

Strategies for the Service Market Place

Paul McKee¹, Steve Taylor², Mike SurrIDGE², Richard Lowe², and Carmelo Ragusa³

¹ BT, Adastral Park Martlesham Heath, Ipswich IP5 3RE, UK
paul.mckee@bt.com

² University of Southampton IT Innovation Centre, 2 Venture Road, Chilworth, Southampton
SO16 7NP, UK

{sjt,ms,rl}@it-innovation.soton.ac.uk

³ Multimedia & Distributed Systems Lab (MDSLab), Department of Mathematics,
University of Messina, Contrada di Dio 98166 S. Agata – Messina, Italy
cragusa@unime.it

Abstract. We describe a number of strategies for a future service oriented market place. We describe the SLA's role within the service framework, and how it enables customers to make value judgements regarding the quality of a service. We also discuss the complexity of too much choice from both the customer and provider points of view, and advocate a "discrete offer" approach. We discuss the "cost of negotiation" and argue that it must be carefully balanced with the cost, value and risk of the offering being negotiated for. We add to the negotiation analysis with presentation and discussion of some results showing a simulated Grid market place and show that it is possible for service providers to deny themselves work through attempting to offer a high quality guaranteed service.

Keywords: Service Oriented Architecture, Service Oriented Infrastructure, Grid Computing, Business, Service Level Agreement.

1 Introduction

Businesses are faced with increasing pressure to introduce products and services in increasingly shorter time scales and at reducing costs. This has led to great interest in technologies such as service oriented architectures (SOA) and its supporting service oriented infrastructure (SOI), grid computing, and the new range of software as a service (SAAS) offerings. Using such technologies it is hoped that businesses will be able to rapidly create new applications, either for use internally or for sale, from services that may be assembled and executed rapidly and economically. To achieve maximum economic benefit it is essential that the services are reused and that, as the technology matures, fewer are developed from scratch. This all leads to the supposition that there will be a market place for applications and services that will have a global reach and potential consumers of services will have to choose from such a global marketplace. In this paper we present some observations about the structure and characteristics of such a service marketplace and suggest strategies that we believe will lead to greater user satisfaction.

For the purposes of this paper, we define a service as: “the doing of work for the benefit of others”. We use this generic definition so as to concentrate on the business aspects of the provision of service rather than a technical definition (for example a Web Service). By definition services are intangible, and may not even exist before they are required. This means that the potential purchaser has limited opportunity to assess the quality of a service before use. Hence there is a need for the consumer to determine the nature of a service and whether it will live up to their expectations. This may be achieved through the use of service level agreements (SLAs), described next.

2 The Role of SLAs

Within the 6th Framework IST project NextGRID we have proposed the use of bipartite (two-party) SLAs to describe both the functional and non-functional aspects of a service to allow consumers to make informed choices. An SLA is an agreement (as denoted by its name) and a contract for service is likely to be based on it, so the SLA describes to what each of its signatories is obliged to do in order to be compliant with the agreement (see [8]). The SLAs under development in NextGRID (also see [13]) are significantly different from what could be regarded as conventional SLAs in two main ways:

- NextGRID SLAs are expressed in terms of *customer business benefit*; and
- NextGRID SLAs may impose *interdependent obligations* on either party.

These differences are discussed in detail below.

We believe it is essential that the SLA have a strong end user benefit component, and that services being offered should be described in terms of their impact at the business level of the consumer, and not in terms of the technology supplied. The marketing benefits of this should be obvious – it is much easier to convince a financial decision maker of the benefits of a service that guarantees one of his key business processes using agreed metrics rather than one that guarantees the availability of, for example, a number of servers. We believe that the strong linkage to business impact benefits both the customer and provider, and this strong business linkage in the SLA structure is discussed in a previous publication [16]. The customer can then assign and manage costs transparently within their business and the service provider retains the flexibility in operations to strive to reduce the cost of delivery and ultimately the price paid by the customer. This view contrasts with the use of SLAs in the more established academic Grid community where the SLAs described in [2], [6] and [20] focus on resources (e.g. computers, network bandwidth, storage devices) rather than services. This is of course appropriate for a community of experienced users running often experimental codes. However, in order to address the needs of business (as opposed to technical) users the value of the service must be articulated at the appropriate – business – level rather than the resource level. The customer should not be concerned with the resources required to provide the service, just that the service exists and provides clear business benefit – management of resources should remain the domain of the service provider. Buco et al [1] concur with this view: “...*The customer need not know the implementation details of the provider’s service level management (SLM) processes...*”. Buyya et al [2] propose a “Grid Economy” in

which the focus is on the resources, low-level computational elements such as CPU, network and storage. Our position is that the value to the customer is at a higher level than the hardware. For example, application codes contain much more customer value than hardware, as they embody functionality that can solve customers' problems.

The second major characteristic of a NextGRID style SLA is that it recognises the fact that in a future service oriented marketplace application delivery will be collaborative and that SLAs need to recognise the impact of this collaboration. We advocate the notion of *interdependent obligations* that describe necessary pre-conditions in order for the SLA to be valid, and that some obligations will be dependent on others. For example, if a customer wishes to outsource payroll processing and asks a service provider to guarantee that the payroll processing will always be completed by the 29th of any month, it is entirely reasonable that the service provider place a pre-condition on the customer requiring that data be available 24hrs before the delivery deadline. The WSLA specification [19] also denotes the use of obligations, but does not determine that they may be dependent relationships, and how these relationships are dependent. Through the use of obligations the NextGRID SLA clearly sets out the expectations and obligations of both parties involved, and we believe is an essential component of a future electronic market place for services.

3 The Dynamic Service Market Place and Choice

Those who market products and services often assume that the more choice given to a customer, the more likely they are to find something that satisfies their requirements. At first glance it seems that this assumption is valid in a future service-based market place. We believe however that trying to offer a wide range of choices may turn out to be less successful than expected, and even counterproductive.

Consider the role of a service provider in such a market. In order to deliver a service against an SLA the service provider will need to characterise the performance of the service in various operational conditions and at all combinations of the various quality parameters the customers may choose. For this, the service provider must produce automated management tools to ensure performance is maintained for all permutations. This is of course extremely difficult, expensive and slow. The customer will see increased costs to pay for these developments, and a reduction in the number of new and innovative services due to the complexity of the task. Although limiting the number of available products may seem counterintuitive, the advantages for controlling costs may be significant. Gottfredson and Aspinall [10] discuss the notion of the "innovation fulcrum" and its impact on a number of manufacturing activities. They describe how careful selection of options and minimisation of complexity leads to a more efficient business and that each business has an optimal number of products. If a business offers too few products they may miss out on customers, while offering too many products leads to greater complexity and unacceptably high management costs. A balance between adequate market coverage and limitation of complexity must be sought.

As a service consumer the problems associated with too much choice may be equally severe. Faced with a service that contains a large number of potentially variable terms the end user is unlikely to be aware of the trade offs associated with

combinations of the various settings, and they may unwittingly specify a set of parameters that significantly increases the cost of the service without increasing the quality. Indeed the availability of a large number of parameters may in fact obscure which parameters are in fact important to them – as discussed by Twing [18] it is essential that SLA focus on critical service items that drive the business, and have clear financial implications. We regard this as further evidence supporting the need for SLAs to be described in terms of business benefits to the consumer. As the end user is a person, other non-technical problems may be observed. Iyengar and Lepper [14] have shown that when a customer has too much choice they do not always see this as a benefit and often they become less likely to buy anything at all, and the wider range of options often leads to dissatisfaction with the final choice. Customers often resort to very simple and non-optimal strategies to choose in such conditions, and this raises an interesting question for automated systems in that if the customer is unable to make the choice effectively in person, how can they program an automated system to achieve the best result for them?

4 Discrete Offers – The Supermarket Approach

If we consider the 20th century approach to shopping much of the complexity has been removed from commodity purchases by some degree of standardisation. Many household purchases have standard pack sizes that facilitate both automated production (thereby reducing costs) and efficient comparison between products by the consumer. It is fairly easy for the customer to make a value judgement between two packets of different branded soap powder, for example. In addition, the prices of the goods would greatly increase if the customer had total freedom to specify their features. For these reasons, within the NextGRID project we have advocated the use of a “discrete offer” system for services. It should be noted that by “discrete offer” we do not mean one single offer, but that every service provider should offer services with fixed quantised service levels, each constituting a separate, “discrete” offering of their service. Customers may then have the confidence that the combination(s) of quality parameters offered will at least work, and that the service provider has created a suitable management infrastructure for each level of the service, hopefully leading to reduced costs.

We recognise the fact that the commoditised approach will not be suitable for every situation and that there will be a continuum stretching from the simple purchasing of commodity services through to more complex fully negotiated contracts. The deciding factor governing which strategy is appropriate will be the cost incurred in the negotiation process versus the cost of the delivered service and this is discussed in the next section.

5 The Cost of Negotiation

The cost of negotiation comprises a significant part of what have been termed “transaction costs”, the roots of which were proposed in 1937 by Coase [5]. Transaction costs are (financial and otherwise) costs associated with the transaction

itself (for example costs associated with the process of procurement), not the item or service that is the subject of the transaction. We believe that the cost of negotiation must be carefully managed. Negotiation cost must be insignificant compared to the cost of the item or service being negotiated for. It is only worth spending much time, money and effort negotiating when the cost, value and risk associated with the item or service being negotiated for are significant.

In line with the discussion in Section 0, the supermarket approach has no negotiation. The customer chooses from a set of discrete offers and pays the labelled price. This is appropriate because the vast majority of real-world supermarket items are low-cost and commoditised. Conversely, a high-value, high-cost, complex service contract will have a great deal of negotiation, as the end item is extremely expensive and the risks of not carefully working out the contact between the customer and provider are significant. In the real world, a contract for the construction of a football stadium would not go ahead without careful negotiation from all parties concerned. The costs, value and risks are too great to leave negotiation to chance.

Applied to the Grid, there are a number of automated “discover, find and bind” protocols and demonstration systems (for example those described in [7] and [17]) some of which use software agents to perform the negotiation to gain access to the services being negotiated for. Whilst these are certainly useful, we assert that there must be a limit determined by the cost, risk and value of the negotiated item, above which humans must get involved. An organisation may be perfectly happy for software systems to take responsibility for low-value agreements, but humans must get involved when the stakes get higher, so that they can make more advanced decisions and take responsibility for their actions on behalf of their organisation. The greater the level of risk and investment, the greater the need there is for human decision making and responsibility. You cannot sue a computer because it made a promise your organisation cannot deliver!

There are a number of factors that determine the total cost of negotiation, and measures that can be taken to control costs arising from each.

1. There is a significant time cost to the customer of searching for suppliers. This implies that customers should search for suppliers infrequently, and record any approved suppliers in an “approved supplier list”.
2. Many negotiation protocols use an “invite-tender” approach. This is where an invitation to tender is issued by a customer and providers respond to it. There is usually a significant delay in the issuance of the invitation and the delivery of the tenders. The delays make this a high-cost approach, and in general this should only be used for high-cost bespoke services. It is useful for building long-term supply relationships e.g. when deciding who to include in an approved supplier list (see above).
3. The customer will have difficulty in comparing the value from offers from different providers – these are not likely to be directly comparable and thus the customer has to compare apples and oranges to determine which suits their needs best. This cost can be controlled by standardising the terms used in discrete offers – an approach used successfully by the supermarkets (see Section 4).

4. As discussed in Section 0, there is a significant risk to customer and provider of giving the customer too much freedom in choosing the service profile – what they choose may not be workable or be too costly to provide. Providers should restrict customers to an appropriate number of discrete offers.
5. There is a considerable time and risk cost to both customer and provider of iterative negotiation. Offers and counter-offers need to be computed and responded to, and there is the significant risk that there is no guarantee that iterative negotiation will converge. Providers can avoid this cost by making only discrete offers, allowing customers to choose between them but not to negotiate over their terms.
6. There is a considerable financial cost of any human interaction (decision making, responsibility, etc). Humans are much more costly than computers, and cannot handle large numbers of negotiations. Service providers and ideally also consumers should automate their procedures for managing offers and making agreements.

In conclusion, providers and customers need to analyse the costs and potential risks of any negotiation they enter into and appropriately target resources. Depending on the situation, too much or too little attention to negotiation and customer choice can be costly or dangerous.

6 Self Denial of Service

In this section we briefly describe simulation results of an initial study into the behaviour of a “Grid market place”. The overall aim of this work was to arrive at a means of simulating or modelling a Grid market place, given that there was no real-world information regarding how different actors operated in such a situation.

This work builds on earlier work in the GEMSS project [9]. This investigated use of the WSLA schema [19] for SLAs along with FIPA negotiation protocols [15], in a business to business Grid middleware based on GRIA [11] (the current GRIA software downloadable from [11] uses the discrete offer approach and does not employ FIPA protocols). The GEMSS middleware was deployed and evaluated using only a small number of sites. Here, we aimed to find out how such a B2B approach would work on a larger scale, with full market competition between larger numbers of sites.

Before starting the work, three options were considered for understanding Grid marketplace behaviour: direct measurement, analytical modelling and simulation. We chose a simulation based approach, as it was deemed to be the most flexible. Measurement was obviously not possible, and it was deemed that simulation allowed more “what-if” types of analysis to be performed than modelling.

We based the simulator code on the GridSim toolkit [3], [12] with appropriate extensions. There are two main actors in the simulation – users and service providers. The business goal for users is to get some (compute-based) work performed within a deadline. Hence, for the users we added a negotiation capability with service providers. The users issue requests for work consisting of a *workload* and a *deadline* by which the workload needs to be processed, evaluate responses from providers, and select the most appropriate based on a utility function. The workload is specified in

units relative to a hypothetical “standard machine”, and all actors in the simulation use the same metric and units to represent the demands of users and the capacity of providers. When evaluating offers from providers, the users differentiate between the offers using a *cost-time* selection strategy. The cost-time algorithm is described in [4], and has price as the major factor but run-time is taken into account where providers cost about the same.

The business goal of the service provider is to provide compute services that process the users’ workloads within the required timeframes. For the service provider we added two major extensions. Firstly, we added the capability to negotiate with users. The providers receive requests and make offers, containing the commitment to perform work for the user and a price. The offers expire after a specific time (denoted as O_p – the *offer validity period*), and users may only accept them within this time. Secondly, we eliminated the disclosure of status by service providers. In the original GridSim, the users are aware of the internal resource utilisation of the service providers. In a market situation, this is clearly not desirable for the service provider, as it gives the user an advantage in any negotiation. Here, the only information users are aware of is whether a provider can complete the work inside a given deadline, and what it will cost the user. There is no communication of the resources used by the provider, only whether the provider can satisfy the user’s requirements.

The scenario chosen models the following situation. In a future large scale dynamic market place it will be trivial for potential customers to request quotations from large numbers of suppliers at frequent intervals in the hope of spotting price trends, or windows of opportunity for low cost operation. It is easy to imagine automated software continually evaluating the providers for bargains. If the service providers commit their resources to offers when making those offers, so as to provide a high-quality, “guaranteed” service, they will be exposed to the significant risk of self-denial of service through committing resources for offers that are not taken up.

We chose a population of 80 users and 20 service providers for this scenario. The ratio between the number of clients and service providers is based on the assumption that in a marketplace of this type (SME clients and larger service providers) there will be significantly more clients than service providers. The population used here was dictated by the need to have enough providers to allow some variation between them, but not so many users that the simulation was impracticably slow.

The service provider population comprises different types, categorized in terms of the number of processing units, the price, and the speed of the processing units:

- **Size:** Big (9 units), Medium (3 units) and Small (1 unit);
- **Speed:** Fast (20 times a “standard” e.g. client machine) and Slow (10 times a standard machine); and
- **Price:** varies between 2½ cents and 7½ cents per hour on a ‘standard’ machine.

To create service providers’ profiles, all combinations of the above categories were initially considered. We then eliminated the non-realistic ones – for example we can see that a small, low priced service provider is not realistic because it will not create enough revenue to survive. We also excluded slow, low capacity, expensive service providers, since they are very unlikely to succeed in a competitive market. The service providers’ distribution thus reflects a marketplace in which large service

providers exploit the economies of scale to reduce their costs and therefore attract the majority of the business. However, we have included some medium and small providers as well to study their impact on the market. The final service providers' profiles are shown in Table 1.

Table 1. Service providers' profiles & parameters

Size - Speed - Price	Number of service providers	Number of processors	Processor rate (multiple of standard processor unit)	Price per processor - hour	Equivalent standard processor - hour price
Big - Slow - Low	3	9	10	€0.25	€0.025
Big - Fast - Low	2	9	20	€0.50	€0.025
Medium - Slow - Low	3	3	10	€0.25	€0.025
Medium - Fast - Low	1	3	20	€0.50	€0.025
Big - Slow - Very High	1	9	10	€0.75	€0.075
Big - Fast - High	2	9	20	€1.00	€0.050
Medium - Fast - High	3	3	20	€1.00	€0.050
Small - Fast - High	5	1	20	€1.00	€0.050

We used one key principal for service providers' operation. When making an offer, the service providers commit resources sufficient to process the workload requested by the user and by the user's deadline. The offer is valid for the duration of O_p . After this time, any resources committed for offers that are not accepted are released and can be offered again. Until then, resources committed to an offer may not be used to back up a new offer, thus ensuring that the service providers do not over-commit their resources.

The simulations were used to investigate the effect of varying O_p on the distribution of work across different providers. Figure 1 shows the situation where O_p is 3456 seconds and shows the amount of outsourced work processed by each of our population of 20 service providers (described in Table 1). In Figure 1, the service providers are denoted by a number and code:

SIZE{Big|Medium|Small} SPEED{Slow|Fast} PRICE{Low|High|V.High}

Due to the users' *cost-time* minimisation strategy (minimise cost then use time to sort any equal-cost providers) we expect the cheapest service providers to get most work. The fastest should also achieve the highest utilisation, since their resources are committed for the least time, and they should be most often in a position to make new offers. Figure 1 shows this effect: SP01 to SP05 inclusive are highly utilized, and within this set the fast providers (SP04 and SP05) are more highly utilised.

Naively, we would expect the work to be shared amongst all the low-priced providers, but we are seeing here that SP09 gets very little and SP06-SP08 get none, even though they are all low-priced. Yet some expensive providers (SP11 and SP12) are winning a great deal of work (they are above 80% utilised).




Fig. 1. Utilisation for different service provider types for $O_p = 3456$ seconds

This is due to the relative agility of these service providers. When providers make offers, they must commit resources for as long as the offers are valid. Hence the low-capacity (medium size and below) service providers commit all their resources much sooner than the larger ones and as a consequence cannot make as many offers as higher capacity providers. Thus the lower capacity providers are considerably less agile and unable to compete, even though their low price should make them attractive to customers.

If we reduce O_p we get a significant change in the work distribution. Figure 2 shows the results of a simulation where O_p was reduced to 1440 seconds.




Fig. 2. Utilisation for different service provider types for $O_p = 1440$ seconds

The major change is that now SP09 (medium size, fast and low-price) is winning much more work (it is now over 90% utilised). SP09 was previously prevented from winning much work because its resources were fully committed, and by the time they became free the more expensive but more agile SP11-SP12 had won most of the work. SP09 is now only committing its resources for 1440 seconds, and is thus able to make more offers and win much more work. The expensive providers SP11 and SP12 still do quite well, but lose out whenever SP09 is in a position to compete with them.

In Figure 3, O_p is further reduced, this time to 720 seconds, and we see another shift in the work distribution following the trend of the lower capacity workers winning work.




Fig. 3. Utilisation for different service provider types for $O_p = 720$ seconds

SP06-08 are now starting to win against SP11 and SP12. Even though they have very low capacity, the further reduction of O_p means that they can make enough offers in order to stand a good chance of some getting accepted.

Note that low capacity providers become more viable with shorter offer validity periods, even though the more expensive but more agile competitors also use the shorter offer validity period. We did not simulate different O_p for each provider, but it is clear that there will be some advantage in cutting it relative to competitors, though only to the point where customers find it too short.

Note also that shortening O_p further does not further improve the situation for the low-agility providers. By optimising the simulation it became possible to reduce O_p to only 5 seconds. Figure 4 shows that although SP06 to SP08 do a little better, they still cannot achieve the same utilisation as the more agile and expensive SP11 and SP12.




Fig. 4. Utilisation for different service provider types for Op = 5 seconds

We conclude that there is a significant business risk to the commitment of resources when making offers, as it makes a provider highly sensitive to the duration of the validity of the offer. The problem for the service providers is in the relationship between their provisioning and negotiation strategies – do they overbook their resources in the hope that some of their offers get turned down, or do they price their service accounting for the offers that are not taken up?

This is common in resource-intensive real business. It is well-known that airlines routinely overbook their planes to account for “no-shows”, and the terms of the contract they have with their customers states that the customers paying the cheaper prices (usually in economy class) may not get a seat on the plane (this is known as “bumping”).

It seems clear that the best way to operate services is by targeting a range of customer groups with a range of different service offers that provide different levels of guarantee (sometimes none at all) at different prices. These would have to be encoded in machine readable SLAs so negotiation can be automated on both sides. The SLAs will need to be much more sophisticated than hitherto, as there will be times when services are not taken up, or are taken up but then not delivered. This also implies a more sophisticated approach to manage the services and resources, in which failure to deliver service is an option, and decisions will be needed as to which services should be maintained. Clearly, there is a need to reduce complexity as much as possible, which suggests that the discrete offer approach should be used by service providers. It is also clear that commitment of resources during negotiation will add to the complexity, and that negotiation protocols should not depend on this.

7 Conclusions

In this paper, we have considered strategies for service oriented future markets. This work has come out of discussion, research and thinking into what is currently called

the “Grid”, and looks like it will evolve into a service oriented market, and hence the emphasis on services in this paper.

We have discussed the difficulty consumers have in judging the quality of an intangible an un-experienced service. The well-known solution is the service level agreement (SLA), and we describe the SLA’s role within the service framework. Our chief findings are twofold. Firstly, an SLA and service providers’ offerings must use terms understandable to and clearly having benefit for customers. Secondly, an SLA should contain interdependent obligations that *all* its signatories (not just the service provider) agree to.

We have advocated a “discrete offer” approach for service providers. We argue that if customers are given too much choice, they are likely to get overwhelmed by the permutations available to them, and may produce unworkable combinations of service factors. In addition, it has been argued that there is an “innovation fulcrum” - an optimum point between adequate market coverage and enough simplicity and manageability.

We have discussed the cost of negotiation and argue that it must be carefully balanced with the cost, value and risk of the offering being negotiated for. We have taken this further using simulations of a Grid market place, which show that it is possible for service providers to deny themselves work through attempting to offer a high quality guaranteed service. This illustrates that ‘transaction costs’ are not limited to the actual cost of a transaction process, but may also need to take account of the impact on business agility and opportunity costs. We conclude by determining that there should be different levels of service related to price targeting different customer needs, and presented to customers as discrete offers, with minimal negotiation and minimal pre-commitment of resources until an SLA is agreed.

Acknowledgments. The authors acknowledge the funding of the European Commission, and the work reported here was conducted within the NextGRID (The Next Generation Grid) Integrated Project, Contract Number 511563. This paper expresses the opinions of the authors and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this paper.

References

1. Buco, M.J., Chang, R.N., Luan, L.Z., Ward, C., Wolf, J.L., Yu, P.S.: Utility computing SLA management based upon business objectives. *IBM Systems Journal* 43(1), 159–178 (2004)
2. Buyya, R., Abramson, D., Venugopal, S.: The grid economy. In: *Proceedings of the IEEE*, vol. 93(3), pp. 698–714 (2005)
3. Buyya, R., Murshed, M.: GridSim: a toolkit for the modeling and simulation of distributed resource management and scheduling for Grid computing. *Concurrency and Computation: Practice and Experience* 14(13-15), 1175–1220 (2002)
4. Buyya, R., Murshed, M., Abramson, D.: A Deadline and Budget Constrained Cost-Time Optimization Algorithm for Scheduling Task Farming Applications on Global Grids. In: *Proceedings of the 2002 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA’02)*, June 24 - 27, 2002, Las Vegas, USA (2002)

5. Coase, R.: The Nature of the Firm. *Economica* 4(16), 386–405 (1937)
6. Czajkowski, K., Foster, I., Kesselman, C.: Agreement-Based Resource Management. *Proceedings of the IEEE* 93(3), 631–643 (2005)
7. Czajkowski, K., Foster, I., Kesselman, C., Sander, C., Tuecke, S.: SNAP: A Protocol for Negotiating Service Level Agreements and Coordinating Resource Management in Distributed Systems. In: Feitelson, D.G., Rudolph, L., Schwiegelshohn, U. (eds.) *JSSPP 2002*. LNCS, vol. 2537, pp. 153–183. Springer, Heidelberg (2002)
8. Davis, Jr., T.H., Fitzgerald, P.K.: Deconstructing Service Level Agreements. *New York Law Journal* (March 4, 2002)
9. GEMSS project, <http://www.gemss.de/>
10. Gottfredson, M., Aspinall, K.: Innovation Versus Complexity: What IS Too Much of a Good Thing? *Harvard Business Review* 83(11), 62–71 (2005)
11. GRIA: www.gria.org
12. GridSim toolkit, <http://www.buyya.com/gridsim/>
13. Hasselmeyer, P., Qu, C., Schubert, L., Koller, B., Wieder, P.: Towards Autonomous Brokered SLA Negotiation. In: Cunningham, P., Cunningham, M. (eds.) *Exploiting the Knowledge Economy: Issues, Applications, Case Studies*, IOS Press, Amsterdam (2006)
14. Iyengar, S.S., Lepper, M.R.: When choice is demotivating: Can one desire too much of a good thing? *Journal of Personality and Social Psychology* 79(6), 995–1006 (2000), <http://www.columbia.edu/~ss957/whenchoiceabstract.htm>
15. Middleton, S.E., Surridge, M., Benkner, S., Engelbrecht, G.: Quality of Service Negotiation for Commercial Medical Grid Services. In: *Journal of Grid Computing*, Springer, Heidelberg (to appear, 2007)
16. Mitchell, B., Mckee, P.: SLAs A Key Commercial Tool. In: Cunningham, P., Cunningham, M. (eds.) *Innovation and the Knowledge Economy: Issues, Applications, Case Studies*, IOS Press, Amsterdam (2006)
17. Shen, W., Li, Y., Ghenniwa, H., Wang, C.: Adaptive Negotiation for Agent-Based Grid Computing. In: Falcone, R., Barber, S., Korba, L., Singh, M.P. (eds.) *AAMAS 2002*. LNCS (LNAI), vol. 2631, pp. 32–36. Springer, Heidelberg (2003)
18. Twing, D.: Are you savvy about SLA negotiations? *Network World* (May 10, 2005)
19. WSLA: <http://www.research.ibm.com/wsla/>
20. Yeo, C.S., Buyya, R.: Service Level Agreement based Allocation of Cluster Resources: Handling Penalty to Enhance Utility. In: *Proceedings of the 7th IEEE International Conference on Cluster Computing (Cluster 2005)*, IEEE Computer Society Press, Los Alamitos (2005)