by truthfully reporting $\hat{c}_i = c_i$. To this end, let $E[u_i(\rho^*(\langle \hat{c}_i, \hat{c}_{-i} \rangle, F))|c_i]$ be service provider i 's expected utility when all other service providers report \hat{c}_{-i} , provider

²The duration functions may be obtained from past or shared experiences, for example from using a trust or reputation system, or simply given by the provider.

³Note that, in this case, this is also the *expected transfer*, since the payment does not depend on the actual outcome.

i reports \hat{c}_i and its actual cost is c_i . Then:

$$\begin{aligned}
& E[u_i(\rho^*(\langle \hat{c}_i, \hat{c}_{-i} \rangle, F)) | c_i] = \\
& \tau_i - c_i \cdot (1 - \text{Prob}(X_{\rho^*(\langle \hat{c}_i, \hat{c}_{-i} \rangle, F)} \leq t_i)) \\
& = E[w_{-i}(\rho^*(\langle \hat{c}_i, \hat{c}_{-i} \rangle, F)) | \hat{c}_{-i}, F] - E[w(\rho^*(\hat{c}_{-i}, F_{-i}))] \\
& \quad - c_i \cdot (1 - \text{Prob}(X_{\rho^*(\langle \hat{c}_i, \hat{c}_{-i} \rangle, F)} \leq t_i)) \\
& = E[w(\rho^*(\langle \hat{c}_i, \hat{c}_{-i} \rangle, F)) | c_i, \hat{c}_{-i}, F] - E[w(\rho^*(\hat{c}_{-i}, F_{-i}))] \quad (8)
\end{aligned}$$

Since $\rho^*(\hat{c}, F)$ is, by definition, the procurement strategy which maximises $E[w]$ given reports \hat{c} , provider i can optimise the first term from its perspective by reporting $\hat{c}_i = c_i$. As for the second term in Equation 8, provider i has no influence on this term, no matter what the revealed cost, since this is based on a procurement strategy where provider i is excluded. Therefore, the service provider is best off revealing its true cost if it wishes to maximise its expected utility, irrespective of the reports of other agents. That is, the mechanism is incentive compatible in *dominant strategies* (i.e., strategy-proof). In addition, note that the expected utility is always positive, and therefore, the mechanism is also individually rational.

While this mechanism displays the desired properties, it has two weaknesses. First, the mechanism is individually rational in *expectation* only, and not *post-execution* individually rational. That is, for particular instances, upon executing the procurement strategy, the incurred costs may be greater than the transfers, resulting in a negative utility for the provider. Furthermore, upon learning this, a provider may *decommit* (i.e., refuse to attempt to execute the task or delay the task indefinitely) and instead forgo the transfers. This is because the transfers are calculated in expectation, and are not based on what occurs in practice.

The second weakness which arises is the computational burden it places on the consumer. The consumer must compute the optimal procurement strategy when considering all providers as candidates, and then the optimal procurement strategy as each provider is removed from consideration. This problem is further exacerbated by the fact that the consumer has limited computational power to start with; that is why it is procuring services from the providers. While an approximate algorithm using a greedy heuristic was proposed for handling settings with large number of providers (as described in Section 2.2), it has been well established that many mechanisms, including VCG mechanisms like ours, may not be incentive compatible if the outcome selected is sub-optimal and does not maximise social welfare [8].⁴ Therefore, heuristics and approximation algorithms must be carefully designed in order to ensure that the mechanism maintains the desired strategic properties.

To this end, in the next section, we investigate alternative solutions which address both the computational overhead and the incentive to decommit, while maintaining incentive compatibility. We also now consider settings where both the costs and durations of services are private information.

3.3 Unknown Costs, Unknown Distributions

In this section, we relax the assumption that the duration distributions are known, and consider mechanisms which need to elicit both the distributions as well as the costs. To this end, we first show that the VCG mechanism no longer exhibits our desired properties and, in particular, providers have an incentive to misreport their duration functions. We then introduce a modified mechanism, which we refer to as the *Execution-Contingent VCG* mechanism, where the transfers are contingent on the actual execution of the procurement strategy and on whether or not the task succeeded. We show

⁴Intuitively, this is because an agent can misreport its information in order to try and manipulate the approximation in its favour.

that this mechanism is incentive compatible and individually rational, and also that providers no longer have an incentive to decommit (addressing the first problem identified in section 3.2).

3.3.1 Failure of the VCG Mechanism

Consider the VCG mechanism, introduced in Section 3.2, with the modification that each provider, i , is asked to report both its cost, \hat{c}_i , and its duration probability, \hat{F}_i . The transfers for this mechanism are calculated as follows:

$$\tau_i = E[w_{-i}(\rho^*(\hat{c}, \hat{F})) | \hat{c}_{-i}, \hat{F}] - E[w(\rho^*(\hat{c}_{-i}, \hat{F}_{-i}))] \quad (9)$$

As before, provider i has no influence on the second term of the transfer function, since this is the expected social welfare that would have been achieved if provider i had not participated in the mechanism in the first place. However, we now show, by example, that a provider can improve its transfer by misreporting F_i in such a way that the first term of the transfer is increased, thus resulting in higher expected utility for the provider.

EXAMPLE 1. *For the sake of simplicity, suppose that provider i only misreports its duration distribution and that all other providers report truthfully. Also suppose that $\rho^*(c, F) = \langle (i, 0) \rangle$. That is, given the true F_i , the optimal procurement strategy is to only invoke provider i and to do so without delay. Now, suppose there exists an alternative distribution, F'_i , such that $\rho^*(c, (F'_i, F_{-i})) = \rho^*(c, F) = \langle (i, 0) \rangle$ (i.e., the strategy remains unchanged) and $F'_i(D) > F_i(D)$ (i.e., the probability of success by the deadline is higher). Clearly, since an increasing probability of success increases the consumer's utility, $E[w(\rho^*(c, (F'_i, F_{-i})))] > E[w(\rho^*(c, F))]$. It also holds that $E[w_{-i}(\rho^*(c, (F'_i, F_{-i})))] > E[w_{-i}(\rho^*(c, F))]$, and so the transfer τ_i is increased when reporting F'_i instead of F_i . Due to the fact that reporting F'_i has no impact on the probability of being invoked (i.e., the allocation remains unchanged), provider i is better off doing so.*

In the above example, the providers have an incentive to misreport their distributions as this will increase the *perceived* expected utility of other agents in the system, and thereby increase the perceived expected social welfare. This, in turn, leads to an increase in the transfers. In this particular case, the provider was able to increase the perceived expected utility of the consumer by increasing the probability of success. However, it is equally possible to construct examples that increase the expected utility of other providers.

Technically, the VCG mechanism fails here because the expected utility of an agent (either the consumer or one of the providers) depends not only on the procurement strategy, but also on the private information of other agents in the system (in this case, the information about the duration functions). Such settings are known as settings with *interdependent types* [6], and it has been shown that, in general, in situations where agents have interdependent types, it is impossible to design a mechanism which ensures that the chosen outcome maximises social welfare and is incentive compatible in dominant strategies (see, for example, [5]). Therefore, we need to make a concession on one of these properties and so, to this end, we introduce a mechanism that still maximises social welfare, but uses a slightly weaker solution concept, the *ex-post equilibrium*. We make this particular choice, because we believe it is still a very natural solution concept, and, in fact, it is often regarded as a practical solution concept [1].

3.3.2 Execution-Contingent VCG

In this section, we introduce a modification of the VCG mechanism, where the transfers made to the service providers are *contingent* on the outcome of the execution of the procurement strategy.

We show that this modification results in a mechanism which is able to elicit both the costs and the duration distributions from the service providers.

As before, each service provider, i , is asked to report its cost, \hat{c}_i , and its duration distribution, \hat{F}_i . The consumer then finds the optimal procurement strategy, $\rho^*(\hat{c}, \hat{F})$, and upon completion of execution and once the outcome is known (i.e., whether or not the task was completed successfully before the deadline and which providers were invoked), the transfers of the providers are determined. Let \mathcal{I}_ρ denote the set of invoked providers. Then:

$$\tau_i = \begin{cases} V - \sum_{j \in \mathcal{I}_{\rho^*(\hat{c}, \hat{F})} \setminus \{i\}} c_j - E[w(\rho^*(\hat{c}_{-i}, \hat{F}_{-i}))] & \text{if succeeded} \\ - \sum_{j \in \mathcal{I}_{\rho^*(\hat{c}, \hat{F})} \setminus \{i\}} c_j - E[w(\rho^*(\hat{c}_{-i}, \hat{F}_{-i}))] & \text{otherwise} \end{cases}$$

Note that, when the task fails, transfers are actually negative (meaning that the providers have to pay a penalty). However, to ensure individual rationality, payments need to be positive in expectation. To this end, we calculate the *expected transfers* (i.e., before execution of the procurement strategy) as in Equation 7 by taking into account the probability of success and the probability of a provider being invoked (and thus incurring the cost), which gives:

$$E[\tau_i] = E[w_{-i}(\rho^*(\hat{c}, \hat{F})) | \hat{c}_{-i}, F] - E[w(\rho^*(\hat{c}_{-i}, \hat{F}_{-i}))] \quad (10)$$

To see how Equation 10 differs from the previous mechanism, note the subtle but important difference between the first term of this equation and the same term of the transfers for the regular VCG (given by Equation 9). Whereas the expected transfers in Equation 9 are calculated based on *reported* distribution functions, these are now based on what actually happens, which corresponds to the true distribution functions (hence the conditioning on the true distribution F). At the same time, however, the optimal procurement strategy is determined *before* execution, and therefore this is still calculated based on the reported distribution functions (in contrast to the case where we assume complete information about the distribution functions — see Equation 6). Hence, the Execution-Contingent VCG mechanism is not incentive compatible in *dominant strategies* (we omit a formal demonstration due to space restrictions). We will, instead, try to achieve the following, weaker notion of incentive compatibility:

DEFINITION 1 (EX POST INCENTIVE COMPATIBILITY).

A mechanism is ex post incentive compatible, if, for each provider i with c_i and F_i , and for all possible cost functions and duration distributions of other providers, c_{-i} and F_{-i} , and for all $\hat{c}_i \neq c_i$ and $\hat{F}_i \neq F_i$,

$$E[u_i(\rho^*(\langle c_i, c_{-i} \rangle, \langle F_i, F_{-i} \rangle))] \geq E[u_i(\rho^*(\langle \hat{c}_i, c_{-i} \rangle, \langle \hat{F}_i, F_{-i} \rangle))].$$

In words, a mechanism is ex post incentive compatible, if, when all service providers but i report their cost and duration distributions truthfully, then no matter what this revealed information is, provider i maximises its expected utility by truthfully reporting its own cost and duration distributions. This is a weaker notion of incentive compatibility than incentive compatibility in dominant strategies, since truthtelling by provider i relies on all other providers also reporting their information truthfully. However, it is stronger than Bayesian incentive compatibility, because it does not depend on prior knowledge of the other providers' private information and because truthtelling is a Nash equilibrium, even when types are revealed after the allocation. Hence, it is often regarded as a realistic solution concept in the mechanism design literature (see, for example, [1] for a detailed discussion). We now show that:

THEOREM 1. *The Execution-Contingent VCG mechanism is: (1) ex post incentive compatible, and (2) individually rational.*

PROOF. Assume that all service providers but i truthfully report their costs and duration distributions. That is, they report c_{-i} and F_{-i} . Then, if provider i reports \hat{c}_i and \hat{F}_i , its expected utility is:

$$\begin{aligned} E[u_i(\rho^*(\langle \hat{c}_i, c_{-i} \rangle, \langle \hat{F}_i, F_{-i} \rangle))] &= \\ E[w_{-i}(\rho^*(\langle \hat{c}_i, c_{-i} \rangle, \langle \hat{F}_i, F_{-i} \rangle)) | c_{-i}, F] &- E[w(\rho^*(c_{-i}, F_{-i}))] \\ - c_i \cdot \left(1 - \text{Prob}\left(X_{\rho^*(\langle \hat{c}_i, c_{-i} \rangle, \langle \hat{F}_i, F_{-i} \rangle)} \leq t_i\right)\right) & \\ = E[w(\rho^*(\langle \hat{c}_i, c_{-i} \rangle, \langle \hat{F}_i, F_{-i} \rangle)) | c, F] &- E[w(\rho^*(c_{-i}, F_{-i}))] \end{aligned}$$

First, we note that the second term on the RHS is independent of provider i 's reported cost and duration distribution. Thus, there is nothing that provider i can do to change this value, given the reports of the other providers. Secondly, the first term of the RHS is computed *after the execution of procurement strategy ρ^** . While the selection of ρ^* depends on the reported cost and duration probabilities, the actual outcome upon execution depends on the true distribution durations. As a result, note that:

$$\begin{aligned} E[w(\rho^*(\langle c_i, c_{-i} \rangle, \langle F_i, F_{-i} \rangle)) | c, F] &\geq \\ E[w(\rho^*(\langle \hat{c}_i, c_{-i} \rangle, \langle \hat{F}_i, F_{-i} \rangle)) | c, F] & \end{aligned}$$

by definition of ρ^* . Thus, if all other providers truthfully report their costs and duration distributions, provider i is also best off revealing its information truthfully, since this will result in the mechanism selecting the procurement strategy which optimises the social welfare in expectation. This, in turn, leads to the expected utility maximisation of provider i .

The Execution-Contingent VCG mechanism is also individually rational, since $E[w(\rho^*(c, F))] \geq E[w(\rho^*(c_{-i}, F_{-i}))]$, implying that $E[u_i(\rho^*(c, F))] \geq 0$. \square

In Section 3.2, we identified two main drawbacks of the VCG mechanism: the computational requirements to calculate the optimal strategy, and the fact that service providers may have an incentive to decommit or delay execution. Interestingly, the latter is no longer a problem when using the Execution-Contingent VCG, despite the possibility that the post-execution utility of the provider may become negative. This is because the utility is calculated based on what actually happened, and any increase in the post-execution social welfare results in the same increase in transfers. Therefore, there is no need to impose additional penalties or a deposit to enforce the schedule. Instead, providers are always incentivised to execute the task and to start at the scheduled time. The first issue relating to the computational overhead of the standard VCG still arises with the Execution-Contingent VCG, as it also needs to compute the optimal procurement strategy. To address this problem, in the next part, we investigate how we can approximate the optimal solution, while maintaining the properties of the mechanism.

3.4 Approximate Mechanism Design

As mentioned in Section 2.2, computing the optimal procurement strategy becomes intractable as the number of available providers increases, and this is of particular importance in our domain, as the consumer has limited computational resources. However, as discussed in Section 3.2, replacing the optimal procurement strategy with a sub-optimal one obtained through use of a heuristic or approximation algorithm, can destroy the incentive properties of the underlying mechanism.

Given this, there are several ways to address the computational problem. In [12], for example, alternative, simpler mechanisms were introduced for a similar setting which required less computation on the part of the consumer. However, these mechanisms relied on complete information about the distribution functions, and cannot easily be extended to a setting where this information needs

to be elicited. Furthermore, these mechanisms resulted in a low efficiency (when no further knowledge about the providers was assumed, the efficiency ranged between 84 and 86 percent of the optimal). For these reasons, we now propose an alternative approach for reducing the computational burden on the consumer.

To this end, we note that Nisan and Ronen showed that it is sometimes possible to have mechanisms which knowingly use sub-optimal outcomes [8]. They proposed that instead of changing the algorithm for finding the optimal outcome (in our case, the optimal procurement strategy), one could restrict the *set of possible outcomes*, and then run the optimal algorithm on this restricted set.⁵ In the following, we apply this approach to our procurement problem, and show that the Execution-Contingent VCG mechanism is incentive compatible for appropriately restricted outcome spaces. Furthermore, we show that this approximation admits a *polynomial-time* solution to calculate the optimal (within the space of allowable outcomes) procurement strategy.

3.4.1 EC-VCG Approximation

Let η denote the maximum number of service providers that can be selected as part of a procurement strategy, and let $\mathcal{P}_\eta = \{\rho \in \mathcal{P} : |\rho| \leq \eta\}$ represent the *reduced set* of strategies. When $\eta = 1$, the reduced set of strategies contains only procurement strategies with no redundancy, while if $\eta = m$, then \mathcal{P}_η is the full procurement strategy space. We propose applying the Execution-Contingent VCG mechanism, but selecting only procurement strategies from the set \mathcal{P}_η . We can show that this restricted version of the Execution-Contingent VCG mechanism is incentive compatible and individually rational under the assumption that η is chosen without using information about the providers (i.e., before the mechanism starts).

THEOREM 2. *For any $1 \leq \eta \leq m$, if the allocation is given by $\operatorname{argmax}_{\rho \in \mathcal{P}_\eta} E[w(\rho|\hat{c}, \hat{F})]$, the Execution-Contingent VCG mechanism with the reduced set of procurement strategies, \mathcal{P}_η , is incentive compatible and individually rational.*

PROOF. Since η is independent of any of reports, no provider can increase the social welfare, and hence its transfers, by misreporting. Therefore, the proof follows directly from Theorem 1. \square

In the following section, we consider the computational properties of finding an optimal strategy in the restricted solution space.

3.4.2 Polynomial-Time Solution

We now show that, once the parameter η is set, then finding the optimal procurement strategy in \mathcal{P}_η becomes polynomial in the number of possible service providers, m . We illustrate this by considering two different scenarios. First, we consider the situation where the optimal invocation times for providers can be found analytically, given a set of providers and their ordering in the procurement strategy (such settings were discussed in Sections 2.2). The problem of finding the optimal procurement strategy then reduces to the problem of finding the optimal ordering of the providers, among all sets of η providers. This is equivalent to searching through all possible ordered subsets of set M of size η , which has size $m!/(m-\eta)!$. Once the optimal procurement strategy is found, then the transfers for all providers must be computed. If a provider is not in the optimal procurement strategy then their transfer is automatically set to zero, and thus we only need to explicitly compute the transfers for at most η providers in the optimal procurement strategy. The overall cost of this is $\eta \cdot (m-1)!/(m-\eta)!$, which results in the overall complexity of $O(m^\eta)$ for running the mechanism.

⁵In particular, if an algorithm is *maximal-in-range*, then VCG-based mechanisms, applied to the restricted problem, are incentive compatible. See [8] for details.

In the case that finding the optimal procurement strategy does not allow for a closed-form analytical solution, another approach is to *discretise time*. Let T denote the total number of discrete time slots before the deadline D .⁶ Now, given an ordered set of candidate providers, each of these providers has at most T possible invocation times, except for the first provider who should be invoked immediately (as mentioned in Section 2.2, it is always optimal to invoke the first provider with no delay). Since there are at most η candidate providers, finding the optimal invocation times therefore requires searching through less than $T^{\eta-1}$ combinations. Together with finding the optimal set of ordered candidate providers, this results in a time complexity of $O(m^\eta \cdot T^{\eta-1})$.

While these approximations have a desirable computational complexity, they may result in sub-optimal solutions. In the following, we analyse more formally how far they can be from the optimal in the worst case.

3.4.3 Worst-Case Performance

As our approximation restricts the set of solutions, it can yield a solution that is significantly worse than the optimal — especially when the optimal strategy, ρ^* , contains many more than η providers. In fact, we show that it can be arbitrarily far from the optimal.

THEOREM 3. *For any $\eta \geq 1$, there exists an m , F and c , such that the ratio between the expected social welfare of the optimal solution and the approximate solution is at least b , for any $b \geq 1$. That is:*

$$\frac{E[w(\rho^*(c, F))]}{\operatorname{argmax}_{\rho \in \mathcal{P}_\eta} E[w(\rho|c, F)]} \geq b \quad (11)$$

PROOF. We prove this by showing how to choose m , F and c , so that the above holds. For simplicity, we assume here that durations of different providers are independent, as shown in Equation 2. For all i , we let $c_i = 0$ and $F_i(D) = 1 - f$, with $1 > f \geq (1 - \frac{1}{b})^{\frac{1}{\eta}}$. Clearly, as providers are free in this example, it is always optimal to invoke all available providers at time $t = 0$. Given this, we can now choose m , so that Equation 11 holds:

$$\begin{aligned} b &\leq \frac{E[w(\rho^*(c, F))]}{\operatorname{argmax}_{\rho \in \mathcal{P}_\eta} E[w(\rho|\hat{c}, \hat{F})]} \\ \Leftrightarrow b &\leq \frac{(1-f^m) \cdot V}{(1-f^\eta) \cdot V} \\ \Leftrightarrow m &\geq \frac{\ln(1-b \cdot (1-f^\eta))}{\ln(f)} \end{aligned} \quad (12)$$

Due to our initial constraints for f , this can always be satisfied. \square

To conclude, we have shown that, through limiting the number of possible outcomes, we can obtain a solution which is polynomial in the number of service providers, while maintaining the desired properties of the mechanism. However, we have also demonstrated that, in general, this approximation can be arbitrarily far from the optimal. Nevertheless, we believe that, in realistic settings, where providers are generally costly, the benefit of increased redundancy diminishes with the number of providers. Therefore, an approximate solution may often be close to the optimal, even when η is chosen to be significantly less than m . To this end, in the next section, we evaluate our approximation empirically.

4. EMPIRICAL EVALUATION

In this empirical evaluation, we are primarily interested in the effect of varying the approximation parameter η on the overall expected social welfare. Although we showed that the approximation can

⁶Note that time slots are not required to be equally spaced. However, it is important that they are set *independent* of the providers' reports. Otherwise, the incentive properties may no longer hold.

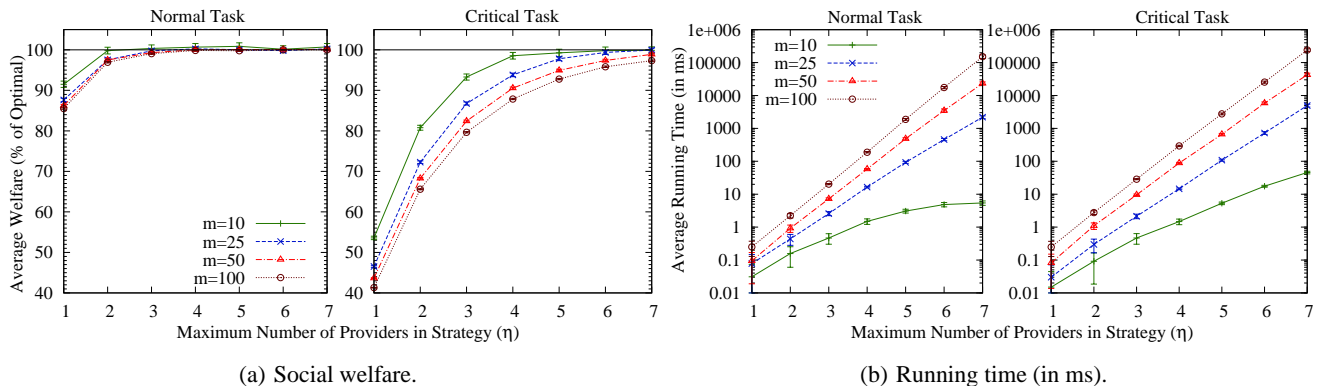


Figure 1: Performance of approximate mechanism.

be arbitrarily far from the optimal, we believe that it can achieve a good performance in realistic environments, where only a few redundant providers can often yield a high expected utility. In particular, if the approximate mechanism achieves a social welfare that is close to the optimal even for small values of η , then the approximation is clearly useful in large environments where finding an optimal solution would be intractable.

In addition to the social welfare, we will also investigate the running time of our proposed mechanism, to verify that it is feasible even when there many providers, and we will briefly discuss how the choice of η affects the consumer’s utility. We start by describing our experimental setup, and then consider all performance metrics.

4.1 Experimental Setup

In our experiments, we assume that the durations of providers are independently and exponentially distributed. For each experimental run, we generate a provider randomly by drawing its cost, c_i , and rate parameter, λ_i , independently and uniformly at random from the interval $[0, 1]$. To consider different environments, we examine two separate scenarios — where the consumer has a *normal* task with a low value $V_{\text{normal}} = 2$ and a long deadline $D_{\text{normal}} = 2$; and a more challenging one, where the consumer has a *critical* task with a high value $V_{\text{critical}} = 8$ and a short deadline $D_{\text{critical}} = 0.5$. These two scenarios were chosen as representative of the general trends our mechanism displays in different environments. In particular, in the first scenario only a few providers will be included in the optimal solution (due to the relatively low uncertainty in meeting the long deadline), while, in the second scenario, there are often ten or more providers in the optimal strategy, to ensure a high success probability despite the short deadline. We also consider environments of different sizes, with $m \in \{10, 25, 50, 100\}$. Finally, in order to ensure statistical significance, we repeat all experiments 1000 times and show 95% confidence intervals.

In the following, we start by examining the expected social welfare achieved by our approximate mechanism.

4.2 Social Welfare

Figure 1(a) shows the expected social welfare that is obtained by our approximate mechanism, as a percentage of the optimal.⁷ The trends shown here are promising, indicating that even a low η can result in a good overall performance. In particular, in the *normal* setting, when $\eta = 1$, the expected social welfare is already 86 – 92% of the optimal, and when $\eta = 2$, this increases to 97 – 100%

⁷For $m = 10$, this is obtained by solving the problem optimally. For larger settings, solving this optimally is infeasible, and so we compare the mechanism to the heuristic approach described in Section 2.2. This has been shown to be near-optimal [12].

(depending on m). Beyond this, the approximation quickly starts to perform as well as the optimal. This is not entirely surprising, because the long deadline and low value result in a relatively low benefit in using redundancy. However, the number of providers in the optimal solution is still higher than 2 in most cases (e.g., when $m = 50$, the near-optimal solution obtained by the heuristic algorithm contains more than 5 providers on average).

In the *critical* setting, we note that a slightly higher η is needed to perform close to the optimal. Here, $\eta = 3$ results in 80 – 93% of the optimal, $\eta = 4$ leads to 88 – 99% and by $\eta = 5$, 93 – 99% is reached (again, depending on m). For $\eta = 7$, the approximation achieves 97% or more in all settings. This is due to the fact that now a higher level of redundancy is required to complete the task successfully within the deadline. We also note that, as in the *normal* setting, a higher m leads to a slightly lower relative performance for the same η . This is because settings with more providers inherently offer better opportunities for redundancy.

Overall, the results are highly promising. In all cases, a relatively low value for η , compared to m , leads to a social welfare that is close to the optimal (or the near-optimal heuristic for larger m). This indicates that our approximate mechanism is feasible for complex environments where computing the optimal allocation would be intractable. Next, we examine how quickly a solution for this approximation is found in practice.

4.3 Running Time

To investigate the running time of our approximate mechanism in practice, we record the total time needed to compute the optimal allocation and all transfers for the settings described above. In more detail, we use a Java implementation of the branch-and-bound algorithm described in Section 2.2, discarding solutions with more than η providers, and we record its run-time on an Intel Pentium Core 2 Quad 2.83 GHz with 4GB RAM and running Windows 7.

The results are shown in Figure 1(b) and display some encouraging trends, indicating that the run-time of the approximate mechanism is small for most values of η . In fact, in all but two of the settings tested here, the mechanism completes in less than a minute. For example, even for the *critical* task with $m = 100$ and $\eta = 5$, the allocation and all transfers are calculated in 2.8 seconds on average (achieving 93% of the optimal). Moreover, for smaller settings, the mechanism often takes less than a second to complete. The longest average run-time recorded is just over 4 minutes ($m = 100$, $\eta = 7$ for the critical task, which achieves 97% of the optimal).

Although we do not discuss larger settings here for reasons of space, we note that the mechanism is still feasible when there are several hundreds of providers. For example, when $m = 200$ for the *critical* task with $\eta = 5$, it finds a solution in 10 seconds. When $m = 400$, this increases to 58 seconds.




Figure 2: Consumer’s utility (horizontal lines indicate consumer’s utility in welfare-maximising procurement strategy).

Finally, we now turn our attention to the consumer’s utility when using our approximate mechanism.

4.4 Consumer’s Utility

So far, we have been concerned mostly with the social welfare obtained by our mechanism. This is because it indicates how well the available providers are used to complete the task. However, it is also interesting to consider the consumer’s utility in this setting, as this is typically lower than the social welfare. This is due to the fact that the consumer pays providers more than their costs, in order to incentivise them to be truthful (in the mechanism design literature, this is commonly referred to as the *information rent*).

To quantify this loss in utility for the consumer, Figure 2 shows the consumer’s expected utility, as a percentage of the optimal social welfare (or near-optimal for larger m). For reference, we also plot the consumer’s expected utility in the optimal (social welfare-maximising) allocation. We note several interesting trends here. First, the utility lost as information rent can be significant when m is relatively low. For example, when $m = 10$, then even in the optimal solution, the consumer achieves only 69% of the social welfare (in the case of the *critical* task). This is because the payments to the agents are relatively high, as each makes a large marginal contribution (the difference between the optimal social welfare and the social welfare when removing the agent from the system, which is how the payment is calculated). However, as m increases, there are more likely to be several high-quality providers, thus leading to lower payments. For example, when $m = 50$ in the same setting, then the consumer obtains 87% of the social welfare.

The second interesting trend in the figures is that the consumer generally comes close to its optimal utility with lower values for η than when considering the social welfare. For example, when $\eta = 4$ for the *critical* task and $m = 50$, the consumer already achieves 99% of its utility in the best allocation (while only obtaining 91% of the social welfare). This means that when the consumer is primarily interested in maximising its own utility, rather than the social welfare, even fewer providers are required to achieve a high performance using the approximate mechanism.

Finally, we note that sometimes the consumer’s utility is slightly higher for certain values of η than they are in the allocation which maximises social welfare (for example, when $m = 10$ and $\eta = 3$ for the *critical* task). This is because, by selecting fewer providers, the marginal contribution of each selected provider decreases. Hence, the consumer may profit from choosing η strategically.

5. CONCLUSIONS

In this paper, we considered a common service procurement setting where services have uncertain execution times. Here, a service consumer may decide to redundantly procure multiple services, in

order to complete a given task within its deadline. However, to do this effectively, it needs to know the providers’ costs and success probabilities over time, which are typically only known by the providers. In this context, we proposed a novel mechanism that incentivises providers to reveal this private information by paying them a transfer that is conditional on the actual task completion time. We show formally that the mechanism is ex post incentive compatible and individually rational.

However, the mechanism relies on finding an optimal procurement strategy, which can be infeasible in large environments. To address this, we also proposed a polynomial-time approximate mechanism, which retains the same economic properties, but typically finds a solution in seconds or less, even when there are hundreds of potential providers. The performance of this approximate mechanism can be balanced explicitly with its computational requirements by setting the maximum number of services to procure redundantly. We showed empirically that a small value for this parameter (around 5) is typically sufficient, even when there are many providers, resulting in a social welfare that is often equivalent to the optimal in normal settings (and achieves 97% or more of the optimal in particularly challenging scenarios).

In future work, we plan to consider settings with multiple interdependent tasks that are part of complex workflows. Such settings are common in realistic service-oriented systems, but also present a more complicated computational problem. We will also examine scenarios with multiple consumers and more dynamism, where tasks and agents arrive or depart over time. Again, these occur frequently in practice and necessitate the use of more complex two-sided market mechanisms.

6. ACKNOWLEDGMENTS

This research was undertaken as part of the ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Systems) project, which is jointly funded by BAE Systems and EPSRC (EP/C548051/1), and as part of the EPSRC funded project on Market-Based Control (GR/T10664/01).

7. REFERENCES

- [1] D. Bergemann and S. Morris. Ex post implementation. *Games and Economic Behavior*, 63(2):527–566, 2008.
- [2] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the grid: Enabling scalable virtual organizations. *International Journal of High Performance Computing Applications*, 15(3):200–222, 2001.
- [3] E. H. Gerding, A. Rogers, K. Larson, and N. R. Jennings. Mechanism design for task procurement with flexible quality of service. In *Proceedings AAMAS08*, pages 1259–1260, 2009.
- [4] B. Hayes. Cloud computing. *Communications of the ACM*, 51(7):9–11, 2008.
- [5] P. Jehiel and B. Moldovanu. Efficient design with interdependent valuations. *Econometrica*, 69(5):1237–1259, 2001.
- [6] V. Krishna. *Auction Theory*. Academic Press, 2002.
- [7] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [8] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29:19–47, 2007.
- [9] R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799, 2008.
- [10] S. D. Ramchurn, C. Mezzetti, A. Giovannucci, J. A. Rodriguez-Aguilar, R. K. Dash, and N. R. Jennings. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research*, 35:119–159, 2009.
- [11] M. P. Singh and M. N. Huhns. *Service-Oriented Computing : Semantics, Processes, Agents*. John Wiley & Sons, Inc., 2005.
- [12] S. Stein, E. Gerding, A. Rogers, K. Larson, and N. R. Jennings. Flexible procurement of services with uncertain durations using redundancy. In *Proceedings IJCAI-09*, pages 292–298, 2009.