

DYNAMIC SLA NEGOTIATION BASED ON WS-AGREEMENT

Antoine Pichot

Alcatel-Lucent, Route De Villejust, 91620 Nozay, France

Oliver Wäldrich, Wolfgang Ziegler

Fraunhofer Institute SCAI, Department of Bioinformatics, 53754 Sankt Augustin, Germany

Philipp Wieder

IT and Media Center, University of Dortmund, 44221 Dortmund, Germany

Keywords: Dynamic Service Level Agreement, Negotiation, Web Services Agreement, Commit Protocol.

Abstract: In Grid, e-Science and e-Business environments Service Level Agreements (SLAs) are often used to establish frameworks for the delivery of services between service providers and the home organisations of the customers. While this high level SLAs, e.g. framework contracts, govern the relationship between the parties, it is desirable for the end-user to have dedicated quality of service (QoS) also for individual services like the orchestration of resources necessary for composed services. Grid level scheduling services typically are responsible for the orchestration and coordination of resources in the Grid, which requires the Grid level scheduler to coordinate resource management systems probably located in different domains. As the site autonomy has to be respected negotiating with the individual sites is the only way to achieve the intended coordination. Electronic SLAs emerged as a new way to negotiate and manage usage of resources in the Grid and are already adopted by a number of management systems. Therefore, it is natural to look for ways to adopt SLAs for Grid level scheduling. To realise this approach we need efficient and flexible protocols supporting dynamic negotiation and creation of SLAs. In this paper we propose and discuss ongoing work on extensions to the WS-Agreement protocol addressing these issues.

1 INTRODUCTION

A Service Level Agreement (SLA) is a contract between a service provider and its customer that describes the service, terms, guarantees, responsibilities and service-level to be provided. In Grids, e-Science and e-Business environments, SLAs are often used to establish service-delivery frameworks between a service provider and a service consumer domain. Such "high-level" SLAs usually define the framework of service provisioning and overall Quality of Service (QoS). However, it is desirable and sometimes even essential for the end-user to have dedicated QoS also for individual services, e.g. in a case of a high-priority service request.

In this article we focus on electronic SLAs dynamically negotiated and created by software programs on behalf of end-users. The use case we consider is resource provisioning, i.e. an SLA-based service provision implies access to various kinds of resources, like

computing nodes, network connections, storage areas or any combination of these.

Resource consumption varies over time and is sometimes dependent on the successful completion of previous tasks. An orchestrator communicates on behalf of customers end-users with several local resource managers to negotiate and create dynamic SLAs. In the rest of the document, for reasons of clarity, we will limit the problem scope to use-cases where computing and network resources are needed, and to a Grid scheduler as orchestrator that has to negotiate, select and schedule resources in order to execute a user's job and fulfil its requirements. As we will see, co-ordinating the access to multiple resources at the same time requires specific protocol features that negotiation and agreement protocols do not necessarily have. Fig. 1 describes the different steps to be performed by a Grid scheduler to co-allocate

resources: resource filtering, SLA negotiation, SLA creation, and job submission. This article presents ongoing work in the Grid Resource Allocation Agreement Protocol working group (GRAAP-WG, <<https://forge.gridforum.org/projects/graap-wg/>>) of the Open Grid Forum. The group is now working on protocol extensions for the WS-Agreement (Andrieux et al., 2007) specification supporting the SLA negotiation and creation phases.

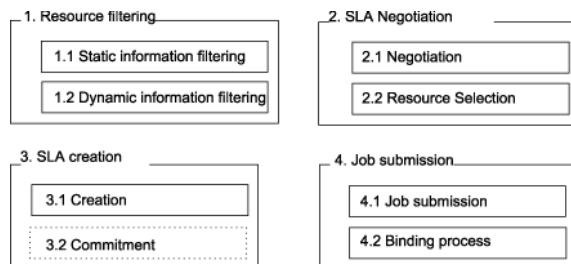


Figure 1: Resource selection & reservation.

In the following section we introduce work related to negotiation of SLAs. Section 3 describes SLA negotiation and creation problems, how distributed databases' commit protocol shed light to this problem. Section 4 describes how these problems can be solved using the Web Services Agreement protocol and Section 5 concludes the paper.

2 RELATED WORK

While WS-Agreement, since May 2007 the proposed recommendation of the Open Grid Forum for a language and a protocol to create SLAs, was becoming stable over the last two years, the interest in using SLAs for resource negotiation was growing constantly. One indicator is the number of projects and developments around resource management and scheduling considering or already using SLAs. In the following we briefly discuss the different approaches.

Some of the projects, e.g. CATNETS (Joita and Rana, 2006), openCCS (Keller, 2007), AgentScape (Mobach et al., 2006), the Community Scheduler Framework (CSF, <<http://sourceforge.net/projects/gcsf/>>), and the VIOLA MetaScheduling Service (MSS) (Wieder et al., 2005), are using WS-Agreement to create the SLAs. Some plan to use WS-Agreement, e.g. the Grid Resource Management System (GRMS, <<http://www.gridge.org/>>) and GridWay (<<http://www.gridway.org/>>) are in the planning phase for integrating WS-Agreement. Finally, some have not yet decided which technology to use or

are not using a web services stack at all: eNanos (<<http://www.bsc.es/grid/enanos/>>), Grid superscalar (Badia et al., 2006) and the Highly-Available Resource Co-allocator (HARC) (MacLaren, 2007).

CATNETS uses a market oriented mechanism of offers and bids to determine the price of a resource and to create an SLA based on the result, thus no negotiation is used. openCCS, AgentScape and MSS are Grid schedulers and use negotiation to refine offers and requests in order to create SLAs. In case of MSS, this includes negotiation with several resource providers to co-allocate multiple resources or support workflow execution across multiple resources. As WS-Agreement does not include a protocol for negotiating the terms of an SLA, the three approaches currently use proprietary extensions of WS-Agreement for the negotiation. GRMS and GridWay currently do plan not use WS-Agreement for SLAs between provider and customer, but plan to use WS-Agreement for the communication between multiple co-operating Grid schedulers. This will allow to negotiate usage of resources that are not part of the environment of the the initially accessed Grid scheduler. Scenarios and the corresponding requirements for the protocol are developed in the Grid Scheduling Architecture research group (GSA-RG, <<https://forge.gridforum.org/projects/gsa-wg/>>) of the Open Grid Forum. The developers of the eNanos Grid resource broker and the Grid superscalar programming environment are currently evaluating the technology for the integration of SLAs into their tools. In contrast, the HARC developers decided not to use a web services stack and do the negotiation based on plain XML-messages over HTTPS.

Moreover, the GRAAP Working Group has started work on extending WS-Agreement to support negotiation of agreements for use cases where the simple step creation are not sufficient. In this paper we discuss some general problems of negotiation in distributed environments and present first ideas how negotiation could be implemented on top of WS-Agreement.

3 NEGOTIATION, CREATION AND COMMIT PROTOCOLS

In this section negotiation is briefly discussed followed by a presentation of commit protocols in distributed databases in Section 3.1 and commit protocols for distributed resource management systems in Section 3.2.

To run a job that requires several resources, like networking and computational resources, managed by

different resource management systems (RMS), several steps must be performed by a Grid scheduler. Upon receipt of the job request, the scheduler starts the first phase: resource filtering based on static information and dynamic information. Static information does not change over time: number of CPUs, operating system, location, etc. Dynamic information changes over time: availability, load, etc. The second phase is the negotiation process and results in the selection of resources that can satisfy the job request. The third phase is the SLA creation phase concluded by the commitment of all service providers (or local RMS) involved leading to an reservation of the negotiated resources as described in the SLA. The last phase is the job submission followed by the execution.

Negotiation is a widely studied topic and there are numerous publications addressing different aspects, e.g. (Shakun, 2002) is a general purpose negotiation journal, (Briquet and de Marneffe, 2006) is a survey about negotiation in distributed resource management systems, while (Kuo et al., 2006b) and (Kuo et al., 2006a) discuss aspects of service negotiation in the Grid. In our context and in the simplest case, a user's job has to be executed and the Grid scheduler has to select between different target systems. If all systems are identical and only one parameter influences the selection, i.e. price, this case is similar to a typical business negotiation between one buyer and several sellers. An auctioning mechanism like the ones described in (Buyya, 2002) can be used. Of course, we take the point of view of an end user, if we look at things from a resource provider's point of view, we have several jobs that compete for one resource, i.e. several buyers and one seller. If we look at the scheduler's point of view, we have many jobs that compete for several resources, i.e. many buyers and many sellers. Buyya (Buyya, 2002) (page 36) also surveyed several distributed resource management systems based on price.

Automatic negotiation of SLAs is a complex and time consuming process (Jennings et al., 2001; Shen et al., 2002; Green, 2004), when even two users have to find an agreement on multiple criteria. Imagine how difficult the problem becomes when multiple entities have to reach an agreement (Czajkowski et al., 2002). When at least two resources are needed at the same time to run a job, e.g. a network connection and a processing resource, several steps have to be performed before reaching an agreement between the resource providers and the consumer. Green (Green, 2004) cites mainly two frameworks for automatic negotiation: ontologies and web services. According to him automated negotiation has three main considerations: The negotiation protocol, the negotiation ob-

jects and the decision-making models. He considers two options existing in order to achieve this type of negotiation. One option is for the originating agent to negotiate separately with each Autonomous System (AS) along each potential path to ensure that an end-to-end path is available. The dominant choice however, is to use a cascaded approach where each AS is responsible for the entire path downstream of itself. This approach enhances agent autonomy as it is only responsible for its immediate links. The autonomy of the cascaded approach struggles however with the issue of price. In a cascading scenario an intelligent agent would need to know the utility functions of all the downstream domains if the best price combination is to be determined, which is private information. In contrast, in this paper we limited the scope to protocols that permit the negotiation of agreements between two parties based on WS-Agreement rather than tackling the full complexity of automated negotiation. These bilateral agreements might then be combined into one single agreement.

3.1 Commit Protocols for Distributed Databases

Distributed transactional systems have been widely studied. One of their objectives is to propagate a consistent state across several systems, in a way that at any time all systems can show a consistent state to users. The consistent state or consistent view maintains and propagates between systems a logical coherent state. To provide crash recovery, several operations are logically grouped into transactions. Those transactions permit the change from one consistent view to another. For instance, you do not credit a bank account if you have not debited another bank account. However, these are two independent operations. A bank's distributed database system must group these two operations in one transaction. Thus it permits the change from one consistent state "before the transfer" to another "after the transfer". Database state changes are visible by other users once a transaction is committed to the system. In distributed systems, each transaction can impact several different systems not co-located. Thus distributed database experts have developed commit protocols (Bhargava, 1987; Kohler, 1981; Oszu and Valduriez, 1991). As Skeen described in (Skeen, 1981), "The processing of a single transaction is viewed as follows. At some time during its execution, a commit point is reached where the site decides to commit or to abort the transaction. A commit is an unconditional guarantee to execute the transaction to completion, even in the event of multiple failures. Similarly, an abort is an uncondi-

tional guarantee to “back out” the transaction so that none of its results persist. If a failure occurs before the commit point is reached, then immediately upon recovering the site will abort the transaction. Commit and abort are irreversible.”

When a user needs to make a change in a distributed database, a co-ordinator will propagate this change on all systems. As Skeen explains, upon receipt of a change request the co-ordinator forwards it to all distributed systems. Upon the change request receipt, all slaves go to the wait state. Then they can decide whether or not to accept this change, and send their response. The co-ordinator collects all responses to the change request, if one of them is negative, it goes in the abort state and sends an “abort” to all systems, if all responses are positive, then the co-ordinator goes in the commit state, and sends a “Commit” to all systems. Upon receiving a “Commit” (respectively “Abort”) all systems must commit (respectively “Abort”) the change request. Fig. 2 (left) represents a slave’s two phase commit protocol finite state machine (FSM). This process is the two phase commit process, supported by a two phase commit protocol.

The problem of this process is that in case of system failure. It’s impossible to know whether the transaction was committed or aborted. The wait state leads to both commit and abort state. For instance, when the co-ordinator fails after having sent a “commit” to some slaves but not all, the remaining slaves are blocked since they cannot know whether the transaction should be aborted or cancelled.

To provide crash recovery, and avoid blocking problems, Skeen introduced a three phase commit protocol. He added an intermediary state before the commitment as shown in Fig. 2 (middle). This state corresponds to a prepare to commit. It’s impossible to jump from this state to an abort state. He proved that if a state transition was possible between the prepare and the abort state, the protocol would be blocking. As a consequence, from any state on the slave’s finite state machine it is possible to determine whether the transaction should be committed or aborted in case of failure. In case of failure a slave in the “Wait” state must abort, while a slave in the “Prepare” state must commit.

3.2 Commit Protocols for Distributed Resource Management Systems

In an environment with distributed RMS providing guarantees on resource usage, a Grid scheduler may create SLAs with its users. In a co-allocation use case, this SLA takes into account several resources

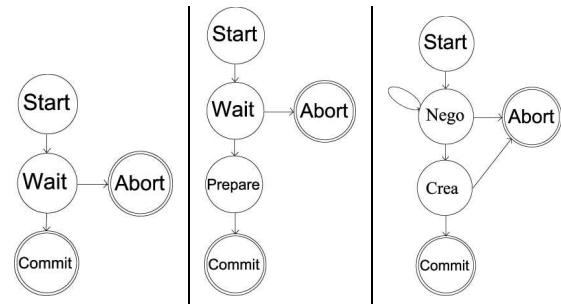


Figure 2: Two phase commit slave’s FSM (left), three phase commit slave’s FSM (middle), and SLA negotiation and creation resource provider’s FSM (right).

coming from several resource providers. With each independent resource provider a bilateral SLA has to be negotiated and created. A Grid scheduler has to create these bilateral SLAs on behalf of its users. For instance, in VIOLA, users may request network and computational resources with a dedicated QoS. The Grid scheduler has to orchestrate the individual reservation of network and computational resources. These two reservations are realised as two bilateral SLAs. The essence of distributed databases’ commit protocol is the transaction: a group of individual operations linked logically. In a distributed resource management system, co-allocation requires multiple bilateral SLAs. For a user or a particular service requiring multiple resources, either all of the individual bilateral SLA must be created, or none. The user SLA creation process is a transaction composed of multiple bilateral SLA creation.

Before reaching an agreement, two steps must be performed: negotiation and creation. The negotiation process can involve all resource providers. Its results are input to a resource provider selection process. When two resources are needed, e.g. network and computing, even if the negotiation involves many compute resource providers, only one computational resource will be selected. For many resources offered, the negotiation process does not lead to an SLA creation process. This is the main reason why negotiation must neither obligate the provider nor the consumer of the SLA. However, the SLA negotiation and creation process should minimise the number of discarded agreement creation requests when it has been previously negotiated. This should occur only when there is a race condition: when two or more users are competing simultaneously for the same resource at the same time. The separation of agreement negotiation and agreement creation process and minimising the number of discarded agreement creations after negotiation are conflicting objectives.

One way to observe atomicity of the SLA creation

is to use a transaction and to rely on a two phase commit protocol. Once resources have been negotiated, the orchestrator starts the SLA creation process by sending an SLA creation request to the selected resource providers. Then each resource provider responds to the request with yes or a counter offer. If all providers agree, the orchestrator sends a commit reservation to all systems. Upon receipt of this message, the reservation is committed and the SLA created. Fig. 2 (right) shows this process.

When the resource provider receives an SLA negotiation offer, its state changes from “Start” to “Nego”. It then answers the negotiation offer by either accepting it or making a counter offer. In case of a counter offer, it stays in the “Nego” state. It can also abort the negotiation and proceed to the “Abort” state. Once the orchestrator decides to start the SLA creation process, upon receipt of the SLA creation request, the resource provider’s state changes to the “Crea” state. It stays there if it accepts the reservation otherwise it goes to the “Abort” state. The final “Commit” state is reached when it receives a “Commit” message from the orchestrator and that resources are reserved and made unavailable to the rest of the world. As mentioned above, this simple two phase commit scenario can lead to a race condition during the SLA creation process. While the resource provider is in the “Crea” state, other users see the previous consistent state where resources are still available. To prevent this, the “Crea” state could imply “locking” resources thus providing a pre-reservation for the transaction lifetime. This prevents other users from reserving the same resource at the same time. In case of a lock request, second users’ transaction must wait for the lock to be released.

Although the FSMs shown in the middle and on the right-hand side of Fig. 2 look similar, we cannot say that the SLA negotiation and creation process is a three phase commit. It is a blocking protocol as described by Skeen (Skeen, 1981). And it does not provide any guarantees against crashes. One could still imagine a non blocking SLA creation protocol relying on a three phase commit providing crash recovery. It will not be discussed in this article.

4 SLA NEGOTIATION AND CREATION WITH WS-AGREEMENT

In order to co-allocate different types of resources and/or resources from different domains, MSS has to negotiate SLAs for the required resources. The eas-

iest way of SLA negotiation is a one step process, where the context, subject and constraints of the negotiation problem are defined. The WS-Agreement protocol natively supports this kind of negotiation by the `getResourceProperties` method. This method returns a set of agreement templates representing acceptable agreement offers for an agreement provider. These agreement templates only provide hints on agreement offers which might be accepted by an agreement provider. They do not guarantee the agreement will be accepted. An agreement template defines one or more services that are specified by their *Service Description Terms* (SDT), their *Service Property Terms* (SPT), and their *Guarantee Terms* (GT). Additionally an agreement provider can constrain the possible values within the SDTs, SPTs, and GTs by defining appropriate creation constraints within the templates.




Figure 3: WS-Agreement one step negotiation.

The creation constraints in agreement template can be static or dynamic. Typical examples of a static creation constraints are the minimum and maximum numbers of CPU, nodes, or memory. As these are properties of computing systems that are not likely to change frequently agreement templates that only contain static information usually are not restricted in their lifetime.

Agreement templates can also contain more dynamic information. Such dynamic information can be used to e.g. restrict the guaranteed execution time of a given service based on the current resource availability. Since the availability of resources is likely to change frequently, templates that contain such dynamic components have a short lifetime. A Grid scheduler can use these dynamic templates to efficiently find suitable time slots in order to e.g. co-allocate resources.

However, it is not always desired to expose availability information, or sometimes it is even not possible to do this in a convenient way. A typical example for this is the creation of an SLA in the network domain. Here, it is simply not possible to include the availability information for all possible network paths in a domain within one single SLA template. This would make the templates far too complex and therefore practically unusable. Therefore, the ef-

efficient agreement on time constraints in SLAs in only one phase is simply not feasible in this case. More advanced multi-step negotiations are needed to solve this problem.

4.1 Negotiation of Agreement Templates

Negotiation requires an iterative process between the parties involved. To rely on WS-Agreement and minimise the extensions to the proposed standard, we suggest not to negotiate SLAs but to negotiate and refine the templates that can be used to create an SLA. Here, our focus is on the bilateral negotiation of agreement templates.

In the following scenario we describe how an agreement initiator (e.g. the Grid scheduler) negotiates agreement templates with two agreement providers (e.g. a network scheduler and a CPU scheduler). We propose a simple offer/counter offer model. In order to use this model in the WS-Agreement protocol, we propose a new function *negotiateTemplate*. This function takes one template as input (offer), and returns zero or more templates (counter offer). The negotiation itself is an iterative process. In the following scenario we describe a simple negotiation process. During the negotiation process we call the agreement initiator 'negotiation initiator'. Accordingly we refer to the agreement providers as 'negotiation responders'.

1. Initialisation of the negotiation process

First, the negotiation initiator initialises the process by querying a set of SLA templates from agreement providers. To do so, it sends a standard WS-Agreement message, *getResourceProperty* request, to agreement providers (not shown in Fig. 4). From this templates, the initiator chooses the most suitable one as a starting point for the negotiation process. This template defines the context of the subsequent iterations. All subsequent offers must refer to this agreement template. This is required in order to enable an agreement provider to validate the creation constraints of the original template during the negotiation process, and therefore the validity of an offer.

2. Negotiation of the template

After the negotiation initiator has chosen an agreement template, it will create a new agreement template based on the chosen one. The new created template must contain a reference to the originating template within its context. Furthermore, the agreement initiator may adjust the content of the new created template, namely the content of the service description terms, the service

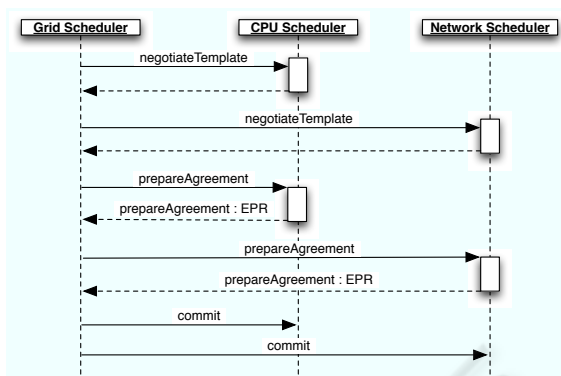


Figure 4: Extended WS-Agreement SLA negotiation.

property terms, and the guarantee terms. These changes must be done according to the creation constraints defined in the original template. Additionally, the negotiation initiator may also include creation constraints within the new created template. These constraints provide hints for the negotiation responder, within which limits the negotiation initiator is willing to create an agreement. For instance, the initial CPU scheduler template can contain “any number of 2GHz x586 CPU between 5pm to 6pm”. And the initiator can request “at least 5 1GHz x586 CPU anytime”. After the initiator created the new agreement template according to its requirements, the template is send to responders via a *negotiateTemplate* message (as shown in Fig. 4)

When the responder has received a *negotiateTemplate* message, it must first check the validity of the input document (refined template). This step includes (i) retrieving the original agreement template that was used to create the input document, (ii) validating the structure of the input document with respect to the originating template, and (iii) validating the changes of the content in the input document with respect to the creation constraints defined in the originating template.

Once this is done, the agreement provider now checks whether the service defined in the request could be provided or not. In our example, it’s only then that the CPU scheduler decides that 5 1GHz x586 CPUs can be provided. If the service can be provided, it just returns the agreement template to the client, indicating that an offer based on that template will potentially be accepted. Otherwise, the provider employs some strategy to create reasonable counter offers. During this process the agreement provider should take into account the constraints of the negotiation initiator. Counter offers are basically a set of new agreement tem-

plates that base on the template received from the negotiation initiator. The relationship between dynamic created templates and original ones must be reflected by updating the context of the new templates accordingly. After creating the counter offers the provider sends them back to the negotiation initiator (*negotiateTemplate* response).

3. *Post-processing of the templates*

After the negotiation initiator received the counter offers from the negotiation responder, it checks whether one or more meets its requirements. If there is no such template, the initiator can either stop the negotiation process, or start again from step 1. If there is an applicable template, the initiator validates whether there is need for an additional negotiation step or not. If yes, the initiator uses the selected template and proceeds with step 2, otherwise the selected template is used to create a new SLA.

4.2 SLA Creation

After the negotiation of an agreement template acceptable for both parties, the initiator needs to create the agreement. At this point, a problem similar to the transaction problem of distributed database systems arises. The goal of a Grid scheduler is to create a set of SLAs with different resource providers in order to provide co-allocation. Therefore, the scheduler first negotiates a set of templates with the providers, which identify the possible provisioning times of the required resources. However, we must not forget that templates only provide hints of what SLAs an agreement provider might accept. There is no guarantee associated with a template. This means that we are in need of a strategy to create all SLAs or none. In principle there are two major strategies to achieve this:

1. *to use transactions to create the SLAs, or*
2. *to create each SLA within one step, applying policies to the SLA.*

The usage of transaction mechanisms to create distributed SLAs, namely the usage of the two phase commit protocol, was already discussed in this paper. Since there is no support for two phase commit in WS-Agreement today, we need to extend the proposed standard to address this problem. This process has been started recently in the OGF working group that created WS-Agreement. A solution consists in adding a type of agreement that must be created in two phases: the first phase is a creation of the agreement triggered by a new *prepareAgreement* message and the second with a new non-standard *Commit* message as shown in Fig. 4.

The other approach is to create an SLA in one step using today's WS-Agreement functionalities, cancellation mechanisms and incentives. In order to realise this, we need to investigate the content of an SLA. On one hand, an SLA describes the service and its properties. On the other hand, it specifies the guarantees for a specific service. In a co-allocation scenario, where a Grid scheduler uses SLAs to co-ordinate e.g. network and computational resources, it employs execution guarantees in order to assure that the different services are provided at the same time. These guarantees may also include costs that are associated with the service if it is provided successfully, as well as penalties that arise when a guarantee is violated. However, an SLA might be prematurely terminated by the agreement initiator, before the service is actually provided. In fact, this is a cancellation of an SLA. When a service provider guarantees a certain execution time for a service, this normally comprises resource reservations. Therefore, the resource provider wants to prevent the termination of an existing SLA. This can be achieved by including a basic payment within the SLA. The basic payment is potentially a very small amount of money that is even charged if the SLA is terminated by the agreement initiator before the service was actually provided. It is therefore a termination penalty and represents the costs for the overhead produced by the resource reservation. In order to enable the Grid scheduler to efficiently negotiate and create SLAs, there could be a certain time period in which the SLA can be terminated without penalty. The duration of this period can dynamically be specified during the negotiation process. The Agreement provider could use a certain trust index in order to determine the maximum length of this period. This offers a feasible solution for the orchestration of multiple resources using the current one-step SLA creation of WS-Agreement.

5 CONCLUSIONS

In this paper we discussed basic functions for resource orchestration in Grids, namely mechanisms to negotiate and create Service Level Agreements using WS-Agreement. SLAs are a basic building block for Grid resource orchestration and distributed resource management. We have shown how a bilateral WS-Agreement based negotiation process is used to dynamically negotiate SLA templates. We proposed an extension of the WS-Agreement protocol, *NegotiateTemplate*, in order to support a simple offer/counter-offer model.

The second relevant part of the resource orchestra-

tion process is the creation of distributed SLAs. We have discussed two different strategies to co-allocate SLAs in Grids. One using a two phase commit with a WS-Agreement *Commit* extension and one using a single phase commit with SLA cancellation and incentives.

ACKNOWLEDGEMENTS

Some of the work reported in this paper is funded by the German Federal Ministry of Education and Research through the VIOLA project (grant #01AK605L), by the European Commission's IST programme through the CoreGRID Network of Excellence (grant #004265), and by the French Ministry of Industry through the CARRIOCAS project.

REFERENCES

- Andrieux, A. et al. (2007). Web Services Agreement Specification (WS-Agreement). Grid Forum Document GFD.107, Open Grid Forum.
- Badia, R., Sirvent, R., Labarta, J., and Perez, J. M. (2006). Programming the GRID: An Imperative Language-based Approach. In Di Martino, B., Dongarra, J., Hoisie, A., Yang, L. T., and Zima, H., editors, *Engineering The Grid: Status and Perspective*, chapter 12. American Scientific Publishers.
- Bhargava, B. (1987). *Concurrency and Reliability in Distributed Database Systems*. Van Nostrand Reinhold.
- Briquet, C. and de Marneffe, P.-A. (2006). Grid resource negotiation: survey with a machine learning perspective. In *PDCN'06: Proceedings of the 24th IASTED international conference on Parallel and distributed computing and networks*, pages 17–22, Anaheim, CA, USA. ACTA Press.
- Buyya, R. (2002). *Economic-based Distributed Resource Management and Scheduling for Grid Computing*, PhD Thesis. Monash University, Melbourne, Australia.
- Czajkowski, K., Foster, I., Kesselman, C., Sander, V., and Tuecke, S. (2002). SNAP : A protocol for negotiation of service level agreements and coordinated resource management in distributed systems. In *Proceedings of the 8th Workshop on Job Scheduling Strategies for Parallel Processing*, pages 153–183, Edinburgh, Scotland.
- Green, L. (2004). Service level negotiation in a heterogeneous telecommunication environment. In *Proceeding International Conference on Computing, Communications and Control Technologies (CCCT04)*, Austin, TX, USA.
- Jennings, N., Faratin, P., Lomuscio, A., Parsons, S., Sierra, C., and Wooldridge, M. (2001). Automated negotiation: Prospects, methods and challenges. *Group Decision and Negotiation*, 10(2):199–215.
- Joita, L. and Rana, O. (2006). WS-Agreement Use in CAT-NETS. Technical report, School of Computer Science and Welsh eScience Centre, Cardiff, UK.
- Keller, A. (2007). openCCS: Computing Center Software. Technical report, Paderborn Center for Parallel Computing, Paderborn, Germany.
- Kohler, W. (1981). A Survey of Techniques for Synchronization and Recovery in Decentralized Computer Systems. *ACM Computing Surveys*, 13(2):148–183.
- Kuo, D., Parkin, M., and Brooke, J. (2006a). A Framework & Negotiation Protocol for Service Contracts. In *Proceedings of the 2006 IEEE International Conference on Services Computing (SCC 2006)*, pages 253–256.
- Kuo, D., Parkin, M., and Brooke, J. (2006b). Negotiating contracts on the grid. In *Exploiting the Knowledge Economy - Issues, Applications, Case Studies, Volume 3, Proceedings of the eChallenges 2006 (e-2006) Conference*, Amsterdam, The Netherlands. IOS Press.
- MacLaren, J. (2007). Co-allocation of Compute and Network resources using HARC. In *Proceedings of "Lighting the Blue Touchpaper for UK e-Science: closing conference of ESLEA Project" (PoS(ESLEA)016)*.
- Mobach, D., Overeinder, B., and Brazier, F. (2006). A WS-Agreement Based Resource Negotiation Framework for Mobile Agents. *Scalable Computing: Practice and Experience*, 7 (1):23 – 36.
- Oszu, M. and Valduriez, P. (1991). *Principles of Distributed Database Systems*. Prentice Hall.
- Shakun, M., editor (2002). *Group Decision and Negotiation*. Springer Netherlands.
- Shen, W., Ghenniwa, H. H., and Wang, C. (2002). Adaptive Negotiation for Agent-Based Grid Computing. In *Proceedings of AAMAS2002 workshop on agentcities: Challenges in Open Agent Environments*, pages 32–36, Bologna, Italy.
- Skeen, D. (1981). Nonblocking Commit Protocols. In *Proceedings of ACM SIGMOD Int'l Conf. Management of Data*, pages 133–142.
- Wieder, P., Wäldrich, O., and Ziegler, W. (2005). A meta-scheduling service for co-allocating arbitrary types of resources. In *Proceedings of the 6th International Conference, Parallel Processing and Applied Mathematics, PPAM 2005*, volume 3911 of LNCS, pages 782 – 791, Poznan, Poland. Springer.