Incentive Compatible Mechanisms for Resource Procurement in Computational Grids with Rational Resource Providers

Hastagiri Prakash and Y. Narahari
Electronic Commerce Laboratory
Department of Computer Science and Automation
Indian Institute of Science, Bangalore - 560 012
hastagiri,hari@csa.iisc.ernet.in

Abstract— **In a computational grid, the presence of grid resource providers who are *rational and intelligent* could lead to an overall degradation in the efficiency of the grid. In this paper, we design incentive compatible *grid resource procurement mechanisms* which ensure that the efficiency of the grid is not affected by the rational behavior of resource providers. In particular, we offer three elegant *incentive compatible* mechanisms for this purpose: (1) *G-DSIC* (Grid-Dominant Strategy Incentive Compatible) mechanism (2) *G-BIC* (Grid-Bayesian Nash Incentive Compatible) mechanism (3) *G-OPT* (Grid-Optimal) mechanism which minimizes the cost to the grid user, satisfying at the same time, (a) *Bayesian incentive compatibility* and (b) *individual rationality*. We evaluate the relative merits and demerits of the above three mechanisms using game theoretical analysis and numerical experiments.**

I. Introduction

Grid computing can be defined as flexible, secure, coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations [4]. Grids enable the sharing, selection, and aggregation of a wide variety of resources including supercomputers, storage systems, data sources, and specialized devices that are geographically distributed and owned by different organizations for solving large-scale computational and data intensive problems in science, engineering, and commerce.

The resources in a grid include computational resources, storage systems, I/O devices, scientific instruments, computer networks, etc. These resources are geographically distributed and owned by different organizations. The sharing of these resources is subject to different administrative policies. Grid resource management is defined as the process of identifying requirements, matching resources to applications, allocating those resources, and scheduling and monitoring grid resources over time in order to run grid applications as efficiently as possible [16]. Procurement of appropriate resources required for executing a complex computational job from a pool of resource providers is an important aspect of grid resource management and is called *resource procurement*.

A. Background for the Problem

In a computational grid, a typical user would want to submit a job that can be split into tasks that execute independently on various nodes on the grid. Specifically, the *parameter sweep type jobs* consist of a number of independent, identical tasks which can all be run in parallel. Consider a grid user who has a parameter sweep type job to be submitted to the grid for computation. The grid user could conduct a procurement auction to procure resources for running the tasks contained in the job. With resource providers who are ready to offer their resources for running these tasks, the grid user would like to minimize the total cost of executing the overall job on hand. The grid user therefore has to come up with an allocation scheme that minimizes the total cost of using the available resources required for executing all the tasks. In order to solve this optimization problem, the grid user should know the true values of the costs of using the resources.

However, the grid resource providers may not be willing to provide the required information in a truthful way. This is because, the resource providers are in general *rational* and *intelligent*. A grid resource provider is *rational* in the game theoretic sense of making decisions consistently in pursuit of his own objectives. Each resource provider's objective is to maximize the expected value of his own payoff measured in some utility scale. Note that *selfishness* or *self-interest* is an important implication of rationality. Each resource provider is *intelligent* in the game theoretic sense of knowing everything about the underlying game that a game theorist knows and he can make any inferences about the game that a game theorist can make. In particular, each resource provider is *strategic*, that is, he takes into account his knowledge or expectation of behavior of other agents. He is capable of doing the required computations. In this paper, our objective is to study the grid resource procurement problem when the resource providers are rational and intelligent. When there is no ambiguity, we use the word *rational* to mean both rationality and intelligence.

B. Relevant Work in Game Theoretic Modeling of Grids

Given the above setting, it has remained an interesting idea to study the behavior of these resource providers in the presence of economic incentives that would make the resource providers voluntarily participate in the grid and perform to their maximum capacity to fulfill the overall goals of the grid and increase its efficiency. Several researchers [3], [17], [19] have studied both market-based and auction-based schemes for achieving the economic incentives. However, the introduction of economic incentives tend to induce rational and revenue-maximizing individuals to alter their bids or prices in order to increase their revenue. In such a setting where we have to ensure truth-elicitation, game theory and mechanism design [15], [13] become immediately applicable. There are a few recent research initiatives which look into the question of looking at these game theoretic issues in grid resource allocation through the glass of mechanism design theory; examples of these include:

- analyzing loss in efficiency due to the presence of rational agents [9]; and
- modeling the load balancing problem as a game and developing a mechanism for it [7].

In [7], the authors investigate the problem of designing mechanisms for resource allocation problem in grids. They consider the problem of incentive compatibility but the model they consider is severely limited because of their usage of a single dimensional type. A slightly different approach to the grid resource management problem is using cooperative games. In [6], the authors formulate the static load balancing problem as a cooperative game between the resource providers and the grid user. Utilizing the Nash bargaining solution, the authors derive an optimal allocation scheme.

Existing work has mostly addressed the resource procurement problem without a comprehensive modeling of the rationality of the grid resource providers. We attempt to do this in this paper by offering three elegant resource procurement mechanisms which satisfactorily address strategic bidding and rational behavior of the resource providers.

C. Contributions and Outline

We consider the resource procurement problem from the viewpoint of a grid user and design three different mechanisms as solutions to the resource procurement problem in computational grids with rational and intelligent resource providers. The mechanisms we design are:

- *G-DSIC* (Grid-Dominant Strategy Incentive Compatible) mechanism which guarantees that truthful bidding is a best response for each resource provider, irrespective of what the other resource providers bid
- *G-BIC* (Grid-Bayesian Nash Incentive Compatible) mechanism which only guarantees that truthful bidding is a best response for each resource provider whenever all other resource providers also bid truthfully
- *G-OPT* (Grid-Optimal) mechanism which minimizes the cost to the grid user, satisfying at the same time, (1) *Bayesian incentive compatibility* (which guarantees that truthful bidding is a best response for each resource provider whenever all other resource providers also bid truthfully) and (2) *individual rationality* (which guarantees that the resource providers have non-negative payoffs if they participate in the bidding process).

We evaluate the relative merits and demerits of the above three mechanisms using game theoretical analysis and numerical experiments. The mechanisms developed in this paper are in the context of *parameter sweep* type of jobs, which consist of multiple homogeneous and independent tasks. We believe the use of the mechanisms proposed transcends beyond parameter sweep type of jobs and in general, the proposed mechanisms could be extended to provide a robust way of procuring resources in a computational grid where the resource providers exhibit rational and strategic behavior.

The rest of this paper is organized as follows. In Section 2, we describe a game theoretic model for the grid resource procurement problem and bring out the relevance of mechanism design. Section 3 explores the design of a dominant strategy incentive compatible mechanism, *G-DSIC*, for the resource procurement problem. We develop, in Section 4, the *G-BIC* mechanism, which is a Bayesian incentive compatible resource procurement mechanism. In Section 5, we design a mechanism which achieves both incentive compatibility and individual rationality while at the same time minimizes the cost of procurement for the grid user. We call this mechanism *G-OPT*. In Section 6, we analyze the performance of the three mechanisms proposed using detailed simulation experiments.

II. A Game Theoretic Model and Mechanism Design Formulation

Consider a grid user who has a job that he is about to submit to the grid for computation. The nature of the job is of *parameter sweep type*, that is, the job consists of a number of tasks all of which can be run independently. We denote the number of tasks comprising the main job by m. That is, the user can split the job into m tasks. These tasks can run independently and there is no dependency between any subset of these m tasks. The user conducts a procurement auction to procure resources for running these m tasks.

The user announces on the grid information service, certain conditions and requirements of both the hardware and software resources required for executing the tasks making up his job. This information is received by all the prospective resource providers. Based on the requirements and conditions, resource providers decide whether or not they want to bid for executing the tasks. Once the user receives these bids, he aggregates them and based on his requirements attempts to find a cost minimizing allocation of resources for executing all the tasks. He selects up to a maximum of m resource providers out of n bidders on whose resources the tasks will be executed.

In our setting, once the grid user provides the information about the tasks and calls for the bids, the resource providers respond with a two-dimensional bid $(\hat{c}, \hat{q})$ where $\hat{c}$ represents the unit cost of executing the tasks and $\hat{q}$ is the capacity of the resource provider (maximum number of tasks that the resource provider can handle). The grid user accrues a benefit $R(q)$ from completing q tasks of the job. In our model, we assume an all-or-nothing payoff to the grid user. That is, the grid user receives zero payoff if only a fraction of the m tasks are completed and receives a constant payoff R only when all the tasks are completed. This constant payoff is the *value* that the grid user attaches to the job. This reflects the real-world quite accurately. Each resource provider i has constant marginal execution cost $c_i \in [\underline{c}, \bar{c}] \subset [0, \infty]$ and also

n	Number of resource providers
N	Set of resource providers $\{1,2,\ldots,n\}$
i	Index for resource providers
m	Number of identical tasks which a grid user wishes to get executed on the grid
c_i	True value of unit cost of executing a task for resource provider i
q_i	True capacity (maximum number of tasks that can be executed) of resource provider i
$\hat{c}_i$	Unit cost of executing a task announced by resource provider i
$\hat{q}_i$	Capacity announced by resource provider i
b_i	$= (c_i, q_i)$ - True type of resource provider i
$\hat{b}_i$	$= (\hat{c}_i, \hat{q}_i)$ - Bid announced by resource provider i
B_i	Set of possible types or bids of resource provider i

TABLE I

NOTATION USED FOR THE MODEL

a maximum capacity $q_i \in [\underline{q}, \bar{q}] \subset [0, \infty]$. F_i denotes the joint distribution of the marginal cost c_i and execution capacity q_i. We work with a joint distribution in order to allow for the cost and quantity values to be correlated. We assume that the resource providers are symmetric in the sense that they all have the same joint distribution function of the marginal cost and execution capacity, with the ranges over which c_i and q_i vary being the same for all resource providers. Note that the actual values of (c_i, q_i) are known only to bidder i.

Since the resource providers are rational and intelligent, they will try to maximize their profit by possibly providing false bids of their valuations. The job of designing an incentive compatible mechanism addresses exactly this issue. The mechanism should make it optimal for the resource providers to bid their true valuations. In our mechanism, the true type of each bidder is represented by $b_i = (c_i, q_i)$ and each bid is represented by $\hat{b}_i = (\hat{c}_i, \hat{q}_i)$. The set of possible true types (and the bids) for resource provider i is denoted by B_i. Note that in our setting $B_i = [\underline{c}, \bar{c}] \times [\underline{q}, \bar{q}]$. Throughout this paper, we also adopt the notational convention of b_{-i} for $(b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$ and B_{-i} for $B_1 \times \ldots \times B_{i-1} \times B_{i+1} \times \ldots \times B_n$. Thus a type profile b is also represented as (b_i, b_{-i}).

Table I summarizes the notation that we use in this paper. Now, one may recall from mechanism design theory in quasi-linear environments [13], [5] that a grid resource allocation mechanism, or any auction for that matter, consists of:

1) an **allocation function** $x : B \to \mathscr{R}^n_+$ which specifies the number of tasks to be allocated to each of the bidders, and
2) **a payment function** $t : B \to \mathscr{R}^n$ which specifies the amount of money transferred from the grid user to the resource providers.

In this paper, we seek to to design a mechanism that would ensure truth elicitation from all the resource providers. Once true values are obtained, the allocation scheme is just an optimization algorithm that minimizes the cost to the grid user. This truthful elicitation can be done in the setting of either dominant strategy equilibrium or Bayesian Nash equilibrium.

III. G-DSIC: A DOMINANT STRATEGY INCENTIVE COMPATIBLE MECHANISM FOR GRID RESOURCE PROCUREMENT

In the case of dominant strategy implementation, truth revelation becomes a best response strategy for each resource provider irrespective of what the other resource providers reveal. The *G-DSIC* mechanism implements the social choice function $f(b) = (x(\hat{b}), t_1(\hat{b}), \ldots, t_n(\hat{b}))$, $\forall \hat{b} \in B$ truthfully in dominant strategies. Here $k(.)$, and $t_i(.)$, $\forall i \in N$ are interpreted in the following way. $x(\hat{b})$ is the allocation rule that represents the number of tasks each resource provider has to execute, given the type profile $\hat{b}$ of the nodes. The payment vector $(t_1(\hat{b}), \ldots, t_n(\hat{b}))$ gives the payments received by the resource providers, given the bid profile $\hat{b}$. For any $i \in N$, if $t_i(\hat{b}) > 0$, then the interpretation is that i receives some positive amount and if $t_i(\hat{b}) < 0$, then the interpretation is that i pays some positive amount.

As per our setting, the resource providers submit their bids and we have to decide the allocation and payments based on the bids. We use the VCG (Vickrey-Clarke-Groves) mechanism [13], [5] to design an auction which achieves truth elicitation in dominant strategies. The appeal behind using a dominant strategy incentive compatible mechanism such as the VCG mechanism is as follows. Such a mechanism will make revealing the true costs a best response strategy for every resource provider regardless of the strategies adopted by the others. Also, the complexity of bidding logic for each resource provider is trivial since he does not have to do any computation to come up with a bidding strategy. Such a mechanism is also very robust since no assumptions need to be made about the strategic behavior of the rational resource providers.

We first define the notion of a *critical resource provider*. This notion is required because the structure of the G-DSIC mechanism demands that the allocation problem be solvable with each resource provider excluded from the set of bidders.

Definition 1: A critical resource provider in a resource procurement scenario is a resource provider without whose presence, the job cannot be completed by any subset of the remaining resource providers and whose presence in some subset of winning resource providers makes the job executable. That is, the sum of the capacities of all the other resource providers is less than the number of tasks m.

The above definition is equivalent to saying that a resource provider i is critical if and only if, $\sum_{j \in N} q_j \geq m$ and $\sum_{j \in N, j \neq i} q_j < m$. The structure of the *G-DSIC* mechanism requires that there be no critical resource provider in the system. The allocation and payment rules for the *G-DSIC* mechanism are the same as that for the VCG mechanisms [13], [5].

Since the G-DSIC scheme is based upon the VCG mechanism, it carries over both the advantages and disadvantages of the VCG mechanism. Following are the desirable properties of the *G-DSIC* mechanism:

1) The mechanism is dominant strategy incentive compatible. This ensures that the resource providers find it optimal to bid their true types irrespective of what the other providers bid. Also, the strategic complexity of the bidders is trivialized since they do not have to worry about the behavior of the other providers and do not have to compute any estimation of their types.
2) The mechanism is individually rational, that is, the resource providers are guaranteed to get a non-negative payoff by participating in the auction.
3) The mechanism is allocatively efficient, that is, the allocation is such that the the total value of the providers is maximized.

For a proof of the above properties, the reader is referred to [18]. We now list two important limitations of the G-DSIC scheme.

1) There should be no critical resource provider. This makes the mechanism unusable at certain times even if there are enough resources with the set of resource providers to complete the job. For example, consider a grid user who has a job with $m = 10$. Now if there are 3 resource providers each of whose capacity valuation is 4. Now in this case each of them is a critical resource provider and the job cannot be completed by the other two alone. In this case, we cannot apply the *G-DSIC* mechanism though the job can be completed by the three of them together.
2) VCG mechanisms, on which the *G-DSIC* protocol is based, are *not budget balanced* [13], [5]. Lack of budget balance is a well known limitation of VCG mechanisms. This means that the *G-DSIC* auction may not sustain itself and could require monetary endowment from an external agency.

The above two limitations motivate us to relax the strong property of dominant strategy incentive compatibility to the much weaker Bayesian incentive compatibility property. This leads to the *G-BIC* mechanism which is discussed next.

IV. G-BIC: A Bayesian Nash Incentive Compatible Mechanism for Grid Resource Procurement

As we have seen in Section 3, there are two important drawbacks of the *G-DSIC* mechanism which limit its applicability. Also, the cost of procurement tends to be high when using the *G-DSIC* mechanism. As a solution to these problems, we design a *Bayesian Nash incentive compatible* mechanism (*G-BIC*) for the grid resource procurement problem. The basic model for this section is the same as the one discussed in Section 3.

A. The G-BIC Mechanism

The dAGVA mechanism, as any other mechanism with quasi-linear utilities, consists of an allocation function and a payment function. In this section, we design a grid resource procurement mechanism based on the dAGVA mechanism. This Bayesian incentive compatible mechanism implements the social choice function (SCF) $f(b) = (k(b), t_1(b), \ldots, t_n(b))$, $\forall b \in B$ corresponding to the job allocation problem in computational grids. Here $k(b)$, and $t_i(b)$, $\forall i \in N$ are interpreted as in Section 3. The allocation rule in the *G-BIC* mechanism is the same as the one for *G-DSIC*. The payment rule follows that of the well known dAGVA mechanism (d'Aspremont and G*é*rard-Varet-Arrow) mechanism [13], [5]. The *G-BIC* mechanism has the following advantages compared to the *G-DSIC* mechanism.

1) The *G-BIC* mechanism will work even in the presence of critical resource providers.
2) The mechanism is budget balanced, hence it is a self-sustaining auction that does not need any external agency to support it monetarily.
3) The payments that the *G-BIC* mechanism yield are typically lower compared to that of *G-DSIC*. This implies that the cost of resource procurement is lower.

Though *G-BIC* overcomes both the limitations of the *G-DSIC* mechanism, it suffers from an important limitation, namely, lack of individual rationality. This implies that some of the resource providers could end up with negative payoffs. This motivates us to look into mechanisms that satisfy both Bayesian incentive compatibility and individual rationality and leads us in a natural way to the *G-OPT* mechanism. This is discussed next.

V. G-OPT: An Optimal Auction Mechanism for Resource Procurement in Computational Grids

The motivation for the *G-OPT* mechanism is provided by the need to come up with a mechanism that minimizes the cost of resource procurement ensuring at the same time that two important properties - (1) Bayesian incentive compatibility and (2) individual rationality are satisfied.

A. Related Work

We discuss here some of the important contributions in the area of optimal mechanism design and also some contemporary results. The problem of designing an optimal mechanism was first studied by Myerson [14]. Myerson considers the setting of a seller trying to sell a single object to one of several possible buyers. Further, the selling game is in a setting of incomplete information where each entity (buyer or seller) is privy only to their own private valuations of the object and not to the valuations of the others (independent private values model). The author introduces a new kind of uncertainty called the preference uncertainty which stems from personal preferences of the agents rather than quality or information uncertainty. The author characterizes all feasible auction mechanisms and derives the allocation rule and payment function for the optimal auction mechanism. This body of work is crucially used in developing the *G-OPT* mechanism in this paper.

Myerson's work was specifically in the setting of single unit single item auctions; this can be easily extended to multi-unit auctions with unit demand. But problems arise when the unit-demand assumption is relaxed. We move into a setting of multi-dimensional type information which makes truth elicitation non-trivial. Several initiatives have addressed this problem, albeit under some restrictive assumptions.

Several attempts have assumed, for example, that even though the seller is selling multiple units (or even objects), the type information of the entities is still one dimensional [20], [1], [2]. The problem with such an assumption is that, even though the model may fit certain specialized requirements, it is generally not applicable to the case of multi-unit demand which makes the type information multi-dimensional.

In [12], the authors provide a unique characterization of the optimal auction problem as a problem of finding an extreme point of a feasible set. They show for well-behaved distributions of the buyers' valuations that virtually any extreme point of the feasible set maximizes the seller's revenue. They also provide an algebraic procedure to check if a given mechanism is an extreme point of the feasible set.

Malakhov *et al* [11] and [10] present a new approach to the problem of designing optimal auctions. They use a network flow approach and model the objective of designing an optimal auction to one of finding a shortest path on a lattice. These were the only efforts which handle the continuous case problem of optimal mechanism design in a discrete setting.

The model that we adopt in this section is based on the work by Kumar *et al* [8]. In this effort, the authors consider a problem of optimal multi-unit procurement and characterize the optimal auction. They also devise a one-shot *get-your-bid* procurement auction for the model they devise.

B. Initial Setup

We would be using a model that is similar to the one discussed in Section 2, with some additional assumptions to enable us to design the required mechanism. Table II describes the notation. In this section, we use the uppercase of a letter to denote the expectation of the component represented by the letter rather than the actual realization. Thus, for instance, $T_i(\hat{c}_i, \hat{q}_i)$, is an expectation of the value $t_i(\hat{c}_i, \hat{q}_i)$.

We first define some notions that would be of use to us in the development of the *G-OPT* mechanism. The offered expected surplus gives us a measure of the surplus that is on offer to a resource provider under a certain mechanism. That is, it characterizes the expected value of the profit to a resource provider, were he to participate in the mechanism.

Definition 2: For a resource procurement mechanism (x,t), we define the *offered expected surplus* as follows.

$$\rho_i(\hat{c}_i, \hat{q}_i) = T_i(\hat{c}_i, \hat{q}_i) - \hat{c}_i X_i(\hat{c}_i, \hat{q}_i)$$

The offered expected surplus when the resource provider i bids $(\hat{c}_i, \hat{q}_i)$ is a measure of the expected transfer payment. The expected surplus $\pi_i(\hat{c}_i, \hat{q}_i)$ of resource provider i when the bid is $(\hat{c}_i, \hat{q}_i)$ is given by

$$\pi_i(\hat{c}_i, \hat{q}_i) = T_i(\hat{c}_i, \hat{q}_i) - c_i X_i(\hat{c}_i, \hat{q}_i) = \rho_i(\hat{c}_i, \hat{q}_i) + (\hat{c}_i - c_i) X_i(\hat{c}_i, \hat{q}_i)$$

The true surplus π_i is equal to the offered surplus only if the mechanism is incentive compatible.

TABLE II

NOTATION USED FOR THE G-OPT MECHANISM

m	number of tasks contained in the job of the grid user
n	number of resource providers
N	$\{1,\ldots,n\}$ - set of resource providers
i	index for resource providers
R	total revenue to the user by completing m tasks
c_i	actual cost of executing a task for provider i
q_i	actual capacity of (maximum number of tasks that can be executed by) resource provider i
b_i	actual type of resource provider i $b_i = (c_i, q_i)$
$\hat{c}_i$	reported cost of task execution for provider i
$\hat{q}_i$	reported capacity of the provider i
$\hat{b}_i$	reported type of provider i $\hat{b}_i = (\hat{c}_i, \hat{q}_i)$
C_i	set of unit cost values of provider i
Q_i	set of capacity values of provider i
B_i	set of types of provider i $B_i = [\underline{c}, \bar{c}] \times [\underline{q}, \bar{q}]$
b	a type profile of the providers, represented by $(b_1,\ldots,b_n)$
$\hat{b}$	a bid profile of providers, represented by $(\hat{b}_1,\ldots,\hat{b}_n)$
B	set of profiles of the types of the providers represented by $B_1 \times B_2 \times \ldots B_n$
b_{-i}	a profile of types without provider i, represented by $(b_1,\ldots,b_{i-1},b_{i+1},\ldots,b_n)$
$\hat{b}_{-i}$	a profile of types without provider i, represented by $(\hat{b}_1,\ldots,\hat{b}_{i-1},\hat{b}_{i+1},\ldots,\hat{b}_n)$
Θ_{-i}	set of profiles of the types without provider i, represented by $B_1 \times \ldots \times B_{i-1} \times B_{i+1} \times \ldots \times B_n$
$f(.)$	social choice function
$x(.)$	allocation rule of $f(.)$
$X_i(.)$	Expected value of allocation to provider i
$t_i(.)$	payment to a provider i
$T_i(.)$	expected payment to a provider i

Following the idea from [14], we use a specific function called the *virtual cost function*, which is then used to rank the resource providers. The use of this virtual cost function in place of the actual cost bids, increases the expected profit of the grid user possibly at the cost of efficiency.

Definition 3: We define *virtual cost* as

$$H_i(c_i, q_i) = c_i + \frac{F_i(c_i|q_i)}{f_i(c_i|q_i)}$$

This virtual cost parameter is similar to the virtual value parameter used in [14], except that here we allow for the cost and the quantity values to be correlated.

Assumptions: The following assumptions are required. While these assumptions are not so strong as to limit the application of the results presented, they simplify the model enough to derive certain results which would not otherwise be achievable.

- *Assumption 1:* For all $i = 1,2,\ldots,n$, the joint density function $f_i(c_i, q_i)$ is completely defined for all $c_i \in [[\underline{c}, \bar{c}]$ and $q_i \in [\underline{q}, \bar{q}]$, so that the conditional density function $f_i(c_i|q_i)$ has full support.

- *Assumption 2:* For all $1 = 1,2,...,n$, the virtual cost function $H_i(c_i,q_i)$ defined in (3) is non-decreasing in c_i and non-Increasing in q_i. It is to be noted that this assumption imposes a condition on the relation between the cost and quantity of each resource provider. This is similar to the regularity assumption in [14]. It holds especially when the cost and the quantity are independent of each other.

C. The G-OPT Mechanism

As explained in Section 3, a grid resource procurement mechanism with quasi-linear utilities consists of:

1) an *allocation function* $x : B \rightarrow \mathscr{R}_+^n$ which specifies the number of tasks to be allocated to each of the bidders, and
2) a *payment function* $t : B \rightarrow \mathscr{R}^n$ which specifies the amount of money transferred from the grid user to the resource providers.

As auction designers, we are looking for the allocation function x and the transfer function t that minimizes the expected payment to be made by the grid user. Now the total expected profit which we are trying to maximize for the user is

$$\pi(x,t) = \mathbf{E}_b[R - \sum_{i=1}^n t_i(b)]$$

subject to:

1) *Feasibility*: $x_i(b) \leq q_i \ \ \forall i = 1,\ldots,n$ and $b \in B$
2) *Individual Rationality*: The expected interim payoff for each bidder is non-negative. That is, $\pi_i(b_i) = \mathbf{E}_{b_{-i}}[t_i(b) - c_i x_i(b)] \geq 0$
3) *Bayesian Incentive Compatibility*: The bidders must be induced by the mechanism to truthfully bid their valuations; in particular, we require that truth revelation be optimal for each resource provider, provided the other resource providers bid truthfully. That is, $E_{b_{-i}}[t_i(b_i,b_{-i}) - c_i x_i(b_i,b_{-i})] \geq$ $E_{b_{-i}}[t_i(\hat{b}_i,b_{-i}) - c_i x_i(\hat{b}_i,b_{-i})], \forall i \in N, \forall b_i \in B_i \forall \hat{b}_i \in B_i$

We receive a cost minimizing mechanism for the grid user as solution to this optimization problem. Any mechanism which satisfies the constraints given above and achieves a minimum cost for the grid user is said to be an optimal auction on the lines of Myerson [14].

The problem of designing an optimal mechanism for the model we had discussed in Section 2 is different from the classic optimal auction presented in [14] in the sense that the demand is not of unit quantity. The grid user has a finite demand and allocates possibly multiple tasks to each resource provider. It is this multidimensional type of the resource providers that make this an interesting and non-trivial problem. We not only have to consider the cost but also maximum quantity as private information. This is non-trivial in the sense that the resource providers are not of unit capacity type or un-capacitated. This bi-dimensional type profile of the resource providers makes the application of some traditional optimizing schemes unsuitable.

1) Characterizing the Optimal Solution: We are now ready to characterize the optimal procurement scheme for the grid resource procurement problem. We provide certain characterizations of all incentive compatible and individually rational mechanisms. From this set of the mechanisms, the one that minimizes the procurement cost for the grid user is the optimal mechanism that we are looking for. We now present an important proposition which characterizes the set of all incentive compatible and individually rational mechanisms.

Proposition 1: For the resource procurement problem under discussion, a feasible allocation rule x is Bayesian incentive compatible and individually rational if the expected allocation $X_i(c_i,q_i)$ is non-increasing in cost valuation c_i and the offered surplus $\rho_i(\hat{c}_i,\hat{q}_i)$ is of the form

$$\rho_i(\hat{c}_i,\hat{q}_i) = \rho_i(\bar{c},\hat{q}) + \int_{\hat{c}_i}^{\bar{c}} X_i(y,\hat{q}_i)dy \qquad (1)$$

.

Also $\rho_i(\hat{c}_i,\hat{q}_i)$ must be non-negative and non-decreasing in $\hat{q}_i$ for all $\hat{c}_i \in [\underline{c},\bar{c}]$.
The reader is referred to [18] for a proof of the above proposition.

D. The G-OPT Allocation and Payment Rules

In this section, we present the allocation rule and the payment function for the *G-OPT* mechanism. The first thing to note before we proceed with the design of the mechanism is that in order to minimize the cost of procurement of resources, the mechanism should be such that the additional cost paid by the user as informational rent must be minimized as much as possible.

Here we present an allocation rule and a payment function which we claim is an optimal mechanism for the grid resource procurement problem. Before we proceed, recall the definition of the virtual cost function

$$H_i(c_i,q_i) = c_i + \frac{F_i(c_i|q_i)}{f_i(c_i|q_i)}$$

We would be using the virtual cost function rather than the actual costs in the optimal mechanism. We would be using these $H_i(c_i,q_i)$ values to compute the assignment vector. We rank the resource providers on the basis of their $H_i(c_i,q_i)$ values and keep allocating tasks to them at their full capacity in increasing order of H_i values until all the tasks are allocated. Let $[i]$ denote the resource provider with i^{th} lowest $H_i(c_i,q_i)$ value and $[\bar{i}]$ be the resource provider such that $[\bar{i}]$ satisfies,

$$\sum_{j=[1]}^{[\bar{i}]-1} q_{[j]} < m \text{ and}$$

$$\sum_{j=[1]}^{[\bar{i}]} q_{[j]} \geq m$$

Then the allocation function is given by

$$x_{[i]} = \begin{cases} \hat{q}_{[i]}, & [i] < [\bar{i}] \\ m - \sum_{j=[1]}^{[\bar{i}]-1} \hat{q}_{[j]}, & [i] < [\bar{i}] \\ 0, & otherwise \end{cases} \qquad (2)$$

With this allocation function, it is easy to see that $X_i(c_i, q_i)$ is non-increasing in c_i and non-decreasing in q_i. Given this allocation rule, we present the payment function that would make this an optimal auction for the resource procurement problem.

$$t_i(\hat{b}) = c_i x_i(\hat{b}) + \int_{c_i}^{\bar{c}} X_i(y, \hat{q}_i) dy \tag{3}$$

We now present some of the properties of the *G-OPT* mechanism that we have designed.

Proposition 2: The allocation mechanism defined by (2) has the following properties.

1) $X_i(c_i, q_i)$ (i.e $E[x_i((c_i, q_i), b_{-i})]$) is non-increasing in c_i for all fixed $q_i \in [\underline{q}, \bar{q}]$.
2) $X_i(c_i, q_i)$ (i.e $E[x_i((c_i, q_i), b_{-i})]$) is non-decreasing in q_i for all fixed $c_i \in [\underline{c}, \bar{c}]$.

Proposition 3: The *G-OPT* mechanism with allocation function, x defined by 2 and the payment function t as defined by 3 is *Bayesian incentive compatible*, *individually rational* and cost minimizing.

VI. Simulation Experiments and Results

In this section, we present the results of some numerical experiments to get an insight into the performance and properties of the three mechanisms. It is logical to compare the performance of the proposed mechanisms against centralized protocols involving benevolent (non-rational) grid resource providers. For this purpose, we define a naive, centralized resource procurement algorithm for benchmarking purposes.

A. A Naive Benchmark Algorithm

The naive algorithm assumes that the bidders are non-strategic and benevolent. That is to say, the central design authority has complete information about all the bidders and the bidders are prepared to perform to their maximum capacity. It is difficult to see the naive algorithm implemented in practice; we only use it as a benchmark the performance of the proposed mechanisms. The results produced by the naive auction algorithm provide a measure of the average cost of procurement when all the resource providers are truthful and benevolent and further are controlled centrally by a single authority. When implementing the naive auction algorithm, we obviously need not be concerned about incentive compatibility or individual rationality. This is the algorithm that will be followed in a centralized setting where the resource providers are not self-interested and are all working to achieve a globally optimal allocation. The ratio of the procurement cost of a designated mechanism to the procurement cost of the naive auction mechanism could be called its *competitive ratio*.

B. Simulation Scenarios

We conducted our simulations with two different settings, and for different sets of resource providers. In each setting, the valuations of the resource providers were randomly generated. We studied the performance of the mechanisms by varying the number of resource providers. Each data point that is plotted is obtained as the average of 1000 replications. This large number of replications is required as the valuations of the resource providers were randomly generated and also the sample space was large. The two different settings for our experimentation are:

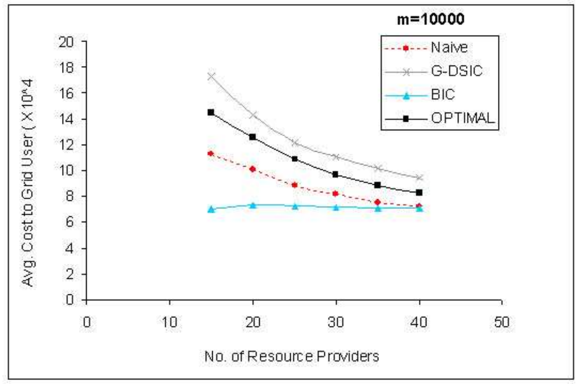

Fig. 1. Comparison of G-OPT, G-DSIC, and G-BIC for Scenario 1

- *Set 1:* In this set, the number of tasks in the main job is set as 10000. The cost valuations, c_i, are uniformly distributed in the interval $[5, 25]$. We choose a discrete distribution of the q_i values, $q_i = \{500, 1000, 1500, 2000\}$ with probabilities $f_i(q_i) = 0.25, 0.25, 0.25, 0.25$. This is a slightly uncommon setting in the grid context as there are 25% of the resource providers who are capable of doing 2000 tasks. Nevertheless, this setting helps us gain an insight into the working of the mechanisms under extreme conditions.
- *Set 2:* This is a more realistic setting, with m again being 10000. The cost valuations c_i are also the same as in the previous set. That is, they are distributed uniformly in the interval $[5, 25]$. In this set, the q_i values are randomly drawn from the set $q_i = 100, 250, 500, 1000, 2000$ with probabilities $f_i(q_i) = 0.2, 0.2, 0.3, 0.2, 0.1$. This scenario captures the setting of a general computational grid quite realistically. In a real life grid scenario, there will be some resource providers who are of medium capacity, a few resource providers who are of very high capacity and a few resource providers who are of lower capacities.

Figures 1 and 2 provide a comparison of performance of the proposed mechanisms. It can be noted that the average cost of procurement decreases with increase in the number of resource providers. This can be explained as follows. The general reduction is procurement cost due to increased number of resource providers is simply a result of availability of higher number of low cost resource providers. It is to be noted that the cost values are uniformly distributed. Hence, the higher the number of resource providers, increasingly larger number of resource providers will have lower cost valuations. Also note that the optimal mechanism completely

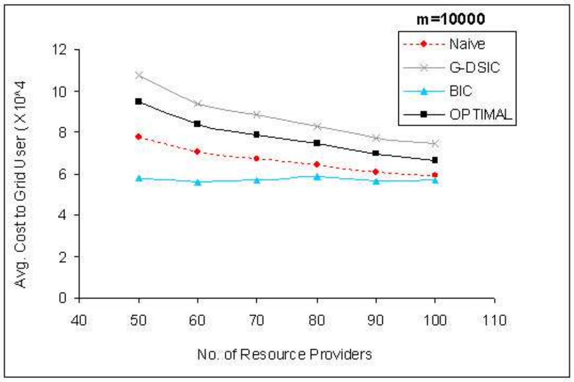

Fig. 2. Comparison of G-OPT, G-DSIC, and G-BIC for Scenario 2

outperforms the *G-DSIC* mechanism in both the settings. An interesting observation is that if participation of all the resource providers is guaranteed explicitly, the procurement costs due to the *G-BIC* mechanism is lower than those of the *G-OPT* mechanism. This is because, although both the mechanisms offer Bayesian incentive compatibility, the costs in *G-OPT* mechanism are higher in order to achieve the implicit guarantee of individual rationality.

VII. Summary and Future Work

We set out to develop resource procurement mechanisms with desirable incentive properties for the procurement of resources for running a parameter sweep job on a computational grid. In Section 3, we presented the first incentive compatible mechanism *G-DSIC*, which was a dominant strategy incentive compatible mechanism. In Section 4, we looked at a weaker form of incentive compatibility and and designed a Bayesian Nash incentive compatible mechanism, *G-BIC*. In Section 5, we identified the problems of both *G-DSIC* and *G-BIC* and came up with the design of a cost minimizing grid resource procurement mechanism, *G-OPT*. In Section 6, we compared the performance of the three proposed mechanisms using simulation experiments.

Some interesting future research directions are:

- Finding strong bounds for mechanisms in terms of the cost of procurement
- Relaxing the assumptions on which the optimal scheme is based and developing a mechanism that will carry forward the same nice properties of the *G-OPT* mechanism.
- Considering other job types apart from parameter sweep and developing robust mechanisms for resource procurement for these jobs
- Generalizing the results obtained to grids more general than computational grids.
- Considering the use of combinatorial auctions and combinatorial exchanges as a model for studying the problem of grid resource procurement.

References

[1] L. Ausubel and P. Crampton. Demand reduction and inefficiency in multi-unit auctions. Technical Report Mimeo, University of Maryland, College Park, MD, 1995.

[2] F. Branco. Multiple unit auctions of an indivisible good. *Economic Theory*, 8:77–101, 1996.

[3] R. Buyya, D. Abramson, J. Giddy, and H. Stockinger. Economic models for resource management and scheduling in Grid computing. *Concurrency and Computation: Practice and Experience*, 14(13):1507–1542, November-December 2002.

[4] I. Foster, C. Kesselman, and S. Tuecke. The Anatomy of the Grid: Enabling Scalable Virtual Organizations. *International Journal of Supercomputer Applications*, 15(3):200–222, Fall 2001.

[5] D. Garg and Y. Narahari. Foundations of Mechanism Design. Technical report, Department of Computer Science and Automation, Indian Institue of Science, Bangalore, India, October 2006.

[6] D. Grosu, A.T. Chronopoulos, and M. Leung. Load balancing in distributed systems: An approach using cooperative games. In *Proceedings of the International Parallel and Distributed Processing Symposium*, 2002.

[7] D. Grosu and A. Das. Auction-based Resource Allocation Protocols in Grids. In *Proceedings of the 16th IASTED International Conference on Parallel and Distributed Computing and Systems*, pages 20–27, November 2004.

[8] A. Kumar and G. Iyengar. Optimal procurement auctions for divisible goods with capacitated suppliers. Technical Report TR-2006-01, Columbia University, 2006.

[9] Y. Kwok, S. Song, and K. Hwang. Selfish grid computing: Game theoretic modeling and nas performance results. In *IEEE/ACM CCGRID 2005*, 2003.

[10] A. Malakhov and R.V. Vohra. An optimal auction for capacitated bidders - a network perspective. Technical report, http://www.kellogg.northwestern.edu/faculty/vohra/htm/res.htm, 2005.

[11] A. Malakhov and R.V. Vohra. Single and multi-dimensional optimal auctions - a network approach. Technical report, http://www.kellogg.northwestern.edu/faculty/vohra/htm/res.htm, 2005.

[12] A.M. Manelli and D.R. Vincent. Multidimensional mechanism design: Revenue maximization and the multiple-good monolpoly. August 2004.

[13] A.S. Mas-Colell, M.D. Whinston, and J.R. Green. *Microeconomic Theory*. Oxford University Press, New York, 1995.

[14] R. B. Myerson. Optimal auction design. *Math. Operations Res.*, 6(1):58–73, Feb. 1981.

[15] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, Cambridge, Massachusetts, 1997.

[16] J. Nabrzyski, J. M. Schopf, and J. Weglarz. *Grid Resource Management: State of the Art and Future Trends*. Kluwer Academic Publishers, 2003.

[17] N. Nisan, S. London, O. Regev, and N. Camiel. Globally Distributed Computation over the Internet - The POPCORN Project. In *International Conference on Distributed Computing Systems*, pages 592–601, May, 1998.

[18] Hastagiri Prakash. A Mechanism design Approach to Resource Procurement in Computational Grids with Rational Resource Providers. Technical report, Master of Science Dissertation,, Department of Computer Science and Automation, Indian Institue of Science, Bangalore, India, October 2006.

[19] Wolski R., Plank J.S., Brevik J., and Bryan T. G-commerce: Market formulations controlling resource allocation on the computational grid. In *AAMAS 2002*, 2002.

[20] R. Wilson. Auctions of shares. *Quaterly Journal of Economics*, 94:675–689, 1979.