

Maximizing Liquidity in Cloud Markets through Standardization of Computational Resources

Ivan Breskovic,^{*} Ivona Brandic,^{*} and Jörn Altmann[†]

^{*}Distributed Systems Group, Institute of Information Systems, Vienna University of Technology, Austria
{breskovic, ivona}@infosys.tuwien.ac.at

[†]TEMEP, Department of Industrial Engineering, College of Engineering, Seoul National University, South Korea
jorn.altmann@acm.org

Abstract—Low liquidity in cloud markets can result in market instability and inefficiency, preventing the successful implementation of ubiquitous computing on demand. To circumvent this issue, it has been suggested to channel demand and supply into a limited number of standardized services. These standardized services can even be automatically adapted to user requirements with the goal of continuously improving market performance. In this paper, we focus on answering how many standardized services should be placed in the market. This work is based on a new definition of liquidity for cloud resources, which in turn has been derived from liquidity definitions of financial markets. Using a simulation framework, we evaluate our method for estimating the optimal quantity of standardized services with respect to market liquidity and demonstrate the benefits of this approach in terms of increase in market efficiency and decrease in users’ cost of participation in the market. The methods presented in this paper have the potential to be applied in other electronic markets as well.

Keywords—Service level agreement, electronic markets, cloud economics, autonomic computing, standardized goods, market modeling, market liquidity, IT services

I. INTRODUCTION

Common to utility computing, cloud computing and grid computing is that they offer computational resources (e.g., software, hardware and computing platform services) in a manner similar to utilities, such as water, electricity and telephony, without regard to where the services are hosted or how they are delivered [4], [15]. However, the current cloud market is fragmented and static, hindering the paradigm’s ability to fulfill its promise of ubiquitous computing on demand and as a commodity. In order to address this issue, electronic markets for trading and/or allocating computational resources (i.e., grid and cloud services) have been proposed [1], [38].

In cloud markets, computational services are described by service parameters and quality of service (QoS) objectives. Each setting of a service, i.e., configuration of parameters, desired values of QoS objectives, and combinations of these elements, defines a new service (also called “computational resource”). Due to a great number of available settings, markets suffer from a vast *diversity and heterogeneity* of computational resources.

In addition to this, computational resources are characterized by a high *dynamism* and a high *fungibility*. Resource dynamism is a result of a large resource variability, a dynamic user base, diverse user behavior in the market, and the

emergence of new actors and new actor types. This results in changing demand and supply. Fungibility of computational resources describes the fact that computational resources can be substituted by other, relatively similar resources without violating end-user expectations. In particular, resource fungibility allows minor differences in specifications of service parameters and service objectives that define new resources.

Dynamism of resources may often have negative effects on market liquidity of a good, i.e., the ability to easily and quickly sell or purchase a good without causing a significant movement in its price. With respect to the cloud market and its extensive variety of computational resources (goods), the probability to find a resource that matches a buyer’s requirement is relatively low and, therefore, drives potential users away. Furthermore, finding a matching provider’s service in such an environment is a costly task, as a buyer must check numerous offerings before finding the closest one to their needs. In our previous work, the means for counteracting this problem was to channel demand and supply into a limited number of standardized services [9], [11], [12]. For this, we proposed automated creation, selection and adaptation of standardized computational resource specifications based on the current demand and supply. By continuously adapting the offering of standardized resources, the marketplace successfully addresses the issue of changing demand and supply and, as demonstrated in [8], increases the homogeneity of computational resources while decreasing the cost of searching for resources.

Standardization of computational resources has the goal of improving market liquidity. This goal is important as it promotes active participation in a market and helps ensuring market efficiency and stability. This, in turn, makes the market more resilient to external shocks, making it more competitive and attractive to both buyers and sellers. However, demonstrating the effects of standardized products on liquidity is far from trivial.

The first objective of this paper is to make the first step towards quantifying liquidity of cloud markets with respect to the resources traded between the participants. For this, we analyze existing approaches for assessing liquidity in financial markets and map the common financial liquidity measures to the computational resource markets. The challenge is to adapt the common measures of liquidity so that they can capture the liquidity of a set of differentiated, but partly substitutable

goods. Using our derived method, we demonstrate the effects of standardized computational resources on market liquidity.

The second objective of the paper is to demonstrate how the ability to assess liquidity of cloud markets opens the possibility of adjusting the product offerings in the market until the maximum point of liquidity is reached. To achieve this objective, we take liquidity as the key performance indicator for adapting resources and find the “ideal” quantity of standardized computational resources so that liquidity is maximized. In order to formalize the process of finding this “ideal configuration”, we introduce a new method. The importance of this method is manifold: it can be used to determine the optimal quantity of standardized resources in electronic markets, but can also help service providers to decide on the services they should offer in the market in order to have the highest probability to find buyers for their services and, therefore, increase their profit and the overall market revenue. This method will fundamentally help the cloud computing market to mature, stabilize and grow.

To summarize, the main contributions of this paper are: (1) an analysis of methods for measuring market liquidity in financial markets and modifying these methods to fit computational goods; (2) a demonstration of the benefits of standardizing computational resources in cloud markets in terms of market liquidity; and (3) the introduction of a method for estimating the optimal quantity of standardized computational resources with respect to market liquidity.

The remainder of the paper is organized as follows. Section II discusses the state-of-the-art in research on electronic markets. Section III presents a short overview of our approach of standardizing computational resources in cloud markets and highlights the importance of liquidity in computing resource markets. In Section IV, we introduce the liquidity measures applied in financial markets and derive a small set of liquidity measures for cloud computing markets. Section V presents a case study and demonstrates the positive effects of resource standardization on market liquidity. A discussion about the usefulness of our approach is given in Section VI. It also introduces the method for estimating the “ideal” quantity of standardized resources. Finally, Section VII concludes the paper.

II. CLOUD MARKETS IN RESEARCH

Several research projects have discussed the implementation of system resource markets [1], [14], [31]–[33], [38]. GRACE developed a market architecture for grid markets and outlined a market mechanism, while the good itself (i.e., computing resource) has not been defined [14]. Moreover, the process of creating agreements between consumers and providers has not been addressed. The SORMA project also considered open grid markets [32], [33]. They identified several market requirements, such as allocative efficiency, budget-balance, truthfulness, and individual rationality [32]. However, they have not considered that a market can only function efficiently with a sufficiently large liquidity. In MACE [39], an abstract computing resource has been defined that can be traded.

However, a detailed specification of a good has not been given. GridEcon proposed a commodity market for cloud computing services [1], [38]. Although an explicit service level agreement for standardized cloud services [37], the cloud service requirements, and the requirements for trading have been defined and specified, the issue of adaptation of standardized goods has not been addressed. In the work on cloud computing value chains [31], many important issues of electronic markets (e.g., improved cloud pricing and licensing models) are discussed. However, while the diversity of virtualized resources was mentioned implicitly, the effect this diversity can have on the market has not been addressed.

Currently, cloud market platforms assume static markets in terms of good specification, user base and participation. However, to enable the flexibility promised by the cloud computing paradigm, such systems have to be adaptive and sustainable. Moreover, in order to enable resource-efficient utilization of on-demand resources, autonomic (self-* [28]) capabilities are essential for the creation of market platforms. Although some scientific works address technical issues to make systems autonomic, such as the development of negotiation protocols to make cloud services self-adaptive [7], or consider using autonomic service management frameworks [19], [29], [34], they do not take economic methodologies into account. Research on autonomic systems focusing on economic methods and considerations is in its early stage. For example, Pardoe et al. propose mechanisms that are able to adaptively adjust their parameters based on the past behavior of participants [35]. Another example is the self-organizing resource allocation mechanism for dynamic application layer networks [43]. The largest step towards this vision, however, is proposed by Breskovic et al. [10], [16], who use the autonomic MAPE loop in the context of electronic markets to automatically adapt the market platform to changed environmental conditions based upon a given concept of “performance”. As possible form of adaptations, the authors specify market’s institutional (i.e., economical) properties (e.g., pricing rules and allocation policies) and properties and the underlying software/hardware infrastructure. However, they do not consider issues such as resource specification and adaptation depending on demand and supply, which is a crucial element of autonomic marketplaces.

Resource allocation, provisioning, QoS-based service selection, and negotiation in grids and clouds have been the subject of many research studies. Various resource-efficient and economically beneficent allocation techniques and methodologies have already been proposed to address these issues. These methodologies include game theoretical approaches [3], [45], stochastic programming [17], bio-inspired mechanisms [20], [41], auction-based algorithms [25], [36], [46] and agent-based approaches [21], [40], [44]. Most of these works perform QoS service selection and resource allocation based on some service performance indicators and economic indicators, including wait time optimization, utilization maximization and economic wastage minimization. However, none of the works identified the importance of market liquidity nor performed

any kind of adaptation of resources traded in cloud markets.

Although market liquidity is commonly used in financial literature as one of the fundamental measures of market attractiveness, efficiency and activity, it is virtually impossible to give one definition of liquidity that covers all of its aspects and fits to all market scenarios. Due to this complexity, measuring market liquidity is far from trivial. However, there are several common measures for approximation of liquidity in market literature with spread measures (e.g., quoted, effective and realized spreads) [6], [27], market depth [5], and immediacy of matching [18] being the most prominent examples. Nevertheless, although they are applicable to the markets trading financial assets, they cannot be directly applied to the computing resource markets due to the fact that the current computing resource market is built of an extremely high number of heterogeneous (i.e., quality differentiated) goods. The challenge is to adapt the common measures of liquidity so that they can capture liquidity of a set of differentiated, but partly substitutable goods. With such a definition, it will be possible to create a computing commodity market, which trades a limited number of standardized goods. A detailed discussion on common liquidity measures in financial markets and their possible application to computing resource markets will be given in Section IV.

Note that besides liquidity, there are several other prominent and important dimensions of market quality [47] such as activity [5], [26] and information model [23]. Activity is measured by indicators such as the daily trading volume and the average daily trade size. The information model explains how information is translated into market prices through order flow. Liquidity, however, is a better measure for the attractiveness of a market as it indicates the ability to quickly trade large size of goods for low cost [24].

III. STANDARDIZED RESOURCES IN ELECTRONIC MARKETS: AN OVERVIEW

In order to position the work in this paper, we shortly summarize our existing work on automatic adaptation of standardized computational resources in cloud markets and identify the next steps towards achieving the vision of self-adaptive cloud resources. We also give a short motivation and summary of the contributions of this paper.

Requirements of services in cloud markets are usually negotiated by means of Service Level Agreements (SLAs). They are binding electronic contracts signed between service buyers and service sellers that formally specify end-user expectations. Before signing legally binding documents, buyers and sellers express the requirements of their services through the use of SLA templates. In our vision of adaptive electronic markets, we differentiate between two types of SLA templates: (1) *private* SLA templates, which are used to specify buyers' requirements and sellers' offerings for services and which they create manually when submitting service offerings and requirements to the market; and (2) *public* SLA templates, which are created automatically by the market platform and

represent standardized services that can be traded in the market [9], [11], [12].

Today's traditional electronic marketplaces support the trading of differentiated services: a buyer's requirement is compared to all sellers' offerings to find the best matching service. This process is often inefficient due to the market dynamism and a large diversity in resources, and requires buyers and sellers to invest a large effort to find the best matching service offerings. On the contrary, in our vision of electronic markets, demand and supply are channeled through a limited number of standardized resources. Buyers and sellers choose between the available standardized services (described by public SLA templates) that describe services closest to their needs. Due to the significant decrease in the quantity of trading artifacts, the effort of finding the best-fitting offering on the market is remarkably lower. However, buyers' and sellers' satisfaction with the limited choice of services is, naturally, lower as well. To minimize this effect, appropriate adaptation methods are applied that constantly and efficiently adjust standardized services to the new market environment and always reflect the requirements of market participants. As it is shown in [8], the tradeoff between buyers' and sellers' satisfaction (i.e., their utility) and cost for finding a trading partner is in the "standardized approach" significantly better than in the "differentiated approach".

In our approach, standardized resources are autonomically adapted using the traditional MAPE loop [28], which is defined as follows. The *Monitoring* component monitors the institutional performance indicators (i.e., properties concerning the economic implementation and behavior of the market platform) to determine the overall "market performance". The properties measured by the monitoring component include, for example, market liquidity, number of active buyers and sellers, revenue, and variety in resource types traded on the market. The *Analysis* component analyzes the monitored data and determines whether publicly available standardized services should be adapted in order to improve the performance of the measured market performance indicators. If an adaptation should be executed, the analysis component inspects demand and supply, i.e., users' private SLA templates, and selects new properties that the new standardized services (i.e., public SLA templates) should reflect. In order to achieve this, clustering algorithms are utilized to group similar user requirements and adaption methods are used to select the preferred properties by each of the clusters of users. Depending on the results of the analysis phase, the *Planning* component determines the set of adaptation actions to modify the specifications of standardized services and/or create new service offerings (i.e., public SLA templates). Finally, the *Execution* component executes the adaption, submits the newly created public SLA templates to the market, and notifies the users of the changes.

In our previous works [8], [11]–[13], we discussed the three of the four adaptation steps: analysis, planning, and execution. In [11], we applied and compared several clustering algorithms for grouping similar requirements of market participants. Using a simulation environment, we identified the

k -means algorithm as the best performing in terms of several evaluation criteria, such as the minimization of participants' costs of utilizing newly created public SLA templates and the maximization of SLA isolation (i.e., the average difference between newly generated public SLA templates). In [12], [30], we demonstrated several methods for creation of public SLA templates based on the private SLA templates of a group of market participants, where the maximum method [30] proved to be the most efficient. Finally, in [13], we introduced an automated method for a cost-efficient creation, adaptation, and utilization of standardized resources as the first step towards optimization of the adaptation process.

The first phase of the autonomic loop (i.e., monitoring) was in our earlier work simplified as it only considered buyers' and sellers' utility (i.e., satisfaction with the standardized services) and the cost that the approach incurred to them. In this paper, we discuss this component in more detail and investigate the methods for measuring market liquidity as the key indicator of market performance. Besides defining methods for assessing liquidity, we quantify the impact of standardizing resources in cloud markets on market liquidity and demonstrate how liquidity measures can be used to determine the "optimal" quantity of standardized resources in the market.

IV. MEASURING LIQUIDITY IN ELECTRONIC MARKETS

A. The meaning and importance of market liquidity

Market liquidity is an important measure of market quality and a concept which is commonly used in financial markets, but can be applied to other types of markets as well. In its essence, it describes how easy it is to trade a certain volume of the considered good. A market is liquid when it has a high level of trading activity, where one can buy and sell with the minimum price deviation.

The essential characteristic of a liquid market is that there are sufficiently many ready and willing buyers and sellers at all times. Market liquidity also depends on the ease with which market participants can carry out transactions. Thus, other things being equal, lower transaction costs contribute to higher market liquidity. In particular, if transaction costs and the costs of the participation in the market are high, the gap between the effective price received by the seller and that paid by the buyer of a service will be large and it will be difficult to match sell and buy orders [22]. Furthermore, if participation costs are high enough to constitute an entry barrier, the market will attract fewer dealers and investors, also lowering trading activity and, consequentially, market liquidity.

In order to work efficiently and to guarantee market stability, a marketplace should have a sufficiently high liquidity. In order for a market to be deeply liquid, a quick, simple and inexpensive exchange of products between buyers and sellers needs to be possible. In markets with a high variety of resource types, as it is the case with cloud markets, this means that it is necessary to ensure a large likelihood of finding a seller's offering for every buyer's requirement and vice versa.

B. Liquidity in financial markets

Due to its complex definition, measuring market liquidity is not a trivial task. Many factors affect liquidity, including institutional factors such as securities law, the regulation and supervision of dealers, and accounting rules. Equally, environmental factors such as the macroeconomic situation and changes play a role. Consideration of all factors that affect liquidity and formulation of possible measures for its quantification have been discussed in various research works with the focus on financial markets [5], [6], [18], [27], [47]. However, the common conclusion of those works is that the highly abstract definition of market liquidity means that it cannot be expressed as an aggregate value. Instead, there exist several standard measures that serve as proxies for its assessment. The most common measures include bid-ask spreads [2], market depth [5], and immediacy [18]. In the following, we describe how each of these measures can successfully capture at least one of the perspectives of market liquidity.

Bid-ask spread denotes the amount by which an ask, i.e., a seller's offered price, exceeds a bid, i.e., a buyer's requested price. The bid-ask spread essentially measures the difference in price between the highest price that a buyer is willing to pay for a product and the lowest price for which a seller is willing to sell it. Large bid-ask spreads indicate high buying possibilities of the buyers: their request prices are higher than the prices offered by the sellers, which results in more numerous trades in the market. Since the large bid-ask spread points to a high trading dynamism, it also leads to the conclusion that the market liquidity rises proportionally to the increase of its value.

Market depth measures the volume of goods traded in the market, i.e., the units that can be sold or bought for a given price impact. Particularly, market depth refers to the maximum size of a trade for any given bid/ask spread. A market may be considered deeply liquid if there are ready and willing buyers and sellers in large quantities, which is directly related to the concept of market depth as a large number of market traders and service offerings as well as a well-designed allocation mechanism result in a large trading volume in a time period. This suggests that a high market depth implies that the assets can be easily purchased or sold. Therefore, high market depth indicates high market liquidity.

The third common approximate measure of market liquidity, **immediacy**, refers to the time needed to successfully trade a certain amount of a product at a prescribed cost. Essentially, immediacy can be measured as the time passed between the submission of a requirement for a service to a market and the allocation (i.e., a match) between the buyer's requirement and a seller's offering. Depending on the actor, it is possible to differentiate buyer's immediacy from seller's immediacy. Small immediacy characterizes a small time needed to close a trade and indicates a liquid market.

C. Liquidity in cloud markets

The presented measures for approximation of market liquidity are commonly used in *financial* literature to measure

liquidity of monetary assets. In the case of cloud markets, we are interested in the liquidity of the market itself, where potentially many heterogeneous goods are traded. Due to the significantly different market organization in terms of goods, definition of liquidity must be modified and adapted to this setting. However, this is not a trivial task. For this reason, instead of defining a complete liquidity model for electronic markets, we will focus only on those aspects that may be affected by the standardization of computational resources in this paper.

One of the key factors of market liquidity is the prices of goods traded in the market: liquidity is strongly affected by buyers' bids and sellers' asks for goods and strongly depends on the market's allocation mechanism as well as the pricing methods used. However, from the perspective of the standardization of goods in electronic markets, we are interested only in the explicit impact of the quantity and structure of the goods on market liquidity. For this reason, the effects that the standardization of services may have on the prices in the market are out of scope of this paper. To achieve this, we simplify the definition of liquidity and assume a static user behavior in terms of pricing. Namely, we assume that the standardization of goods does not affect the bidding strategies of market participants: they are willing to bid for the standardized goods with the same prices as for the differentiated goods. Note that this assumption would most probably not hold in the real-life markets for several reasons. For example, users' satisfaction with the standardized goods may be lower than with the exact goods they need, which would result in the lower bidding prices. On the other hand, the positive impacts on market liquidity and participation costs (which will be demonstrated later in the paper) would have positive impacts on the bidding prices. However, the assumption of "static pricing" provides a simplified view on market liquidity and allows us to avoid uncertainty about the real cause of the change of market prices.

To quantify the impact of the product standardization on market quality, we consider the definitions of the common liquidity measures. Due to the simplification of the assessment model, we are not interested in the bid-ask spread, as it only depends on the current market prices. However, we are interested in the other two standard measures: market depth and immediacy, which we modify to the *overall market depth* and the *search cost*.

1) *Overall market depth*: Similarly as in financial markets, we use market depth to indicate the number of matches between requirements and offers during the trading time. In financial markets, depth points to the trading volume of one asset. In the computing resource market characterized by the heterogeneity of services, however, market depth can be seen as the cumulative value for all goods in the market. To differentiate between these measures, we use the term *overall market depth* to indicate the cumulative trading volume in computing resource markets. Details on how the overall market depth can be measured in adaptive cloud markets will be presented in Section V-A.

2) *Search cost*: In its original definition, immediacy strictly represents the *time* needed to successfully trade a product in the market. It is presented in time units and is defined for every single asset in the market. Although it is a valuable indicator of market liquidity, it is hard to strictly associate it to the variety of resource types in the market, as many factors (e.g., performance of market mechanisms and pricing algorithms, as well as various exogenous factors) may affect immediacy. To avoid these conflicts, we consider immediacy in its broad form: the effort needed to be invested in order to find a trading partner. In our context, this effort describes the number of comparisons between a buyer's (or a seller's) requirement and sellers' (or buyers') offerings in the market until the most suitable service in the market is found. The effort is, hence, associated to the search of a fitting service offering and is, for this reason, termed search cost. The search cost is of a particular importance in those markets in which goods are purchased and (re)sold very often. Considering cloud computing, this is the case in markets in which computational resources are rented on a short term (e.g., Amazon EC2 Spot Instances¹). Note that the search cost strictly correlates to the immediacy, since a large search cost is always a result of more numerous execution steps in the market, which requires more computation time.

Having the two liquidity measures in mind, we conclude that the goal of increasing market liquidity can be achieved by increasing the overall market depth while reducing the buyers' and sellers' search cost. Due to the inversely proportional relation between these values, this goal can be additionally expressed as maximization of the aggregate liquidity measure

$$lqdt_a = \text{overall market depth} / \text{search cost}. \quad (1)$$

Note that Eq. (1) does not present a "final" and "unique" measure of liquidity, i.e., a measure that depicts market liquidity independently from the overall market depth and the search cost. In particular, as it will be demonstrated in Section V, all presented liquidity measures are equally important and cannot be interpreted autonomously, i.e., without considering other measures simultaneously.

V. CASE STUDY

A set of realistic and effective measures for quantification of liquidity in cloud markets may bring enormous benefits in various research areas, including design and implementation of market platforms, allocation mechanisms and pricing, as well as market assessment and autonomic market adaptation. In this paper, we focus on its two fundamental applications with respect to the standardization of computational resources. First, we study the impact of standardized services to market liquidity and compare it to the liquidity of electronic markets with numerous differentiated and heterogenous services. Second, we explore the possibility of using market liquidity as the main market performance indicator to determine the quantity and quality of standardized services. To perform this analysis,

¹<http://aws.amazon.com/ec2/spot-instances/>

we define a case study in a simulated market environment described in [11] for evaluating the trade of both differentiated and standardized products, and examine market liquidity of the market in the given trading scenarios.

A. Simulation environment and testbed

In our earlier works, we introduced a simulation framework [11] and a testbed [8] for evaluating the approach of standardization of computational resources from the perspective of the cost of creating and utilizing standardized services. For the sake of consistency, in this paper we adopt the previously established simulation testbed, which we shortly summarize here.

The simulation process is started by a random generation of buyers' and sellers' private SLA templates containing a fixed amount of 4 SLA parameters and 4 associated service level objectives (SLOs). The desired parameter values specified in users' SLOs are given in form of ranges of real numbers. For example, an SLO value for an SLA parameter *ErrorRate* may be [0,1]%, stating that any value between 0% and 1% is acceptable for the user. The values of SLOs are created randomly, but with a predefined width of the value range, which is given as a percentage of the maximum possible SLO value range. Simulation settings are presented in Table I.

TABLE I
SIMULATION SETTINGS

Parameter	Value
No. of market users	$200 < n < 15000$
Portion of buyers in the number of users	50%
Portion of sellers in the number of users	50%
No. of parameters in SLA templates	4
Width of the SLO value range	10%
Method to cluster users' preferences	<i>k</i> -means [11]
Method to adapt standardized services	Maximum method [30]
No. of services required/offered by one user	1

After creating users' private SLA templates, the simulation of trade is performed. As previously mentioned, we separately simulate the trade of differentiated and standardized services.

The trade of **differentiated services** is started by buyers who "manually" search for appropriate trading partners. Hereby, buyers iterate through all sellers' offerings and compare them to their requirements by considering the values of service objectives. A match between a buyer's requirement and a seller's offering is found if the intersection of the SLO value ranges of all SLA parameters is not an empty set. Once a buyer finds a match, the search is stopped and the buyer is ready to establish a trade.

The trade of **standardized services** is started by the *k*-means clustering algorithm that is applied to group similar requirements and offers based on the SLO values from users' SLA templates [11]. The adaptation method named "maximum method" [30] is applied and a set of new public SLA templates is created (or the existing set updated). As described in [8], unlike private SLA templates, a public SLA template defines single parameter values for SLOs instead of ranges of acceptable values. Once the adaptation cycle is finished,

buyers and sellers iterate through the newly created public SLA templates and check whether the new products have the required specification. Specifications of the templates match only if the SLO values of all SLA parameters from a public SLA template are inside the value ranges specified by their private SLA templates. Note that only those users who have a matching specification are able to be part of a trade. Finally, buyers are matched with the randomly chosen sellers from the set of users who opt for the same public SLA template. Note that finding a matching public SLA template does not guarantee an allocation with another trading party, as it may happen that the demand and the supply are not balanced.

In both the differentiated and the standardized approaches, the overall market depth is measured as the accumulative trading volume, i.e., the total number of buyers and sellers who have received an allocation. On the other hand, search cost is measured as the number of comparisons between SLA templates that the users have to perform in order to find a suiting trading partner. These are the comparisons between the buyers' and the sellers' private SLA templates in case of the differentiated approach, and between users' (buyers' and sellers') private SLA templates and the public SLA templates in case of the standardized approach.

For the sake of simplicity, each market participant can have either the role of a buyer or a seller, but not both at the same time. Furthermore, each participant wishes to sell or purchase only one service during the whole simulation period. Therefore, once an allocation occurs (i.e., when a match between a seller's offer and a buyer's requirement is found), both the buyer and the seller are removed from the list of users who have not yet received an allocation. Although the latter limitation may contradict the scalability requirements of the cloud computing paradigm that promise virtually unlimited resources, it helps in simplification of the simulation model without affecting the approach of standardizing computational resources.

Regarding the credibility of the simulated approach with respect to the random creation of SLAs, note that the real-world "production" SLAs are currently very limited: they are used only for describing infrastructure services (i.e., in the infrastructure-as-a-service business model), while for the other models such as platform-as-a-service (PaaS) and software-as-a-service (SaaS) they are still not utilized. The SLA parameters contained by the SLAs used in our experiments are simple modifications of common production SLAs used in the IaaS model. However, it is noteworthy that the real-world SLAs are more complex than the simulated SLAs since they contain more SLA parameters and more differences in their definitions. It is also important to note that the motivation for our approach of standardizing computational resources gets on its significance with the additional complexity of users' SLAs. Therefore, we believe that our approach would demonstrate even better results with the real-world production SLAs than in the simulated environment. This analysis is, however, out of scope of this paper and will be examined in detail in our future work.

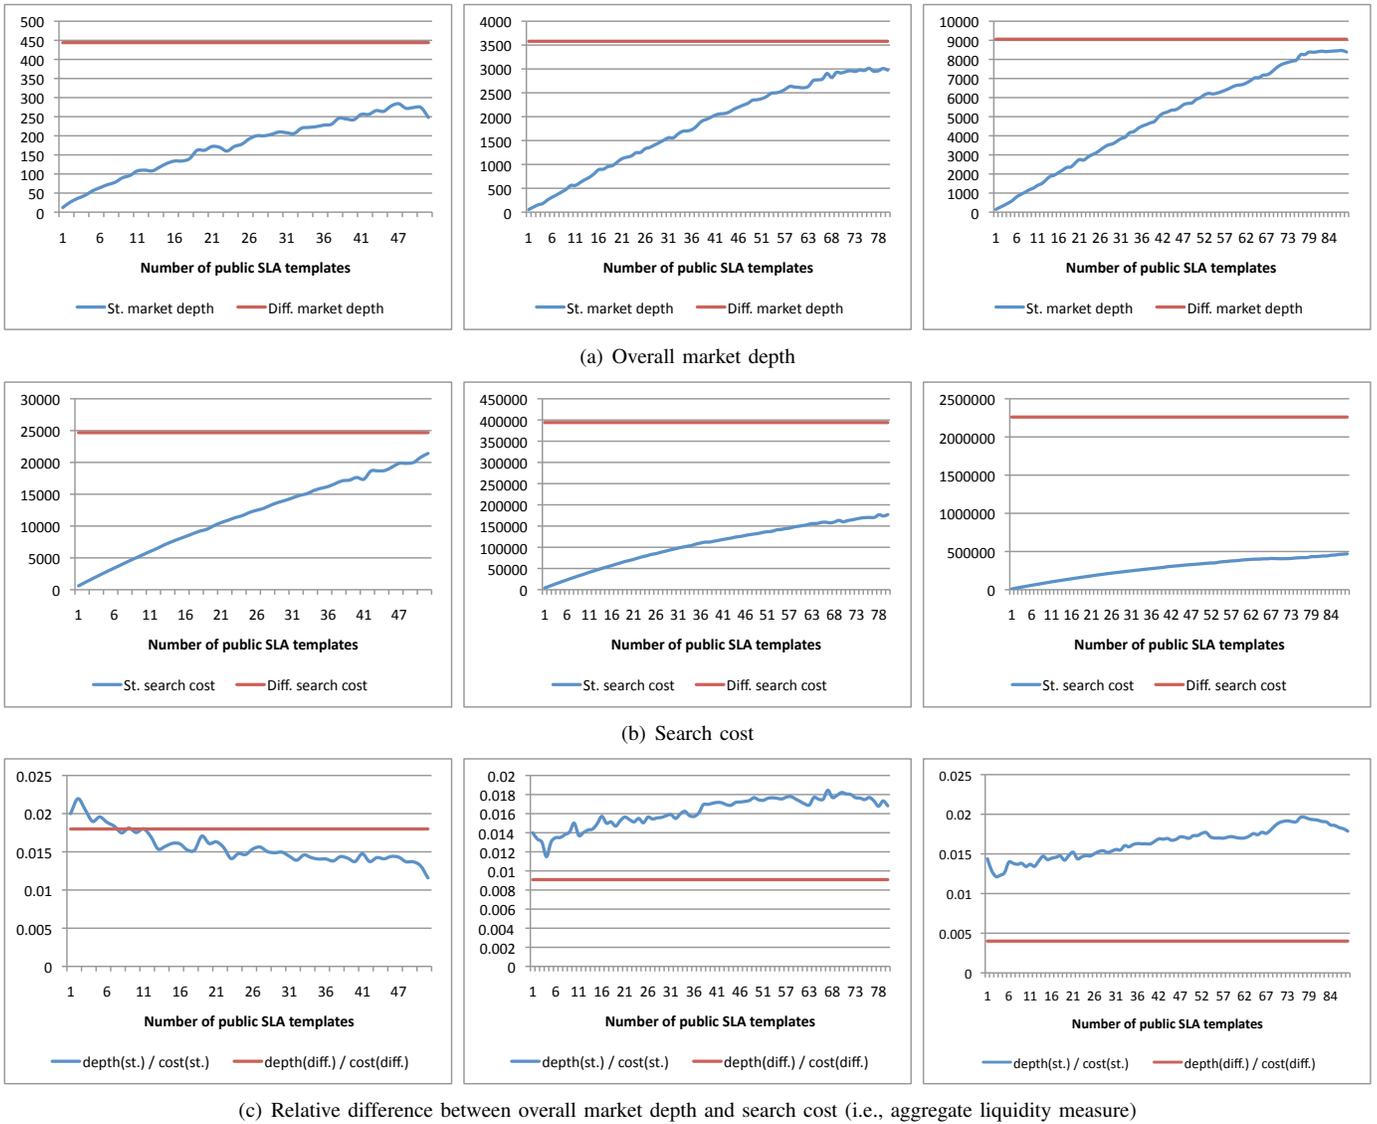


Fig. 1. Simulation results for 600, 4000, and 10000 market participants

B. Evaluation results

To begin our analysis of the impact of standardized products on market liquidity, we simulate a trade of exclusively differentiated products as well as a trade of standardized products in various environments, out of which we here discuss three: with 600, 4000, and 10000 market participants trading. In each of the scenarios, demand and supply are evenly distributed, i.e., 50% of traders are buyers and the remaining 50% sellers. Note that the latter assumption does not hold in the real-world scenarios as cloud consumers currently significantly outnumber cloud providers. However, this assumption increases market activity and improves visibility of effects of resource standardization on market liquidity. In our future work, we will investigate whether and to which extent the ratio between demand and supply affects the evaluation results.

Fig. 1 presents the simulation results. It contains 9 graphs arranged in 3 columns and 3 rows. Each of the rows presents

one of the liquidity measures: Fig. 1(a) depicts the overall market depth, Fig. 1(b) depicts the search cost, and Fig. 1(c) depicts their relative difference (also called “the aggregate liquidity measure”). In each of the rows, the left-hand graph presents the trade of 600 market participants, the middle graph the trade of 4000 participants, and the right-hand graph the trade of 10000 market participants. The horizontal axis depicts the number of created standardized products (i.e., public SLA templates) and the vertical axis depicts the result values of liquidity measures. Note that the change in the number of standardized products is the only change in the market condition that we consider. This change, of course, does not effect the trade of the differentiated goods. Nevertheless, in order to simplify the comparison between the two trading approaches, Fig. 1 depicts the continuous but constant values of liquidity measures for the “differentiated approach”.

1) *Overall market depth*: We begin our discussion on simulation results by considering overall market depth depicted in Fig. 1(a). The graphs in this figure present the expected dominance of the differentiated products over the standardized products in terms of the number of successful allocations of requirements and offers for services. Naturally, due to the high variety in resource types when trading differentiated products, probability to find an offer similar to a user's requirement is significantly higher. Regarding the standardized approach, the graphs obviously demonstrate that the overall market depth grows with the increasing number of standardized resource types in the market. In order to achieve the same amount of overall market depth as with the differentiated resources, the standardized approach should create a very large number of standardized resources. This, however, means that each standardized resource would be approximately equal to one private SLA template in the market and would only slightly differ from the differentiated approach. Moreover, since the number of standardized resources cannot be larger than the number of differentiated resources, we conclude that the standardized approach will always achieve a lower or equal value of the overall market depth when compared to the differentiated approach.

When comparing the overall market depth for the differentiated and the standardized products, an interesting result occurs. As shown in Fig. 1(a), the depth for the standardized approach, although always lower than for the differentiated approach, significantly rises with the number of market traders. Considering the values of the left-hand graph, we conclude that the standardized approach achieves the maximum of 64% of the depth value of the differentiated approach when there are 600 active traders in the market. This value is achieved with 47 standardized products, which is 6.4 times less than the number of products in the differentiated market. When the number of market participants is increased to 10000 (the right-hand graph), the standardized approach achieves up to 93.4% of the overall market depth of the differentiated approach, although the number of standardized resources increased only slightly. In this case, when the maximum depth is achieved there are 86 standardized resources in the market which is almost 58 times less than the number of differentiated products. The main reason for this behavior is diversity of resource types which is reasonably constant, no matter the number of active traders. Namely, with only a small number of traders, the number of different resource types in the market is large and the number of users requiring or offering one resource type is low. On the contrary, with a sufficiently large number of traders, the number of different resource types is only moderately higher, but with more users requiring a same resource type. This means that the number of standardized products needed to keep the overall market depth stable grows slowly with the number of market participants.

The discovery of the relatively constant diversity in resource types may lead to a conclusion that a fixed amount of standardized products is needed to achieve a certain level of market efficiency with the number of traders playing no role

in its determination. This is, however, not the case for several reasons. First, as depicted in Fig. 1(a), after a certain amount of standardized products is created, the growth of the overall market depth with every new standardized products decelerates because the currently existing standardized products already reflect the needs of most of the traders. Secondly, as it will be soon explained, this growth is not sufficiently high to cover the expenses of the introduction of more standardized products.

2) *Search cost*: Fig. 1(b) presents the effort needed to find a trading partner in the simulated market environment. On the contrary to the perspective of overall market depth, the differentiated approach is significantly inferior to the standardized approach when considering the search cost. For the differentiated goods, buyers must iterate through active sellers' offerings until they find a matching service. This means that the maximum search cost is

$$cost_{max}^{diff.} = no. buyers \times no. sellers. \quad (2)$$

In the standardized approach, all users (buyers and sellers) iterate through public SLA templates, which means that the maximum search cost is

$$cost_{max}^{st.} = (no. buyers + no. sellers) \times no. st. resources. \quad (3)$$

Since the number of buyers and sellers is always remarkably larger than the number of standardized resources, the effort needed in the differentiated approach is always larger when compared to the standardized approach. However, the realistic cost values are usually notably lower than the theoretical maximum in both the differentiated and the standardized approach for two reasons. Firstly, users usually find required services before iterating through the whole list of available resources. Secondly, in the differentiated approach, once an allocation occurs, the buyer's request and the seller's offer are not considered in future iterations, which means that the number of active (i.e., non-allocated) users is reduced. Nevertheless, search cost grows with the number of market participants and the quantity of standardized products. Since sellers appear in the market much faster and in larger quantities than the standardized resources, the positive effects of the standardization becomes greatly obvious with the increasing number of market participants.

In the differentiated approach, the growth of the search cost slows down with the increase in the number of active market traders. For instance, when there are 600 traders in the market, the measured cost is 27% of the maximum possible cost (Eq. (2)). On the contrary, when there are 10000 market traders, the value of the search cost is only 9% of the maximum cost. This is due to a large number of traders, which means that they can easily find a trading partner (due to the diversity in offerings) and have to perform a smaller number of search iterations. Similarly, the increase in the number of standardized products slows down the rise of the search cost in the standardized approach since users can easily find an appropriate service.

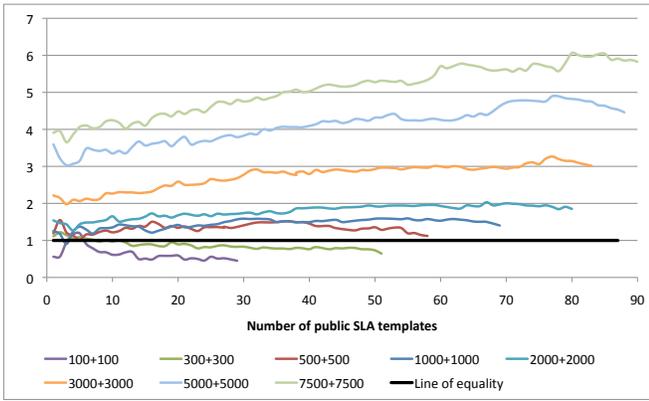


Fig. 2. Standardized vs. differentiated approach in terms of aggregate liquidity measure

3) *Aggregate liquidity measure*: The conclusion that can be drawn from the previous analysis of standardized vs. differentiated approach in terms of overall market depth and search cost is that the standardization of services facilitates the search for the appropriate services (i.e., it minimizes the cost necessary to perform that action), but provides a lower chance to find a match for a certain service requirement or an offering. It stays unclear, however, which of the two approaches creates a better trading environment, as each of them provides a better performance from one of the two aspects of liquidity. The doubt that remains is, therefore, whether the lower overall market depth of the standardized approach compensates for the low search cost and whether the standardization improves the overall market liquidity. To answer this question, we use the aggregate liquidity measure (Eq. (1)). Market liquidity is improved if the depth is increased while the effort is reduced. Therefore, the market platform has the goal of maximizing the relative difference between these two values.

Fig. 1(c) depicts the values of the aggregate liquidity measure and shows that the positive effects of the standardization become present only with a sufficiently high number of market participants. In a market with a limited amount of traders, a high diversity in resource types is distributed among a low number of market users and, as already explained, many standardized products are needed to achieve even a moderate overall market depth. At the same time, a large amount of standardized products increases the cost of searching for the most appropriate service offering. With the increase in the number of buyers and sellers, every new standardized product brings more to the overall market depth, but increases the search cost only slightly. Therefore, the more participants are in the market, the more benefits the standardization brings.

The question of how large the positive effect of the standardization to the market liquidity is can be answered by observing the results depicted in Fig. 2. The figure compares the two approaches by looking at the relative difference between the aggregate liquidity measures of the approaches, i.e.,

$$l = \frac{depth^{st.}/cost^{st.}}{depth^{diff.}/cost^{diff.}} = \frac{depth^{st.}}{depth^{diff.}} \times \frac{cost^{diff.}}{cost^{st.}}. \quad (4)$$

Essentially, Fig. 2 shows how many times the value of the aggregate liquidity measure in the standardized approach is higher than the value in the differentiated approach. The graph presents the measured values for the various number of market participants (in the legend depicted in the form of *no. buyers + no. sellers*) and the “line of equality”, which represents the value in which the two trading approach behave equally in terms of market liquidity (i.e., when the value of the aggregate liquidity measure is 1).

Fig. 2 emphasizes the limited performance of the standardized approach with the low number of traders, but also its great outperformance when the number of traders is sufficiently high. In the demonstrated scenario, the standardized approach achieves up to 6 times higher amount of “aggregate liquidity” with 15000 market participants, which is achieved by creating only 80 standardized products. This result is important and noteworthy as we demonstrated, using a simulation scenario, that the standardization of goods in small markets may even hurt the market efficiency and stability. On the contrary, it brings enormous savings and benefits in a market where the demand and the supply are sufficiently high.

VI. ESTIMATING THE “IDEAL” NUMBER OF STANDARDIZED RESOURCES

In the previous section, we demonstrated the positive effects of the standardization of computational resources on market liquidity in (simulated) electronic markets. In this section, we continue this study and analyze the possibilities of using methods for approximation of liquidity to maximize the benefits of the standardized approach. In particular, we look into automatic estimation of the number of standardized products that, when introduced, increase the aggregate liquidity measure to its maximum point. Since market liquidity is the main performance indicator for the “benefit”, the number of products is “ideal” when the liquidity is maximized. Finding this number is, therefore, a matter of finding the market setting in which a single liquidity measure achieves its maximal value. However, as already described in Section IV, there is no universal measure of market liquidity. Overall market depth and search cost cannot be taken as the only measures of liquidity as they both must be considered in order to balance low search cost and high overall market depth. The aggregate liquidity measure (i.e., the relative difference between the depth and the cost) may help to find the “optimal” point, but cannot be taken into consideration individually either. To demonstrate this with an example, in the left-hand graph of Fig. 1(c) it seems that, considering the value of the aggregate liquidity measure, the standardized approach outperforms the differentiated approach when there are only a few standardized products in the market. This, however, does not hold, as a user’s probability to find a trading partner at this point is almost nonexistent. A market with such a low matching probability has almost no benefits for buyers and sellers who would, in this case, almost certainly leave the market. Therefore, despite extremely low search costs that increase the aggregate liquidity measure, these results cannot be taken into consideration. However, as the aggregate

liquidity measure is the closest we can get to the universal indicator of market liquidity, we address this issue by taking only those scenarios into consideration in which the overall market depth of the standardized approach reaches at least 50% of the value achieved by the differentiated approach. This step ensures at least a moderately satisfying outcome to the users. For the cases in which overall market depth is lower than this predefined threshold, we conclude that the standardization does not pay off and that only differentiated products should be traded.

Fig. 3(a) depicts the number of standardized products created for a certain number of buyers and sellers in the market when the aggregate liquidity measure is maximized. Due to considering only the values where the overall market depth is sufficiently high, the scenarios with the lower number are not represented. The figure presents the market behavior already discussed in Section V: the “ideal” number of standardized products increases with a significantly slower pace than the number of traders. Moreover, the growth slows down with the increase in the number of traders. Therefore, after a certain number of standardized products are created, there are not many benefits of introducing additional products, no matter the number of market participants.

The usefulness of the market behavior depicted in Fig. 3(a) is limited, as it can only be used to estimate the “ideal” number of standardized products in a low number of situations: when the number of sellers and buyers in the market is high enough, the introduction of additional products does not have a large effect on market liquidity. As already discussed, in this case the “ideal” number of standardized products stays constant. To avoid this limitation, we consider the same results from a different perspective. Fig. 3(b) depicts the number of buyers and sellers per one standardized product in the market. The linear behavior of the graph shows that the diversity of resource types rises only slowly with the number of market participants, but the number of users using the same resource type increases. Therefore, in a market with more traders, the number of standardized products is not high, but the number of users per one product is. This measure does not only illustrate market behavior but may also help to estimate the “ideal” number of standardized products. In particular, the linear growth presented in Fig. 3(b) can be estimated with the following equation:

$$no. users per resource = 0.12 \times no. users + 13.17 \quad (5)$$

The prediction probability is larger than 99% (i.e., with an R^2 value larger than 0.99). In statistics, the coefficient of determination R^2 is the proportion of variability in a data set that is accounted for by the statistical model ([42], pp. 187, 287). It provides a measure of how well future outcomes are likely to be predicted by the model. Note that in the given equation, *no. users* represents the sum of the number of buyers and the number of sellers in the market.

With respect to Eq. (5), the “ideal” number of standardized

resources can be estimated using the following equation:

$$no. resources = \left\lfloor \frac{no. users}{0.12 \times no. users + 13.17} \right\rfloor \quad (6)$$

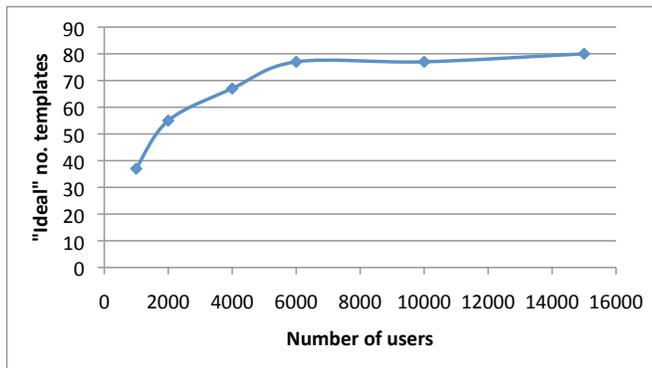
Note, however, that this result is valid only for the given trading scenario. The estimation function and the certainty depend on the demand and supply, i.e., the diversity in resource types in the market, and differ in other (real-case) scenarios. However, having a (reasonable) assumption that the diversity of users’ requirements is relatively limited even in the real implementations of electronic markets, the same estimation method may be used in those environments, but with Eq. 5 adapted to the observed demand and supply.

If properly used, the presented method, i.e., the appropriately adapted Eq. 6, can be used to efficiently estimate the “ideal” number of standardized resources for every market situation. Namely, after having enough data to build the estimation function with a sufficiently high prediction probability, it is not necessary to check the varying number of standardized resources in order to find the “ideal” number, but it is possible to use the equation to quickly compute it. In order to allow a particular certainty, another method can be applied to confirm that indeed the number with the maximum aggregate liquidity measure has been selected. Namely, as the simulation scenario presented in Section V-B has demonstrated, after the maximal aggregate liquidity value has been reached, the dynamics of the growth of overall market depth suddenly change: even if a larger quantity of standardized products is introduced, the depth stops rising or even loses its value (cf. Fig. 1(a)). Therefore, it is possible to check whether any close number of standardized resources larger than the chosen “ideal” number increases the overall market depth and, thus, validate the estimated quantity.

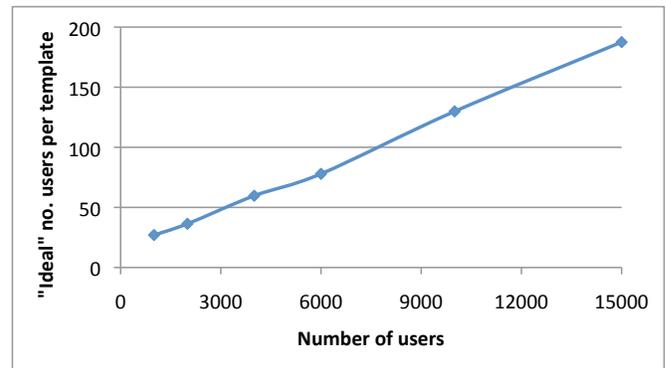
VII. CONCLUSION AND FUTURE WORK

Despite many attempts, an implementation of an efficient and dynamic market model for trading computational services is still missing. Due to the broad resource variability and still low number of market participants, the lack of standardized computational services can cause low market liquidity and reduce the attractiveness of the market platform. Addressing this issue may decide whether (or when) cloud computing will finally make traditional vendor relationships obsolete.

In this paper, we emphasized the need for quantifying liquidity in electronic markets and derived a set of liquidity measures from the most notable literature on assessment of financial markets. Using the derived “liquidity model”, we demonstrated substantial benefits that the standardization of computational resources brings in terms of market liquidity. Furthermore, we used the measure of liquidity to determine the “optimal setting” of a market, i.e., a quantity of standardized products that should be created so that the liquidity is maximized. The usefulness of the derived method is manifold. First, it improves liquidity and, therefore, market activity and attractiveness, resulting in more numerous trades and potentially attracting more market participants. Second, this



(a) "Ideal" number of SLA templates



(b) "Ideal" number of users per SLA template

Fig. 3. Estimating the "ideal" number of standardized goods

method can be used by service providers to decide on the set of services they should offer so that the probability to find a buyer is maximized, therefore increasing providers' profit and the market revenue.

In our future work, in addition to market liquidity, we will consider other measures of market quality and analyze the possibility of creating a "(near) optimal market setting", in which all of the measures are balanced and maximized. Furthermore, we will consider more complicated SLA specifications, in order to demonstrate the benefits of resource standardization in the real-world trading scenarios.

REFERENCES

- [1] J. Altmann, C. Courcoubetis, and M. Risch, "A marketplace and its market mechanism for trading commoditized computing resources," *Annals of Telecommunications*, vol. 65, pp. 653–667, 2010.
- [2] Y. Amihud and H. Mendelson, "Asset pricing and the bid-ask spread," *Journal of Financial Economics*, vol. 17, no. 2, pp. 223 – 249, 1986.
- [3] D. Ardagna, B. Panicucci, and M. Passacantando, "A game theoretic formulation of the service provisioning problem in cloud systems," in *Proceedings of the 20th international conference on World wide web*, ser. WWW '11. New York, NY, USA: ACM, 2011, pp. 177–186.
- [4] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia, "Above the clouds: A Berkeley view of cloud computing," EECS Department, University of California, Berkeley, Tech. Rep., Feb. 2009.
- [5] M. J. Barclay, T. Hendershott, and D. T. McCormick, "Competition among trading venues: Information and trading on electronic communications networks," *The Journal of Finance*, vol. 58, no. 6, pp. 2637–2665, 2003.
- [6] E. Boehmer, G. Saar, and L. Yu, "Lifting the veil: An analysis of pre-trade transparency at the nyse," *Journal of Finance*, vol. 60, no. 2, pp. 783–815, Apr. 2005.
- [7] I. Brandic, D. Music, and S. Dustdar, "Service mediation and negotiation bootstrapping as first achievements towards self-adaptable grid and cloud services," in *Grids meet Autonomic Computing Workshop 2009. In conjunction with the 6th International Conference on Autonomic Computing and Communications*. Barcelona, Spain: ACM, 2009, pp. 1–8.
- [8] I. Breskovic, J. Altmann, and I. Brandic, "Creating standardized products for electronic markets," *Future Generation Computer Systems*, 2012.
- [9] I. Breskovic and I. Brandic, "Increasing the popularity of cloud computing by improving its market performance," *Awareness Magazine: Self-Awareness in Autonomic Systems*, 2011.
- [10] I. Breskovic, C. Haas, S. Caton, and I. Brandic, "Towards self-awareness in cloud markets: A monitoring methodology," in *Proceedings of the 9th IEEE International Conference on Dependable, Autonomic and Secure Computing*, ser. DASC '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 81–88.
- [11] I. Breskovic, M. Maurer, V. C. Emeakaroha, I. Brandic, and J. Altmann, "Towards autonomic market management in cloud computing infrastructures," in *Proceedings of the 1st International Conference on Cloud Computing and Services Science*, ser. CLOSER'11, 2011, pp. 24–34.
- [12] —, "Achieving market liquidity through autonomic cloud market management," in *Cloud Computing and Services Science*, ser. Service Science: Research and Innovations in the Service Economy, I. Ivanov, M. van Sinderen, and B. Shishkov, Eds. Springer New York, 2012, pp. 91–107.
- [13] I. Breskovic, M. Maurer, V. C. Emeakaroha, I. Brandic, and S. Dustdar, "Cost-efficient utilization of public sla templates in autonomic cloud markets," in *Proceedings of the 4th IEEE International Conference on Utility and Cloud Computing*, ser. UCC '11. Washington DC, USA: IEEE Computer Society, 2011, pp. 229–236.
- [14] R. Buyya, D. Abramson, and J. Giddy, "A case for economy grid architecture for service oriented grid computing," *Parallel and Distributed Processing Symposium*, vol. 2, 2001.
- [15] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," *Future Generation Computer Systems*, vol. 25, pp. 599 – 616, 2009.
- [16] S. Caton, I. Breskovic, and I. Brandic, "A conceptual framework for simulating autonomic cloud markets," in *Proceedings of the 3rd International Conference on Cloud Computing*, ser. CloudComp '11. Washington DC, USA: IEEE Computer Society, 2012.
- [17] S. Chaisiri, R. Kaewpuang, B.-S. Lee, and D. Niyato, "Cost minimization for provisioning virtual servers in amazon elastic compute cloud," in *Proceedings of the 2011 IEEE 19th Annual International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, ser. MASCOTS '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 85–95.
- [18] G. C. Chako, J. W. Jurek, and E. Stafford, "The price of immediacy," *The Journal of Finance*, vol. 63, no. 3, pp. 1253–1290, 2008.
- [19] Y. Cheng, A. Leon-Garcia, and I. Foster, "Toward an autonomic service management framework: A holistic vision of SOA, AON, and autonomic computing," *Communications Magazine, IEEE*, vol. 46, no. 5, pp. 138–146, May 2008.
- [20] L. Chimakurthi and M. K. SD, "Power efficient resource allocation for clouds using ant colony framework," *CoRR*, vol. abs/1102.2608, 2011.
- [21] K. P. Clark, M. Warnier, and F. M. T. Brazier, "An intelligent cloud resource allocation service - agent-based automated cloud resource allocation using micro-agreement," in *Proceedings of the 2nd International Conference on Cloud Computing and Services Science*, F. Leymann, I. Ivanov, M. van Sinderen, and T. Shan, Eds. SciTePress, 2012, pp. 37–45.
- [22] Committee on the Global Financial System, "How should we design deep and liquid markets? The case of government securities," in *CGFS Working Group Reports*, no. 11, Oct. 1999. [Online]. Available: <http://www.bis.org/publ/cgfs13.pdf>
- [23] L. R. Glosten and P. R. Milgrom, "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Northwestern University, Center for Mathematical Studies in Economics and Management Science, Discussion Papers 570*, Aug. 1983.

- [24] L. Harris, *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press, 2002.
- [25] L. He and T. R. Ioerger, "Forming resource-sharing coalitions: a distributed resource allocation mechanism for self-interested agents in computational grids," in *Proceedings of the 2005 ACM symposium on Applied computing*, ser. SAC '05. New York, NY, USA: ACM, 2005, pp. 84–91.
- [26] T. Hendershott, C. M. Jones, and A. J. Menkveld, "Does algorithmic trading improve liquidity?" *The Journal of Finance*, vol. 66, no. 1, pp. 1–33, 2011.
- [27] R. D. Huang and H. R. Stoll, "Dealer versus auction markets: A paired comparison of execution costs on nasdaq and the nyse," *Journal of Financial Economics*, vol. 41, no. 3, pp. 313–357, Jul. 1996.
- [28] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, pp. 41–50, Jan. 2003.
- [29] C. Lee and J. Suzuki, "An autonomic adaptation mechanism for decentralized grid applications," in *Consumer Communications and Networking Conference, 2006. CCNC 2006. 3rd IEEE*, vol. 1, Jan. 2006, pp. 583–589.
- [30] M. Maurer, V. C. Emeakaroha, I. Brandic, and J. Altmann, "Cost-benefit analysis of an SLA mapping approach for defining standardized cloud computing goods," *Future Generation Computing Systems*, vol. 28, no. 1, pp. 39–47, 2012.
- [31] A. B. Mohammed, J. Altmann, and J. Hwang, "Cloud computing value chains: Understanding businesses and value creation in the cloud," in *Economic Models and Algorithms for Distributed Systems*, ser. Autonomic Systems, D. Neumann, M. Baker, J. Altmann, and O. Rana, Eds. Birkhäuser Basel, 2010, pp. 187–208.
- [32] D. Neumann, J. Stösser, and C. Weinhardt, "Bridging the adoption gap—developing a roadmap for trading in grids," *Electronic Markets*, vol. 18, pp. 65–74, Feb. 2008.
- [33] J. Nimis, A. Anandasivam, N. Borissov, G. Smith, D. Neumann, N. Wirstm, E. Rosenberg, and M. Villa, "SORMA: Business cases for an open grid market: Concept and implementation," in *Grid Economics and Business Models*, ser. Lecture Notes in Computer Science. Springer Berlin / Heidelberg, 2008, vol. 5206, pp. 173–184.
- [34] W. H. Oyen and S. A. Deloach, "Towards a systematic approach for designing autonomic systems," *Web Intelligence and Agent Systems*, vol. 8, pp. 79–97, Jan. 2010.
- [35] D. Pardoe, P. Stone, M. Saar-Tsechansky, and K. Tomak, "Adaptive mechanism design: a metalearning approach," in *Proceedings of the 8th international conference on Electronic commerce*, ser. ICEC '06. ACM, 2006, pp. 92–102.
- [36] F. Pascual, K. Rzdca, and D. Trystram, "Cooperation in multi-organization scheduling," *Concurrency and Computation: Practice and Experience*, vol. 21, no. 7, pp. 905–921, May 2009.
- [37] M. Risch and J. Altmann, "Enabling open cloud markets through WS-agreement extensions," in *Grids and Service-Oriented Architectures for Service Level Agreements*, P. Wieder, R. Yahyapour, and W. Ziegler, Eds. Springer US, 2010, pp. 105–117.
- [38] M. Risch, J. Altmann, L. Guo, A. Fleming, and C. Courcoubetis, "The gridecon platform: A business scenario testbed for commercial cloud services," in *Proceedings of the 6th International Workshop on Grid Economics and Business Models*, ser. GECON '09. Springer-Verlag, 2009, pp. 46–59.
- [39] B. Schnizler, D. Neumann, D. Veit, and C. Weinhardt, "Trading grid services - a multi-attribute combinatorial approach," *European Journal of Operational Research*, vol. 187, no. 3, pp. 943–961, Jun. 2008.
- [40] K. M. Sim, "Towards complex negotiation for cloud economy," in *Proceedings of the 5th international conference on Advances in Grid and Pervasive Computing*, ser. GPC'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 395–406.
- [41] G. Singh, C. Kesselman, and E. Deelman, "Application-level resource provisioning on the grid," in *Proceedings of the Second IEEE International Conference on e-Science and Grid Computing*, ser. E-SCIENCE '06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 83–.
- [42] R. G. D. Steel and J. H. Torrie, *Principles and Procedures of Statistics*. McGraw-Hill New York, 1960.
- [43] W. Streitberger and T. Eymann, "A simulation of an economic, self-organising resource allocation approach for application layer networks," *Computer Networks*, vol. 53, pp. 1760–1770, Jul. 2009.
- [44] H. Tianfield, "Towards agent based grid resource management," in *IEEE International Symposium on Cluster Computing and the Grid*, vol. 1, May 2005, pp. 590–597.
- [45] G. Wei, A. V. Vasilakos, Y. Zheng, and N. Xiong, "A game-theoretic method of fair resource allocation for cloud computing services," *The Journal of Supercomputing*, vol. 54, no. 2, pp. 252–269, Nov. 2010. [Online]. Available: <http://dx.doi.org/10.1007/s11227-009-0318-1>
- [46] G. Yong, Y. Li, Z. Wei-ming, S. Ji-chang, and W. Chang-ying, "Methods for resource allocation via agent coalition formation in grid computing systems," in *Proceedings of the 2003 IEEE International Conference on Robotics, Intelligent Systems and Signal Processing*, vol. 1, Oct. 2003, pp. 295–300.
- [47] S. S. Zhang, M. Wagener, A. Storckenmaier, and C. Weinhardt, "The quality of electronic markets," in *Proceedings of the 2011 44th Hawaii International Conference on System Sciences*, ser. HICSS '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 1–10.