How to trade Electronic Services?
Current Status and Open Questions

Benjamin Blau, Carsten Block, Jochen Stößer
Universität Karlsruhe (TH), Germany
{blau; block; stoesser}@iism.uni-karlsruhe.de

Abstract

Electronic Services become more and more important for our daily life. News and communication services are among the most prominent examples that drastically transformed the way we keep ourselves informed and relate to each other. But new application areas for electronic services such as grid computing, security and surveillance, ambient assisted living, or intelligent facility management especially with focus on energy optimization are already emerging. All these services are tradeable goods meant to add value to our daily life and thus come at a price. This is why in this paper we focus on research done and future research on mechanisms for trading services. Our main claim is that different types of services need different types of trading mechanisms. Thus, at first we classify electronic services into different categories and then describe suitbale trading mechanisms for each of these categories. Despite the work done already, a lot of additional research has to be accomplished in this particular field. We therefore conclude our paper with a roadmap for future research on how to trade services.

Keywords: Markets, Auctions, Electronic Services, Service Classification

---

1 Introduction

The dynamic allocation and pricing of IT services has become an important application domain for market mechanisms, in particular for auctions. As IT services are becoming increasingly standardized, there is a trend towards the dynamic sourcing of IT services from specialized service providers that economize on scale and scope (Rappa, 2004). Amazon's Elastic Compute Cloud (`http://aws.amazon.com/ec2`), Sun Microsystem's network.com (`http://www.network.com`) and Salesforce's force.com (`http://www.force.com`) are prominent precursors of this trend. The potential benefits of computing services are straightforward: lower fixed costs for hardware, software licenses and maintenance, less energy consumption for electricity and cooling, state-of-the-art services, and a focus on key competencies, processes and products.

Furthermore we claim that in future different types of services such as those for ambient assisted living, facility management, energy optimization, or surveillance and security will be combined and executed on unified generic operation platforms. These platforms serve as a technical foundation for services to run on and mainly provide common access to sensor and actor infrastructures as well as to communication facilities while enforcing authentication and authorization constraints. Already today a first generation of service market places such as `http://www.strikeiron.com` is established where different types of services (ranging from simple IP-lookups to complex creditworthiness checks) with different business models are offered at different prices to different types of consumers. We believe that more of these market places will evolve and that they will become ever more integrated into our daily life. Thus sometimes in future we will probably be able to decide spontaneously which entertainment, information, communication, energy optimization or other services to consume sitting in our living rooms, which are then equipped with new human computer interfaces that seamlessly integrate into our daily life.

This is the reason why in this paper we aim at describing the *current status* in the design of market mechanisms that we see as key-enablers for the efficient trading of such services. Our main claim is, that different types of market mechanisms have to be developed for different types of services. Moreover, we want to highlight important – and from our point of view – not yet resolved issues in the context of service trading, which may serve as a *roadmap for future research.*

This paper is structured as follows. In Section 2 we describe a basic classification scheme for distinguishing three different types and layers of IT services in a *Service Decomposition Model.* In Section 3 we summarize the most important requirements that arise when designing market mechanisms for such service settings and we briefly outline three different scenarios for the application of such service trading mechanisms. At the core of this paper, in Section 4 we propose a roadmap for future research in the trading of services to increase awareness and stir up discussion on this topic.

2 Classification of Services

In this section we give a thorough classification of groups of services that share common characteristics from a technical and economic perspective as depicted in Figure 1. Our Service Decomposition Model is based on a classification by Blau & Schnizler (2008). Our model distinguishes three different service layers: Utility Services, Elementary Services and Complex Services.

2.1 Utility Services

Utility Services reflect a vision where (IT) services can be accessed dynamically in analogy to electricity and water: "Utility computing is the on-demand delivery of infrastructure, applications, and business processes in a security-rich, shared, scalable, and standards-based computer environment over the Internet for a fee. Customers will tap into IT resources – and pay for them – as easily as they now get their electricity or water." (Rappa, 2004) According to Rappa, utilities are characterized by necessity, reliability, ease of use, fluctuating utilization patterns, and economies of scale. Rappa suggests to base pricing in utility computing on metering usage (also coined "pay-what-you-use" or "pay-as-you-go"), as is the case with classic utilities such as water, telephone and Internet access. With the fast rise of energy prices, the meaning of utility services is even extended back to the roots where the name originally came from: Chase *et al.* (2001) describe how basic computing services in hosting centers need to be managed explicitly taking into account energy consumption as a relevant optimization criterion. Bianchini & Rajamony (2004) describe how "heterogeneous server clusters can be made more efficient by conserving power and energy while exploiting information from the service level, such as request priorities established by service-level agreements" while Moore *et al.* (2008) propose even temperature aware computing solutions for data centers.

One can easily see that for most of the utility services a separate and independent trading would lead to inefficient results as in such a case, individuals could for example be allocated with computing resources on one market without being able to acquire the necessary energy resources on the other or vice versa. Furthermore if services that make a difference on the outcome are not priced properly, inefficient allocations are inevitable. If for example waste heat produced by data centers would have a market price it would no longer be disposed to the atmosphere. Instead low prices for waste heat would attract businesses that are in need of heat and that are willing to pay a certain price for it as long as it remains below their own reservation cost. But even if not such businesses exist, in a concrete case, one could still think of using the waste heat as input for absorption coolers that are able to turn at least parts of this (valuable) waste heat into (valuable) cooling energy, which in turn could be used within the data center again.

However, even in this metered model, prices are temporarily static and do not fully reflect the dynamics of demand and supply. Moreover, setting appropriate prices is a complex task for utility service providers. This is where we think (auction) markets in combination with electronic bidding agents should come into play as the former are an efficient instrument of determining prices based on demand and supply while the latter provide the level of automation

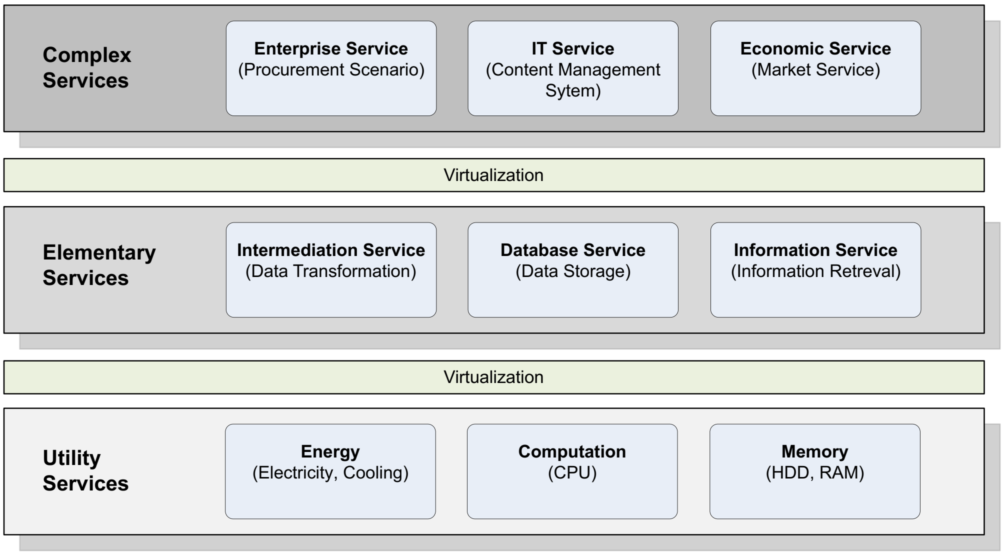

Figure 1: Service Decomposition Model

needed for a wide adoption in practice.

2.2 Elementary Services

Elementary Services provide basic functionality such as virtualization of physical resources or intermediation services. They can be fully described by well-defined interfaces consisting of simple attributes such as throughput, response time or reliability that are specified through well-established standards such as WSDL. Therefore elementary services yield clear semantics of input and output capabilities. From an economic perspective a lot of research has been done in the field of market mechanisms for trading homogenous services in different environments such as the Web or in Grids. Here market mechanisms are deemed promising since they induce service requesters to make more efficient use of scarce resources (e.g. distributing demand across time if prices are high) (Lai, 2005). Moreover, resource owners have an incentive to contribute to Grids in return for the market price.

2.3 Complex Services

Services in this layer are characterized by high degree of specialization and heterogeneity. Description of complex services in a standardized manner can hardly be accomplished using well-established formalisms and therefore demands for ontology-based description frameworks as introduced by Blau *et al.* (2008); Lamparter *et al.* (2007). Complex Services facilitate a vast variety of elementary services combined into a network topology that is shaped by service configurations, interrelations and dependencies. Functionality of multiple sub-services offered by different decentralized providers contributes to value-added complex services fulfilling an overall goal. An example for a complex service is an enterprise service as part of a software-as-a-service business model that offers the realization of a complete business scenario consisting of interdependent processes such as procurement or service order processing. Another example is a video surveillance service, which requires sub-services for capturing, analyzing, and storing video streams (`http://www.sorma-project.eu`). Research with respect to description frameworks and market mechanisms for trading complex services is still in its infancy and consequently needs thorough investigation.

3 Requirement Analysis

In this section we provide an overview of requirements that trading mechanisms for different types of services commonly have to fulfill. These requirement are divided into two different categories, (i) domain-specific requirements (based on empirical evidence) that target practical and applicability issues and (ii) economic requirements that address desirable mechanism properties and theoretical aspects (cf. Parkes (2001); Jackson (2003)). Subsequently, for each of the service layers depicted in Figure 1, we shortly describe an exemplary market mechanism used to allocate and price services from the respective layer and highlight the most important requirements for each scenario.

3.1 Domain-specific requirements

Requirement 1 ***Computational tractability:*** *The market outcome (allocations and prices) need to be determined in polynomial runtime in the size of the market input, that is the number of service requests and offers.*

Requirement 2 ***Combinatorics:*** *Service requesters oftentimes need combinations of services. Only obtaining a subset of such a combination is of no value and reduces the likelihood of efficient allocations. Reducing the "exposure" risk for bidders is thus desirable (Rothkopf* et al.*, 1998).*

Requirement 3 ***Time constraints:*** *Service requesters and providers must be able to specify time constraints, e.g. to support advance reservation and service level guarantees. Furthermore, with time constraints in place, different prices for the the same service offered at different times might evolve. This sets an incentive to service consumers to avoid high price time slots where possible while service providers are stimulated to provide more of their services during these times. Some studies e.g. in the utility business indicate that such a time aware trading of services can lead to load shifting bahavior and thus a better leveling of the service production and consumption overall (Strapp* et al.*, 2007).*

Requirement 4 ***Technical feasibility*** *Market allocations have to be compliant with technical boundaries. Capacity limitations of (computer or electricity) grids may for example render a market allocation where demand and supply from opposite sites of the network are matched infeasible as it would overstretch existing network capacities. Another example are ramp-up times for machines that have to be taken into account when finding market allocations. The simplest approach to avoid technical difficulties is to internalize these constraints into the market, but usually the computational complexity of the market mechanisms then increases drastically so that the problem is only shifted but not eliminated.*

Requirement 5 ***Incomplete Information:*** *When designing markets we usually assume that all participants have all relevant information available to make optimal decisions. In practice this is usually not the case. The provisioning of real time prices or real time resource monitoring usually impose serious technical problems so that technical monitoring limitations have to be taken into account already during the design phase of the mechanisms.*

Requirement 6 ***Task Automation:*** *While allocating services through markets can be shown to be economically efficient, it will also be a tedious task to service consumers and producers. Thus automating the trading of services by means of electronic agents is desirable in order to increase acceptance and thus adoption in practice.*

3.2 Economic requirements

Requirement 7 ***Allocative efficiency:*** *The market should maximize the system's overall value by allocating the most valuable service requests to the most cost-efficient providers.*

Requirement 8 ***Budget-balance:*** *The market must be self-sustainable in that it does not need to be subsidized by outside payments. The payments from the demand-side of the market must cover the payments to the service providers.*

Requirement 9 ***Individual rationality:*** *Market participants must not suffer any loss from participating in the market.*

Requirement 10 ***Incentive compatibility:*** *In order to be able to maximize the* true *allocative efficiency, service requesters and providers must be induced to truthful reports of their characteristics.*

3.3 Customization and Pricing of Complex Services

Addressing the issues that come along with customization and pricing of complex services, Blau *et al.* (2008) propose an ontology framework that facilitates the design and description process of complex services. This process results in a graph of subservice instances that together form a complex service instance (R2) as depicted in Figure 2.

In order to provide an overall functionality through a complex service, adequate services from each functional cluster have to be allocated. Each service offer is fully specified through a set of attributes and internal costs that the provider has to bear for a service invocation by a certain predecessor service. Consequently every feasible path from source to sink within the network yields a valid instantiation of the complex service.
For instance, a finance service that computes the risk of a portfolio facilitates sub-services that provide computing, storage and information retrieval functionality. Hence, each functional cluster contains subservice offers that provide the same type of functionality. Each sub-service provider has to bear internal costs for service invocation which are depend on the predecessor service. The service requester wants to buy a combination of sub-services that satisfies the overall functionality and yields the highest utility based on her preferences.
Based on such a network Blau et al. design a path auction that allows service providers to announce prices for invocation and usage of their services. The path auction is designed in a way that requesters can specify multiattribute utility functions that not only depend on the price but also on non-functional service characteristics. The utility accounts for different aggregation schemes of attributes such as *and*, *average*, *minimum*, *maximum*, *sum*. In such a multiattribute setting the payment scheme induces truthful revelation (R10) of the reservation prices and all announced attribute values of the sub-service providers. This ensures that the – in terms of requesters utility – most efficient complex service satisfying (R7) and (R9) will be identified and subsequently selected. Due to the high complexity in combinatorial auctions, the allocation problem is NP-hard and therefore not computational feasible which is a major drawback especially for online mechanisms. Proposed multiattribute path auction implicitly defines feasible combinations of sub-services within the network topology. This representation of suitable bundles reduces the combinatorial complexity and consequently enables a solution of the allocation problem in polynomial time which meets (R1).

3.4 Market-based Scheduling of Elementary Computing Services

Grid computing is a computing paradigm where utility services such as CPU and memory are shared across administrative boundaries, e.g. between enterprises and / or scientific computing centers. Organizations only need to accommodate the basic load on local resources, which leads to lower hard- and software expenses and requires less energy for electricity and cooling. Scheduling becomes a key challenge in this setting due to its inter-organizational character and dynamic demand and supply. Market mechanisms are deemed promising to lead to efficient resource allocations by explicitly targeting these characteristics and by providing incentives to contribute idle resources to the grid in return for the market price. Most computing services are elementary services since they consist of well-defined resources such as CPU, memory and bandwidth. The key requirements in this environment are computational tractability (R1), combinatorics (R2), and a budget-balanced, individually rational market (R8, R9). Interactive applications require the *timely allocation* of *bundles of resources* (e.g. CPU *and* memory). The market must support trading by both resource providers and requesters, and must be self-sustainable. The contribution of Stößer *et al.* (2007) is the proposal of a greedy, market-based scheduling heuristic which achieves this distinct trade-off: it is designed so as to obtain an approximately efficient allocation schedule at low computational cost while accounting for strategic, self-interested users in a heterogeneous environment.

Users ususally cannot be expected to continuously monitor the dynamic market situation and the requirements of the application which is to be executed remotely. Hence, a key issue for future research will be the design and implementation of "intelligent" tools that assist the users in interacting with the market (cf. MacKie-Mason & Wellman (2006)).

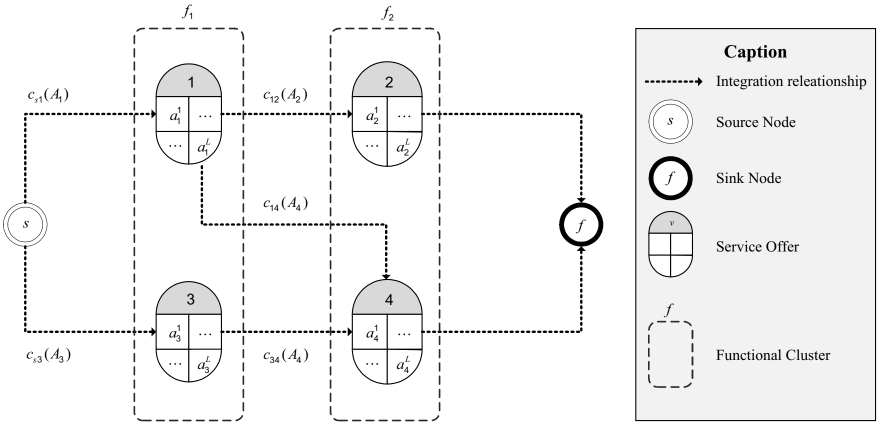

Figure 2: Formal Model of a Complex Service

3.5 Coallocation of utility services

For a long time utility services were considered to be ubiquitously available and cheap. This perception has changed recently when energy prices started to rise strongly (OECD, 2007). Increasing prices make energy a relevant (cost) factor for the provisioning of utility services such as CPU power. Sun & Lee (2006) found in an empirical study that (i) data centers overall are highly intensive energy consuming areas, and (ii) that up to 70% of the total energy consumption were not devoted to computing but to heating, ventilating, air conditioning (HVAC), lighting, and uninterruptable power supply (UPS), which is inline with findings reported by Ziff Davis (2005). Greenberg *et al.* (2006) state that *"the energy used by a typical rack of state-of-the art servers, drawing 20 kilowatts of power at 10 cents per kWh, uses more than $17,000 per year in electricity."* and a recent report by EPA (2007) found the energy used by by U.S. servers and data centers to be *"about 61 billion kilowatt-hours (kWh) in 2006 (1.5 percent of total U.S. electricity consumption) for a total electricity cost of about $4.5 billion.*

Overall it becomes obvious that trading mechanisms for utility services should not be focused on IT components like CPU power, disk space, or RAM only, but in future have to take real utilities such as power, cooling, or heating into account as well in order to achieve the goal of providing a given service level at minimum overall cost. The key requirements for trading utility services are thus technical feasibility (R4), combinatorics (R2), time constraints (R3), computational tractability (R1), and – for a wide adoption – automation (R6). As with the other services, the market should be designed to be budget-balanced, individually rational and allocative efficient (R8, R9, R7) in order to ensure the overall efficiency of this trading scheme. The contribution of Block *et al.* (2008) is a mechanism for trading basic utility services in small scale grid infrastructures, explicitly taking technical constraints such as minimum and maximum load levels or bundling requirements (in particular for cogenerated power and heat resources) into account.

4 Roadmap for future research

All examples described in the previous section meet only subsets of the aforementioned requirements, thus in this section we describe some of the issues that have to be taken into account when trading services but that have not been properly researched and integrated into existing trading mechanisms.

4.1 Market Concatenation

As introduced in Section 3, there can be multiple market mechanisms both across the layers of our service decomposition model as well as within one layer. Up to now, each of these mechanisms has been investigated in isolation. But in practice, these mechanisms may be closely intertwined. For example, a mechanism for the pricing of a complex service may depend on mechanisms for basic and / or utility services. An interesting question for future research may thus be to model and investigate these dependencies: How does strategic behavior change if a user acts on multiple markets at the same time? And how does the computational complexity of the markets on one service layer affect the tractability of the markets on another service layer?

4.2 Pricing of Complex Services

The heterogenic character of complex services as discussed in Section 2 implies many challenges when it comes to provision and price determination. Especially in a distributed environment, where decentralized providers contribute to the functionality of complex services, prices are difficult to determine. Self-interested participants in a value creation network that forms a complex service follow and adapt strategies that maximize their individual utility without willingly contributing to the overall goal. Hence, incentive schemes implemented by adequate mechanisms have to be designed addressing these issues. Furthermore the strategic behavior of service providers and requesters in such a market requires intense investigation to understand the implications of mechanism design decisions.

4.3 Requirement Engineering

In Section 3, we introduced both technical and economic requirements which a mechanism for service trading should ideally satisfy. Historically, mechanism design research has mainly focused on the economic properties and produced various impossibility and possibility results, such as the prominent Gibbard-Satterthwaite and Myerson-Satterthwaite impossibility theorems and Groves mechanisms (Parkes, 2001). In the last decade, the area of *algorithmic* mechanism design introduced domain-specific requirements, in particular focusing on combinatorial auctions and computational aspects as well as the interrelation with the economic requirements (cf. Nisan *et al.* (2007)).

We see at least two important questions with respect to the increasing complexity of the problem space in designing market mechanisms:

- How should the domain-dependent and the economic properties be balanced? Oftentimes a trade-off between various requirements is inherent to the problem. Real-life markets cannot abstract from the technical peculiarities of their application domain and thus must explicitly cope with all constraints no matter if they are of technical and economical nature.
- In previous research, the focus was often on incentive issues. But is this really the most important requirement, given the bounded rationality and irrationality of users and the complexity of their (real-life) strategy space?

4.4 Preference Elicitation and Automated Bidding

While market mechanisms exhibit compelling features, two important building blocks are missing in both theory and practice which hampers their use: *preference elicitation* and *automated trading.* It is a complex burden for both users and providers to (i) assess their true valuation for a certain service and combination of services and to (ii) efficiently communicate this valuation to the market.

Probably the largest body of research on preference elicitation in *auction-based* systems stems from the domain of combinatorial auctions (cf. (Conen & Sandholm, 2001; Zinkevich *et al.*, 2003; Parkes, 2005; Nisan & Segal, 2005)). However, this previous work focuses on a separate issue: If users know their valuation, but communication between the users and the market is costly, how to efficiently query users for their valuations given the specific structure of the underlying allocation problem. Different from this literature, we use the term preference elicitation to denote the users' problem of determining their true valuation, i.e. questions such as "What am I willing to pay for using a server with application X, a dual-core processor and 2 GB of memory for one hour?" or "How much compensation do I expect if I agree to postpone a certain part of my planned energy consumption to a later point in time accepting that a certain machine has to suspend operation for that amount of time?". There is currently not much research available in this area, which is surprising as it is a prerequisite for any market-based approach.

MacKie-Mason & Wellman (2006) study the automation of the user-market interaction by means of trading agents. By equipping users with such (at least partially) automated tools, the communication with the market can be drastically simplified since human users do not constantly need to monitor the market outcome and update their requests. One prominent outcome of this research is the TAC trading agent competition (`http://www.sics.se/tac/`) where research teams compete in designing trading agents for a specific market mechanism.

Another approach worthwhile some further research would be the combination of existing preference elicitation techniques. One could for example try to forecast energy consumption profiles using statistical methods like exponential smoothing or regression models and try to further enhance these results through conjoint analysis. Popular techniques for estimating a user's *valuation* for a certain good or service are Conjoint Analyses (Luce & Tukey, 1964; Green & Rao, 1971) and the Analytical Hierarchy Process (Saaty, 1990). The main problem with these approaches is that they quickly become infeasible with increasing numbers of service attributes and attribute values. This is exacerbated by the large number of transactions especially in markets for elementary services.

In summary, hitherto there has been neither research that tries to combine the issues of preference elicitation and automated trading from the user's perspective nor research that tries to find suitable (graphical) representations of the basic problem which supports users in explaining their preferences easily.

4.5 Reputation Mechanisms

Research in the design of market mechanisms assumes the ex post compliance of market participants; after the mechanism has determined the allocation and the resulting pricing, market participants are assumed to adhere to the market's decisions. In reality, however, this is not self-evident due to moral hazard problems. Consequently, there is a need for reputation mechanisms which induce the participants to "truthful" ex post behavior after the market has cleared (Resnick *et al.*, 2000; Bolton *et al.*, 2002; Dellarocas, 2003), as opposed to the market mechanism which aims at inducing truthful bidding. However, analogously to the mechanism design problem, another layer of complexity is added by the characteristics of the application domain at hand. In computational grids, for instance, when a job fails or a wrong result is returned to the user, it is hard to detect whether this was due to intentional misbehavior of the resource provider or due to technical reasons which are neither controlled by the user nor the provider, programming errors of the user etc. An important challenge will thus be the design of reputation mechanisms which are tailored towards the specifics of the application scenario and which need to be intertwined with the design of the market mechanism to ensure "truthful" overall behavior to avoid market failure.

References

Bianchini, R. & Rajamony, R. (2004). Power and energy management for server systems. *Computer*, **37**, 68–76.

Blau, B. & Schnizler, B. (2008). Description languages and mechanisms for trading service objects in grid markets. In M. Bichler, T. Hess, H. Krcmar, U. Lechner, F. Matthes, A. Picot, B. Speitkamp & P. Wolf, eds., *Multikonferenz Wirtschaftsinformatik, MKWI 2008, München, 26.2.2008 - 28.2.2008, Proceedings*, GITO-Verlag 2008 Berlin.

Blau, B., Neumann, D., Weinhardt, C. & Michalk, W. (2008). Provisioning of service mashup topologies. In *ECIS-16*.

Block, C., Neumann, D. & Weinhardt, C. (2008). A market mechanism for energy allocation in micro-chp grids. In *HICSS-41*.

Bolton, G., Katok, E. & Ockenfels, A. (2002). How effective are online reputation mechanisms? Discussion Papers on Strategic Interaction 2002-25, Max Planck Institute of Economics, Strategic Interaction Group.

Chase, J.S., Anderson, D.C., Thakar, P.N., Vahdat, A.M. & Doyle, R.P. (2001). Managing energy and server resources in hosting centers. *SIGOPS Oper. Syst. Rev.*, **35**, 103–116.

Conen, W. & Sandholm, T. (2001). Preference elicitation in combinatorial auctions. *Proceedings of the 3rd ACM conference on Electronic Commerce*, 256–259.

Dellarocas, C. (2003). The digitization of word-of-mouth: Promise and challenges of online feedback mechanisms. *Management Science*, **49**, 1407–1424.

EPA (2007). Epa report to congress on server and data center energy efficiency. Tech. rep., U.S. Environmental Protection Agency.

Green, P. & Rao, V. (1971). Conjoint Measurement for Quantifying Judgmental Data. *Journal of Marketing Research*, **8**, 355–363.

Greenberg, S., Mills, E., Tschudi, B., Rumsey, P. & Myatt., B. (2006). Best practices for data centers: Results from benchmarking 22 data centers. In *Proceedings of the 2006 ACEEE Summer Study on Energy Efficiency in Buildings.*.

Jackson, M.O. (2003). Mechanism theory. In U. Derigs, ed., *Encyclopedia of Life Support Systems*, EOLSS Publishers.

Lai, K. (2005). Markets are dead, long live markets. *ACM SIGecom Exchanges*, **5**, 1–10.

Lamparter, S., Ankolekar, A., Grimm, S. & Studer, R. (2007). Preference-based Selection of Highly Configurable Web Services. In *Proceedings of the 16th Int. World Wide Web Conference (WWW'07)*, 1013–1022, Banff, Canada.

Luce, R. & Tukey, J. (1964). Simultaneous conjoint measurement: A new type of fundamental measurement. *Journal of Mathematical Psychology*, **1**, 1–27.

MacKie-Mason, J. & Wellman, M. (2006). Automated Markets and Trading Agents. In *Handbook of Computational Economics, Volume 2: Agent-Based Modeling*, North Holland.

Moore, J., Sharma, R., Shih, R., Chase, J., Patel, C. & Ranganathan, P. (2008). Going beyond cpus: The potential of temperature-aware solutions for the data center. Whitepaper, Hewlett Packard Labs.

Nisan, N. & Segal, I. (2005). Exponential communication inefficiency of demand queries. *Proceedings of the 10th conference on Theoretical aspects of rationality and knowledge*, 158–164.

Nisan, N., Roughgarden, T., Tardos, E. & Vazirani, V.V. (2007). *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.

OECD (2007). World energy outlook 2007. Tech. rep., International Energy Agency.

Parkes, D. (2001). *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*. Ph.D. thesis, University of Pennsylvania.

Parkes, D. (2005). Auction design with costly preference elicitation. *Annals of Mathematics and Artificial Intelligence*, **44**, 269–302.

Rappa, M. (2004). The utility business model and the future of computing services. *IBM Systems Journal*, **43**, 32–42.

Resnick, P., Zeckhauser, B., Friedman, E. & Kuwabara, K. (2000). Reputation systems. *Communications of the ACM*, **12**, 45–48.

Rothkopf, M.H., Pekeč, A. & Harstad, R.M. (1998). Computationally manageable combinational auctions. *Management Science*, **44**, 1131–1147.

Saaty, T. (1990). *Multicriteria Decision Making - The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. RWS Publications.

Stößer, J., Neumann, D. & Anandasivam, A. (2007). A Truthful Heuristic for Efficient Scheduling in Network-Centric Grid OS. In *ECIS-15*.

Strapp, J., King, C. & Talbott, S. (2007). Ontario energy board smart price pilot final report. Tech. rep., Ontario Energy Board.

Sun, H.S. & Lee, S.E. (2006). Case study of data centers' energy performance. *Energy and Buildings*, **38**, 522–533.

Ziff Davis (2005). Power consumption and cooling in the data center: A survey. Tech. rep., Ziff Davis Media.

Zinkevich, M., Blum, A. & Sandholm, T. (2003). On polynomial-time preference elicitation with value queries. *Proceedings of the 4th ACM conference on Electronic commerce*, 176–185.