



SEOUL NATIONAL UNIVERSITY
College of Engineering
Technology Management, Economics and Policy Program

Using SLA Mapping to Increase Market Liquidity

**Marcel Risch, Ivona Brandic,
and Jörn Altmann**

TEMEP Discussion Paper No. 2009:37

서울대학교

공과대학, 기술경영경제정책대학원과정
151-742 서울시 관악구 신림동 산 56-1

Technology Management, Economics and Policy Program
College of Engineering, Seoul National University
San 56-1, Shillim-Dong, Kwanak-Gu, Seoul 151-742, KOREA
Phone: ++82-2-880-9140, Fax: ++82-2-880-8389

Technology Management, Economics and Policy Program (TEMEP) is a graduate program at Seoul National University. It includes both M.Sc. and Ph.D. programs, which are integrated into the graduate programs at the College of Engineering.

The *TEMEP Discussion Paper Series* is intended to serve as an outlet for publishing research about theoretical, methodological, and applied aspects of industrial economics, especially those related to the institute's areas of specialization, namely management of technology, information and communication technology, telecommunication, services, health industries, energy industries, as well as infrastructures for development. In particular, paper submissions are welcome, which analyze current technology and industry related issues, discuss their implications, and suggest possible alternative policies. The objective is to gain insights into important policy issues and to acquire a balanced viewpoint of policy making, technology management, and economics. It will enable us to identify the problems in industries accurately and to come up with optimal and effective guidelines. In addition, another important aim of this series is to facilitate communication with external research institutes, individual researchers, and policy makers worldwide.

Research results disseminated by TEMEP may include views on policies for information and communication technologies, technology management, telecommunication, energy, health, and development economics, but the institute itself takes no institutional policy position. If appropriate, the policy views of the institute are disseminated in separate policy briefs. Thus, any opinion expressed in the TEMEP Discussion Paper Series is those of the author(s) and not necessarily the opinion of the institute.

Finally, the current editor of this series would like to thank Prof. Dr. Almas Heshmati for establishing this working paper series in January 2009 and for guiding its development during the course of the first 55 issues. In particular, Prof. Heshmati provided invaluable contributions for setting up the local IT infrastructure, registering the series with RePEc, disseminating the working papers, and setting the quality requirements.

Jörn Altmann, Editor

Office: 37-305

Technology Management, Economics and Policy Program

College of Engineering

Seoul National University

599 Gwanak-Ro, Gwanak-Gu

Seoul 151-742

South-Korea

Phone: +82-70-76787-6676

Fax: +1-501-641-5384

E-mail: jorn.altmann@acm.org

Using SLA Mapping to Increase Market Liquidity

Marcel Risch¹, Ivona Brandic², Jörn Altmann¹

Department of Industrial Engineering

College of Engineering

Seoul National University

Seoul, South-Korea

marcel.risch@temep.snu.ac.kr, ivona@infosys.tuwien.ac.at, jorn.altmann@acm.org

First Version 2009-11-22, Revised Version 2010-04-15

Abstract: Research into computing resource markets has mainly considered the allocative fairness of market mechanisms. It has not been discussed how a large variety of resource types influences the market liquidity. Markets containing large numbers of buyers and sellers for heterogeneous resources suffer from a low likelihood of matching offers and requests. Traders therefore have the high risk of not being able to trade resources. We suggest a solution that derives SLA templates from a large number of heterogeneous SLAs in the market and, by using these templates instead of the original SLAs, facilitates SLA mapping. The usefulness of this approach is demonstrated through simulation results and a comparison with an alternative approach, in which SLAs are predefined.

Keywords: Service level agreements, Cloud computing, commodity goods, SLA matching, market liquidity, market mechanisms, SLA matching.

JEL Classification Numbers: C15, C63, D40, L86, M15, M21

1. Introduction

When developing markets for Cloud resources, it should be considered that the Cloud resources are designed to be liquid. Illiquid goods (which include rare goods and differentiated goods) have a higher risk that they cannot be sold or purchased when needed, driving market participants away [19]. Therefore, markets can only function efficiently with a sufficient number of market participants.

Creating such a market with a large number of Cloud resource traders is far from trivial. Firstly, consumers will only join an open Cloud market, if they are able to find what they need quickly. Secondly, providers will only join the market if they can be fairly certain that the resources they are trying to sell will be sold for the right price. Should either of these conditions not be met, providers and consumers will not participate in the market.

Open Cloud markets face an unusual challenge in this instance. The widespread use of virtualization enables resource providers to create a wide range of tradable resource types. At the same time, resource consumers can also specify their needs very precisely. In such a case, this lack of liquidity will make the market unattractive to consumers and providers.

In this paper, we demonstrate the problem (low liquidity for each resource type) caused by a large number of resource definitions. To counteract this problem, we introduce an approach, based on SLA mapping, which ensures sufficient liquidity in the market. These SLA mapping techniques not only simplify the search for similar offers but also allow us to derive public SLA templates from all existing offerings. These SLA mappings map parts of a consumer-defined SLA document to a public SLA template. The purpose of SLA mappings is twofold: Firstly, users may discover services with less effort and secondly, based on predefined learning functions and accumulated SLA mappings; the proposed SLA mapping approach facilitates user-driven definitions of public SLA templates, increasing the chances of matching of offers in the future.

Summarizing, the contributions of this paper are: (1) the demonstration of the problem caused by a large number of resource types; (2) an approach to overcome the shortcomings of an open market by introducing an SLA mapping approach; and (3) a first evaluation of the proposed SLA mapping approach considering different strengths and weaknesses.

2. Computing Resource Markets

The research into resource markets can be divided into two groups, when looking at their attempts of describing the tradable good. The first group does not define goods at all, while the second group focuses on one aspect of a computing resource only. However, neither group discusses the liquidity of goods in Cloud computing markets.

The first group consists to a large extent of early Grid market designs [1-3, 8]. One such example can be seen in [1], where Buyya et al. describe the entities of a Grid market. However, the question of how the market mechanism can handle a multi-dimensional good has not been answered. Similarly, GRACE [2] is a market architecture for Grid markets and outlined a market mechanism without defining the good “computing resource” or considering that consumers and providers have to agree to the resource specifications.

The second group of Grid market research has simplified the computing resource good. When developing the MACE exchange [8], the authors recognize the importance of developing a definition for the tradable good and abstract computing resources into services which can be traded. However, a detailed specification of a computing resource service has not been given and hence its effects on market liquidity cannot be assessed. The Spawn market was envisioned to work with CPU time slices [9]. While matching demand and supply is trivial, the fact that the resources have been reduced to a single component is not realistic, as the CPU slice requirements depend on the CPU vendor due to different instruction sets. Lastly, the Tycoon market was developed before virtualization tools (e.g. Xen [11]) became widely used [10]. Initially, it worked with basic computing cycles and it was planned to extend this market by making use of virtualization. However, it seems that the effort has been discontinued. The SORMA project focused on fairness and efficient resource allocation [3-6]. The project identified several requirements for Grid markets [7]. However, the requirement analysis has not considered that a market can only function efficiently, if there is sufficient liquidity of goods.

Overall, much of the computing resource market research has worked with either simplified definitions of the tradable good or without defining the good at all. This means that the issue of maintaining a sufficiently high liquidity (i.e. the likelihood of matching bids and asks) in such a market has not been addressed.

Despite some lacking research, a large number of commercial Cloud providers have entered the utility computing market, offering various types of services. On the one hand, there are resource providers, such as Amazon (e.g. EC2 [12]) and Tsunamic Technologies [13], who provide computing resources. Next, there are providers, who not only sell their own resources but also their own software services, such as Google Apps and Salesforce.com [14-15]. Furthermore, there are companies that attempt to run a mixed approach, i.e. they allow users to create their own services but also offer services themselves., e.g. Sun N1 Grid [16], Salesforce.com, and Microsoft Azure [17]. Most of these providers have in common that they only sell a single type of resources (or, in the case of Amazon, resource types which are based on a basic instance [18]). This limited number of different resource types enables a market creation, since all demand is channeled towards very few resource types.

3. The Modified Preferential Attachment Rule

Virtualization makes the use of Cloud computing environments more interesting, since virtualization allows for portability of machine images. For a Cloud market, however, the introduction of virtualization introduces the problem that providers can define their own resource types to differentiate themselves from other providers. The impact of this behavior on liquidity is shown with an example of a double auction market mechanism. This liquidity issue also exists for other market mechanisms.

3.1 Double Auction

There have been a number of proposals for using double auctions for Grid markets [21-23], since double auctions have shown their suitability for non-storable commodities, such as electricity [24]. Contracts (i.e. service level agreements) in the double auction market, that we sketch here, are traded as follows: First, traders send their bids and asks to the market. Second, using the continuous double auction matching procedure, the market attempts to match bids and asks. Once a match is made, both parties receive a contract (i.e. signed SLA) with each other's information. The advantages of using such a contract lie in the fact that simple bids and asks suffice for trading resources.

3.2 Matching Performance for a Double Auction

We have developed a simulation for a double auction market, for which we defined a range of tradable resource types. Any trader could decide at the start of the day, whether resources had to be purchased or could be sold. The demand generator created a normally distributed demand across all traders. The simulation was repeated with various numbers of traders and various numbers of resources types. Each simulation allowed traders to interact over a period of 500 simulated trading days. Using these parameters, we measured the percentage of matched bids to the total number of bids. This measure was chosen, since the absolute trading volume only has meaning, if it can be compared to historical data. Our results are shown in Figure 1.




Figure 1. Matching Probability for Bids.

The top line shows the likelihood of matching a bid in a market with 5000 resource types. Initially, the matching probability is low, since the supply and the demand are spread out over many resource types and, hence, it is less likely that a buyer and seller will meet. However, the likelihood increases, as the number of traders in the market grows. To determine how many traders would be needed to achieve a match probability of 75%, we performed a linear regression on the available data and found that about 39,800 traders would be needed. The grey line indicates the likelihood of matching a bid in a market with 10,000 resource types. Again, we can see that the probability is quite low but increases slowly with the number of traders. Using linear regression, we found that about 46,400 traders are needed to achieve a 75% match probability in this market. The dashed line indicates the likelihood of matching a bid in a market where each parameter of a resource can be set by traders, leading to 276,480 possible resource permutations. It can be seen that this probability is quite low and rises very slowly. A linear regression analysis indicated

that about 33 million traders would be required to achieve a 75% matching probability for bids. In general, these results show that, for a market to function properly, a sufficiently large base of traders must exist. Since large numbers of continuously acting traders are uncommon in current resource markets, our results indicate that a market, which sells fewer resource types, has higher liquidity. Next, we will show how resource types can be homogenized through SLA mapping, to get less resource types.

4. SLA Mapping

In this section, we explain the SLA mapping approach used to obtain public SLA templates, focusing on the lifecycle of SLA templates, management of SLA mappings, and SLA transformations used for the realization of SLA mappings. Finally, we introduce a marketplace, in which SLA mappings can be used to limit the number of resource types and, consequently, increase liquidity in the computing resource market. Note that automated trading is not required, purchases can be made manually.

4.1 The importance of SLAs in markets

Based on the outcome of the analysis of the liquidity problems in markets, we are faced with an interesting research challenge. On the one hand, to exploit the potential of open markets, a large number of traders is necessary. On the other hand, the large number of traders inflates the variety of resources available. Without some form of SLA matching or some way of limiting resource diversity, the set of available computing resources in a market would grow, spreading supply and demand across a wide range of resources, thereby reducing the liquidity of each resource type.

Some current adaptive SLA matching mechanisms are based on OWL, DAML-S [25-26]. Another semantic technology work (by Oldham et al.) describe a framework for semantic matching of SLAs based on WSDL-S and OWL [16]. Dobson et al. present a unified quality of service (QoS) ontology, which is applicable to specific scenarios such as QoS-based Web services selection [24]. Ardagana et al. present an autonomic Grid architecture with mechanisms to dynamically reconfigure service center infrastructures to fulfill varying QoS requirements [27]. Koller et al. discuss autonomous QoS management, using a proxy-like approach for exploiting SLAs to define certain QoS parameters that a service has to maintain during its interaction with a specific customer [28].

None of the presented approaches address the issues of the open market, where traders meet on demand. In most existing approaches, traders have to agree either on specific ontologies [26-27] or have to belong to a specific portal [28]. None of the discussed approaches handle semi-automatic definitions of SLA mappings, allowing negotiations between inconsistent SLA templates. Also, none of the presented approaches allow user-driven definitions of publicly available SLA templates.

4.2 The SLA Template Lifecycle

The following figure illustrates the lifecycle of SLA templates. In particular, it shows how public SLA templates are generated and managed.




Figure 2. SLA Template Lifecycle.

Figure 2 shows how an initially generated SLA template (step 1), which may not necessarily reflect the needs of users from a specific domain, can be used to derive realistic templates for a new specific domain. As indicated through step 1, we assume that for specific domains, specific SLA templates are initially generated. These generated SLA templates are published in the public registry (step 2). Thereafter, SLA mappings are defined manually by users (step 3). At the same time, learning functions for the adaptation of these public SLA templates are defined. A learning function determines how often a specific SLA mapping had to be used during a predefined time period before the modified SLA template becomes a template stored in the repository (step 5). In such an environment, the application of SLA mappings can also be done automatically, as described in Section 4.3 (step 4).

4.3 Managing SLAs

Figure 3 shows the architecture for managing SLA mappings. The registry comprises different SLA templates where each of them represents a specific application domain, such as medicine or telecommunication.




Figure 3. Template Registry




Figure 4. SLA Transformation

Service providers may assign their service to a particular public SLA template (step 1 in Figure 3) and, if necessary, assign SLA mappings afterwards making modifications to the original template (step 2). Service consumers may search for the services using meta-data and search terms (step 3). After finding appropriate services, each service consumer may define mappings (i.e. modifications) to a posted template (step 4). Besides, public SLA templates should be updated frequently according to predefined adaptation rules to reflect

the actual SLAs used by traders (step 5). To reduce the mapping cost for traders though, SLA mappings should also be executed automatically and public SLA templates should be defined and deleted dynamically.

4.4 SLA Mapping

To explain the SLA mapping, we describe how the transformation from modified public SLAs to the final SLA template is conducted. Figure 4 depicts a scenario for defining XSL transformations of SLAs expressed in the Web Service Level Agreements (WSLAs) specification language.

Public SLA templates are published in a searchable registry, from which traders may download and compare them with their local SLA template. If any differences are discovered, traders may write rules for transforming (XSL transformation) the local WSLA template to the remote template. The rules are stored in the database and can be applied at runtime to the remote WSLA template.

Figure 4 shows how negotiations can be performed using SLA transformation. Each of the two parties generates a WSLA. The locally generated WSLA plus the rules defining transformations from the local WSLA to the remote WSLA deliver a WSLA, which is compliant with the remote WSLA. Reversing this process, the remote WSLA plus the rules defining transformations from the remote to the local WSLA deliver a WSLA, which is compliant to the local WSLA. Thus, the negotiation may be done between non-matching WSLAs in both directions: from service consumer to service provider and vice versa. Thus, both parties may match on a publicly available WSLA template. To facilitate matching, SLAs are transferred into a canonical form as described in [29].

For example, in a double auction market, traders generate their own WSLA plus the rules for transforming their local WSLA to the market WSLA (i.e. the WSLA of the auction marketplace). The market has a learning function which can determine which of the WSLAs best matches the requirements of the traders. Based on this result, new public WSLAs can be created or old ones can be removed.

4.5 Application of the public SLA template in double auction markets

In a double auction market environment, the market participants trade using public SLA templates, which describe the traded computing resources, the software running on these resources, the Terms of Use, and the price [30]. In Figure 5, a graphical representation of such a tradable public SLA template is shown.

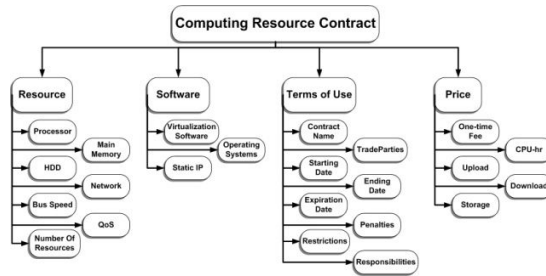


Figure 5. A Sample SLA Template Obtained from a Learning Function.

A market trading only a small number of public SLA templates encourages consumers to map their demand to the existing SLA templates, and providers to map their supply to the public SLA templates. However, once the market is operating, provisions can be made to add additional SLA templates (as new resources become available) and to remove SLA templates. Two rules are suggested to address this issue: First, a new tradable SLA is to be added, if the learning function determines that there is sufficient level of demand for this type of SLA template. Second, the removal of SLA templates depends on the trading volume: If the traded volume falls beneath a threshold for a certain time period, the SLA template will be removed.

Trading of SLA templates has to take into account that specific starting times and ending times are parameters of the templates. Standardization of starting times and usage durations provides one solution. Standardization for usage durations is fairly trivial, since it will become apparent that certain resources are commonly used for certain durations and less popular for other durations. Focusing on these popular periods will be sufficient for successful standardization. The learning function described previously can be used to determine popular starting and ending times.

5. Discussion and Validation

The following table gives an overview of the advantages and disadvantages of the SLA Mapping approach compared to an approach, in which the SLAs are predefined.

Table 1. Strengths and Weaknesses of SLA Mappings and Standardized SLAs.

Approach	Strengths	Weaknesses
SLA Mapping	<ul style="list-style-type: none"> - No explicit resource limitations - Flexible SLA usage - Increased liquidity - User-driven approach 	<ul style="list-style-type: none"> - SLA translations are time intensive - System abuse by submitting fake SLA mappings
Standardized SLA Templates	<ul style="list-style-type: none"> - Clearly defined resources - Adaptability of Standardized SLAs - Increased liquidity 	<ul style="list-style-type: none"> - Standardized SLAs may prevent users with unusual resource profiles to join the market

The advantage of the SLA mapping approach is that users face few limitations if they determine their resource needs. They can define their SLA templates flexibly. The SLA transformation then determines those public SLA templates that minimize the differences

to all user-defined SLAs. However, the SLA translations are time-intensive, which can be a problem if the SLAs are large. Furthermore, users can submit fake SLA mappings which have to be detected by a learning function.

5.1. Consequences of few resource types

Figure 7 shows the gains that can be achieved by transforming many heterogeneous SLAs into a few public SLA templates, which can be traded in a double auction. This figure was obtained using the same simulation environment, which was used to demonstrate the liquidity problems earlier.




Figure 7. Matching Probability for Bids with Fewer SLA Templates

The top line shows the likelihood of matching a bid for a market with one tradable SLA template. The likelihood is high (75%), since supply and demand is channeled to a single resource type. The next lower line shows the likelihood of a bid being matched for a market with 100 tradable SLA templates. Using linear regression, we calculated that a market with about 12,000 traders would have a matching probability of 75%. The lowest line shows the likelihood of a bid being matched in a market with 1000 tradable SLA templates. Using linear regression, we have calculated that, to reach a matching probability of 75%, 17,000 traders would be needed.

The matching probability of about 75% for a market with a single resource type seems quite low. The main cause of a lower matching probability was that bids had to be satisfied by a single ask, since multi-tier applications usually must be closely co-located. The impact of duration and number of resources has been outlined in [20].

Concluding, double auction markets for computing resources require lots of standardization. Having many similar resource types distributes demand and supply, thereby threatening the liquidity. If resources are standardized, supply and demand can be channeled to fewer resources. Despite the reduction in overall utility caused by standardizing, standardization will increase liquidity. Our proposed SLA mapping approach creates standardized goods that can be traded in computing markets.

6. Conclusion

We have demonstrated that a market trading diverse computing resources can suffer from a reduced market liquidity. Since this problem makes a market extremely unattractive to

traders, we have developed an approach to ensure that the market liquidity remains high. Our approach allows finding public SLA templates that are closest to the SLA templates defined by the user, therefore meeting the needs of users best. This approach is very flexible, since it adapts to the changing needs of the users. We also described the rules established within our approach to ensure that new SLA templates can be added and old SLA templates can be removed. We believe, based on the simulation results and the comparison with an alternative approach, that our approach will prove to be valuable in the development of an open Cloud market and will ensure sufficient liquidity to attract traders.

References

- [1] R. Buyya, S. Vazhkudai, "Compute Power Market: Towards a Market-Oriented Grid," First IEEE International Symposium on Cluster Computing and the Grid, CCGrid, pp.574, 2001.
- [2] R. Buyya, D. Abramson, J. Giddy. An Economy Grid Architecture for Service-Oriented Grid Computing. 10th IEEE International Heterogeneous Computing Workshop, HCW 2001. IEEE Computer Society Press: Los Alamitos, CA, 2001.
- [3] The SORMA project, <http://sorma-project.org/>, 2009.
- [4] R. Nou, F. Julia, J. Guitart, J. Torres. Dynamic Resource Provisioning for Self-adaptive Dynamic Heterogeneous workloads in SMP Hosting Platforms. International Conference on e-Business, 2008.
- [5] L. Amar, A. Mualem, J. Stoesser. On the Importance of Migration for Fairness in Online Grid Markets. International Conference on Autonomous Agents and Multiagent Systems, 2008.
- [6] L. Amar, A. Barak, E. Levy, M. Okun. An On-line Algorithm for Fair-Share Node Allocations in a Cluster. IEEE International Symposium on Cluster Computing and the Grid, CCGRID, 2007.
- [7] D. Neumann, J. Stoesser, C. Weinhardt. Bridging the Grid Adoption Gap – Developing a Roadmap for Trading Grids. Bled eConference, Merging and Emerging Technologies, Processes, and Institutions, 2007.
- [8] B. Schnizler, D. Neumann, D. Veit, C. Weinhardt. Trading Grid Services - A Multi-Attribute Combinatorial Approach. European Journal of Operational Research, vol.187, no.3. pp.943-961, 2008.
- [9] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, W. S. Stornetta. Spawn: A Distributed Computational Economy. IEEE Transactions on Software Engineering, vol.18, no.2, pp.103-117, February 1992.
- [10] K. Lai, L. Rasmusson, E. Adar, L. Zhang, B. A. Huberman. Tycoon: An Implementation of a Distributed, Market-Based Resource Allocation System. Multiagent Grid Systems, vol.1, no.3, pp.169-182, August 2005.

- [11] XenSource, Inc., <http://xen.org/>, 2009.
- [12] Amazon Elastic Compute Cloud (Amazon EC2), <http://aws.amazon.com/ec2/>, 2009..
- [13] Tsunamic Technologies Inc., <http://www.clusterondemand.com/>, 2008.
- [14] Google Apps, <http://www.google.com/apps/>, March 2009.
- [15] Salesforce.com, <http://www.salesforce.com>, March 2009.
- [16] Sun Grid, <http://www.sun.com/service/sungrid/index.jsp>, 2009.
- [17] Microsoft Azure, <http://www.microsoft.com/windowsazure/>, 2009.
- [18] Amazon EC2 Instance Types, <http://aws.amazon.com/ec2/instance-types/>, 2008.
- [19] P. A. Samuelson, W. D. Nordhaus. Economics. McGraw-Hill/Irwin, 18 edition, ISBN 0072872055, July 2004.
- [20] J. Altmann, C. Courcoubetis, G. D. Stamoulis, M. Dramitinos, T. Rayna, M. Risch, C. Bannink. GridEcon - A Market Place for Computing Resources. GECON 2008, Workshop on Grid Economics and Business Models, Springer LNCS, August 2008.
- [21] C. Weng, X. Lu, G. Xue, Q. Deng, M. Li. A Double Auction Mechanism for Resource Allocation on Grid Computing Systems, In GCC, pp.269, 2004.
- [22] U. Kant, D. Grosu. Double Auction Protocols for Resource Allocation in Grids. Proceedings of the International Conference on Information Technology: Coding and Computing (ITCC 2005), vol.1, no.1. IEEE Computer Society, Washington, DC, pp.366-371, April 2005.
- [23] C. Weng, M. Li, X. Lu. Grid Resource Management Based on Economic Mechanisms. Journal of Supercomputing. vol.42, no.2, pp.181-199, November 2007
- [24] European Energy Exchange, <http://www.eex.com/en/>, 2008.
- [25] G. Dobson, A. Sanchez-Macian. Towards Unified QoS/SLA Ontologies. Proceedings of the IEEE Services Computing Workshops, SCW 2006, pp.18-22, Chicago, Illinois, USA September 2006.
- [26] N. Oldham, K. Verma, A. P. Sheth, and F. Hakimpour. Semantic WS-Agreement Partner Selection. Proceedings of the 15th International Conference on World Wide Web, WWW 2006, Edinburgh, Scotland, UK, May 2006.
- [27] D. Ardagna, G. Giunta, N. Ingra, R. Mirandola, B. Pernici. QoS-Driven Web Services Selection in Autonomic Grid Environments. International Conference Grid Computing, High Performance and Distributed Applications, GADA 2006. France, November 2006.
- [28] B. Koller, L. Schubert. Towards Autonomous SLA Management Using a Proxy-Like Approach. Multiagent Grid Systems. vol.3, no.3, IOS Press, The Netherlands, 2007.
- [29] I. Brandic, D. Music, S. Dustdar. Service Mediation and Negotiation Bootstrapping as First Achievements Towards Self-Adaptable Grid and Cloud Services. Grids Meet Autonomic Computing Workshop, GMAC09. In conjunction with the 6th

International Conference on Autonomic Computing and Communications Barcelona, Spain, June 2009.

- [30] M. Risch, J. Altmann. Enabling Open Cloud Markets Through WS-Agreement Extensions. Service Level Agreements in Grids Workshop, in conjunction with GRID2009, CoreGRID Springer Series, Banff, Canada, October 2009.

TEMEP Discussion Papers

- 2009-01: Tai-Yoo Kim, Almas Heshmati and Jihyun Park, “Perspectives on the decelerating agricultural society”.
- 2009-02: Yunhee Kim, Jeong-Dong Lee and Almas Heshmati, “Analysis of Pay Inequality and its Impacts on Growth and Performance in the Korean Manufacturing Industry”.
- 2009-03: Chul-Yong Lee and Jongsu Lee. “Demand Forecasting in the Early Stage of the Technology’s Life Cycle Using Bayesian Update”.
- 2009-04: Jongsu Lee and Chul-Yong Lee, “A Forecasting Model Incorporating Replacement Purchases: Mobile Handsets in South Korea’s Market”.
- 2009-05: Ki H. Kang and Jina Kang, “Revisiting Knowledge Transfer: Effects of Knowledge Characteristics on Organizational Effort for Knowledge Transfer”.
- 2009-06: Ki H. Kang and Jina Kang, “Does partner type matter in R&D collaboration for product innovation?”.
- 2009-07: Ki H. Kang and Jina Kang, “How Do Firms Source External Knowledge for Innovation?: Analyzing Effects of Different Knowledge Sourcing Methods”.
- 2009-08: Ki H. Kang and Jina Kang, “Do external knowledge sourcing methods matter in service innovation?: analysis of South Korean service firms”.
- 2009-09: Gunno Park and Jina Kang, “The effects of teacher firms’ characteristics and student firms’ absorptive capacity on firm performance in technology alliances”.
- 2009-10: Donghyun Oh and Almas Heshmati, “A Sequential Malmquist-Luenberger Productivity Index”.
- 2009-11: Tausch Armo and Almas Heshmati, “Re-Orient? MNC Penetration and Contemporary Shifts in the Global Political Economy”
- 2009-12: Donghyun Oh, Almas Heshmati and Hans Lööf, “Technical Change and Total Factor Productivity Growth for Swedish Manufacturing and Service Industries”
- 2009-13: Tai-Yoo Kim, Almas Heshmati and Jihyun Park, “The Faster Accelerating Knowledge-based Society”
- 2009-14: Junseok Hwang, Jörn Altmann and Kibae Kim, “The Structural Evolution of the Web2.0 Service Network”
- 2009-15: Alireza Abbasi, Jörn Altmann and Junseok Hwang, “Evaluating Scholars Based on Their Academic Collaboration Activities: The RC-Index and CC-Index for Quantifying Collaboration Activities of Researchers and Scientific Communities”

- 2009-16: Dong-hyun Oh, Hans Lööf and Almas Heshmati, “The Icelandic Economy: A Victim of the Financial Crisis or Simply Inefficient?”
- 2009-17: Sungki Lee, Donghyuk Choi and Yeonbae Kim, “Contextual Effects on the Complementarities between R&D Activities: An Empirical Analysis of the Korean Manufacturing Industry”
- 2009-18: Younghoon Kim, Yeonbae Kim and Jeong-Dong Lee, “Corporate Venture Capital and Its Contribution to Intermediate-Goods Firms in South Korea”
- 2009-19: Yunhee Kim, Jae Young Choi and Yeonbae Kim, “Complementarity and Contextuality in the Adoption of Information Systems in Korean Firms”
- 2009-20: Marcel Risch and Jörn Altmann, “Enabling Open Cloud Markets Through WS-Agreement Extensions”
- 2009-21: Almas Heshmati and Rachid El-Rhimaoui, “Effects of Ownership and Market Share on Performance of Mobile Operators in MENA Region”
- 2009-22: Donghyun Oh, Almas Heshmati and Hans Lööf, “Total Factor Productivity of Korean Manufacturing Industries: Comparison of Competing Models with Firm-Level Data”
- 2009-23: Almas Heshmati, Dan Johansson and Carl Magnus Bjuggren, “Effective Corporate Tax Rates and the Size Distribution of Firms”
- 2009-24: Seong-Sang Lee and Yeonbae Kim, “The Impact of Efficiency Parameters on Firms’ Innovative Activities: Evidence from Korean Firm-Level Data”
- 2009-25: Junseok Hwang, Donggook Choi, Jongeun Oh and Yeonbae Kim, “Competition in the Korean Internet Portal Market: Network Effects, Profit, and Market Efficiency”
- 2009-26: Jae Young Choi, Yeonbae Kim, Yungman Jun and Yunhee Kim, “A Multivariate Probit Analysis of Korean Firms’ Information System Adoption: An Empirical Analysis on the Determinants of the Adoption and Complementarity among the Information Systems”
- 2009-27: Flávio Lenz-Cesar and Almas Heshmati, “Determinants of Firms Cooperation in Innovation”
- 2009-28: Firas H. Al-Hammadany and Almas Heshmati, “Determinants of Internet Use in Iraq”
- 2009-29: Jeong-Dong Lee, “Industry Dynamics and Productivity Research”
- 2009-30: Nam TRAN-QUY and Almas Heshmati, “Measurements and Determinants of Pay Inequality and its Impacts on Firms Performance in Vietnam”
- 2009-31: Nils Karlson, Marcus Box and Almas Heshmati, “Generality, State Neutrality and Unemployment in the OECD”

- 2009-32: Chang Seob Kim, Yoonmo Koo, Je-jung Choi, Junhee Hong and Jongsu Lee, "Consumer Preferences for Automobile Energy Efficiency Grades"
- 2009-33: Jongsu Lee, Jae Young Choi and Youngsang Cho, "A Forecast Simulation Analysis of the Next-Generation DVD Market Based on Consumer Preference Data"
- 2009-34: Jungwoo Shin, Chang Seob Kim and Jongsu Lee, "Model for Studying Commodity Bundling with a Focus on Consumer Preference"
- 2009-35: Yuri Park, Hyunnam Kim and Jongsu Lee, "Empirical Analysis of Consumer Preferences for Mobile Phones and Mobile Contents in the Presence of Indirect Network Effects"
- 2009-36: Jörn Altmann and Juthasit Rohitratana, "Software Resource Management Considering the Interrelation between Explicit Cost, Energy Consumption, and Implicit Cost: A Decision Support Model for IT Managers"
- 2009-37: Marcel Risch, Ivona Brandic and Jörn Altmann, "Using SLA Mapping to Increase Market Liquidity"