Resource Management in Grids: Overview and a discussion of a possible approach for an Agent-Based Middleware

Mehrdad Senobari, Michal Drozdowicz, Maria Ganzha, Marcin Paprzycki, Richard Olejnik, Ivan Lirkov, Pavel Telegin, Nasrollah Moghadam Charkari

► To cite this version:

Mehrdad Senobari, Michal Drozdowicz, Maria Ganzha, Marcin Paprzycki, Richard Olejnik, et al.. Resource Management in Grids: Overview and a discussion of a possible approach for an Agent-Based Middleware. Saxe-Coburg Publications, Stirlingshire, UK. PARALLEL, DISTRIBUTED AND GRID COMPUTING FOR ENGINEERING, Saxe-Coburg Publications, Stirlingshire, UK, pp 141-164, 2009, Computational Science, Engineering & Technology Series, ISSN 1759-3158. 10.4203/csets.21.8 . hal-00834409

HAL Id: hal-00834409

<https://hal.science/hal-00834409>

Submitted on 17 Jun 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Resource management in Grids—overview and a discussion of a possible approach for an agent-based middleware

M. Senobari¹, M. Drozdowicz², M. Ganzha², M. Paprzycki^{2,3}, R. Olejnik⁴, I. Lirkov⁵, P. Telegin⁶ and N. M. Charkari¹

¹Tarbiat Modares University, Tehran, Iran ²System Research Institute Polish Academy of Sciences, Warsaw, Poland

³Warsaw Management Academy, Warsaw, Poland

⁴University of Sciences and Technologies of Lille, Lille, France

⁵Institute for Parallel Processing, Bulgarian Academy of Sciences, Sofia, Bulgaria

⁶SuperComputing Center Russian Academy of Sciences, Moscow, Russia

Keywords: Grid computing, agent-based resource management, resource management, job scheduling, meta scheduling.

Resource management and job scheduling are important research issue in computational Grids. When software agents are used as resource managers and brokers in the Grid a number of additional issues and possible approaches materialize. The aim of this chapter is twofold. First, we discuss traditional job scheduling in Grids, and when agents are utilized as Grid middleware. Second, we use this as a context for discussion of how job scheduling can be done in the agent-based system under development.

References

Abstract

Resource management and job scheduling are important research issue in computational Grids. When software agents are used as resource managers and brokers in the Grid a number of additional issues and possible approaches materialize. The aim of this chapter is twofold. First, we discuss traditional job scheduling in Grids, and when agents are utilized as Grid middleware. Second, we use this as a context for discussion of how job scheduling can be done in the agent-based system under development.

Keywords: Grid computing, agent-based resource management, resource management, job scheduling, meta-schedulers

1 Introduction

According to [42], “computational Grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to computational resources.” In general, resources (also named services; see [33]) can include supercomputers, data repositories, clusters, sensors, workstations, programs, or individual PCs. However, in this chapter we are interested in Grid as a collection of computational (rather than data / software) resources, which have to be effectively managed. However, the heterogeneous highly dynamic nature of the Grid makes development of Grid Resource Management Systems (GRMS) a challenge. There have been many projects focused on designing and implementing GRMSs. As it has been argued in [62], the most important issues in the process are: (a) supporting adaptability, extensibility, and scalability, (b) allowing systems with different administrative policies to inter-operate while preserving site autonomy, (c) co-allocating resources, (d) supporting quality of service, and (e) meeting computational cost constraints.

In the case of Grid computing scheduling means allocation (matching) resources to incoming jobs (we use terms task and job interchangeably). While there exists exhaustive work on scheduling in “traditional” parallel and distributed systems, it is rarely applicable to the Grid. The main reason is that they often assume that tasks are completed utilizing homogeneous dedicated resources residing in a single administrative domain [62]. However, the vision of the Grid assumes that it consists of heterogeneous non-dedicated resources residing in multiple administrative domains.

One of interesting approaches to resource management in the Grid is based on software agents. While arguments presented in [41, 68] there are not without critics, we are in the process of developing a GRMS, where agent teams are the basic managerial infrastructure [35]. Thus far we have concentrated our attention on interactions between users and the team leader (the *LMaster* agent) [48], and interactions between *Worker* agents and the actual Grid infrastructure [65]. Currently, we need to address the question: how will the *LMaster* select the *Worker* to execute a given job (for the time being we assume that a single *Worker* will execute a single job). In this context, first, we present an overview of approaches to Grid resource management and scheduling. Second, we discuss how understanding these approaches can help us to solve the load distribution problem within our agent-based infrastructure.

2 Grid resource managers and schedulers

An early survey of Grid resource management systems was compiled in 2002 by Buyya et al. [62]. There, three basic approaches to job scheduling were distinguished: (i) centralized, (ii) hierarchical, and (iii) decentralized. The centralized scheduling is easiest to manage and deploy, but is not well suited for the Grid. Its main disadvantages are lack of scalability and fault-tolerance. In the hierarchical approach, schedulers at higher levels manage larger sets of resources (and utilize schedulers at lower levels). Hierarchical schedulers are utilized, for instance, in the 2K distributed operating system [46], as well as in Darwin [31] and Legion [32] resource management systems. On the other hand, Ninf [53], MOL [45], and Bond [23] projects adopt decentralized scheduling. Note that Grid schedulers usually cannot control Grid resources directly. They mostly work as *brokers* (or *agents*). However, Grid schedulers can also be tightly coupled with applications (application-level scheduling; see [36]).

While Grid heterogeneity includes applications, resources, as well as middleware components and their relations, it is possible to generalize processes involved in Grid scheduling. For instance, in [36] it is argued that Grid scheduling consist of: (1) resource discovering and filtering, (2) resource selecting and scheduling according to certain objectives, and (3) job submission. In [36] we can find a bird-eye view of the structure of the Grid resource management. Here, the Grid Scheduler (GS) receives jobs from Grid Users, performs matchmaking (utilizing an Information Service—IS), and generates a schedule. In other words, result of scheduling is a map between jobs and available resources. Note that in the Grid multiple schedulers may exist in separate administrative domains, and may be organized in different structures. Delivery

of a feasible schedule requires detailed information about (a) properties of the job and (b) of available resources, and (c) current state of the system. Here, the Grid Information Service (GIS) provides system state information. After the schedule is prepared, submitting a task to selected resource(s) and monitoring execution is the responsibility of the Launching and Monitoring (LM) service. Another functionality of the LM is staging necessary executables and input data. Local Resource Manager (LRM) performs the final scheduling and submission inside of a specific domain. The LRM has additional control over the available resources (e.g. it enforces administrative policies over resource usages, and reports the resource status information to the GIS). Let us now look into some details of popular schedulers.

2.1 AppLeS

The AppLeS project [1] was based at UC San Diego and was focused on development of scheduling agents for applications running in Grids. The AppLeS system collects resource information from the Network Weather Service (NWS) running at each computing node, and dispatches tasks to lighter loaded nodes; while scheduling actual execution of applications is local. AppLeS uses other RMSs, e.g. Globus, Legion, or NetSolve, to execute actual jobs (i.e. it can be viewed as a meta-middleware placed above the standard Grid middleware). Each application has embedded AppLeS agents that perform resource scheduling. AppLeS applications utilize templates that define specific computational models and thus allow reuse of application patterns (e.g. parametric and master-slave application templates have been provided). AppLeS facilitates predictive heuristic state estimation model, online rescheduling and fixed application-oriented scheduling policy [26, 62]. According to our best assessment the AppLeS project is no longer active. All of its publications and the WWW site were completed before 2001 and most of its participants have moved out of UC San Diego.

2.2 Nimrod/G

Nimrod/G is a Grid resource broker based on an economy-driven approach to manage resources and schedule jobs. It utilizes services provided by other Grid middleware (e.g. Globus, Legion, Condor), and the GRACE [5] trading mechanisms. More detailed explanation of economical methods used in Nimrod/G is presented in section 3.2. Note that, while at the Nimrod/G site there are references to work completed in 2007, the latest version of this middleware v3.0.1 was released in October, 2005. Thus, to the best of our knowledge, today this project is no longer active.

2.3 OpenPBS

OpenPBS [10, 11] is a simple workload management solution intended for small clusters of dedicated homogeneous nodes. Here, computers are federated into a virtual pool of resources. Workload is scheduled to run within this virtual pool, based on

simple scheduling algorithms. OpenPBS is one of workload managers accessible from the CSF meta-scheduler (see, Section 2.6). However, the last release of the OpenPBS as an independent project (v2.3.16) happened in 2001. At the same time, the PBS Professional, is the commercial product developed and sold by the Altair corporation.

2.4 NetSolve

The NetSolve project [37] was focused on execution of scientific applications in heterogeneous environments, while utilizing different scheduling algorithms for different applications. Job completion time estimation was based on performance and load models, while a dynamic job queue was used for job ordering. Length of this queue was adaptively adjusted based on historical performance data (an example of a system-level scheduling [26]). In addition, mechanisms for scheduling multi-step, data-dependent jobs have been implemented [20]. Recently, the NetSolve project has been extended to Grids through the GridSolve infrastructure [9]. Both projects are active; for instance, a new release of GridSolve software appeared on 2008-12-18.

2.5 Condor

Condor [2] is a high-throughput computing environment that manages large collections of diversely owned machines. It utilizes a centralized scheduler based on the ClassAd matchmaker. To overcome the disadvantages of centralized scheduling, Condor allows the matchmaker (and/or the user) to forward requests to another matchmaker through the gateway flocking mechanism. The Condor project is still under development and has a large community of users.

2.6 Community Scheduler Framework

The Community Scheduler Framework (CSF) is an open source Web Services Resource Framework compliant [15] metascheduler built for the Globus Toolkit [3, 66, 40]. The CSF provides interface and tools for Globus users to create reservations, define scheduling policies and submit jobs to the Grid. CSF functionalities can be extended to utilize other schedulers and support different Grid deployment models. For instance, using CSF allows a single interface access to (i) Load Sharing Facility (LSF); [8]), (ii) OpenPBS, (iii) Condor, and (iv) Sun Grid Engine (SGE [12]). The CSF is the default metascheduler for the Globus Toolkit 4. This indicates that the CSF is not only active, but likely to be developed further (with the development of Globus).

2.7 ADAJ and SOAJA

The aim of the ADAJ (Adaptative Application in Java) and the SOAJA (Service Oriented Adaptive Java Applications) projects is to develop an infrastructure to run applications in “desktop Grids” [57, 49]. ADAJ is a programming and execution envi-

ronment, which contains mechanisms for dynamic re-distribution of components of an applications (in response to load imbalances among Grid nodes). The SOAJA is a WSRF-based ([61]) service oriented extension of ADAJ. It is being developed on top of ADAJ, by adding the webservices layer. Furthermore, in SOAJA adjustments are to be made to make it less dependent on proprietary solutions (e.g. the Enterprise Service Bus and the JavaParty); see [44] for more details. Similarly to ADAJ, SOADAJ is to facilitate workload measurement and component-level load balancing.

3 Approaches to task scheduling in Grids

Job/task scheduling has a long history of research in parallel and distributed computing systems. Obviously, a large number of methods developed in that context have been adapted to Grid scheduling. Since a thorough summary of Grid scheduling methods can be found in [36, 52], we will focus our attention only on economic methods, which fit very well with agent systems (see literature related to utilization of software agents in e-commerce scenarios, for instance [18, 19]).

3.1 Economic models

Metaphorically speaking, computational Grids and the power grid share the same general economic model. In the power grid, we use electricity and pay for usage of a “unit of electricity.” In the computational Grid we are to use computational resources and pay for their usage (possibly, one day a “unit of computing” will be established and globally accepted). Therefore, in the computational Grid we can distinguish producers and consumers (with their objectives) and commodities (resources) that are traded (e.g. programs, data, storage, CPU cycles, etc.). To support this aspect of the Grid (Grid economy), we need infrastructure (algorithms/policies) similar to that in the e-commerce. In other words, we need to be able to establish price of resource, service level agreement (SLA) and its enforcement, secure payment, etc. [26]. Similarly to the case of the (commodity) market, it is possible to utilize transactions based on both bartering and commodity prices. In the case of bartering, exchange of resources takes place (e.g., storage space for CPU time). When pricing is used, price of services should be based on supply, demand and possibly other economic factors.

For the Grid economy to materialize, Grid users need both: (1) to be able to specify resource requirements, and (2) to establish their preferences (e.g. that a given job should be high priority and price is not important). At the same time the Grid infrastructure (acting as a resource broker) has to be able to select resources that meet these requirements. Obviously, what is also needed are effective mechanisms for price / SLA negotiations. Mechanisms of this type have been studied in research in broadly understood e-commerce. For instance, when resource price is considered alone, mechanisms for establishing price equilibrium should be utilized.

According to [26], service providers and users can be functionalized as represented

by Grid Service Providers (GSPs) and Grid Resource Brokers (GRBs). The GSPs deliver Grid enabled resources (e.g. Globus or ADAJ) as well as Grid Trading Services (GTS) to facilitate resource usage negotiations (based on user requests delivered by the GRBs). Interactions between GRBs and GSPs during contract negotiations are mediated by a Grid Market Directory (GMD).

Depending on the mechanism used for establishing transaction details (negotiating the Service Level Agreement, the price, etc.) either GRBs or GSPs can initiate the negotiation. For instance, a GRB may invite proposals from GSPs' and select one that satisfies its requirements (e.g. job will be done before the deadline and within cost constraints). Alternatively, a GSP may invite bids from prospective users and offer its services to the highest bidder. Note that both GSPs and GRBs have constraints to be satisfied and utility functions to be maximized. Let us look into details of selected mechanisms that can be used to negotiate transaction details (for more information, as well as a list of projects that, before 2002, used economic models, see [26]).

3.1.1 Bargaining Mechanism

In the bargaining model, GRBs bargain with GSPs, for instance for lower access price and/or higher usage duration. In the e-commerce literature this model is also known as *iterative bargaining* [58]. Here, GRBs and GSPs have objective functions and negotiate as long as their objectives are met or until it is established that finding an agreement is not possible. Negotiation can involve a single item (e.g. price), or multiple items (e.g. price and deadline). According to [26] this model is particularly useful when market supply-and-demand and service prices are not clearly established.

3.1.2 Tender/Contract-Net Mechanism

Tender/Contract-Net model is one of the popular models for service negotiations (see, also [4]). Here, the GRB announces its requirements (using a specific template) and invites bids from GSPs. Interested GSPs evaluate the announcement and respond by submitting their bids. The GRB awards the contract to the GSP that submitted the best offer. In the case when no (satisfactory) offer is obtained the GRB may adjust and re-submit its call for proposals. The negotiation template may include, among others, addressee, requirements specifications (e.g. Linux, x86arch, and 1024MB memory), task/service description (e.g. a Matlab job), maximum price (optional), bid specification (what should offer contain), expiration time (deadline for receiving bids), etc. Note that the Tender/Contract-Net allows finalizing contracts without bargaining. This simplifies interactions and can improve the efficiency of the system (e.g. it is easy to imagine a very long sequence of iterations taking place in the case of Bargaining).

3.1.3 Auction Model

According to [26] the Auction model involves one-to-many negotiations, between a GPS and multiple GRBs, and reduces negotiation to a single value (i.e., price). Note

that this view is a clear oversimplification of what has been discussed in the e-commerce literature, see for instance [18, 50, 21]. Let, us however accept this view as an approximation of what could be used in this case. Here, the basic process involves: (a) start of the Auction, (b) bid submission, (c) agreement formation, or establishing that agreement cannot be formed (see, also [19]). Most popular forms of Auctions are: English Auction, First-price sealed-bid, Vickrey Auction, and Dutch Auction. However, the literature concerning single and multi-item auctions considers a much broader spectrum of contract negotiation mechanisms.

3.2 Job scheduling in Nimrod/G

To illustrate application of an economic model in Grid resource management, let us look into the Nimrod/G approach. In [27], three economy-based algorithm used in the Nimrod/G are presented (note that neither one of them utilizes the elaborate economic mechanisms presented above):

- Time Minimization—complete job(s) within time and budget constraints,
- Cost Minimization—complete job(s) within time and budget constraints,
- None Minimization—complete job(s) within time and budget constraints.

Let us assume that a task to be completed consists of one or more jobs. When the Time Minimization algorithm is used, the goal is to complete the task as quickly as possible (within the available budget). The key steps of this algorithm are as follows [27]:

1. For each available resource, use information about previously assigned jobs to estimate the completion time for a new job.
2. Sort resources according to the expected completion times.
3. Assign a given job to the resource for which the completion time is the “shortest,” while the cost is less than the remaining budget per job.
4. Repeat until all jobs are assigned.

In the case of the Cost Minimization, the goal is to complete the task as cheaply as possible, while still satisfying the deadline constraint ([27]):

- Sort available resources according to the advertised cost (cheapest first).
- Utilize estimates of job execution times to assign as many jobs as possible to cheapest resources; without exceeding the deadline.

Finally, for the None Minimization algorithm, the goal is to complete the task within the deadline and cost constraints, but no minimization of either is attempted [27]:

- Divide resources in such a way that in each case cost per job is less than budget available per job.
- In the case of the cheaper resources, assign jobs (inverse) proportionally to the estimated job completion time (e.g. “cheap resource” with estimated completion time = 4 received twice as many jobs as resource with estimated completion time = 8).
- In the case of more expensive resources, repeat steps (until all jobs are assigned), but use recalculated budget per job (budget based on money left after cheaper resources have been contracted).

Note that implementation of either one of these three strategies involves taking care of special situations. For instance, during systems startup completion times estimates are unknown, and thus a supplemental strategy has to be used. Similarly, when there are “too many” jobs in the task and they cannot be *all* assigned to available resources (an example of an infeasible schedule), system has to be able to manage such situation.

4 Agent based scheduling systems

Since in our project we utilize software agents as Grid resource managers, let us now focus our attention on approaches to Grid resource scheduling appearing in the context of agent-Grid integration.

4.1 ARMS

ARMS was an agent-based Grid RMS. It used agents for resource advertisement and discovery [28]. It utilized performance prediction mechanisms provided by the Performance Analysis and Characterise Environment (PACE) toolkit [56]. Furthermore, scheduling was focused on QoS requirements; e.g. users had to specify an explicit job execution deadline.

In ARMS homogeneous cooperating agents are organized in a hierarchy, and their goal is to manage and schedule applications over available Grid resources. Each agent acts as a representative for a single Grid resource, while PACE is used to create a hardware characterization template. Next, hardware model and services information is spread across the agent hierarchy. This information is utilized to build the Agent Capability Table (ACT). Each ACT item is a tuple containing the agent ID, information about services as well as performance. ACTs are updated periodically and both pull and push methods are used to maintain them.

When an application is submitted to the system, it includes an associated performance model and requirements related to its execution (e.g. start time, deadline, etc.). Service discovery involves communicating with agents neighboring in the hierarchy (upward and downward). It is claimed that this approach is therefore more scalable when the Grid becomes large.

ARMS agents consist of three layers: (i) communication layer, (ii) coordination layer, and (iii) local management layer. The communication layer acts as an interface to the external world. It receives service advertisements and discovery messages and forwards them to appropriate modules in the coordination layer, which perform matchmaking and scheduling. For the service discovery message, agent tries to find an available Grid resource. This involves utilization of the (a) application model, (b) job requirements, and (c) the PACE engine. Specifically, the expected execution time of a given application on a selected resources is estimated and compared with the requirements. If time constraints are satisfied for one of them, process is completed. Otherwise agent forwards the request to other agents (higher or lower in the hierarchy) to find resource that will satisfy user defined constraints.

In [29], an ant-based self-organizing mechanism is utilized to perform load balancing for batch jobs with no explicit execution deadlines. It is shown that application of such mechanism improves global load balancing in the system (given large enough number of ants). Separately, in [30], scheduling dependent jobs and executing workflows in ARMS was discussed.

It should be noted, that while this approach seems very interesting, work on the PACE framework and the ARMS system is not pursued further since approximately 2003. What is left are only papers reporting results (no agent code pertinent to any part of the system can be found).

4.2 JADE Extensions

In their work, Poggi et al. extended the JADE agent framework to be used in Grid applications [59]. They argued that realizing an agent-Grid integration is possible through: (i) extending Grid middleware to support agent features, or (ii) extending agent-based middleware to support functionalities of the Grid. They follow the second approach by attempting at adding new features to JADE to realize a Grid environment. Considered extensions are mechanisms for: code distribution, reconfiguration, goal delegation, load balancing optimization and QoS definition. To realize these goals new types of agents are proposed. They are to support: (i) rule-based creation and composition of tasks, and (ii) mobility of the code at the task level (i.e., JADE behaviors or rules are exchanged by agents). First, a *Drools agent* is developed, which uses the Drools rule engine, to receive and execute rules (coming from other agents). The *BeanShell agent* receives and executes behaviors coming from other agents. It integrates the BeanShell engine, which allows usage of Java as a scripting language. Here, each rule can contains scripts in its condition, consequence or extractor fields. When a rule is scheduled for execution (its preconditions are satisfied), Drools invokes

the BeanShell interpreter to execute the code contained in the consequence of the rule. To address security issues in the Grid, the JadeS security framework is used.

Work on extending JADE agent framework was originally reported in 2004. Since then the following three papers, in some way related to the subject, have been published [54, 55, 60]. It is thus difficult to see this project as being actively pursued. This is even more so, as there is no Grid related add-on listed among JADE add-on software [6].

4.3 Bond

Bond is a Java-based object-oriented middleware for network computing. Bond was developed to create an infrastructure for a Virtual Laboratory. It was to support scheduling of complex tasks and data annotation for data intensive applications. One of the goals of the Bond system was to facilitate collaborative activities through support for knowledge and workflow management. These functionalities are based on a distributed object system [22] (e.g. to store and process resource information). Resources exchange information using messaging and utilizing the KQML language. Dissemination of resource information is achieved through periodic data pushes. In Bond, mechanism called distributed awareness is used to learn about existence of other agents. Specifically, each node maintains information about nodes it has communicated with, and periodically exchanges it with other agents. In this way, information about existing agents is propagated in the system. Finally, job scheduling is decentralized and utilizes predictive pricing models for state estimation [22].

Again, situation is similar to the two other systems described thus far. All publications related to the Bond system have appeared in the 1999–2003 time frame. Therefore, we have to conclude that the project is no longer active.

4.4 Agent-based Scheduling Framework

Agent-based Scheduling Framework (ASF) is an agent-centered scheduling approach applied to Grid scheduling [63]. ASF is composed of a metascheduler and autonomous agents attached to computing resources. The main idea of the ASF is to reduce the responsibilities of a conventional metascheduler. Specifically, the main issues that the ASF attempts to address are ([63]):

- The workload of the metascheduler grows when the number of computing resources grows.
- If the metacheduler is overloaded, accuracy of scheduling degrades as it is going to be based on incomplete information (not all information can be processed in time to make accurate predictions).
- What is needed is a new approach to managing large numbers of heterogeneous computers; especially, when many domains with various operation policies are

combined in a Grid as a VO (conventional frameworks are not ready to deal with such complex issues).

To achieve its goals the ASF relies on agents that autonomously search for jobs; instead of jobs being assigned to them by the metascheduler. Thus, the ASF is “dual” to conventional metaschedulers, which continuously collect state information of resources under their control, decide about the scheduling policy and *push* jobs to selected resources. In the ASF, instead, each agent discovers jobs that can be processed by its resource and retrieves them from the metascheduler (a *pull* based approach).

A prototype of the ASF has been implemented based on the Globus Toolkit 4. Experimental results of utilization of the ASF metascheduler have been discussed in [63]. It was shown, that for a relatively small environment (a heterogeneous cluster with 3 nodes(!)) utilization of the ASF resulted in an 11% reduction of the total elapsed time of job processing.

4.5 MAGDA

Mobile Agent based Grid Architecture (MAGDA; [16]) is a Java-based mobile agent toolkit designed to overcome some of the limitations of existing Grid middlewares. According to its creators [16], these include:

- Lack of ability to migrate an application from one system to another.
- Low level of abstraction of the heterogeneity of the environment.
- Lack robust fault tolerance.
- Existing information and monitoring frameworks do not scale to the Grid level (or are focused only on specific issues).
- Lack of support of task migration, monitoring and execution (with adequate checkpointing).

MAGDA supports (1) resource discovery, (2) performance monitoring and load balancing, and (3) task execution within the Grid. It has a layered architecture following the Layered Grid Model [43] and thus it should be possible to implement or integrate MAGDA components with other systems that are based on the same model.

In the current release of MAGDA, service discovery is performed by help of Web Services and Web Services technologies (such as UDDI) are used to implement this component. Application-level load balancing is provided by means of a *Coordinator Agent*. This agent manages lists of registered workers and of available free hosts. When an agent is initialized, coordinator updates list of registered agents. It stores information about their state of computation, and their relative speed. In this way, the *Coordinator Agent* can recognize imbalance in the system and ask the most loaded

worker to split its workload, part of which will then be assigned to the less loaded worker.

After migration from Aglets [14] to JADE [6], MAGDA is continually being developed with the aim of maximization of integration with Web Services, workflow management, service orchestration and choreography [17].

5 Job scheduling in the *Agents in Grid* project

Let us now look how issues discussed above can be utilized in the *Agents in Grid* project. In our approach it is assumed that agents work in teams [35]. Each team is managed by the *LMaster* agent. This agent also represents the team to the outside world. In other words, agents representing *Users*, named *LAgents* interact with *LMasters* either to contract job execution [34], or to join the team and to become a *Worker* [48]. Each *LMaster* is supported by an *LMirror* agent, which stores copies of crucial team data (e.g. list of workers, list of contracts, status of job execution, etc.). In the context of this chapter, we are interested in processes that take place when a job is contracted and forwarded to a selected *Worker* to be actually executed.

5.1 Forming the team

To start we need to consider the way that workers join the team. As described in [48], the process consists of the following steps: (a) *User* specifies conditions of joining to its *LAgent*, (b) the *LAgent* interacts with the *Client Information Center*, represented by the *CIC* agent, to obtain list of teams that seek workers satisfying certain conditions (e.g. CPU power, available memory. etc.), (c) upon receiving such list, the *LAgent* eliminates these that are not trustworthy [38], (d) next it utilizes the FIPA Contract Net Protocol [13] to negotiate which team to join (note that, in the context of this chapter, we omit details of contract negotiations; e.g. how the *LMaster* establishes the optimal price; we simply assume that one of bargaining mechanisms described in Section 3.1 is used), (e) negotiations can result in a success (joining a team), or in a failure (no acceptable team was found). Here, we are particularly interested in: (i) description of capabilities of resource(s) represented by the *LAgent*, and (ii) details of the contract. To illustrate them, let us consider the ACL message in Figure 5.1. It could have been sent by the *LAgent* as a call for proposals (within the Contract Net Protocol).

Here, agent (*demi*), representing machine (*barszcz*) with CPU running at 2.6 GHz, 2 Gbytes of RAM and 8 Gbytes of disc space available for Grid applications, is sending a *take-me* message to agent *bruce* (*LMaster* of some team residing on node *tequila*). Agent *demi* is proposing the following contract conditions: availability each day between 22:00 and 7:55 (next day in the morning), and contract duration 14 days.

It is obvious that this is the information that matches what is utilized by the above described schedulers. First, provided are capabilities of the machine (Grid node) represented by agent *demi*. They are needed to be able to estimate job completion time.

```

(cfp
 :sender (agent-identifier :name demi@barszcz:1099/JADE)
 :receiver (agent-identifier :name bruce@tequila:1099/JADE)
 :content
  ((action
   (agent-identifier :name bruce@tequila:1099/JADE)
   (take-me
    :configuration (hardware
                     :cpu 2.6
                     :memory 2048
                     :quota 8000)
    :conditions (condition
                 :availability (frequency
                               :unit (day)
                               :day-time (period
                                           :from 00000000T22000000
                                           :to 00000000T07550000))
                 :contract-duration +00000014T0000000000))
    :language fipa-s10
    :ontology joining-ontology
    :protocol fipa-contract-net
   )
 )
)

```

Figure 1: Sample call for proposals, containing description of available resources and conditions of joining

Second, let us assume that *bruce* “takes” *demi* to be a *Worker* in the team. In this case, contract details allow *bruce* to approximate if a task contracted by its team can be completed while the Grid node that agent *demi* represents will be still available. It would be very bad for the reputation of the “bruce-team” if a job was to be paused until the “demi-node” comes back online, and as a result contract conditions would not be fulfilled. Observe that this form of utilization of available information follows the Nimrod/G-based approach.

5.2 Contracting job execution

The second aspect that needs to be considered is: what happens when *User* would like a job to be executed in our system. As described in [34] the process consists of the following steps: (a) *User* specifies to its *LAgent* acceptable conditions of job execution, (b) the *LAgent* contacts the *CIC* agent to obtain the list of teams that have the required resources, (c) this list is then adjusted on the basis of trust considerations ([38]), (d) as in the previous case, FIPA Contract Net protocol is utilized as a negotiation mechanism to find the best team to execute the task, (e) process can result in a success (finding a team to do the job), or in a failure (no acceptable team is found). Here, of particular interest is available information concerning the job. This information is a part of the Contract Net CFP. To illustrate our initial approach let us present a snippet of our ontology of constraints:

```

:NegotiationSet a owl:Class .
:negotiationParam a owl:ObjectProperty ;
  rdfs:domain :NegotiationSet ;
  rdfs:range NegotiationParam .
:NegotiationParam a owl:Class .

```

```

: paramWeight
  a owl:DatatypeProperty , owl:FunctionalProperty ;
  rdfs:domain :NegotiationParam ;
  rdfs:range xsd:float .

: Cost a owl:Class ;
      rdfs:subClassOf :NegotiationParam .
: costConstraint
  a owl:ObjectProperty , owl:FunctionalProperty ;
  rdfs:domain :Cost ;
  rdfs:range :FloatConstraint .
: costValue
  a owl:DatatypeProperty , owl:FunctionalProperty ;
  rdfs:domain :Cost ;
  rdfs:range xsd:float .
: JobStartTime a owl:Class ;
      rdfs:subClassOf :NegotiationParam .
: jobStartTimeValue
  a owl:DatatypeProperty , owl:FunctionalProperty ;
  rdfs:domain :JobStartTime ;
  rdfs:range xsd:dateTime .
: jobStartTimeConstraint
  a owl:ObjectProperty , owl:FunctionalProperty ;
  rdfs:domain :JobStartTime ;
  rdfs:range :TimeConstraint .
: JobEndTime a owl:Class ;
      rdfs:subClassOf :NegotiationParam .
: jobEndTimeValue
  a owl:DatatypeProperty , owl:FunctionalProperty ;
  rdfs:domain :JobEndTime ;
  rdfs:range xsd:dateTime .
: jobEndTimeConstraint
  a owl:ObjectProperty , owl:FunctionalProperty ;
  rdfs:domain :JobEndTime ;
  rdfs:range :TimeConstraint .

```

Here, we can see that three constraints of job execution were conceptualized: (i) *cost* (as a constraint, this value is private to the *LAgent*, but is also used in the contract), (ii) *job start time*, and (iii) *job end time*. It is obvious, that *job start time* and *job end time* may provide useful information for job schedulers, but this needs to be considered further. Note that, it is possible to use the *job start time* constraint also in the case when there is no estimate as to how long this job is going to take. However, in this case we cannot have the *job end time* specified as well, as there is no feasible way to estimate if this constraint can be satisfied. Furthermore, in this case (*job start time* only), job has to be scheduled on a node with non-stop contract for an extended time (e.g. constant contract). Utilization of the *job end time* requires existence of at least some estimate of the job execution time. This information may originate from: (a) job description (could be provided by the *User* as a part of the CFP), or (b) could be based on past execution times. However, the latter case requires that information about the “type of the job” has to be a part of the CFP. Only then the *LMaster* would be able to estimate job execution time utilizing historical data. These considerations show clearly that the question of description of jobs, resources, constraints has to be revisited. Being aware of this, in another chapter of this book ([39]), we have presented an overview of resource descriptions utilized in various Grid systems, as well as attempts at creating an ontology of the Grid. As a result we have found that the Core Grid ontology is the one that is most likely going to be utilized in our system. However,

we have established also that it will have to be extended to include additional aspects of resource management, brokering and scheduling, necessary for our system (e.g. to deal with various constraints discussed above and support trust management).

Let us now assume that an appropriate Grid ontology has been selected and extended to facilitate robust descriptions of: (1) resources, (2) job characteristics, (3) job execution constraints, and (4) trust related concepts. It should be obvious that such ontology will support advanced contract negotiation scenarios. In this context let us stress that following [42, 33], we believe that one of the key aspects of future Grid computing will be the economic model (involving *Users* who pay for job execution, or are paid for usage of their resources). This being the case, we consider economic mechanisms described in Sections 3.1 and 3.2 as necessary for high-level resource management in the Grid. In this context, we expect that development of a robust Grid ontology will allow us to utilize semantic reasoning and thus opens the Grid up to a large number of autonomic contract negotiations mechanisms considered, among others, in e-commerce research (see, also [18, 64, 67, 51]). However, we will omit the very process of contract negotiations as outside of scope of this chapter.

Now, let us see what happens when a CFP reaches the *LMaster*. The first thing that it has to do is to check job execution constraints. Let us assume that on the basis of available data, the *LMaster* (i) can obtain an estimate of job execution time; or (ii) establishes, that there is a contradiction in constraints—open ended job (no completion time estimate available) combined with the *job end time* constraint; or (iii) finds out that it has to deal with an open ended job. In the case of constraint contradiction, the *LMaster* can send an ACL *REFUSE* message to the CFP originator. In the remaining two cases the *LMaster* can combine the job and constraint information with its knowledge of: (a) all “job execution contracts” that have been signed thus far, (b) current “work contracts” of all team *Workers*, and (c) current status of team *Workers* (which *Workers* are available, what is their workload etc.) to decide whether the proposed job can be executed within specified constraints or not. For instance, the *LMaster* should not contract an open ended job if it does not have a trustworthy *Worker* with non-stop contract. Similarly, it cannot contract a job with deadline in 24 hours, if all of its *Workers* are 100% utilized for the next 36 hours. In such cases, the *LMaster* responds with an ACL *REFUSE* message. If job execution is feasible (i.e, constraints can be satisfied), then the *LMaster* sends a “positive” response to the originator of the CFP. This response may contain a detailed final proposal (e.g. in the case of contract net negotiations) or may be a part of more elaborate contract negotiations.

Recall that we assume that *Worker* agents work within their contract and are paid on its basis and thus do not possess access to the information what price does the *LMaster* charge for their work (however, they may explore offers from multiple teams to find the market price of resources they offer). In this way the sole responsibility for contract negotiation is on the *LMaster*. Obviously, in a different scenario it could be assumed that the *LMaster* would negotiate “subcontracts” with its *Workers* (e.g. use a local Contract Net protocol) and use results of these (sub)negotiations to prepare a response to the *User*. However, we consider this approach unnecessarily complicated

and stay with the model in which the *LMaster* is the only contract negotiator.

5.3 Selecting worker to execute a task

Thus far we have established that, during the process of accepting team members (*Workers*), the *LMaster* obtains a complete description of resources represented by each one of them. Furthermore, for each *Worker* it knows details of its contract. Additionally, during evaluation of the Call for Proposals, the *LMaster* is able to establish that its team is capable of fulfilling the request within specified constraints. Let us now assume that the *LMaster* was able to successfully complete contract negotiations. Now, it has to obtain all information/data necessary to complete the job (including, for instance, needed files or information about their locations; see, [65], for more details). Recall that the *LMaster* is the *only* representative of the team known to the outside world. Here we follow basic tenets of agent system design [69] and reduce the number of possible agent interactions (i.e. *LAgents* representing *Users* do not know individual *Workers*). Assuming that the *LMaster* has all the information ready for the job execution to be initiated, the following two questions arise: (1) which agent should actually execute the job, and (2) when should it receive it.

In general, we have identified three possible responses to the first question. The first one is based on the ASF approach (see, Section 4.4). Here, the *LMaster* acts as a meta-scheduler, while *Workers* contact it to obtain the next task to be executed. While this approach has some potential advantages (e.g. reducing the workload of the *LMaster* and giving *Workers* more autonomy), it does not seem to work well in the case of Grid economy, as it is not able to handle complex SLA's. For instance, let us assume that a special job-contract has been successfully negotiated. This task requires specific resources and immediate commencement of execution. In return, price is substantially higher than a typical one. In this case the *LMaster* cannot wait until the right resource becomes available and requests that job (recall, that in the ASF, agents pull tasks from the meta-scheduler "at their will"). Instead, the *LMaster* should be able to "act" and rearrange work of the team in such a way that the high-priority high-paying job would start executing immediately while using the right resources. This illustrates that the ASF approach is not easy to combine with the economic model.

The second possible approach is conceptually based on [24, 25]. Here the *LMaster* and *Workers* negotiate job execution. Note that, as specified above, we assume that the *LMaster* is the sole contract negotiator, while *Workers* act in a non-competitive way. As a matter of fact, they should support each other, as the success of the team means also their (financial) success. Obviously, this assumption may need to be changed (introducing agent spoilers and/or competitive/selfish behaviors of *Workers*), but this would require reexamination of all trust-focused considerations (see [38]) and thus, for the time being, will not be pursued further. In the negotiation-based approach, when the *LMaster* wins a contract, it could utilize the FIPA Contract Net to inform *Workers* and to find out which of them could execute it. In response to the CFP, *Workers* would inform about their conditions of job execution. Answers received from each *Worker*

could be evaluated on the basis of their resources, contracts and trust to establish which should execute the given job. Here, trust is used to attempt at avoiding (or at least to minimize the impact of) failures of *Workers* (see, [38] for more details). While in [24, 25] improvement of order of 10% resulting from utilizing this method has been reported, we can see also some problems with this approach. Basically, this is the problem discussed in the previous solution. Let us assume that a high-paying job has been negotiated and has to start immediately. In the current scenario it is possible to envision that as a result of internal negotiations this job may be started immediately on an appropriate machine. The selected *Worker* stops executing its current job to work on the special task. However, this situation may have a domino effect. Now, the job that has been executed on that machine may need to be moved to another one (it also has a relatively high priority and a close deadline), and so on. Now, it is extremely difficult to envision how such job movement can be achieved in the situation when each *Worker* is an autonomous entity that may or may not agree to the change. And, even if they do agree, the process of job-alignment may involve large number of additional negotiations, or will actually be realized using the third approach (discussed in the next paragraph). Overall, in this approach, autonomy given to *Workers* may turn against the capability of the team to efficiently complete some jobs, and thus compete in the marketplace.

Finally, the *LMaster* itself can decide about the task assignment without any communication with its *Workers*. In other words, the *LMaster* can act as a meta-scheduler. This being the case, following examples of meta-schedulers described above, the *LMaster* can use, historical performance data for each *Worker*, information about their current load, etc., to establish which resource should execute which job. We can assume here that the *LMaster* not only can decide which *Worker* will execute which job, but also can issue an *order* that a given *Worker* should send its job to another, while the recipient has to accept it. In other words, we functionalize the *LMaster* as an omnipotent meta-scheduler. It is easy to observe that this approach minimizes communication between the *LMaster* and *Workers* and overall programmatic complexity of interactions within the team; only orders are sent to *Workers*, who report their completion. Furthermore, this means that *Workers* do not have to have special reasoning capacities (required in the second approach, in particular). Finally, this approach provides an easy way to deal with jobs that come at various priorities and deadlines (it is the *LMaster* that takes care of them). Unfortunately, we can find also an important disadvantage: requirements placed on the *LMaster* increase considerably as it has to be able to deal with all possible scenarios without any “help” from *Workers*. This, in turn, substantially increases hardware requirements of the node it runs on. This requirement propagates also to the *LMirror*, which has to be as good as the *LMaster* as it may take its place at any moment. Furthermore, the *LMaster* may become the bottleneck of the team (it will not scale with increasing team-size), which was one of the important disadvantages of the centralized scheduling pointed above. However, our proposed system is to be adaptive. Thus, the *LMaster* that cannot handle the load will loose both clients and *Workers*. As a result its team will either decrease in size (to the size it can manage successfully), or disappear completely. We can thus assume that, as

the time passes, each *LMaster* will be able to establish size of the team it can manage. The question that remains open is, will the size of the team be large enough for this approach to be feasible in a long run? However, an answer to this question can be established only experimentally. Summarizing, we believe that the advantages of this approach outweigh its potential disadvantages. We are thus likely to pursue this approach as the first attempt at introducing job scheduling into agent teams in our system.

As far as the second question is concerned (when should the job be transferred to the selected *Worker*), taking into account the dynamic nature of the Grid, there is no easy answer to it. First, recall that we assume that each node can disappear without warning. This means that staging a job at node *X* as soon as it is contracted may result in a wasted effort if this node crashes. However, since the *LMaster* can crash as well, it means that data needs not only to be kept until it is released to the selected *Worker*, but also has to be mirrored by the *LMirror*. We will thus leave this question open, pending further analysis.

5.3.1 Monitoring in the system

One of the important issues in most schedulers is monitoring the workload. In the case of our system not only the workload of *Workers* has to be monitored. More importantly, their very existence needs to be established. Recall that we assume that Grid nodes can disappear without any warning. However, monitoring existence of nodes has been addressed in [47]. There, we have reported how we have implemented a mechanism (based on principles derived from network management) in which the *LMaster* is “pinging” its *Workers* and expects that they respond in time to a certain number of pings. Lack of response is an indication of a problem with the *Worker*.

Let us now assume that the *LMaster* monitors existence of team members. What it needs is information about their workloads. Since Java does not have a direct method to access the load information of the underlying system, an external library had to be used to fulfill this requirement. *Jsysmon* is a Java library (which works both in Linux and Windows) that permits Java applications to access system monitoring information, e.g. the CPU or the Memory usage [7]. This library is used by the *Worker* agent. Specifically, in the *Worker* agent, we have added a new behavior (*UsageReporterBehaviour*), which has been implemented as a subclass of the standard JADE *TickerBehaviour*. At predefined intervals, this behavior collects the load information from the local system (using the *Jsysmon* library commands), and sends it as an ACL message to the *LMaster*. At the same time, the *LMaster* has been extended by adding a workload data collecting behavior (*MonitorStatusBehaviour*), which is a subclass of standard JADE *CyclicBehaviour*, which receives the load messages from team members and stores them for future use.

6 Concluding remarks

The aim of this chapter was three-fold. First, to present a brief overview of Grid resource management techniques found in standard Grid middlewares. Second, we have considered attempts at utilizing software agents as a Grid middleware, and thus as resource brokers and job (meta-)schedulers. Finally, we have discussed how the knowledge gathered in the first two parts of the chapter influences our thinking about job scheduling within our system. We have realized, again, that time has come to infuse our system with a robust Grid ontology (see also, [39]). This ontology will not only allow us to utilize a broad range of possible SLA negotiation mechanisms. It will also provide the scheduler with the necessary information about the resources, the job and its execution constraints. Finally, analysis of available scheduling techniques performed within the context of our system (based heavily on the Grid economy based vision of the nature of Grid computing of the future) pointed out that at this stage we should proceed with making the *LMaster* an omnipotent manager and meta-scheduler, which has total command over *Workers* in its team. We will report on our progress in subsequent publications.

Acknowledgments

Work of the Polish team was in part supported from the “Funds for Science” of the Polish Ministry for Science and Higher Education for years 2008-2011, as a research project (contract number N N516 382434). Collaboration of the Polish and Bulgarian teams is partially supported by the *Parallel and Distributed Computing Practices* grant. Collaboration of Polish and French teams is partially supported by the PICS grant *New Methods for Balancing Loads and Scheduling Jobs in the Grid and Dedicated Systems*. Collaboration of the Polish and Russian teams is partially supported by the *Efficient use of Computational Grids* grant.

References

- [1] “About AppLeS”, <http://www.cs.ucsd.edu/groups/hpcl/apples/hetpubs.html#AppLeS>.
- [2] “Condor Manual”, <http://www.cs.wisc.edu/condor/manual/>.
- [3] “CSF”, http://www.globus.org/grid_software/computation/csf.php.
- [4] “FIPA Contract Net Interaction Protocol Specification”, <http://www.fipa.org/specs/fipa00029/SC00029H.html>.
- [5] “GRACE—GRid seArch & Categorization Engine”, <http://www.grace-ist.org/>.

- [6] “Homepage of the JADE project”, <http://jade.tilab.com/>.
- [7] “Jsysmon—JAVA library for system monitoring”, <http://jsysmon.sourceforge.net/>.
- [8] “Load Sharing Facility”, <http://help.unc.edu/4484>.
- [9] “NetSolve(GridSolve) Overview”, <http://icl.cs.utk.edu/netsolve/overview/index.html>.
- [10] “OpenPBS homepage”, <http://www.openpbs.org/>.
- [11] “PBS Pro homepage”, <http://www.pbsgridworks.com/>.
- [12] “Sun Grid Engine Project”, <http://www.sun.com/software/gridware/>.
- [13] “Welcome to the FIPA”, <http://www.fipa.org/>.
- [14] “Wikipedia: Aglets”, <http://en.wikipedia.org/wiki/Aglets>.
- [15] “The WS-Resource Framework”.
- [16] R. Aversa, B. Di Martino, N. Mazzocca, S. Venticinque, “MAGDA: A Mobile Agent based Grid Architecture”, *Journal of Grid Computing*, 4(4): 395–412, 2006, 10.1007/s10723-006-9049-1.
- [17] R. Aversa, B. Di Martino, N. Mazzocca, S. Venticinque, “A skeleton based programming paradigm for mobile multi-agents on distributed systems and its realization within the MAGDA Mobile Agents platform”, *Mobile Information Systems*, 4(2): 131–146, 2008.
- [18] C. Badica, M. Ganzha, M. Paprzycki, “Implementing Rule-Based Automated Price Negotiation in an Agent System”, *Journal of Universal Computer Science*, 13(2): 244–266, 2007.
- [19] C. Bartolini, C. Preist, N. Jennings, “Architecting for Reuse: A Software Framework for Automated Negotiation”, in *Proceedings of AOSE*, Volume 2585, pages 88–100. Springer, Berlin, 2002.
- [20] M. Beck, H. Casanova, J. Dongarra, T. Moore, J. Plank, F. Berman, R. Wolski, “Logistical quality of service in NetSolve”, *Computer Communications*, 22(11): 1034–1044, 1999.
- [21] H. Benameur, B. Chaib-draa, P. Kropf, “Multi-item auctions for automatic negotiation”, *Journal of Information and Software Technology*, 44: 291–301, 2002.

- [22] L. Boloni, K. Jun, K. Palacz, R. Sion, D. Marinescu, “The Bond Agent System and Applications”, in *Proceedings of the Second International Symposium on Agent Systems and Applications and Fourth International Symposium on Mobile Agents*, Volume 1882, pages 99–112. Springer-Verlag, London, UK, 2000.
- [23] L. Boloni, D. Marinescu, “An object-oriented framework for building collaborative network agents”, in *Intelligent systems and interfaces*, pages 31–64. Kluwer Academic Publishers, Norwell, MA, USA, 2000.
- [24] H.J. Burckert, K. Fischer, G. Vierke, “Holon Transport Scheduling with Teletruck”, *Applied Artificial Intelligence*, 14(7): 697–725, August 2000.
- [25] S. Bussmann, K. Schild, “An Agent-Based Approach to the Control of Flexible Production Systems”, in *Proc. of the 8th IEEE Int. Conf. on Emergent Technologies and Factory Automation (ETFA 2001)*, Volume 2, pages 481–488. IEEE CS Press, Los Alamitos, CA, 2001.
- [26] R. Buyya, D. Abramson, J. Giddy, H. Stockinger, “Economic models for resource management and scheduling in Grid computing”, *Concurrency and Computation: Practice and Experience*, 14(13–15): 1507–1542, 2002.
- [27] R. Buyya, J. Giddy, D. Abramson, “An Evaluation of Economy-based Resource Trading and Scheduling on Computational Power Grids for Parameter weep Applications”, in *Proceedings of the Second Workshop on Active Middleware Services (AMS 2000)*. Kluwer Academic Press, Pittsburgh, USA, August 2000.
- [28] J. Cao, “ARMS: An agent-based resource management system for grid computing”, *Scientific Programming*, 10(2): 135–148, 2002.
- [29] J. Cao, “Self-organizing agents for grid load balancing”, in *Proceedings of the Fifth IEEE/ACM International Workshop on Grid Computing (GRID’04)*, 2004.
- [30] J. Cao, S.A. Jarvis, S. Saini, G.R. Nudd, “GridFlow: Workflow Management for Grid Computing”, in *3rd IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID’03)*, 2003.
- [31] P. Chandra, A. Fisher, C. Kosak, T. Ng, P. Steenkiste, E. Takahashi, H. Zhang, “Darwin: Customizable Resource Management for Value-Added Network Services”, in *Network Protocols, IEEE International Conference on*, pages 177–188. IEEE Computer Society, Los Alamitos, CA, USA, Oct 1998.
- [32] S. Chapin, D. Katramatos, J. Karpovich, A. Grimshaw, “The Legion Resource Management System”, in *Job Scheduling Strategies for Parallel Processing*, pages 162–178. 1999.
- [33] K. Czajkowski, I. Foster, C. Kesselman, “Agreement-based resource management”, Volume 93(3), pages 631–643, 2005.

- [34] M. Dominiak, M. Ganzha, M. Paprzycki, “Selecting grid-agent-team to execute user-job—initial solution”, in *Proceedings of the Conference on Complex, Intelligent and Software Intensive Systems*, pages 249–256. IEEE CS Press, Los Alamitos, CA, 2007.
- [35] M. Dominiak, W. Kuranowski, M. Gawinecki, M. Ganzha, M. Paprzycki, “Utilizing agent teams in grid resource management—preliminary considerations”, in *Proceedings of the IEEE J. V. Atanasoff Conference*, pages 46–51. IEEE CS Press, Los Alamitos, CA, 2006.
- [36] F. Dong, S.G. Akl, “Scheduling Algorithms for Grid Computing: State of the Art and Open Problems”, Technical report, Queen’s University School of Computing, January 2006.
- [37] J. Dongarra, “NetSolve: A network server for solving computational science problems”, *The International Journal of Supercomputer Applications and High Performance Computing*, 11(3): 212–223, 1997.
- [38] M. Drozdowicz, M. Ganzha, W. Kuranowski, M. Paprzycki, I. Alshabani, R. Olejnik, M. Taifour, M. Senobari, I. Lirkov, “Software Agents in ADAJ: Load Balancing in a Distributed Environment”, *Applications of Mathematics in Engineering and Economics’34*, pages 527–540, 2008.
- [39] M. Drozdowicz, M. Ganzha, M. Paprzycki, R. Olejnik, I. Lirkov, P. Telegin, M. Senobari, “Ontologies, Agents and the Grid—an Overview”, in *Proceedings of the PARENG’2009 Conference*, 2009, in press.
- [40] L. Ferreira, M. Batista, S. Fibra, C.Y. Lee, C.A.Q. Silva, J. Almeida, F. Lucchese, N. Keung, *Grid Computing Products and Services*, IBM Redbooks, 2005.
- [41] I. Foster, N. Jennings, C. Kesselman, “Brain Meets Brawn: Why Grid and Agents Need Each Other”, *Autonomous Agents and Multiagent Systems, International Joint Conference on*, 1: 8–15, 2004.
- [42] I. Foster, C. Kesselman, *The Grid: Blueprint for a Future Computing Infrastructure*, Morgan Kaufmann Publishers Inc., 1999.
- [43] I. Foster, C. Kesselman, S. Tuecke, “The Anatomy of the Grid: Enabling Scalable Virtual Organizations”, *International Journal of High Performance Computing Applications*, 15(3): 200–222, 2001.
- [44] M. Ganzha, M. Paprzycki, I. Lirkov, “Trust Management in an Agent-based Grid Resource Brokering System—Preliminary Considerations”, in M. Todorov (Editor), *Applications of Mathematics in Engineering and Economics’33*, Volume 946 of *AIP Conf. Proc.*, pages 35–46. American Institute of Physics, College Park, MD, 2007.

- [45] J. Gehring, A. Streit, “Robust Resource Management for Metacomputers”, in *HPDC*, pages 105–112, 2000, URL citeseer.ist.psu.edu/gehring00robust.html.
- [46] F. Kon, R. Campbell, M. Mickunas, K. Nahrstedt, F. Ballesteros, “2K: A Distributed Operating System for Dynamic Heterogeneous Environments”, in *HPDC’00: Proceedings of the 9th IEEE International Symposium on High Performance Distributed Computing*, pages 201–210. IEEE Computer Society, Washington, DC, USA, 2000, ISBN 0-7695-0783-2.
- [47] W. Kuranowski, M. Ganzha, M. Paprzycki, I. Lirkov, “Supervising Agent Team an Agent-based Grid Resource Brokering System—Initial Solution”, in F. Xhafa, L. Barolli (Editors), *Proceedings of the Conference on Complex, Intelligent and Software Intensive Systems*, pages 321–326. IEEE CS Press, Los Alamitos, CA.
- [48] W. Kuranowski, M. Paprzycki, M. Ganzha, M. Gawinecki, I. Lirkov, S. Margenov, “Agents as resource brokers in grids—forming agent teams”, in *Proceedings of the LSSC Meeting*, Volume 4818 of *LNCS*. Springer, Berlin, 2007.
- [49] E. Laskowski, M. Tudruj, R. Olejnik, B. Toursel, “Bytecode Scheduling of Java Programs with Branches for Desktop Grid”, *Future Generation Computer System (FGCS) -The International Journal of Grid Computing: Theory, methods and Applications*, 23(8): 977–982, November 2007.
- [50] B.Q. Li, J.C. Zeng, M. Wang, G.M. Xia, “A negotiation model through multi-item auction in multi-agent system”, in *Machine Learning and Cybernetics, 2003 International Conference on*, Volume 3, pages 1866–1870, 2003.
- [51] J. Li, R. Yahyapour, “Negotiation Strategies for Grid Scheduling”, in *Advances in Grid and Pervasive Computing*, pages 42–52. 2006.
- [52] Y. Li, Z. Lan, “A Survey of Load Balancing in Grid Computing”, *Computational and Information Science*, pages 280–285, 2005.
- [53] H. Nakada, M. Sato, S. Sekiguchi, “Design and Implementation of Ninf: Towards a Global Computing Infrastructure. Future Generation Computing Systems (Metacomputing Special Issue)”, 1999.
- [54] A. Negri, A. Poggi, M. Tomaiuolo, “Intelligent Task Composition and Allocation through Agents”, pages 255–260, 2005.
- [55] A. Negri, A. Poggi, M. Tomaiuolo, P. Turci, “Dynamic Grid tasks composition and distribution through agents”, *Concurrency and Computation: Practice and Experience*, 18(8): 875–885, 2006.

- [56] G. Nudd, D. Kerbyson, E. Papaefstathiou, S. Perry, J. Harper, D. Wilcox, “Pace—A Toolset for the Performance Prediction of Parallel and Distributed Systems”, *Int. J. High Perform. Comput. Appl.*, 14(3): 228–251, 2000, ISSN 1094-3420.
- [57] R. Olejnik, E. Laskowski, B. Toursel, M. Tudruj, I. Alshabani, “DG-ADAJ: a Java Computing Platform for Desktop Grid”, in K.W. Marian Bubak, Michal Turala (Editor), *Cracow Grid Workshop '05 Proceedings*. Academic Computer Centre CYFRONET AGH, Cracow, Poland, April 2006.
- [58] M. Paprzycki, M. Ganzha, “Adapting Price Negotiations to an E-commerce System Scenario”, in K. Saeed, et.al. (Editors), *Proceedings of the CISIM Conference*, pages 380–386. IEEE CS Press, Los Alamitos, CA, 2007.
- [59] A. Poggi, M. Tomaiuolo, P. Turci, “Extending JADE for Agent Grid Applications”, in *Proceedings of the 13th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises*, 2004.
- [60] A. Poggi, M. Tomaiuolo, P. Turci, “An Agent-Based Service Oriented Architecture”, in *Proceedings of the WOA'07*, pages 157–165, 2007.
- [61] R. Prodan, T. Fahringer, “From Web Services to OGSA: Experiences in Implementing an OGSA-based Grid Application”, *Grid*, 00: 2, 2003, ISBN 0-7695-2026-X.
- [62] K.K. Rajkumar Buyya, M. Maheswaran, “A taxonomy and survey of grid resource management systems for distributed computing”, *Software Practical Experience*, 32(2): 135–164, 2002.
- [63] K. Sakamoto, H. Sato, “A Resource-Oriented Grid Meta-Scheduler Based on Agents”, in *Proceedings of the 25th IASTED International Multi-Conference, Parallel and Distributed Computing and Networks*. Innsbruck, Austria, 2007.
- [64] B. Schnizler, *Resource Allocation in the Grid, A Market Engineering Approach*, PhD thesis, 2007.
- [65] M. Senobari, M. Drozdowicz, M. Paprzycki, W. Kuranowski, M. Ganzha, R. Olejnik, I. Lirkov, “Combining an JADE-agent-based Grid infrastructure with the Globus middlewareInitial Solution”, in M. Mohammadian (Editor), *Proceedings of the CIMCA-IAWITC 2008 Conference*, pages 890–895. IEEE CS Press, Los Alamitos, CA, 2008.
- [66] C. Smith, “Open source metascheduling for Virtual Organizations with the Community Scheduler Framework (CSF)”, 2003, Technical Whitepaper.
- [67] Z. Tan, *Market-Based Grid Resource Allocation Using A Stable Continuous Double Auction*, PhD thesis, 2007.

- [68] H. Tianeld, R. Unland, “Towards self-organization in multi-agent systems and Grid computing”, *Multiagent and Grid Systems*, 1(2): 89–95, 2005.
- [69] M. Wooldridge, *An Introduction to Multiagent Systems*, John Wiley & Sons (Chichester, England), 2002.