

Economic Models for Resource Management and Scheduling in Grid Computing

Rajkumar Buyya, David Abramson, Jonathan Giddy, and Heinz Stockinger[‡]

CRC for Enterprise Distributed Systems Technology
School of Computer Science and Software Engineering
Monash University, Melbourne, Australia

[‡] CMS Experiment, Computing Group
CERN, European Organization for Nuclear Research
CH-1211 Geneva 23, Switzerland

Email: {rajkumar, jon, davida}@csse.monash.edu.au, Heinz.Stockinger@cern.ch

Abstract: The accelerated development in Peer-to-Peer (P2P) and Grid computing has positioned them as promising next generation computing platforms. They enable the creation of *Virtual Enterprises* (VE) for sharing resources distributed across the world. However, resource management, application development and usage models in these environments is a complex undertaking. This is due to the geographic distribution of resources that are owned by different organizations or peers. The resource owners of each of these resources have different usage or access policies and cost models, and varying loads and availability. In order to address complex resource management issues, we have proposed a computational economy framework for resource allocation and for regulating supply and demand in Grid computing environments. This framework provides mechanisms for optimizing resource provider and consumer objective functions through trading and brokering services. In a real world market, there exist various economic models for setting the price of services based on *supply-and-demand* and their value to the user. They include commodity market, posted price, tender and auction models. In this paper, we discuss the use of these models for interaction between Grid components to decide resource service value, and the necessary infrastructure to realize each model. In addition to usual services offered by Grid computing systems, we need an infrastructure to support interaction protocols, allocation mechanisms, currency, secure banking, and enforcement services. We briefly discuss existing technologies that provide some of these services and show their usage in developing the Nimrod-G grid resource broker. Furthermore, we demonstrate the effectiveness of some of the economic models in resource trading and scheduling using the Nimrod/G resource broker with deadline and cost constrained scheduling for two different optimization strategies on the World Wide Grid (WWG) testbed that has resources distributed across five continents.

1. Introduction

Computational Grids and Peer-to-Peer (P2P) computing systems are emerging as a new paradigm for solving large-scale problems in science, engineering, and commerce [1][5]. They enable the creation of *Virtual Enterprises* (VE) for *sharing* and *aggregation* of millions of resources (e.g., SETI@Home [23]) geographically distributed across organizations and administrative domains. They comprise heterogeneous resources (PCs, work-stations, clusters, and supercomputers), fabric management systems (single system image OS, queuing systems, etc.) and policies, and applications (scientific, engineering, and commercial) with varied requirements (CPU, I/O, memory, and/or network intensive). The *producers* (resource owners) and *consumers* (resource users) have different goals, objectives, strategies, and supply-and-demand patterns. More importantly both resources and end-users are geographically distributed with different time zones. In managing such complex environments, traditional approaches to resource management that attempt to optimize system-wide measure of performance cannot be employed. Traditional approaches use centralized policies that need complete state information and a common fabric management policy, or decentralized consensus based policy. Due to the complexity in constructing successful Grid environments, it is impossible to define an acceptable system-wide performance matrix and common fabric management policy [13]. (The concepts discussed in this paper apply to both P2P and Grid systems although we can argue about some of their technical, social, and political differences. However, we use the term Grid for simplicity and brevity).

In [2][3][4][5], we proposed and explored the usage of an economics based paradigm for managing resource allocation in Grid computing environments. The economic approach provided a fair basis in successfully managing decentralization and heterogeneity that is present in human economies. Competitive economic models provide algorithms/policies and tools for resource sharing or allocation in Grid systems. The models can be based on bartering or

prices. In the *bartering-based model*, all participants need to own resources and trade resources by exchanges (e.g., storage space for CPU time). In the *price-based model*, the resources have a price, based on the demand, supply, value, and the wealth in the economic system.

We envision a future in which economically intelligent and economically motivated peer-to-peer and Grid-like software systems will play an important role in distributed service-oriented computing. The resource management systems need to provide mechanisms and tools that facilitate the realization of the goals of both resource owners and users. The resource consumers' need a *utility model*—to allow them to specify resource requirements and preference parameters—and *brokers* to that provide strategies for choosing appropriate resources that meet user requirements. The resource owners need mechanisms for *price generation schemes* to increase system utilization and *protocols* that help them offer competitive services. For the market to be competitive and healthy, coordination mechanisms are required that help the market reach an equilibrium price—the price at which the supply of a service equals the quantity demanded.

Most of the related work in Grid computing dedicated to resource management and scheduling problems adopt a *conventional style* where a scheduling component decides which jobs are to be executed at which site based on certain cost functions (Legion [7], Condor [28], AppLeS [26], Netsolve [27], Punch [25]). Such cost functions are often driven by system-centric parameters that enhance system throughput and utilization rather than improving the utility of application processing. They treat resources as if they all cost the same price and the results of all applications have the same value even though this may not be the case in reality. The end user does not want to pay the highest price but wants to negotiate a particular price based on the demand, value, priority, and available budget. Also, the results of different applications have different value at different times. In an economics approach, the scheduling decision is not done statically by a single scheduling entity but directed by the end users requirements. Whereas a conventional cost model often deals with software and hardware costs for running applications, the economic model primarily charges the end user for services that they consume based on the value they derive from it. Pricing based on the demand of users and the supply of resources is the main driver in the competitive, economic market model. Thus, we stress that a user is in competition with other users and a resource owner with other resource owners.

The main contribution of this paper is to provide economic models, system architecture, and policies for resource management in Grid environments. Currently, the user community and the technology are still rather new and not well accepted and established in commercial settings. However, we believe the Grid can become established in such settings by providing incentive to both consumers and resource owners for being part of the Grid. Since the Grid uses the Internet as a carrier for providing remote services, it is well positioned to create a computational ecology, cooperative problem solving environment, and means for sharing computational and data resources in a seamless manner. Up to now, the idea of using Grids for solving large computationally intensive applications has been more or less restricted to the scientific community. However, even if, in the scientific community, the pricing aspect seems to be of minor importance, funding agencies need to support the hardware and software infrastructure for Grids. Economic models help them manage and evaluate resource allocations to user communities. The system managers may impose quota limitations and value different resources with a different number of tokens [33]. In such environments, resource consumers certainly prefer to use economic driven schedulers to effectively utilize their tokens by using lightly loaded cheaper resources.

2. Players in the Grid Marketplace

The two key players driving the Grid marketplace are: *Grid Service Providers* (GSPs) providing the traditional role of *producers* and *Grid Resource Brokers* (GRBs) representing *consumers*. The Grid computing environments provide the necessary infrastructure including security, information, transparent access to remote resources, and information services that enable us to bring these two entities together [5]. Consumers interact with their own brokers for managing and scheduling their computations on the Grid. The GSPs make their resources Grid enabled by running software systems (such as Globus [6] or Legion [7]) along with *Grid Trading Services* (GTS) to enable resource trading and execution of consumer requests directed through GRBs. The interaction between GRBs and GSPs during resource trading (service cost establishment) is mediated through a *Grid Market Directory* (GMD) (see Figures 1 to 5). They use various economic models or interaction protocols for deciding service access price. These models are discussed in Section 3. Our Grid Architecture for Computational Economy (GRACE), proposed in our earlier work [5], provides detailed information on the role played by GSPs, GRBs, and GMDs and compares them to actual Grid components and implementations.

As in the conventional marketplace, the users' community (GRBs) represents the demand, whereas the resource owners' community (GSPs) represents the supply. We emphasize the user community and how they can influence the

pricing of Grid resources via their brokers. Numerous economic models including microeconomic and macroeconomic principles have been proposed in the literature. Some of the commonly used economic models that can be employed for managing resources environment, include:

- Commodity Market Model
- Posted Price Model
- Bargaining Model
- Tendering/Contract-Net Model
- Auction Model
- Bid-based Proportional Resource Sharing Model
- Community/Coalition/Bartering Model
- Monopoly and Oligopoly

Several research systems (see Table 1) have attempted to apply the concept of computational economy models to database management, CPU cycles, storage, and distributed computing. They include Mariposa [8], Mungi [17], Popcorn [21], Java Market [18], Enhanced MOSIX [19], JaWS [30], Xenoservers [31], D’Agents [32], Rexec/Anemone [22], Spawn [20], Mojo Nation [24], and Nimrod-G [1][2]. These systems have been targeted to manage single or multiple resources and they can be categorized as follows:

- *Single domain computing systems:* Enhanced MOSIX and Rexec/Anemone.
- *Agent-based systems:* Xenoservers and D’Agents.
- *Distributed database management system:* Mariposa
- *Shared storage management system:* Mungi.
- *Web-based distributed systems:* Popcorn, Java Market, and JaWS.
- *Multi-domain distributed Grid system:* Nimrod-G

Many of the resource management systems presented in Table 1 support a single model for resource trading, provide their own programming model, and are implemented as monolithic systems. To overcome these limitations, the modern grid computing systems use a layered architecture. Typically, in a Grid marketplace, the resource owners, and users can use any one or more of these models or even combinations of them in meeting their objectives [5]. Both have their own expectations and strategies for being part of the Grid. The resource consumers adopt the strategy of solving their problems at low cost within a required timeframe. The resource providers adopt the strategy of obtaining best possible return on their investment while trying to maximize their resource utilization by offering a competitive service access cost in order to attract consumers. The resource consumers can choose providers that best meet their requirements. The design and architecture for the development of Grid systems using these economic models has been discussed in the next section.

SYSTEM NAME	ECONOMY MODEL	PLATFORM	REMARKS
Mariposa [8] (UC Berkeley)	Bidding (Tendering/ ContractNet). Pricing based on load and historical info.	Distributed database.	It supports budget-based query processing and storage management.
Mungi [17] (University of New South Wales)	Commodity market (renting storage space that increases as available storage runs low, forcing users to release unneeded storage.)	Storage servers.	It supports storage objects based on bank accounts from which rent is collected for the storage occupied by objects. .
Nimrod-G [1][2] (Monash University)	It supports economy models such as commodity market, spot market, and contract-net for price establishment.	A Grid of distributed computers (PCs, WS, and clusters)	It supports deadline and budget constrained scheduling algorithms for executing task-farming applications on large scale distributed systems depending on their cost, power, and availability and users quality of service requirements.

Popcorn [21] (Hebrew University)	Auction. (Highest bidder gets access to resource and it transfers credits from buyer to the seller account.)	Web browsers. (<i>Popcorn parallel code</i> runs within a browser of CPU cycles seller.)	Popcorn API-based parallel applications need to specify a budget for processing each of its modules.
Java Market [18] (John Hopkins University)	QoS based computational market. (The resource owner receives $f(j, t)$ award for completing f in time t .)	Web browsers. (JavaMarket runs <i>standard Java Applets</i> within a browser).	One can sell CPU cycles by pointing Java-enabled browser to Portal & allow execution of Applets.
Enhanced MOSIX [19] (Hebrew Uni., Israel)	Commodity market (resource cost of each node is known)	Clusters of computers (Linux PCs)	It supports process migration such that overall cost of job execution is kept low.
JaWS [30] (University of Crete)	Bidding (Tendering)	Web browsers	It is similar to Popcorn.
Xenoservers [31] (University of Cambridge)	Bidding (Proportional resource sharing)	Single computer	Accounted execution of untrusted code.
D'Agents [32] (Dartmouth College)	Bidding (Proportional resource sharing)	Single computer or Mobile Agents	Agents bid function is proportional to benefit.
Rexec/Anemone [22] (UC Berkeley)	Bidding/Auction (for proportional resource sharing)	Clusters (A market-based Cluster Batch Queue System)	Users assign utility value to their application and system allocates resources proportionally.
Mojo Nation [24] (Autonomous Zone Industries, CA)	A Credit-based partnership and/or bartering model. (Contributors earn credits by sharing storage and spend them when required)	Network storage.	It is a content-sharing community network. It combines marketplace and bartering approach for file/resource sharing.
Spawn [20] (Xerox PARC)	Second-price Auction (uses sponsorship model for funding money to each task depending on some requirements)	Network on workstations. Each WS executes a single task per time slice	It supports execution of concurrent program expressed in the form of hierarchy of processes that expand and shrink size depending on the resource cost.
Supercomputing center [33] (Computing Services for Academic Research at the University of Manchester)	Commodity market and priority-based model (they charge for CPU, memory, storage, and human support services)	MPPs, Crays, and Clusters, and Storage servers.	Any application can use this service and QoS is proportional to user priority and scheduling mechanisms.

Table 1: Computational economy based resource management systems.

Both GRBs and GSPs can initiate resource trading and participate in the interaction depending on their requirements and objectives. GRBs may invite bids from a number of GSPs and select those that offer the lowest service costs and meet their deadline and budget requirements. Alternatively, GSPs may invite bids in an auction and offer services to the highest bidder as long as its objectives are met. Both GSPs and GRBs have their own utility functions that must be satisfied and maximized. The GRBs perform a cost-benefit analysis depending on the deadline (by which the results are required) and budget available (the amount of money the user is willing to invest for solving the problem). The resource owners decide their pricing based on various factors [5]. They may charge different prices for different users for the same service or it can vary depending on the specific user demands. Resources may have different prices based on environmental influences such as the availability of larger core memory and better communication bandwidth with an outside world.

Grid brokers (note that in a Grid environment each user has his own broker as his agent) may have different goals

(e.g., different deadlines and budgets), and each broker tries to maximize its own good without concern for the global good. However, given adequate information, independent rational self-interest achieves the optimal resource allocation, both for the individual and for the society. This needs to be taken into consideration in building automated negotiation infrastructure. In a *cooperative distributed computing or problem-solving environment* (like cluster computers), the system designers impose an *interaction protocol* (possible actions to take at different points) and a *strategy* (a mapping from one state to another and a way to use the protocol). This model aims for global efficiency as nodes cooperate towards a common goal. On the other hand, in Grid systems, brokers and GSPs are provided with an interaction protocol, but they choose their own private strategy (like in multi-agent systems), which cannot be imposed from outside. Therefore, the negotiation protocols need to be designed assuming a *non-cooperative, strategic* perspective. In this case, the main concern is what social outcomes follow given a protocol, which *guarantees that each broker/GSP's desired local strategy is best for that broker/GSP and hence the broker/GSP will use it.*

Various criteria used for judging effectiveness of a market model are [15]:

- Social welfare (global good of all)
- Pareto efficiency (global perspective)
- Individual rationality (better off by participating in negotiation)
- Stability (mechanisms that cannot be manipulated, i.e., behave in the desired manner)
- Computational efficiency (protocols should not consume too much computation time)
- Distribution and communication efficiency (communication overhead to capture a desirable global solution).

3. Economic Models in a Grid Context

In the previous section we identified a few popular models that are used in human economies. In this section we discuss the use of those economic models and propose architecture for realizing them. The discussion on realizing negotiation protocols based on different economic models is kept as generic as possible. This ensures that our proposed architecture is free from any specific implementation and provides a general framework for any other Grid middleware and tools developers. Particular emphasis will be placed on framework and heuristics that Grid resource brokers (*G-Brokers*) can employ for establishing service price depending on their customers' requirements. The service providers publish their services through the Grid market directory. They use Grid trading services' declarative language for defining cost specification and their objectives such as access price for various users for different times and durations along with possibilities of offering discounts to attract users during off-peak hours. The Grid trading server (GTS) can employ different economic models in providing services. The simplest would be a commodity model wherein the resource owners define pricing strategies including those driven by the demand and resource availability. The GTS can act as auctioneer if Auction-based model is used in deciding the service access price or an external auctioneer service can be used (see Figure 6).

For each of the economic models, firstly, the economic model theory, its parameters and influences are discussed and then a possible solution is given for a current Grid environment and how they can be mapped to existing Grid tools and architectures or what needs to be extended. In the classical economic theory there are different models for specific environmental situations and computing applications. Since the end-user interaction is the main interest of this paper, we point out possible interactions with the broker.

3.1. Commodity Market (Flat or Supply-and-Demand Driven Pricing) Model

In the commodity market model, resource owners specify their service price and charge users according to the amount of resource they consume. The pricing policy can be derived from various parameters and can be *flat* or *variable* depending on the resource *supply and demand*. In general, services are priced in such a way that supply and demand equilibrium is maintained. In the *flat price model*, once pricing is fixed for a certain period, it remains the same irrespective of service quality. It is not significantly influenced by the demand, whereas in a *supply and demand* model prices change very often based on supply and demand changes. In principle, when the demand increases or supply decreases, prices are increased until there exists equilibrium between supply and demand. Pricing schemes in a Commodity Market Model can be based on:

- Flat fee
- Usage Duration (Time)
- Subscription

- Demand and Supply-based [11]

The resource owners publish their prices through yellow pages like Grid market directory (GMD) service (see Figure 1). This is accomplished by defining the price specification that GTS can use for publishing service access price in the market directory. A simple price specification may contain the following parameters.

```

. . .
consumer_id           // this can be same Grid-ID
peak_time_price       // 9am-6pm: office hours on working days
lunch_time_price      // (12.30-2pm)
offpeak_time_price    // (6pm-9am),
discount_when_lightly_loaded // if load is less than 50% at any time
raise_price_high_demand // % raise price if average load is above 50%
price_holiday_time    // during holidays and week ends!
. . .

```

Traditionally, computational services are priced based on their production cost and desired profit margin. However, the consumers' perception of value is based on parameters such as supply and demand for resources, priority and service quality requirements. Therefore, the resource value in Grid economy needs to be defined as a function of many parameters as follows:

- Resource Value = Function (Resource strength, Cost of physical resources, Service overhead, Demand, Value perceived by the user, Preferences);

The last three parameters are difficult to capture from consumers unless they see any benefit in disclosing them and vary from time to time, from one application to another. However, there are consumers who prefer regular access to resources during a particular period of the day. For example, those involved in making regular decisions on supply chain management of goods shipping from inventory to the departmental stores prefer calendar-based guaranteed access, and stable but competitive pricing to resources unlike spot-market based access to services [35]. In this case demand and preferences are clear, pricing policy can be easily negotiated in advance in a competitive and reasonable manner and resource quality of services can be guaranteed through reservation during the required period as agreed in advance.

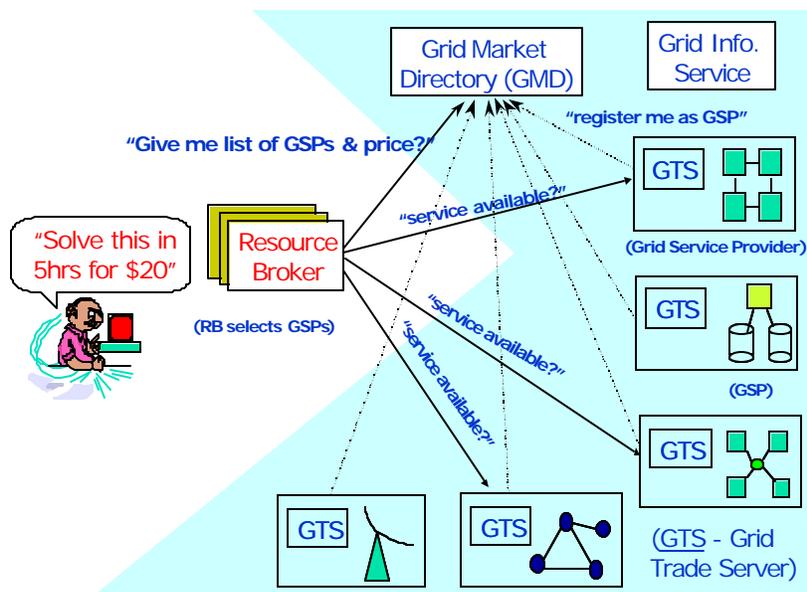


Figure 1: Interaction between GSPs and users in a commodity market Grid for resource trading.

Consumers can be charged for access to various resources including CPU cycles, storage, software, and network. The users compose their application using higher-level Grid programming languages. For example, in our Nimrod problem-solving environment we provide a declarative programming language for composing parameter sweep applications and defining application and user requirements such as deadline and budget [1]. The resource broker (working for the user)

can carry out the following steps for executing applications:

- The broker identifies service providers.
- It identifies suitable resources and establishes their prices (by interacting with GMD and GTS).
- It selects resources that meet its utility function and objectives (lower cost and meet deadline requirements). It uses heuristics and/or historical knowledge base while choosing resources and mapping jobs to them.
- It uses resource services for job processing and issues payments as agreed.

As we are focusing on generic framework, implementation specific details for releasing the above steps are not presented. For example, implementation specific details of our Nimrod/G resource broker [1] [2][4] vary from other related systems.

3.2. Posted Price Model

The posted price model is similar to the commodity market model, except that it advertises special offers (see Figure 2) in order to attract (new) consumers to establish market share or motivate users to consider using cheaper slots. In this case, brokers need not negotiate directly with GSPs for price, but use posted prices as they are generally cheaper compared to regular prices. The posted-price offers will have usage conditions, but they might be attractive for some users. For example, during holiday periods, demand for resources is likely to be limited and GSPs can post tempting offers or prices aiming to attract users to increase resource utilization. The activities that are specifically related to the posted-price model in addition to those related to commodity market model are:

- Grid Service Providers (GSPs) post their special offers and associated conditions etc. in Grid Market Directory.
- Broker looks at GMD to identify if any of these posted services are available and fits its requirements.
- Broker enquires (GSP) for availability of posted services.
- Other steps are similar to those pointed out in commodity market model.

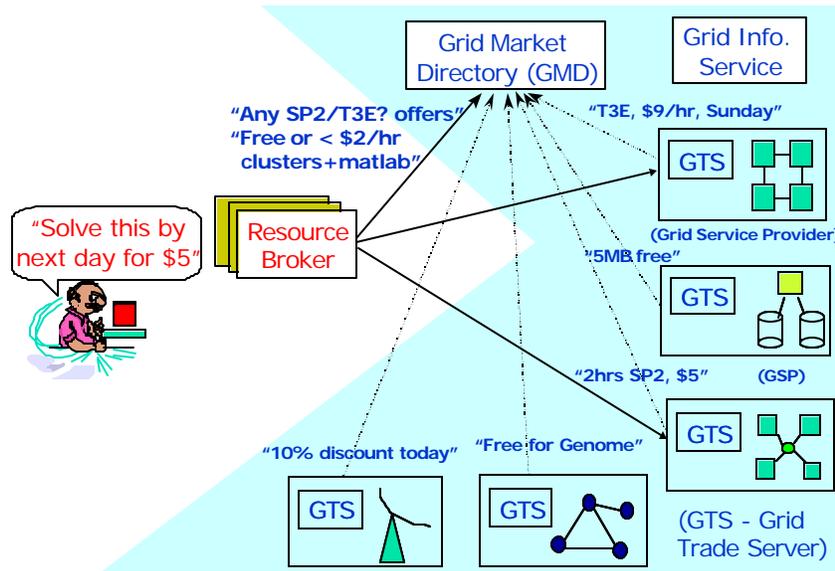


Figure 2: Posted price model and resource trading in a computational market environment.

3.3. Bargaining Model

In the previous models, the brokers pay access prices, which are fixed by GSPs. In the bargaining model, resource brokers bargain with GSPs for lower access price and higher usage duration. Both brokers and GSPs have their own objective functions and they negotiate with each other as long as their objectives are met. The brokers might start with a very low price and GSPs with a higher price. They both negotiate until they reach a mutually agreeable price (see Figure 3) or one of them is not willing to negotiate any further. This negotiation is guided by user requirements (e.g., deadline is too relaxed) and brokers can take risk and negotiate for cheaper prices as much as possible and can discard expensive

machines. This might lead to lower utilization of resources, so GSPs might be willing to reduce the price instead of wasting resource cycles. Brokers and GSPs generally employ this model when market *supply-and-demand* and service prices are not clearly established. The users can negotiate for a lower price with promise of some kind favour or using GSPs services even in the future.

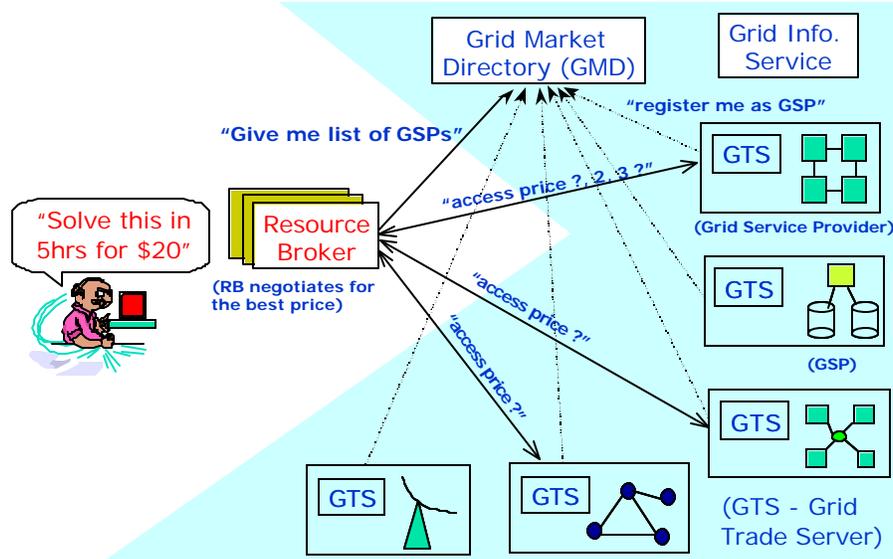


Figure 3: Brokers bargaining for lower access price in their bid for minimizing computational cost.

3.4. Tender/Contract-Net Model

Tender/Contract-Net model is one of the most widely used models for service negotiation in a distributed problem-solving environment [14]. It is modeled on the contracting mechanism used by businesses to govern the exchange of goods and services. It helps in finding an appropriate service provider to work on a given task. Figure 4 illustrates the interaction between brokers and GSPs in their bid to meet their objectives. A user/resource broker asking for a task to be solved is called the *manager* and resource that might be able to solve the task is called potential *contractor*. From a manager's perspective, the process is:

1. Consumer (Broker) announces its requirements (using deal template) and invites bids from GSPs.
2. Interested GSPs evaluate the announcement and respond by submitting their bids.
3. Broker evaluates and awards the contract to the most appropriate GSP(s).
4. The broker and GSP communicate privately and use the resource (R).

The contents of the deal template used for work announcement include, addressee (user), eligibility requirements specifications (for instance, Linux, x86arch, and 128MB memory), task/service abstraction, optional price that the user is willing to invest, bid specification (what should offer contain), expiration time (deadline for receiving bids).

From a contractor's/GSP perspective, the process is:

1. Receive tender announcements/advertisements (say in GMD).
2. Evaluate service capability.
3. Respond with bid.
4. Deliver service if bid is accepted.
5. Report results and bill the broker/user as per the usage and agreed bid.

The advantage of this model is that if the selected GSP is unable to deliver a satisfactory service, the brokers can seek services of other GSPs. This protocol has certain disadvantages. A task might be awarded to a less capable GSP if a more capable GSP is busy at award time. Another limitation is that the GRB manager has no obligation to inform potential contractors that an award has already been made. Sometimes, a manager may not receive bids for several reasons: (a) all potential GSPs are busy with other tasks, (b) a potential GSP is idle but ranks the proposed tender/task

below the other tasks under consideration, (c) no GSPs, even if idle, are capable of offering service (e.g., resource is Windows NT-based, but user wants Linux). To handle such cases, a GRB can request quick response bids to which GSPs respond with messages such as *eligible*, *busy*, *ineligible* or *not interested*. This helps the GRB in making changes to its work plan. For example, the user can change deadline or budget to wait for new GSPs or attract existing GSPs to submit bids.

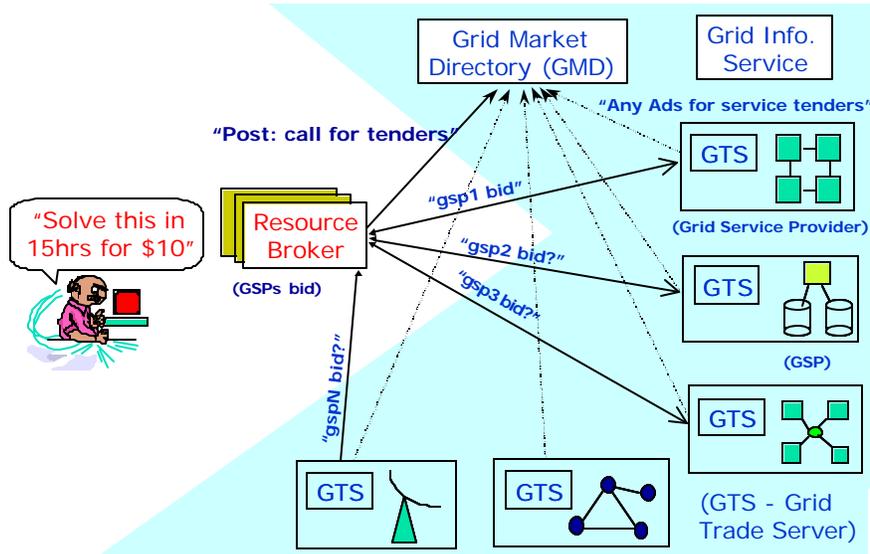


Figure 4: Tender/ContractNet model for resource trading.

The tender model allows directed contracts to be issued without negotiation. The selected GSP responds with an *acceptance* or *refusal* of award. This capability can simplify the protocol and improve the efficiency of certain services.

3.5. Auction Model

The auction model supports one-to-many negotiation, between a service provider (seller) and many consumers (buyers), and reduces negotiation to a single value (i.e., price). The auctioneer sets the rules of auction, acceptable for the consumers and the providers. Auctions basically use market forces to negotiate a clearing price for the service.

In the real world, auctions are used extensively, particularly for selling goods/items within a set duration. The three key players involved in auctions are: resource owners, auctioneers (mediators), and buyers (see Figure 5). Many e-commerce portals such as Amazon.com and eBay.com are serving as mediators (auctioneers). Both buyers' and sellers' roles can also be automated. In a Grid environment, providers can use an auction protocol for deciding service value/price (see Figure 6). The steps involved in the auction process are:

- GSPs announce their services and invite bids.
- Brokers offer their bids (and they can see what other consumers offer if they like - depending on open/closed).
- Step (b) goes on until no one is willing to bid higher price or auctioneer stops if the minimum price line is not met.
- GSP offers service to the one who wins.
- Consumer uses the resource.

Auctions can be conducted as open or closed depending on whether they allow back-and-forth offers and counter offers. The consumer may update the bid and the provider may update the offered sale price. Depending on these parameters, auctions can be classified into four types:

- English Auction (first-price open cry)
- First-price sealed-bid auction
- Vickrey (Second-price sealed-bid) auction [9]
- Dutch Auction
- Double Auction (Continuous)

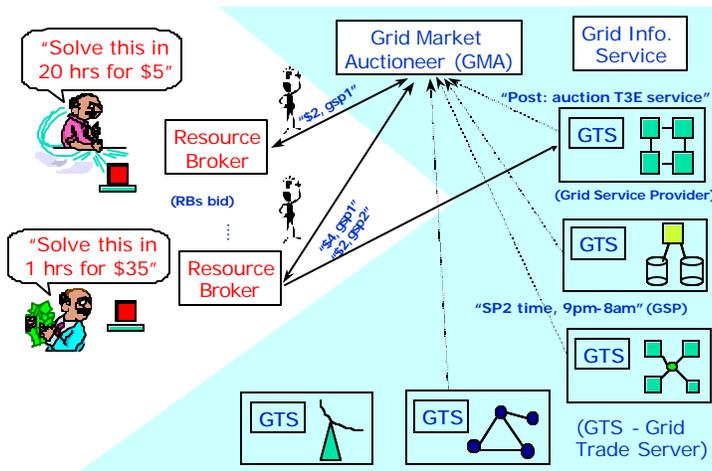


Figure 5: Auctions using external auctioneer.

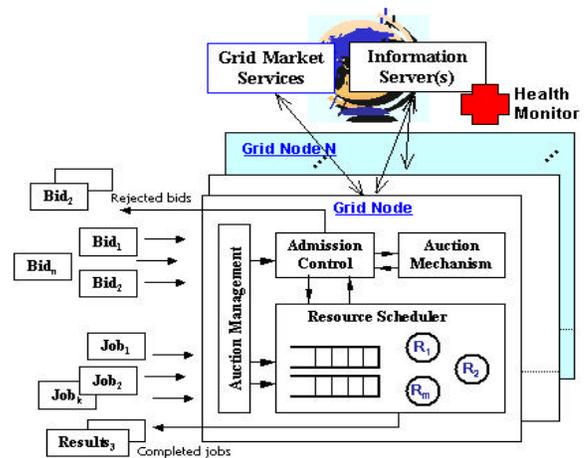


Figure 6: Auction using their own Auctioneer.

English Auction (first-price open cry) — all bidders are free to increase their bids exceeding other offers. When none of the bidders are willing to raise the price anymore, the auction ends, and the highest bidder wins the item at the price of his bid. In this model, the key issue is how GRBs decide how much to bid. A GRB has a private value (as defined by the user) and can have a strategy for a series of bids as a function of its private value and prior estimation of other bidder's valuations, and the past bids of others. The GRB decides the private value depending on the user-defined requirements (mainly deadline and budget that he is willing to invest for solving the problem). In the case of private value English auctions, a GRB's dominant strategy is to always bid a small amount "higher" than the current highest bid, and stop when its private value price is reached. In correlated value auctions, the policies are different and allow the auctioneer to increase the price a constant rate or at the rate he wishes. Those not interested in bidding anymore can openly declare so (open-exit) without re-entry possibility. This information helps other bidders and gives a chance to adjust their valuation.

First-price sealed-bid auction — each bidder submits one bid without knowing the others' bids. The highest bidder wins the item at the price of his bid. In this case a broker bid strategy is a function of the private value and the prior beliefs of other bidders' valuations. The best strategy is bid less than its true valuation and it might still win the bid, but it all depends on what the others bid.

Vickrey (Second-price sealed-bid) auction — each bidder submits one bid without knowing the others' bids. The highest bidder wins the item at the price of the second highest bidder [9]. The implementation architecture and strategies is similar to the ContractNet/Tender model discussed earlier.

Dutch Auction — the auctioneer starts with a high bid/price and continuously lowers the price until one of the bidders takes the item at the current price. It is similar to first-price sealed-bid auction because in both cases bid matters only if it is the highest, and no relevant information is revealed during the auction process. From the broker's bidding strategic point of view, Dutch auction is similar to English (first-price sealed-bid auction). The key different between them is that in English auction bids start with low opening and increase progressively until demand falls whereas, in Dutch auction bids start with low opening and decrease progressively until demand rises to match supply.

The interaction protocols for Dutch auction are as follows: the auction attempts to find market price for a good by starting a price much higher than the expected market value, then progressively reducing the price until one of the buyers accepts the price. The rate of reduction price is up to the auctioneer and they have a reserve price below which not to go. If the auction reduces the price to reserve price with no buyers, the auction terminates. In terms of real time, Dutch auction is much efficient as the auctioneer can decrease the price at a strategic rate and first higher bidder wins. In an Internet wide auction, it is appealing in terms of automating the process wherein all parties can define their strategies for agents that can participate in multiple auctions to optimize their objective functions.

Double Auction — This is one of the most common exchange institutions in the marketplace whose roots go back to ancient Egypt and Mesopotamia [36]. In fact, it is the primary economic model for trading of equities, commodities, and derivatives in stock markets (e.g., NASDAQ). In the double auction model, buy orders (*bids*) and sell orders (*asks*) may be submitted at anytime during the trading period. If at any time there are open bids and asks that match or are compatible in terms of price and requirements (e.g., quantity of goods or shares), a trade is executed immediately. In this auction

orders are ranked highest to lowest to generate demand and supply profiles. From the profiles, the maximum quantity exchanged can be determined by matching *asks* (starting with lowest price and moving up) with demand *bids* (starting with highest price and moving down). Researchers have developed software-based agents mechanisms to automate double auction for stock trading with or without human interaction [37].

The double auction model has high potential for Grid computing. The brokers can easily be enabled to issue *bids* depending on budget, deadline, job complexity, scheduling strategy, and resource characteristics requirements and GSPs can issue *asks* depending on current load and perceived demand, and price constraints. Both orders can be submitted to GMD agents that provide continuous clearance or matching services. Since bids are cleared continuously, both GRBs and GSPs can make instant decision with less computational overhead and complexity.

All the above auctions differ in terms of whether they are performed as open or closed auctions and the offer price for the highest bidder. In open auctions, bidding agents can know bid value of other agents and will have an opportunity to offer competitive bids. In closed auctions, the participants' bids are not disclosed to others. Auctions can suffer from collusion (if bidders coordinate their bid prices so that the bids stay artificially low), deceptive auctioneers in case of Vickrey auction (auctioneer may overstate the second highest bid to the highest bidder unless that bidder can vary it), deceptive bidders, counter speculation, etc.

3.6. Bid-based Proportional Resource Sharing Model

Market-based proportional resource sharing systems are quite popular in cooperative problem-solving environments like clusters (in single administrative domain). In this model, the percentage of resource share allocated to the user application is proportional to the bid value in comparison to other users' bids. The users are allocated credits or tokens, which they can use for having access to resources. The value of each credit depends on the resource demand and the value that other users place on the resource at the time of usage. For example, consider two users wishing to access a resource with similar requirements, but first user is willing to spend 2 tokens and the second user is willing to spend 4 tokens. In this case, the first user gets 1/3 of resource share whereas the second user gets 2/3 of resource share, which is proposal to the value that both users place on the resource for executing their applications.

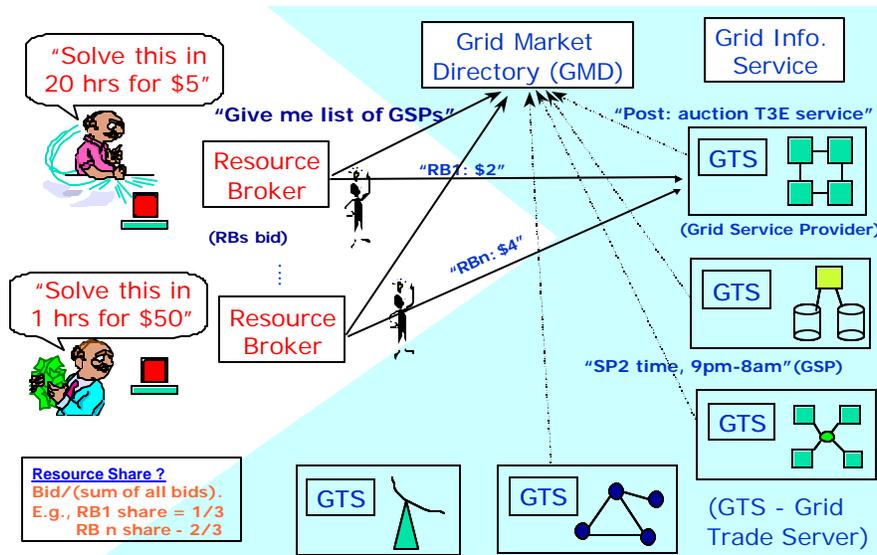


Figure 7: Market-based Proportional Resource Sharing.

This can be a good way of managing a large shared resource in an organization or resource owned by multiple individuals (like multiple departments in a university) can have credit allocation mechanism depending on investment they made. They can specify how much of credit they are willing to offer for running their applications on the resource. For example, a user might specify low credits for non-interactive batch jobs and high credits for interactive jobs with high response times. GSPs can employ this model for offering a QoS for higher price paying customers in a shared resource environment (as shown in Figure 7). Systems such as Rexec/Anemone and Xenoservers, D'Agents CPU market employ proportional resource sharing model in managing resource allocations [5].

3.7. Community/Coalition/Bartering/Share Holders Model

A community of individuals shares each other's resources to create a cooperative computing environment. Those who are contributing their resources to a common pool can get access to that pool. A sophisticated model can also be employed here for deciding how much resources share contributors can get. It can involve credits that one can earn by sharing resource, which can then be used when needed. A system like Mojonation.net employs this model for storage sharing. This model works when those participating in the Grid have to be both service providers and consumers.

3.8. Monopoly/Oligopoly

In the previously mentioned models we have assumed a competitive market where several GSPs and brokers/consumers determine the market price. However, there exist cases where a single GSP dominates the market and is the single provider of a particular service. In economic theory this model is known as a monopoly. Users cannot influence the prices of services and have to choose the service at the price given by the single GSP who monopolized the Grid marketplace. As regards the technical realization of this model, the single site puts the prices into the GMD or information services and brokers consult it without any possibility to negotiate prices.

The competitive markets are one extreme and monopolies are the other extreme. In most of the cases, the market situation is *oligopoly*, which is in between these two extreme cases: a small number of GSPs dominate the market and set the prices.

3.9. Other Influences on Market Prices

We now state more influences on price setting strategies in competitive, international markets. Supply and demand is the most common one but one also has to take into account national borders and different pricing policies within different countries such as taxation, consumer price index, inflation, etc. These factors are not dealt with in this paper, however implementations may need to consider them. There are micro and macro-economic factors that play an important role. One can also neglect them and build a price model on which all the Grid consumers have to agree. So this would correspond to an international market with special rules. Then, a model has to be formed for price changes. What is the factor for that change? Is there a monopoly that can decide what to do? Is the market transparent with optimally adapted prices? These are some of the main questions that need to be answered by GSPs when they decide their prices in an international market. A broker may consult the Grid Information Service to find out where the price for a particular service is minimal. For instance, one country might impose special taxes on a service whereas another country does not.

There are occasions where resources are not valued as per the actual cost of resources and overhead involved in offering services. When new players enter the market, in order to attract customers from the existing GSPs they are likely to offer access at minimal price by under valuing resources. This leads to price wars as GSPs are caught in a price cutting round to compete with each other. Measures such as intervention of *price regulation authorities* can be in place to prevent the market from collapsing or leaving it to the market to consolidate naturally.

4. Economy in a Data Grid Environment

In computational Grid environments, large computational tasks that do not use very large amounts of data are solved. In Data Grid environments [10], large amounts of data are distributed and replicated to several sites all around the globe. Here, efficient access to the data is more important than scheduling computational tasks. When accessing large data stores, the cost for accessing data is important [12]. Can a single user afford to access a third of all data in a Petabyte data store? Certain restrictions and cost functions need to be imposed in order to obtain a good throughput for data access and to provide a fair response time for multiple users. An economic model can help achieve local or global efficiency. Paying higher prices can result in accessing a larger amount of data. How these prices can be mapped to the requirements of scientific users is still an open research issue. However, it is clear that some optimizations and restrictions for data access are required.

4.1. A Case for Economy in a Scientific Data Grid Environment

In this subsection we discuss the possible use of economy in a scientific Data Grid environment, in particular in the DataGrid project [10]. We claim that scheduling based on economic models is different in a conventional "business environment" and in the scientific community. Whereas in the first, a Grid user pays explicitly a certain amount of

money in order to get a service (e.g. pay Euro 100 for running application x in 10 minutes somewhere in the Grid), in the scientific community an explicit payment does not seem to be useful.

As regards the DataGrid project, several possible applications can be identified that require scheduling. Here, we concentrate only on the scheduling of user requests to (replicated) data. Let us define the problem. We assume several concurrent users (in the order of 100), replicated data stores with several Terabytes, up to Petabytes of data each and a finite throughput of data servers at a single site. A site can have several data servers and a hierarchical disk pool with tapes but each single site will have a restriction on the maximum amount of data it can serve at a time. Thus, an optimization problem occurs that tries to optimize the throughput for a large user community. For instance, a single user should not be able to request a Terabyte of data per day and consequently use all the data server resources. A fair scheduling concept is required that allows multiple, concurrent user requests. The problem can be identified as a high throughput problem. A scheduler at a single site has to take care of a fair scheduling concept. This compares to the conventional market model where a single user can increase his/her own throughput by offering to pay higher prices for a service. In the Data Grid environment a single user is interested in analyzing data and does not want to pay money explicitly for the physics analysis job. In other words, data analysis should not be restricted by resource prices since the main focus of physics analysis is to find some new information in large amounts of data rather than making money by selling data resources. However, this needs to be regulated to provide fair and improved quality of service to all users.

Data access in a Data Grid environment needs to be measured, regulated, and it has to be translated into costs, but shall not be expressed in currency units. One possible mapping of access requests to costs is to define the maximum data throughput of a site in the Data Grid. Based on this, a maximum number of tokens can be distributed to the users. For instance, a distributed data server at a site is able to serve 10 TB of data per day and thus 10 000 000 tokens are available and distributed to possible users. By default, a single user may only access as much data as he has tokens. This gives other users a chance to access data. However, the amount of data that they access for a given token needs to be based on parameters such as demand, system load, QoS requested, etc. This helps in better managing resources and improving QoS offered by the system. The users can tradeoff between QoS and tokens.

The local site scheduler has to take care of admission control. Now, a single user might want to access more data than the available tokens. This requires re-distribution and negotiation of tokens. The negotiation step can then be done with the various economic models discussed earlier. Each user gets a certain amount of budget (expressed in tokens) and can be expressed in terms of pre-allocation priorities. For instance, users A and B are allocated X and Y number of tokens. Tokens can be renewed each day based on the current demand and the data server capabilities for serving data. The tokens can be converted into money when we compare it to a commercial environment (like 10 tokens for \$1). The Data Grid system can have different pricing policies for data access at different times to encourage the users to organize and plan themselves depending on their priorities. For example, charging users 10 tokens per MB of data access during the peak time and 6 tokens per MB of data access during the off-peak time can help users to manage and prioritize their work flow effectively. Accessing replicas from different sites may also result in having different prices for different replicas. Such policies lead to an economy in data grid environments. Consequently, tokens are assigned to users depending on the value/importance of users and their work. We conclude that in a scientific Data Grid environment economic models can be applied, but a scheduling instance is still required that controls the overall throughput for a large user community.

4.2. Data Economy

In [12] a cost model for distributed and replicated data over a wide area network is presented. Cost factors for the model are the network, data server and application specific costs. Furthermore, the problem of job execution is discussed under the viewpoint of sending the job to the required data (code mobility) or sending data to a local site and executing the job locally (data mobility). In the economic model, the main focus is on executing a job at any site as long as the cost for job execution is minimal. Based on the previous case study we summarize this topic by introducing the term *Data Economy* that applies to several fields where data is distributed over Grids. The domains include

- Content (like music, books, news papers) sales using techniques such as micro-payments, aggregation, subscription, and subsidy.
- Community or bartering model for content sharing.
- Data and replica sites access in Data Grid environments.

5. Nimrod-G: A Computational Economy driven Grid Resource Broker

In this section, we briefly discuss the Nimrod-G resource broker as an example of a Grid system that uses a computational economy driven architecture for managing resources and scheduling task farming application on a large-scale distributed Grid resources. The Nimrod-G toolkit and resource broker is developed by leveraging services provided by Grid middleware systems such as Globus, Legion, Condor, and the GRACE trading mechanisms. These middleware systems provide a set of low-level protocols for secure and uniform access to remote resources; and services for accessing resources information and storage management. The modular and layered architecture of Nimrod-G is shown in Figure 8. A detailed discussion on the Nimrod system architecture and implementation [1][2], scheduling algorithms [4], and its ability to execute real world applications such as ionization chamber calibration [1] and drug design [34] on the Grid can be found elsewhere.

The Nimrod-G toolkit provides a simple declarative parametric modeling language for expressing parametric experiments. It uses novel resource management and scheduling algorithms based on economic principles. Specifically, it supports user-defined deadline and budget constraints for schedule optimizations and manages the supply and demand for resources in the Grid using a set of resource trading services provided by GRACE (Grid Architecture for Computational Economy) [3]. The GridBank (G-Bank) has been proposed as an infrastructure for managing accounts of resource owners and users along with handling electronic payments. A case for realizing economy Grid for service-oriented computing has been presented our earlier work [5].

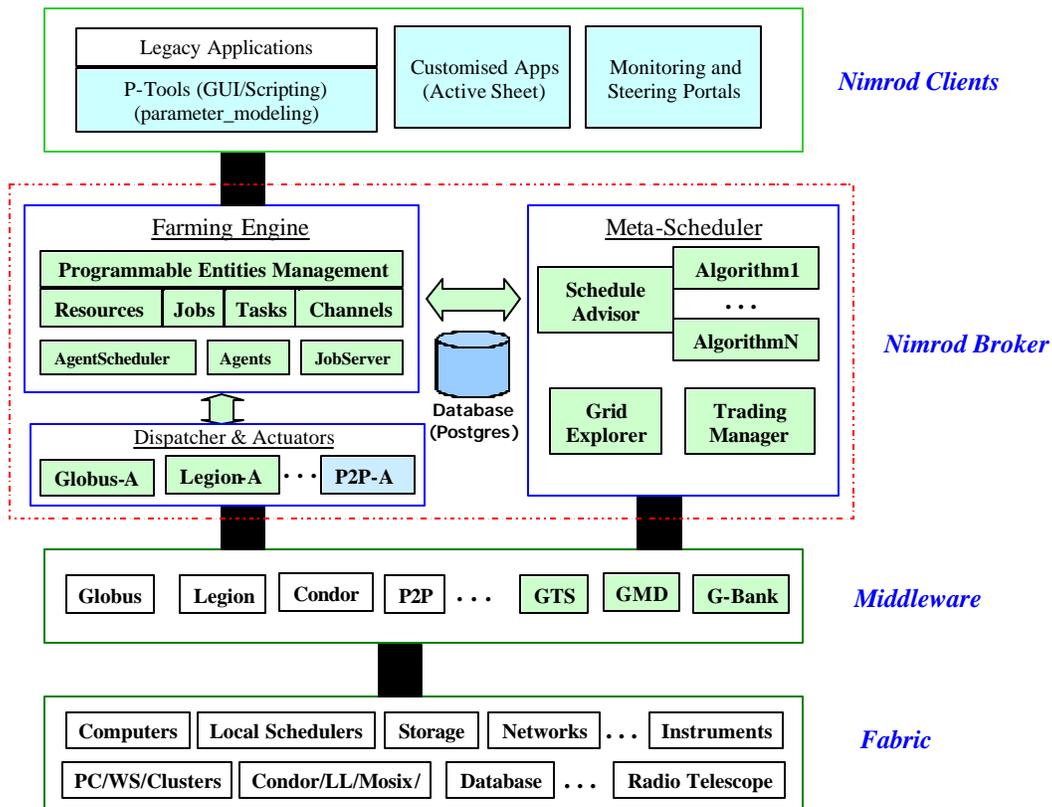


Figure 8: A Layered architecture of Nimrod-G system.

The Nimrod-G Resource broker is responsible for determining the specific requirements that an experiment places on the Grid and performing resource discovery, resource trading, scheduling, deploying jobs on Grid resources, starting and managing job execution, and gathering results back to the home node. The four key components of the Nimrod-G Resource Broker are: the *task farming engine*; the *scheduler* that consists of a grid explorer for resource discovery, a schedule advisor backed with scheduling algorithms, and a resource trading manager; a *dispatcher and actuators* for deploying agents on grid resources, and *agents* for managing the execution of Nimrod-G jobs on grid resources.

The Task Farming Engine (TFE)

The Nimrod-G task-farming engine is a persistent and programmable job control agent that maintains the record of experiment (jobs, parameters, tasks, deadline, and budget), manages, and coordinates with other components. It consists of a database to provide persistence that is accessed through a thin management interface. The farming engine is responsible for the parameterization of the experiment, the actual creation of jobs, the maintenance of job status, and provides the means for interaction between the clients, the schedule advisor, and the dispatcher. The TFE interacts with the scheduler and dispatcher in order to process jobs. It manages the experiment under the direction of schedule advisor, and then instructs the dispatcher to deploy jobs on selected resources and manage their execution.

The TFE maintains the state of an entire experiment and ensures that it is recorded in persistent storage. This allows the experiment to be restarted if root node fails. The TFE exposes interfaces for job, resource, and task management along with the job-to-resource mapping APIs. Accordingly, scheduling policy developers can use these interfaces to implement other schedulers without concern for the complexity of low-level remote execution mechanisms.

The programmable capability of task-farming engine enables “plugging” of user-defined schedulers and customised clients or problem solving environments (e.g., ActiveSheets [29]) in place of the default components. The task-farming engine is a coordination point for processes performing resource trading, scheduling, data and executable staging, remote execution, and result collation.

The Scheduler

The scheduler is responsible for resource discovery, resource trading, resource selection, and job assignment. The resource discovery algorithm interacts with an information service (the MDS in Globus), identifies the list of authorized and available machines, trades for resource *access cost*, and keeps track of resource status information. The resource selection algorithm is responsible for selecting those resources that meet the *deadline and budget* constraints along with optimization requirements. Nimrod-G provides three different scheduling algorithms [4].

The Dispatcher and Actuators

The dispatcher triggers appropriate actuators to deploy agents on grid resources and assign one of the resource-mapped jobs for execution. Even though the schedule advisor creates a schedule for the entire duration based on user requirements, the dispatcher deploys jobs on resources periodically depending on load and number of CPUs that are available. We have implemented different dispatchers and actuators for each different middleware service. For example, a Globus-specific dispatcher is required for Globus resources, and a Legion-specific component for Legion resources.

Agents

Nimrod-G agents are deployed on Grid resources dynamically at runtime depending on the scheduler’s instructions. The agent is responsible for setting up the execution environment on a given resource for a job. It is responsible for transporting the code and data to the machine; starting the execution of the task on the assigned resource and sending results back to the TFE. Since the agent operates on the “far side” of the middleware resource management components, it provides error-detection for the user’s task, sending the information back to the TFE.

The Nimrod-G agent also records the amount of resource consumed during job execution, such as the CPU time and wall clock time. The online measurement of the amount of resource consumed by the job during its execution helps the scheduler evaluate resource performance and change the schedule accordingly. Typically, there is only one type of agent for all mechanisms, irrespective of whether they are fork or queue nodes. However, different agents are required for different middleware systems.

6. Scheduling Experiments

We have performed Deadline and Budget Constrained (DBC) scheduling experiments on the World Wide Grid (WWG) [16] testbed shown in Figure 9. The testbed has computational resources owned by different organizations distributed across five continents: Asia, Australia, Europe, North America, and South America. It contains heterogeneous resources such as PCs, workstations, SMPs, clusters, and vector supercomputers running operating systems such as Linux, Sun Solaris, IBM AIX, SGI IRIX and Compaq Tru64. Further, the systems use a variety of job management systems such as OS-Fork, NQS, Condor, RMS, PBS, and LSF. Most of these resources support secure remote access through the Globus

system and a Linux cluster at Virginia is managed using the Legion system. The Solaris workstation from where this scheduling experiment is performed runs Globus, Legion, and Condor systems along with the Nimrod-G resource broker. The Nimrod-G agents are deployed at runtime on resources for managing the execution of jobs.

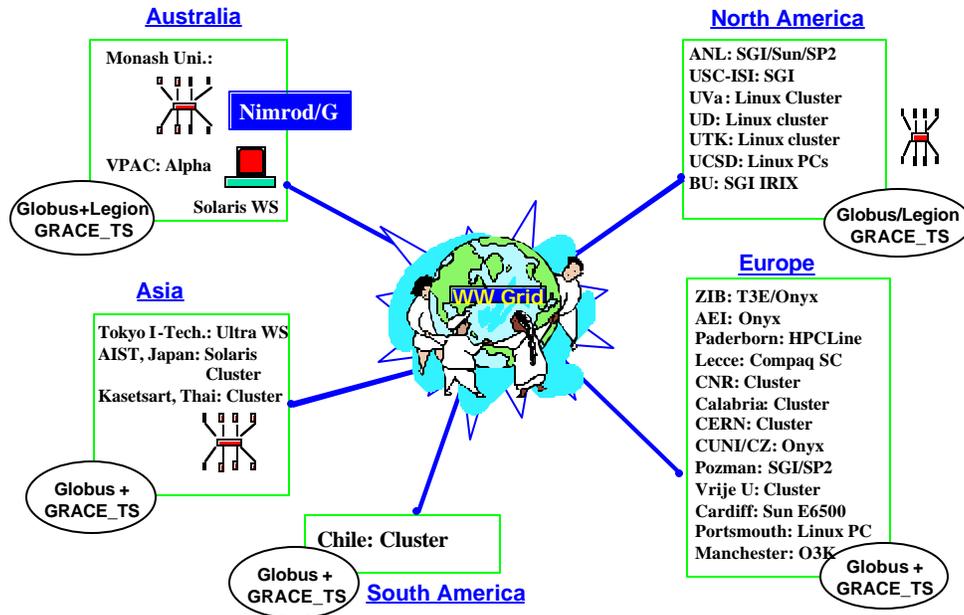


Figure 9: The World Wide Grid (WWG) Testbed.

We have created a hypothetical parameter sweep application (PSA) that executes a CPU intensive program with 200 different parameter scenarios or values. The program *calc* takes two input parameters and saves results into a file named “output”. The first input parameter *angle_degree* represents the value of angle in degree for processing trigonometric functions. The program *calc* needs to be explored for angular values from 1 to 200 degrees. The second parameter *time_base_value* indicates the expected calculation complexity in minutes plus 0 to 60 seconds positive deviation. That means, the program *calc* is modeled to execute for anywhere between 10 to 11 minutes on resources randomly. A plan for modelling this application as a parameter sweep application using the Nimrod-G parameter specification language is shown in Figure 10. The first part defines parameters and the second part defines the task that needs to be performed for each jobs. As the parameter *angle_degree* is defined as a range parameter type with values varying from 1 to 200 in step of 1, it leads to the creation of 200 jobs with 200 different input parameter values. To execute each job on a Grid resource, the Nimrod-G resource broker depending on its scheduling strategy, first copies the program executable to a Grid node, then executes the program, and finally copies results back to the user home node and stores output with job number as file extension.

```
#Parameters Declaration
parameter angle_degree integer range from 1 to 200 step 1;
parameter time_base_value integer default 10;

#Task Definition
task main
  #Copy necessary executables depending on node type
  copy calc.$OS node:calc
  #Execute program with parameter values on remote node
  node:execute ./calc $angle_degree $time_base_value
  #Copy results file to use home node with jobname as extension
  copy node:output ./output.$jobname
endtask
```

Figure 10: Nimrod-G Parameter Specification Plan.

We have conducted two scheduling experiments for a given deadline of 4 hours (with 18 minutes of extension for the second experiment) and budget of 250000 (G\$ or tokens) with different optimization strategies [4]:

1. Optimize for Time: This strategy produces results as early as possible, but before a deadline and within a budget limit.
2. Optimize for Cost: This strategy produces results by deadline, but reduces cost within a budget limit.

In these scheduling experiments, the Nimrod-G resource broker employed the *commodity market* model for establishing a service access price. It used grid resource trading services for establishing connection with the Grid Trader running on resource providers' machines and obtained service prices accordingly. The broker architecture is generic enough to use any of the protocols discussed above for negotiating access to resources and choosing appropriate ones. The access price varies from one consumer to another and from time to time, as defined by the resource owners. Depending on the deadline and the specified budget, the broker develops a plan for assigning jobs to resources. While doing so it does dynamic load profiling to learn the ability of resources to execute jobs. Thus, it adapts itself to the changing resource conditions including failure of resources or jobs on the resource. The heuristics-based scheduling algorithms employed by Nimrod-G broker are presented in our early work [4].

We have used a subset of resources of the WWG testbed in these scheduling experimentations. Table 2 shows resources details such as architecture, location, and access price along with type of Grid middleware systems used in making them Grid enabled. A snapshot of the Nimrod-G monitoring and steering clients taken immediately after the completion of application processing is shown in Figure 11 and 12. The resources used in both experiments are (time/space) shared resources with many other independent users. Hence, they were partially available to us, which changed dynamically depending on other users requirements and priorities. The access price indicated in the table is being established dynamically using GRACE resource trading protocols (commodity market model), but is based on an arbitrary assignment by us for demonstration purposes only.

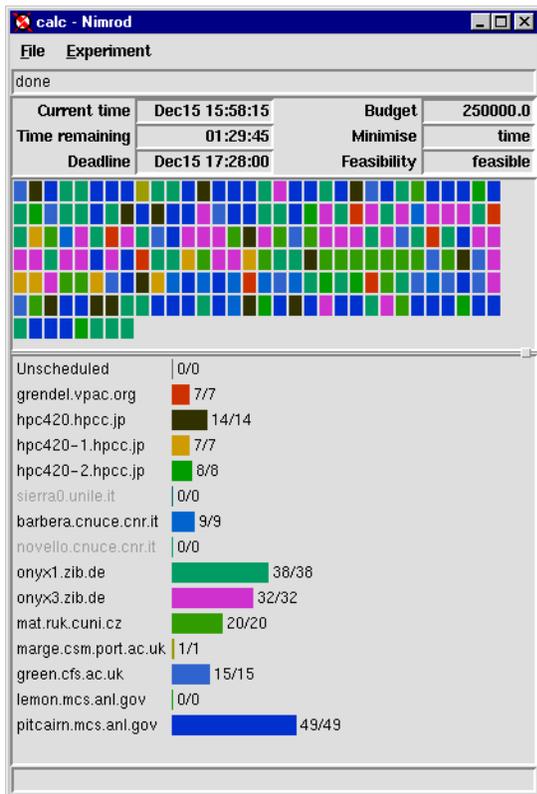


Figure 11: A Snapshot of the Nimrod-G Monitor during “Optimize for Time” Scheduling Experiment.

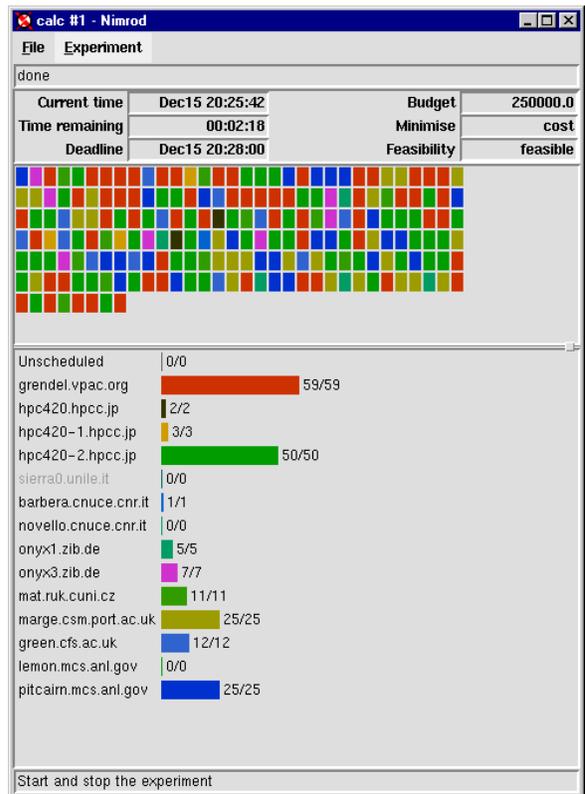


Figure 12: A Snapshot of the Nimrod-G Monitor during “Optimize for Cost” Scheduling Experiment.

Organization & Location	Vendor, Resource Type, # CPU, OS, hostname	Grid Services, Fabric, and Role	Price (G\$/CPU sec.)	Number of Jobs Executed	
				TimeOpt	CostOpt
Monash University, Melbourne, Australia	Sun: Ultra-1, 1 node, bezek.dstc.monash.edu.au	Globus, Nimrod-G, CDB Server, Fork (Master node)	--	--	--
VPAC, Melbourne, Australia	Compaq: Alpha, 4 CPU, OSF1, grendel.vpac.org	Globus, GTS, Fork (Worker node)	1	7	59
AIST, Tokyo, Japan	Sun: Ultra-4, 4 nodes, Solaris, hpc420.hpcc.jp	Globus, GTS, Fork (Worker node)	2	14	2
AIST, Tokyo, Japan	Sun: Ultra-4, 4 nodes, Solaris, hpc420-1.hpcc.jp	Globus, GTS, Fork (Worker node)	1	7	3
AIST, Tokyo, Japan	Sun: Ultra-2, 2 nodes, Solaris, hpc420-2.hpcc.jp	Globus, GTS, Fork (Worker node)	1	8	50
University of Lecce, Italy	Compaq: Alpa cluster, OSF1, sierra0.unile.it	Globus, GTS, RMS (Worker node)	2	0	0
Institute of the Italian National Research Council (CNR), Pisa, Italy	Unknown: Dual CPU PC, Linux, barbera.cnuce.cnr.it	Globus, GTS, Fork (Worker node)	1	9	1
Institute of the Italian National Research Council (CNR), Pisa, Italy	Unknown: Dual CPU PC, Linux, novello.cnuce.cnr.it	Globus, GTS, Fork (Worker node)	1	0	0
Konrad-Zuse-Zentrum Berlin, Berlin, Germany	SGI: Onyx2K, IRIX, 6, onyx1.zib.de	Globus, GTS, Fork (Worker node)	2	38	5
Konrad-Zuse-Zentrum Berlin, Berlin, Germany	SGI: Onyx2K, IRIX, 16 onyx3.zib.de	Globus, GTS, Fork (Worker node)	3	32	7
Charles University, Prague, Czech Republic	SGI: Onyx2K, IRIX, mat.ruk.cuni.cz	Globus, GTS, Fork (Worker node)	2	20	11
University of Portsmouth, UK	Unknown: Dual CPU PC, Linux, marge.csm.port.ac.uk	Globus, GTS, Fork (Worker node)	1	1	25
University of Manchester, UK	SGI: Onyx3K, 512 node, IRIX, green.cfs.ac.uk	Globus, GTS, NQS, (Worker node)	2	15	12
Argonne National Lab, Chicago, USA	SGI: IRIX lemon.mcs.anl.gov	Globus, GTS, Fork (Worker node)	2	0	0
Argonne National Lab, Chicago, USA	Sun: Ultra -8, Solaris, 8, pitcairn.mcs.anl.gov	Globus, GTS, Fork (Worker node)	1	49	25
Total Experiment Cost (G\$)				199968	141869
Time to Finish Expt. (Min.)				150	258

Table 2: The WWG testbed resources used in scheduling experiments, job execution and costing.

6.1. DBC Constrained Time Optimization Scheduling

The first experiment, *Optimize for Time* scheduling, we performed on December 15, 2001 at 13:28:00, Australian Eastern Daylight Saving Time (AEDT), with 4 hours deadline and finished on the same day by 15:58:15. A snapshot of the Nimrod-G monitoring and steering client, taken immediately after the completion of experiment, is shown in Figure 11. This experiment took 2½ hours to finish the processing of all jobs using resources available at that time with an

expense of 199968 G\$. Figure 13 shows the number of jobs in execution on different resources and Figure 14 shows the total number of jobs in execution on the Grid during the experiment execution period. It can be observed during the first hour of deadline, called the *calibration phase*, the broker aggressively consumed resources for processing jobs to bring the experiment to a feasible state.

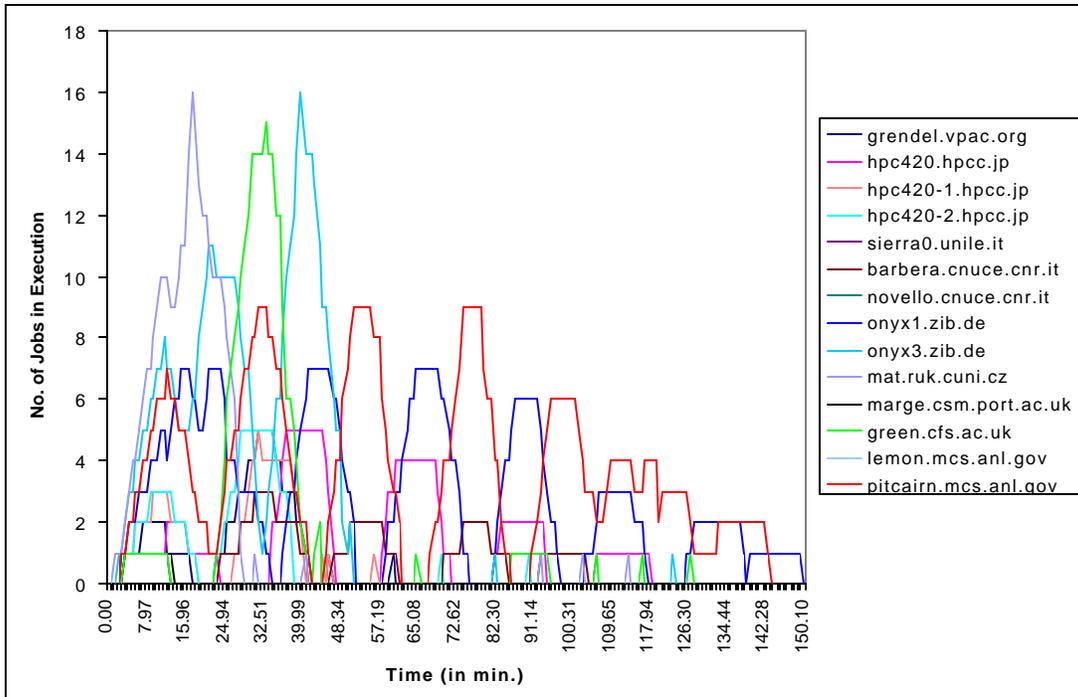


Figure 13: No. of jobs in execution on different Grid resources during DBC time optimization scheduling.

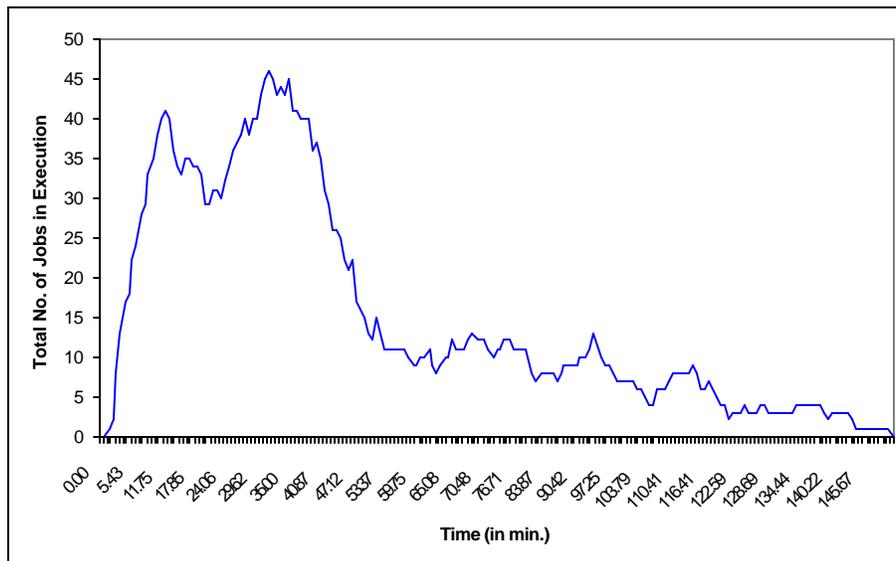


Figure 14: Total No. of jobs in execution on Grid during DBC time optimization scheduling.

Throughout the experiment, two resources, marked in *gray* color in Figure 11, were unavailable (either they were shutdown or their resource information servers had failed). We were unable to schedule any jobs on the first ANL node, lemon.mcs.anl.gov, as that was overloaded with jobs from other users. Figure 15 shows the number of jobs processed on

different resources selected depending on their cost and availability. Figure 16 shows the total number of jobs processed on the Grid. Figures 17 shows the corresponding expenses of processing on different resources and Figure 18 shows the aggregated processing expenses. From the graphs it can be observed that the broker selected resources to ensure that the experiment was completed at the earliest possible time given the current availability of resources and the budget limitations. It continued to use expensive resources even after the calibration phase deepening on the amount of remaining budget. Another interesting pattern to be observed in Figures 15 and 17 is that the amount of budget consumed by a resource is not always in proportion with the consumption amount and jobs processed. The most expensive machine (3G\$/sec.), onyx3.zib.de, was used to process less jobs compared to the other two machines, onyx1.zib.de and pitcairn.mcs.anl.gov, but we had to spend more budget for processing on it.

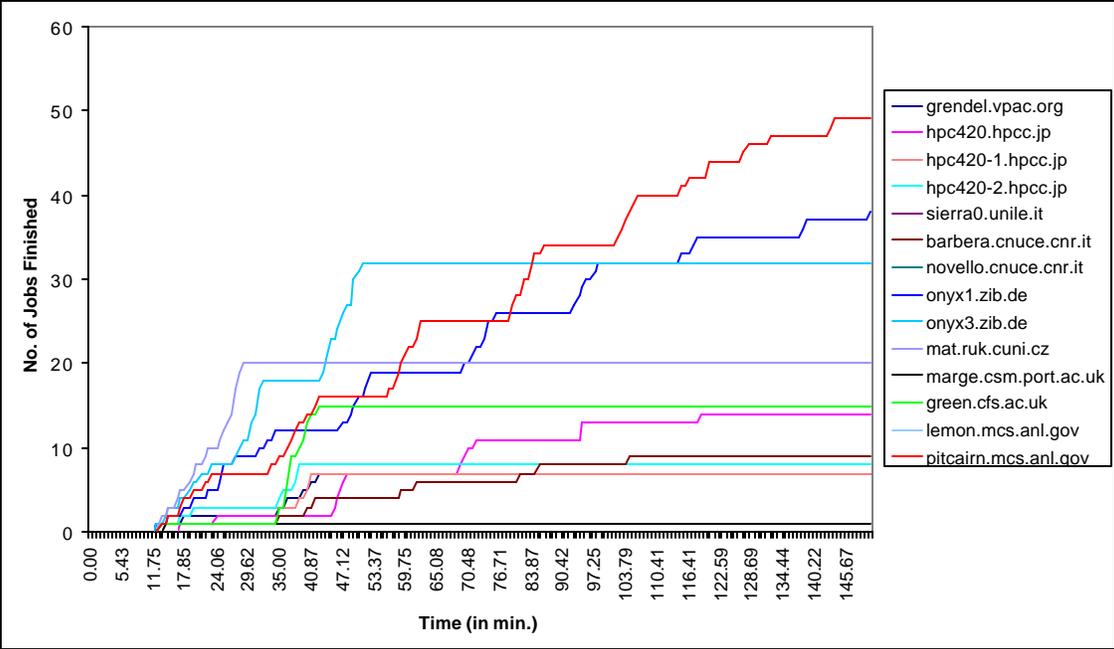


Figure 15: No. of jobs processed on different Grid resources during DBC time optimization scheduling.

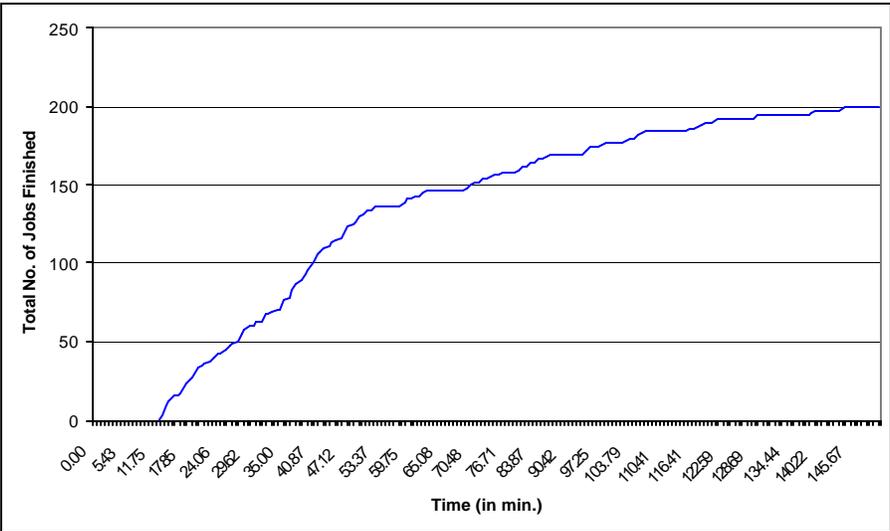


Figure 16: Total no. of jobs processed on Grid resources during DBC time optimization scheduling.

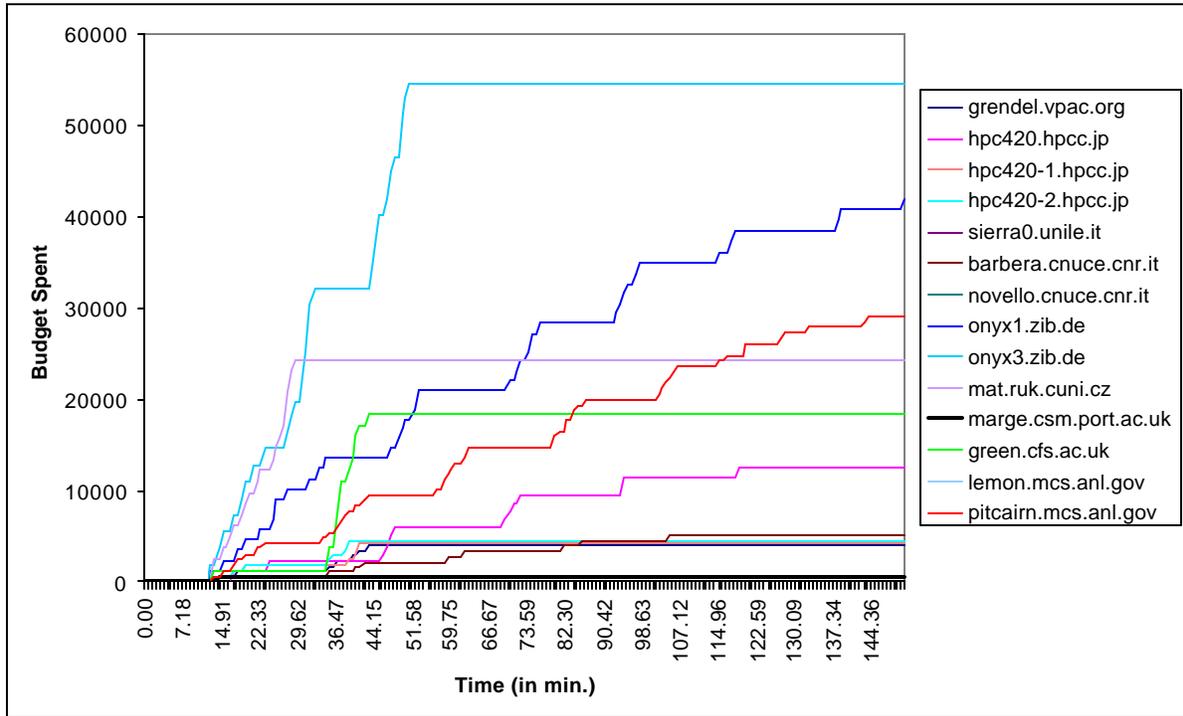


Figure 17: The amount spent on different Grid resources during DBC time optimization scheduling.

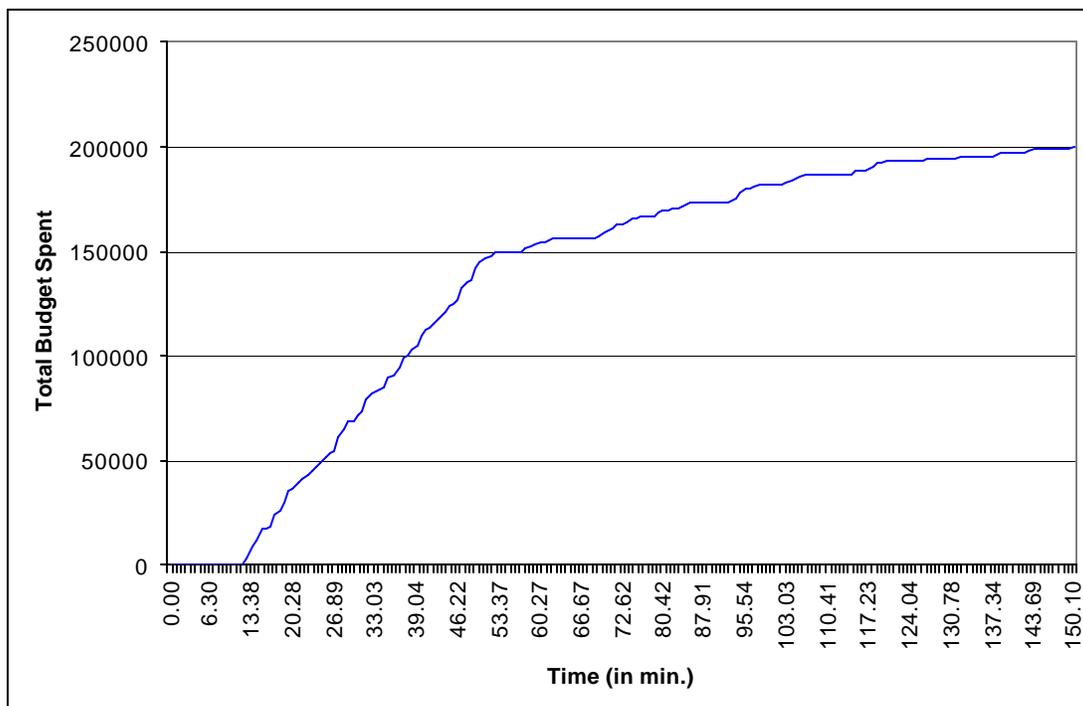


Figure 18: The total amount spent on Grid during DBC time optimization scheduling.

6.2. DBC Constrained Cost Optimization Scheduling

The second experiment, *Optimize for Cost* scheduling, was performed on December 15, 2001 at 16:10:00, Australian Eastern Daylight Saving Time (AEDT), with 4 hours deadline, which has been extended by 20 minute after 3 hours of execution time. It finished on the same day by 20:25:52. A snapshot of the Nimrod-G monitoring and steering client, taken immediately after the completion of experiment, is shown in Figure 12. This experiment took 4 hours and 18 minutes to finish the processing of all jobs using resources available at that time with an expense of 141869 G\$. Even though it took more time compared to the first experiment, it saved 58099 G\$ from the expenses.

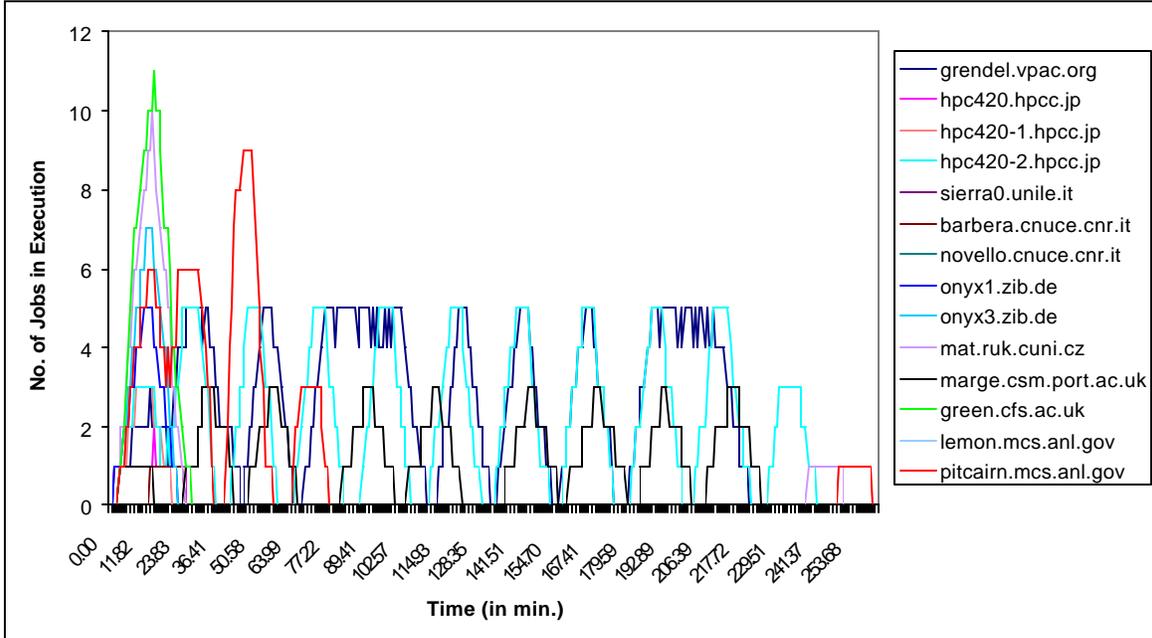


Figure 19: No. of jobs in execution on different Grid resources during DBC cost optimization scheduling.

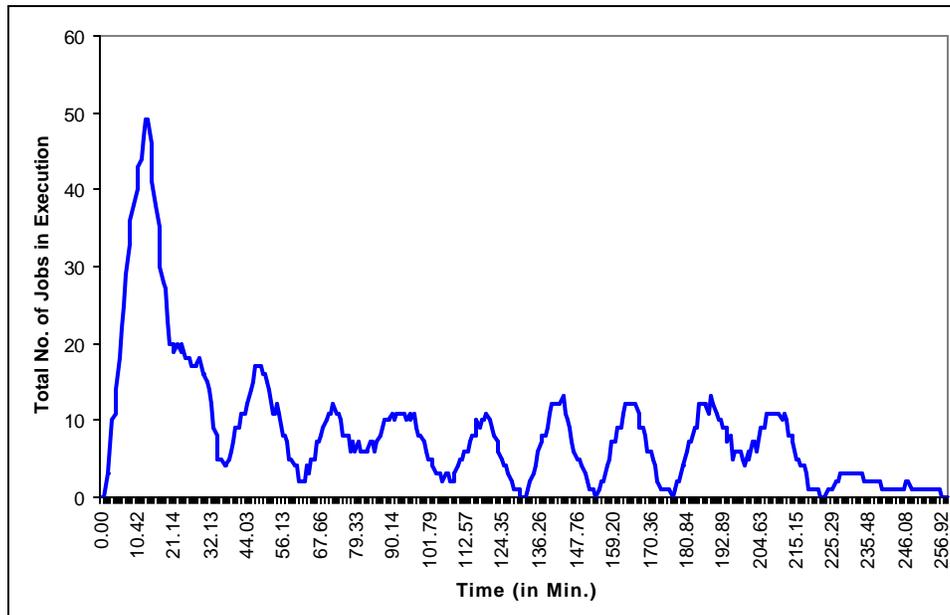


Figure 20: Total No. of jobs in execution on Grid during DBC cost optimization scheduling.

Figure 19 shows the number of jobs in execution on different resources and Figure 20 shows the total number of jobs in execution on the Grid during the experiment execution period. It can be observed during the first half-hour, called the *calibration phase*, the broker aggressively consumed resources for processing jobs to bring the experiment to a feasible state. During the experiment, one Italian resource, marked in *gray* color in Figure 12, was unavailable. As in the first experiment, we were unable to schedule any jobs on the first ANL node, *lemon.mcs.anl.gov*, as that was overloaded with jobs from other users. Figure 21 shows the number of jobs processed on different resources selected depending on their cost and availability and Figure 22 shows the aggregation of jobs processing on the Grid at different times during the experimentation. Figures 23 and 24 show the corresponding amount of budget spent for processing on individual resources and the Grid respectively.

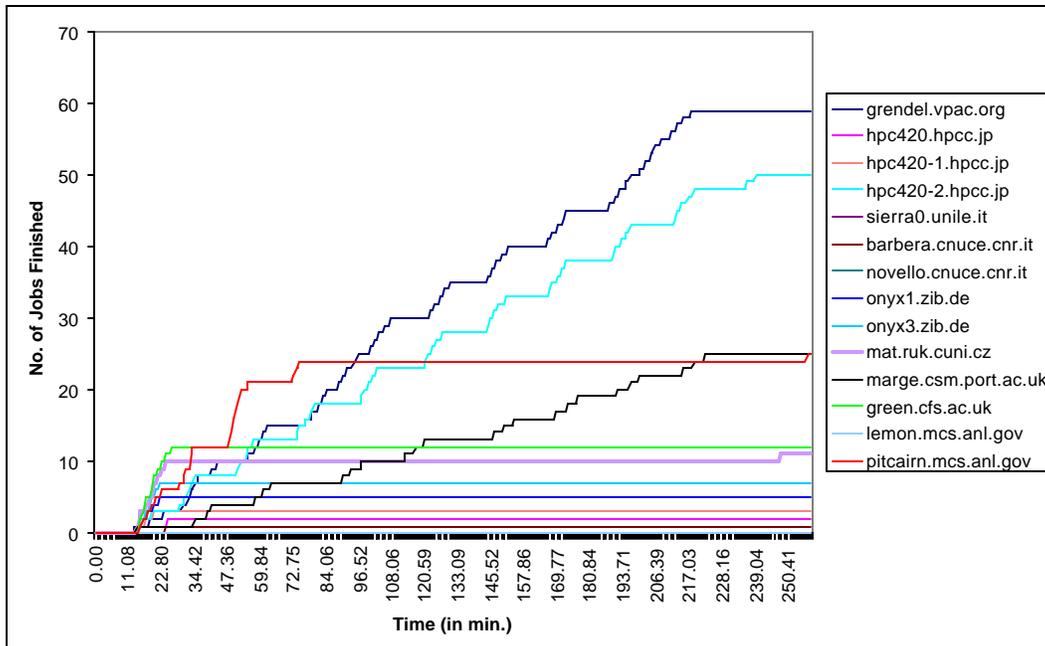


Figure 21: No. of jobs processed on different Grid resources during DBC cost optimization scheduling.

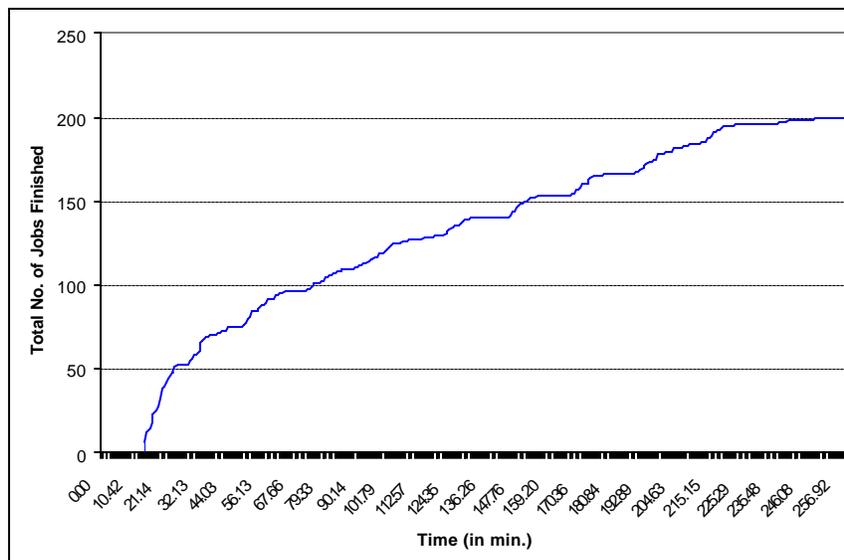


Figure 22: Total No. of jobs processed on Grid during DBC cost optimization scheduling.

From the graphs it can be observed that, after the calibration phase, the broker selected the cheapest and powerful resources extensively to process jobs, as cost minimization was the top priority as long as the deadline could be met. However, it did use a moderately expensive resource, for example, resource `mat.ruk.cuni.cz` costing 2G\$/CPU-sec., to process one job to ensure that deadline was met and this was essential as the availability of the cheapest resources had changed from the forecasted availability. As in the first experiment, it can be observed that the amount of budget consumed by a resource was not always in proportion with the number of jobs processed (see Figures 21 and 23) since we had to spend the higher amount for processing on expensive resources.

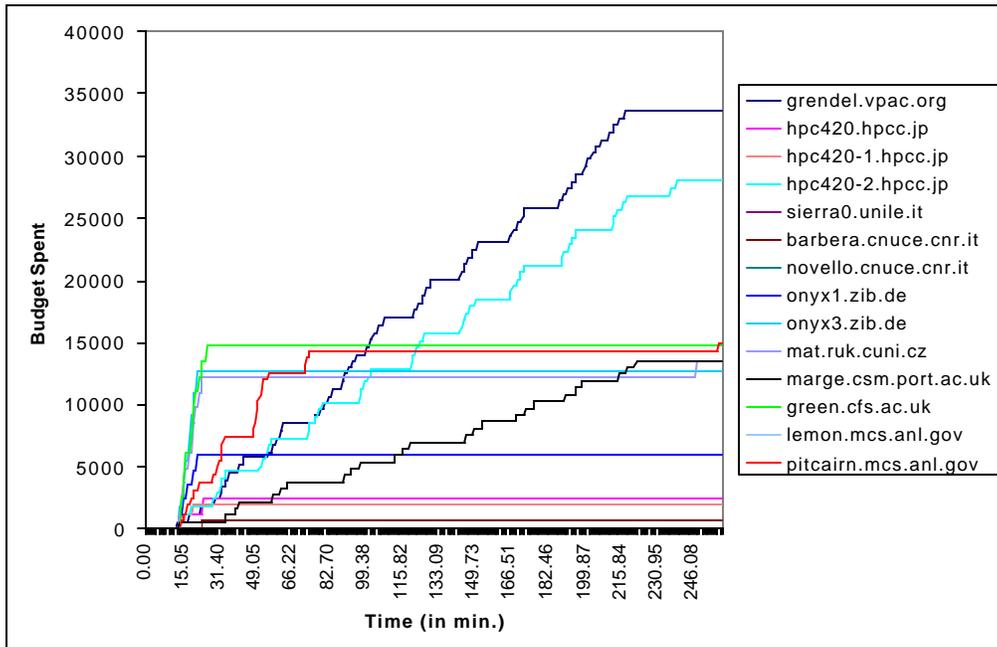


Figure 23: The amount spent on different Grid resources during DBC cost optimization scheduling.

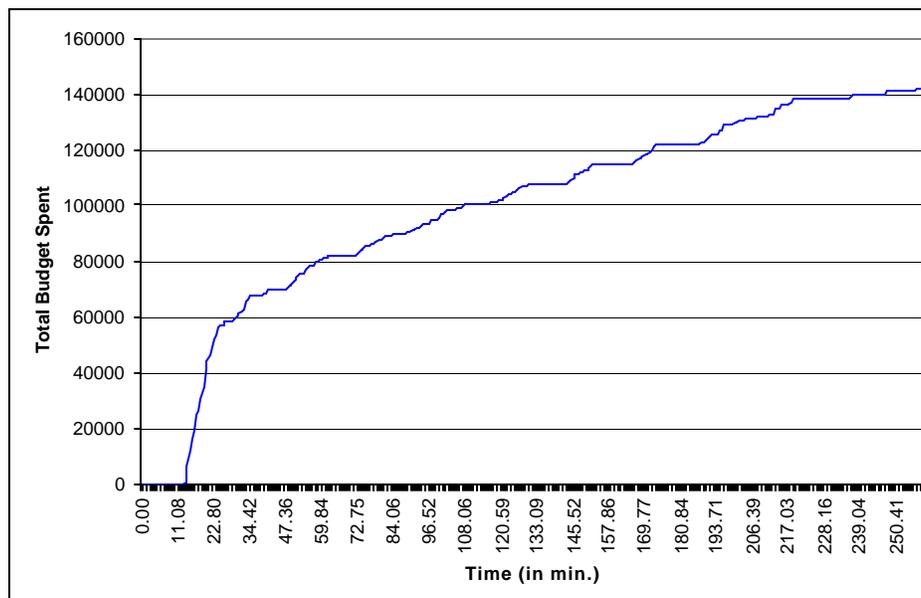


Figure 24: The total amount spent on Grid during DBC cost optimization scheduling.

7. Summary and Conclusion

Computational Grids enable the creation of a virtual computing environment for sharing and aggregation of distributed resources for solving large-scale problems in science, engineering, and commerce. The resources in the Grid are geographically distributed and owned by multiple organizations with different usage and cost policies. They have a large number of self-interested entities (distributed owners and users) with different objectives, priorities and goals that vary from time to time. The management of resources in such a large and distributed environment is a complex task. We proposed the use of computational economy as a metaphor for the management and regulation of supply-and-demand for resources. It provides decentralized resource management capability and is adaptable to changes in the environment and user requirements. It is a scalable, controllable, measurable, and easily understandable policy for management of resources. We discussed several market-based economic models such as commodity market, tenders, and auctions along with architecture and algorithms for their implementation in the Grid computing systems.

We briefly presented a computational economy driven Grid system called Nimrod-G. The Nimrod-G resource broker supports the deadline and budget constrained algorithms for scheduling task-farming applications on large-scale distributed systems. The results of scheduling applications with different quality of service requirements on the World Wide Grid resource demonstrates the power of these economic models in optimal and effective usage of resources. More importantly, it provides mechanisms to trade-off QoS parameters, deadline and computational cost and offers incentive for relaxing their requirements—reduced computational cost for relaxed deadline when timeframe for earliest results delivery is not too critical. We believe, this approach of providing an economic incentive for resource owners to share their resources and resource users to trade-off between the deadline and budget, promotes the Grid as a platform for mainstream computing, which can lead to the emergence of a new service oriented computing industry.

References

- [1] D. Abramson, J. Giddy, and L. Kotler, *High Performance Parametric Modeling with Nimrod/G: Killer Application for the Global Grid?*, Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS 2000), May 1-5, 2000, Cancun, Mexico, IEEE Computer Society (CS) Press, USA, 2000.
- [2] R. Buyya, D. Abramson, and J. Giddy, *Nimrod/G: An Architecture for a Resource Management and Scheduling System in a Global Computational Grid*, Proceedings of the 4th International Conference and Exhibition on High Performance Computing in Asia-Pacific Region (HPC ASIA 2000), May 14-17, 2000, Beijing, China, IEEE CS Press, USA, 2000.
- [3] R. Buyya, D. Abramson, and J. Giddy, *An Economy Driven Resource Management Architecture for Global Computational Power Grids*, Proceedings of the 2000 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000), June 26-29, 2000, Las Vegas, USA, CSREA Press, USA, 2000.
- [4] R. Buyya, J. Giddy, and D. Abramson, *An Evaluation of Economy-based Resource Trading and Scheduling on Computational Power Grids for Parameter Sweep Applications*, Proceedings of the 2nd International Workshop on Active Middleware Services (AMS 2000), August 1, 2000, Pittsburgh, USA, Kluwer Academic Press, USA, 2000.
- [5] R. Buyya, D. Abramson, and J. Giddy, *A Case for Economy Grid Architecture for Service-Oriented Grid Computing*, Proceedings of the International Parallel and Distributed Processing Symposium: 10th IEEE International Heterogeneous Computing Workshop (HCW 2001), April 23, 2001, San Francisco, California, USA, IEEE CS Press, USA, 2001.
- [6] K. Czajkowski, I. Foster, N. Karonis, C. Kesselman, S. Martin, W. Smith, and S. Tuecke, *A Resource Management Architecture for Metacomputing Systems*, Proceedings of the 4th International Workshop on Job Scheduling Strategies for Parallel Processing, March 30, 1998, Orlando, Florida, USA, Springer Verlag Press, Germany, 1998.
- [7] S. Chapin, J. Karpovich, and A. Grimshaw, *The Legion Resource Management System*, Proceedings of the 5th Workshop on Job Scheduling Strategies for Parallel Processing, April 16, 1999, San Juan, Puerto Rico, USA, Springer Verlag Press, Germany, 1999.
- [8] M. Stonebraker, R. Devine, M. Kornacker, W. Litwin, A. Pfeffer, A. Sah, and C. Staelin, *An Economic Paradigm for Query Processing and Data Migration in Mariposa*, Proceedings of the 3rd International Conference on Parallel and Distributed Information Systems, Sept. 28-30, 1994, Austin, Texas, USA, IEEE CS Press, USA, 1994.

- [9] W. Vickrey, *Counter-speculation, auctions, and competitive sealed tenders*, Journal of Finance, Vol. 16, No. 1, pp. 9-37, March 1961.
- [10] W. Hoschek, J. Jaen-Martinez, A. Samar, H. Stockinger, and K. Stockinger, *Data Management in an International Data Grid Project*, Proceedings of the First IEEE/ACM International Workshop on Grid Computing (GRID 2000), Dec. 17, 2000, Bangalore, India, Springer Verlag Press, Germany, 2000.
- [11] L. W. McKnight and J. Boroumand, *Pricing Internet Services: Approaches and Challenges*, IEEE Computer, Vol. 33, No. 2, pp. 128-129, IEEE CS Press, USA, Feb. 2000.
- [12] H. Stockinger, K. Stockinger, E. Schikuta, and I. Willers. *Towards a Cost Model for Distributed and Replicated Data Stores*, Proceedings of the 9th Euromicro Workshop on Parallel and Distributed Processing (PDP 2001), Feb. 7-9, 2001, Italy, IEEE CS Press, USA, 2001.
- [13] D. Ferguson, C. Nikolaou, J. Sairamesh, and Y. Yemini, Economic Models for Allocating Resources in Computer Systems, In *Market-based Control: A Paradigm for Distributed Resource Allocation*, World Scientific Press, Singapore, 1996.
- [14] R. Smith and R. Davis, *The Contract Net Protocol: High Level Communication and Control in a Distributed Problem Solver*, IEEE Transactions on Computers, Vol. C-29, No. 12, pp. 1104-1113, Dec. 1980, IEEE CS Press, USA.
- [15] T. Sandholm, Distributed Rational Decision Making, In *Multi-Agent Systems: A Modern Introduction to Distributed Artificial Intelligence*, G. Weiss, editor, The MIT Press, USA, 2000.
- [16] R. Buyya, *World Wide Grid testbed*, <http://www.buyya.com/ecogrid/wwg/>, June 2001.
- [17] G. Heiser, F. Lam, and S. Russell, *Resource Management in the Mungi Single-Address-Space Operating System*, Proceedings of Australasian Computer Science Conference, February 4-6, 1998, Perth Australia, Springer-Verlag, Singapore, 1998.
- [18] Y. Amir, B. Awerbuch, and R. Sean Borgstrom, *A Cost-Benefit Framework for Online Management of a Metacomputing System*, Proceedings of 1st International Conference on Information and Computational Economy, October 25-28, 1998, Charleston, SC, USA.
- [19] Y. Amir, B. Awerbuch., A. Barak A., S. Borgstrom, and A. Keren, *An Opportunity Cost Approach for Job Assignment in a Scalable Computing Cluster*, IEEE Transactions on Parallel and Distributed Systems, Vol. 11, No. 7, pp. 760-768, IEEE CS Press, USA, July 2000.
- [20] C. Waldspurger, T. Hogg, B. Huberman, J. Kephart, and W. Stornetta, *Spawn: A Distributed Computational Economy*, IEEE Transactions on Software Engineering, Vol. 18, No. 2, pp 103—117, IEEE CS Press, USA, February 1992.
- [21] N. Nisan, S. London, O. Regev, and N. Camiel, *Globally Distributed computation over the Internet - The POPCORN project*, International Conference on Distributed Computing Systems (ICDCS'98), May 26 - 29, 1998, Amsterdam, The Netherlands, IEEE CS Press, USA, 1998.
- [22] B. Chun and D. Culler, *Market-based proportional resource sharing for clusters*, Technical Report CSD-1092, University of California, Berkeley, USA, January 2000.
- [23] W. T. Sullivan, III, D. Werthimer, S. Bowyer, J. Cobb, D. Gedye, and D. Anderson, *A new major SETI project based on Project Serendip data and 100,000 personal computers*, Proceedings of the 5th International Conference on Bioastronomy, 1997. Available at: http://setiathome.ssl.berkeley.edu/woody_paper.html, June 2001.
- [24] Mojo Nation - <http://www.mojonation.net/>, June 2001.
- [25] N. Kapadia and J. Fortes, *PUNCH: An Architecture for Web-Enabled Wide-Area Network-Computing*, Cluster Computing: The Journal of Networks, Software Tools and Applications, Vol. 2, No. 2, pp. 153-164, Baltzer Science Publishers, The Netherlands, Sept. 1999.
- [26] F. Berman and R. Wolski, *The AppLeS Project: A Status Report*, Proceedings of the 8th NEC Research Symposium, Berlin, Germany, May 1997.
- [27] H. Casanova and J. Dongarra, *NetSolve: A Network Server for Solving Computational Science Problems*, International Journal of Supercomputing Applications and High Performance Computing, Vol. 11, No. 3, pp 212-223, Sage Publications, USA, 1997.

- [28] M. Litzkow, M. Livny, and M. Mutka, *Condor - A Hunter of Idle Workstations*, Proceedings of the 8th International Conference of Distributed Computing Systems (ICDCS 1988), January 1988, San Jose, CA, IEEE CS Press, USA, 1988.
- [29] D. Abramson, P. Roe, L. Kotler L, and D. Mather, *ActiveSheets: Super-Computing with Spreadsheets*, Proceedings of the 2001 High Performance Computing Symposium (HPC'01): Advanced Simulation Technologies Conference, April 22-26, 2001, Seattle, Washington, USA.
- [30] S. Lalis and A. Karipidis, *An Open Market-Based Framework for Distributed Computing over the Internet*, Proceedings of the First IEEE/ACM International Workshop on Grid Computing (GRID 2000), Dec. 17, 2000, Bangalore, India, Springer Verlag Press, Germany, 2000.
- [31] D. Reed, I. Pratt, P. Menage, S. Early, and N. Stratford, *Xenoservers; Accounted execution of untrusted code*, Proceedings of the 7th Workshop on Hot Topics in Operating Systems (HotOS-VII), Rio Rico, Arizona, USA, March 28-30, 1999, IEEE CS Press, USA, 1999.
- [32] J. Bredin, D. Kotz, and D. Rus, *Utility Driven Mobile-Agent Scheduling*, Technical Report CS-TR98-331, Dartmouth College, Hanover, NH, USA, October 3, 1998.
- [33] J. Brooke, M. Foster, S. Pickles, K. Taylor, and T. Hewitt, *Mini-Grids: Effective Test-beds for Grid Application*, Proceedings of the First IEEE/ACM International Workshop on Grid Computing (GRID 2000), Dec. 17, 2000, Bangalore, India, Springer Verlag Press, Germany, 2000.
- [34] R. Buyya, K. Branson, J. Giddy, and D. Abramson, *The Virtual Laboratory: Enabling Molecular Modelling for Drug Design on the World Wide Grid*, Technical Report Monash-CSSE 2001-103, Monash University, Melbourne, Australia, Dec. 2001.
- [35] L. Norskog, *A Personal Communication on Economics and Grid Allocation*, Enron Broadband Systems, March 14, 2001. Available at: <http://www.buyya.com/ecogrid/comments/enron.txt>
- [36] K. Reynolds, *The Double Auction*, Agorics, Inc., 1996. <http://www.agorics.com/Library/Auctions/auction6.html>
- [37] R. Das, J. Hanson, J. Kephart, and G. Tesauro, *Agent-Human Interactions in the Continuous Double Auction*, Proceedings of the International Joint Conferences on Artificial Intelligence (IJCAI), August 4-10, 2001, Seattle, Washington, USA.