Melanie Holloway, Dieter Schuller, Ralf Steinmetz: Customized Cloud Service Quality: Approaching Pareto-Efficient Outcomes in Concurrent Multiple-Issue Negotiations. In: Raicu, I.; Rana, O.; Buyya, R.: Proceedings of the 2015 IEEE/ACM 8th International Conference on Utility and Cloud Computing UCC 2015, p. 256-260, IEEE Computer Society, December 2015. ISBN 978-0-7695-5697-0.

Customized Cloud Service Quality: Approaching Pareto-Efficient Outcomes in Concurrent Multiple-Issue Negotiations

Melanie Holloway, Dieter Schuller, and Ralf Steinmetz
Technische Universität Darmstadt, Multimedia Communications Lab (KOM)
Rundeturmstr. 10, 64283 Darmstadt, Germany
Email: firstname.lastname@KOM.tu-darmstadt.de

***Abstract*—To date, cloud service consumers usually cannot obtain customized quality guarantees according to their specific business constraints. However, on-demand service provisioning, a large number of services with multiple quality parameters, and a plethora of consumers prevent manual negotiations that are typically conducted in face-to-face meetings. Therefore, the design and realization of appropriate mechanisms for an automated negotiation of service level agreements plays a major role for cloud service markets to emerge. Moreover, from a business point of view, the negotiating parties also have specific demands on the performance of such mechanisms and want to obtain the best result that is achievable while not revealing their private information. In this paper, we propose a negotiation mechanism for the envisioned scenario, that allows to approach efficient results despite private information, and evaluate its performance.**

I. Introduction

Cloud computing has become a major paradigm of our time by providing configurable computing resources on-demand over the Internet, similar to traditional utilities like electricity or water [1]. However, the high flexibility and minimal management effort when using cloud-based services is accompanied by a loss of control over quality aspects such as performance and security. In order to address this issue, cloud providers offer so-called Service Level Agreements (SLAs). Basically, an SLA is a formal contract between a cloud consumer and a cloud provider aiming to ensure that a certain level of quality is maintained by the cloud provider. To date, these SLAs are mainly static and thus, cannot be customized according to individual business constraints. From a business perspective, more flexible SLAs are required since IT is a key enabler to realize business processes. However, traditional negotiations of individual SLAs, where the negotiating parties settle a contract in a face-to-face meeting, are not feasible when it comes to service provisioning from the cloud in an ad-hoc manner for a large number of consumers. Hence, automated negotiation mechanisms are required.

Our research focuses on a scenario, where a broker negotiates concurrently with multiple providers on behalf of a consumer over multiple issues of a desired service. Since there can be multiple providers in the cloud market offering similar services with heterogeneous properties, consumers will wish to compare their offers first before establishing a contract with a certain provider [2]. Furthermore, rational negotiating parties should not “leave extra money on the table”, but aim at Pareto-optimal solutions [3]. Basically, a Pareto-optimal solution is achieved, if no party could be better off without making the opponent(s) worse off, i.e., such a solution is efficient (e.g., [4]). In general, optimal negotiation results can be determined, if the opponent’s preferences are public information, an assumption often considered in game theory (cf. [5]). However, since the negotiating parties have competing interests, they won’t be willing to disclose any private information to another party. Furthermore, when negotiating over multiple issues with different levels of importance for each issue, trade-offs are possible. That means that by increasing the value of one Quality of Service (QoS) parameter to a certain extent while decreasing the value of another QoS parameter by a certain amount, the overall benefit for one or both negotiating participants can be increased [6]. However, this raises the question which trade-off to select for a proposal, since several offers with different values for the multiple issues exist that exhibit the same benefit for one of the negotiating parties, but vary in the benefit for the opponents. Different negotiation mechanisms have been proposed in several related research works, e.g., [3], [6], [7]. However, to the best of our knowledge, none of them addresses Pareto-efficient outcomes in such a setting. All in all, multiple issues, the possibility for trade-offs, and multiple providers with private information introduce high complexity in the decision process. Therefore, we propose a new negotiation mechanism that allows to obtain efficient results despite incomplete knowledge. Our approach is also applicable to further improve preliminary, inefficient agreements, while obtaining joint gains.

The remainder of this paper is structured as follows. Section 2 describes the basic negotiation model of concurrent multiple-issue negotiations in cloud-based systems. In Section 3, we describe our new negotiation mechanism and Section 4 presents the results of our evaluation. The paper closes with a summary in Section 5.

II. Customized Cloud Service Quality

The considered scenario explores concurrent one-to-many negotiations for a specific service, i.e., a broker negotiates with a set of cloud providers in a bilateral manner. The broker can also offer monitoring services to consumers later on. We have already developed corresponding monitoring solutions in our former works (e.g., [8]) and focus on the negotiation of SLAs in the work at hand. The corresponding generic negotiation model is based on our former work in [2]. A negotiation is conducted over the non-functional parameters of a specific service. A proposal represents a set $X = \{x_1, ..., x_n\}$ of values for n QoS parameters under negotiation. The requirements of a consumer or provider on a service are expressed in terms of constraint intervals, each reflecting the lower and upper bound for a given QoS parameter (cf. [5]) , i.e., $x_i \in [min_i, max_i]$. The initial proposal usually contains either the minimum or maximum value for each issue. Thus, the remaining, corresponding lower or upper bound defines the reserve value. Each constraint interval can be mapped to a utility range from 0 to 1 using a scoring function. In doing so, the utility a negotiating party assigns to a value of issue x_i can be determined using the following scoring function (e.g., [6]):

$$U_i(x_i) = \begin{cases} \dfrac{max_i - x_i}{max_i - min_i} & \text{if } U_i(x_i) \uparrow \text{ as } x_i \downarrow \quad (1a) \\ \dfrac{x_i - min_i}{max_i - min_i} & \text{if } U_i(x_i) \uparrow \text{ as } x_i \downarrow \quad (1b) \end{cases}$$

As already mentioned above, we assume, that each negotiating party also specifies different levels of importance in terms of weights $W = \{w_1, ..., w_n\}$ with $\sum_{i=1}^{n} w_i = 1$ for all n issues. We further assume, that the parties negotiate over several rounds alternating in making a proposal. In each round, the benefit of an offer is determined based on the utility function of a party. Following other research works (e.g., [6]), we consider linear, additive utility functions in the work at hand due to computational simplicity. However, this does not limit the applicability of our negotiation mechanism. The total utility for a given offer can then be calculated as the weighted sum over all $U_i(x_i)$. In a negotiation, the objective of each party is to maximize its total utility. Despite conflicting interests, they aim to reach a mutual agreement. Thus, they have to make concessions during each round. The amount of concession is thereby defined by the strategy each negotiating party applies. Concerning the acceptance of a proposal, it is intuitive that a proposal is acceptable to one party, if the utility gained from that proposal is equal to or higher than the utility obtained from the proposal the party is going to offer in the next round. However, since our scenario considers concurrent negotiations with multiple providers, multiple acceptable proposals may exist. Therefore, a consumer will accept that proposal among all the proposals received from the set of providers, that maximizes the consumer's utility.

III. PARETO-OPTIMIZING NEGOTIATION MECHANISM

A. Concept

From a business point of view, the negotiating parties do not only aim at reaching a mutual agreement, but rather want to obtain the best result that is achievable. However, when negotiating over multiple issues, trade-offs are possible. Consequently, several different offers exist that exhibit the same utility to a proposing party, but vary in the utility for the opponent. Since we focus on an incomplete information setting, a proposing party does not know which of the offers exhibiting the same utility provides a higher benefit to the opponent(s) or represents a Pareto-efficient solution. In order to address that issue, we propose a two-phase meta-strategy (cf. Figure 1a) and a corresponding negotiation protocol (cf. Figure 1b) that enhances the well-known alternate offers protocol.

Fig. 1: Overview of our negotiation approach

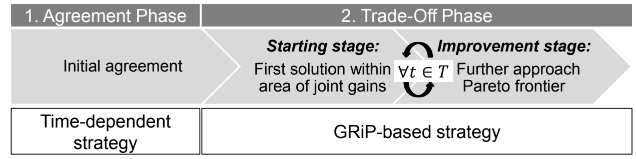

(a) Two-phase meta-strategy

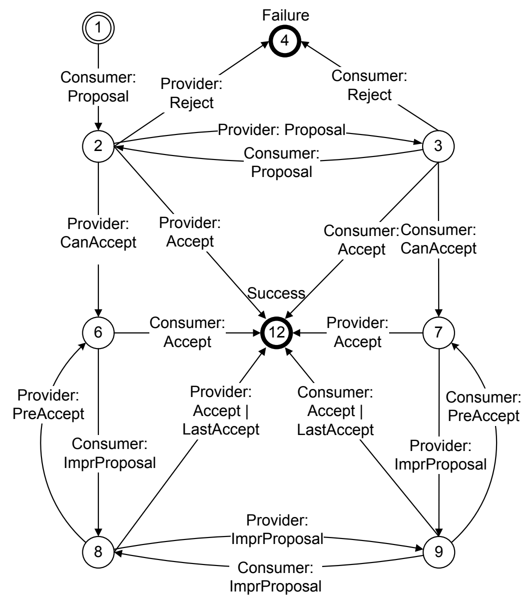

(b) Extended alternate offers negotiation protocol

In the *Agreement Phase*, the negotiating parties alternate in making proposals over multiple rounds in order to approach mutual interests while applying time-dependent strategies. During each round, a party can either withdraw from the negotiation (*Reject* message), accept the current proposal as final agreement (*Accept* message) or accept the current proposal as initial agreement and signal its willingness to improve it (*CanAccept* message). After having obtained an initial agreement with a certain provider, both parties proceed to the second phase, the *Trade-Off Phase*. In that phase, the negotiating parties exchange improvement proposals in order to further improve the current agreement. Having received such an improvement proposal, a negotiating party has four options. The party can accept that improvement proposal and either terminate the negotiation (*Accept* message) or signal its continuing willingness to further improve that last proposal (*PreAccept* message). In both cases, a new agreement is reached based on that improvement proposal. Otherwise, the negotiating party can reject that improvement proposal and either compile a counter improvement proposal or terminate the negotiation (*LastAccept* message). In the latter, the negotiation ends with an agreement based on the last accepted proposal. Hence, the current improvement proposal is discarded. An improvement proposal is only accepted as new agreement if joint gains are achieved. This assumption is valid in a business scenario, since negotiating parties would not enter the second phase, if the possibility exists that there will be only singlesided gains for the opponent. In doing so, the negotiating parties iteratively approach Pareto-efficient outcomes.

The work at hand considers negotiations over the three issues, price, execution time, and availability for simplicity. However, our approach is applicable to an arbitrary number of issues in general. We further assume that trade-offs can be conducted between the price and any of the other QoS attributes. This assumption is intuitive, since providers will charge higher prices for the provisioning of a higher quality resulting from a certain QoS attribute. The same applies for the provisioning of a lower quality level at a lower price. Both cases may lead to a higher utility for both negotiating parties. Consequently, trade-offs can be conducted regarding two different directions. All in all, four different types of trade-offs are possible in our three-issue case. During the *Trade-Off Phase*, the QoS attributes are considered issue-by-issue. For each issue, an iterative search for improvements is conducted based on trade-offs concerning that issue until a stopping criterion is reached. That phase is again split into two different stages: a *Starting Stage* and an *Improvement Stage*. Both stages are iteratively repeated over all types t of possible trade-offs T. The goal of the *Starting Stage* is to determine a first solution point within the area of joint gains and the *Improvement Stage* is used to explore the area of joint gains in order to further approach a Pareto-efficient solution.

B. Negotiation Strategies

1) Agreement Phase: During this phase, we apply time-dependent strategies (cf. [5]) for proposal generation. Time is critical in e-commerce, but other strategies could also be applied, since the first phase only aims to reach a preliminary agreement. At the beginning, the negotiating parties generate their initial proposal with the most preferred values for each issue and make concessions in the next rounds. Concerning the amount of concession, three different classes of tactics are typically distinguished [5]: *conceder*, where great concessions are already made after the beginning, *linear*, where concessions are made in near constant rate, and *boulware*, where concessions are only made shortly before the end.

2) Trade-Off Phase: In order to search for an improvement proposal residing in the area of joint gains, a negotiating party can use its own, current indifference curve as a reference. In general, the preferences at a given point of an indifference curve can be mathematically described by the marginal rate of substitution (MRS) (e.g., [4]). Consequently, a negotiating party has to deviate from its own, current MRS in order to make an improvement proposal. Whether such an improvement proposal leads to a joint improvement or not, depends on the opponent's current MRS. Regarding our scenario, joint improvements can only be found as long as there is a difference between the consumer's MRS_{cons} and the provider's MRS_{prov}. They are equal in case of a Pareto-efficient solution. Thus, joint improvements are only made if one of the following conditions holds:

- $MRS_{prov} > MRS_{cons}$: decrease in quality and price, ΔMRS_{cons} is positive, ΔMRS_{prov} is negative
- $MRS_{prov} < MRS_{cons}$: increase in quality and price, ΔMRS_{cons} is negative, ΔMRS_{prov} is positive

Since we consider an incomplete negotiation setting, we are not aware of the opponent's current MRS and we do not want to reveal the own MRS. Hence, we propose a new negotiation strategy (in the following denoted as GRiP-based strategy) considering greed, risk, and patience of the negotiating parties. In order to apply this strategy, the following parameters must be defined by each negotiating party a:

- $risk\ factor\ R = \{\rho_1^a, ..., \rho_n^a\}$: the percentage of (initial) deviation from the own MRS for each issue
- $greed\ factor\ G = \{\gamma_1^a, ..., \gamma_n^a\}$: the percentage of additional (or subtractive) units for each issue
- $patience\ P = \{\pi_1^a, ..., \pi_n^a\}$: the maximum number of trials for improvement in one direction for each issue
- $cooling\ factor\ C_1 = \left\{\tau_1^{a1}, ..., \tau_n^{a1}\right\}$: the speed of convergence of the risk factor of a given issue to zero
- $cooling\ factor\ C_2 = \left\{\tau_1^{a2}, ..., \tau_n^{a2}\right\}$: the speed of convergence of the risk factor of a given issue to the last observed successful risk factor

Now, in order to generate an improvement proposal, a trade-off is conducted between one of the QoS parameters and the price. In doing so, the new value $x_i^{'}$ for the parameter i that is considered in the current trade-off is determined as follows based on the level of greed γ_i^a of a negotiating party a. Depending on the search direction, the current value x_i of that parameter is either increased or decreased.

$$x_i^{'} = x_i \pm x_i * \gamma_i^a \tag{2}$$

In order to compute the new price value $price^{'}$ of the improvement proposal, $\Delta price$ is determined based on the additional/subtractive units of issue i times the new $MRS_i^{'}$ for that issue. Again, an increase/decrease of the price depends on the search direction.

$$price^{'} = price \pm \Delta price = price \pm \left(x_i * \gamma_i^a * MRS_i^{'}\right) \tag{3}$$

The amount of deviation from the current MRS_i for a given issue i is thereby determined by the risk factor ρ_i^a resulting in the new $MRS_i^{'}$:

$$MRS_i^{'} = |MRS_i| \pm |MRS_i| * \rho_i^a \tag{4}$$

The values for all the other QoS parameters are not changed in the improvement proposal. Based on the GRiP-based strategies, we obtain the negotiation algorithm for the *Trade-Off Phase* as described in the next section.

C. Negotiation Algorithm

During the *Starting Stage* of the *Trade-Off Phase*, a search for a first solution in the area of joint gains is performed for a given issue i and search direction. For this purpose, a negotiating party a tries to achieve a joint improvement with an increase/decrease of ρ_i^a percent of its own, current MRS_i at the beginning. The lower ρ_i^a, the lower is the potential gain in one's own utility and the higher is the potential gain in utility for the opponent. However, the higher ρ_i^a, the higher is the risk that a negotiating party proposes an improvement outside the area of joint gains. Hence, in case that an improvement proposal is not successful, ρ_i^a is (further) reduced according to Equation 5 based on the number f of own unsuccessful improvement proposals so far using cooling factor τ_i^{a1}.

$$\rho^a = \rho_i^a * e^{\left(-\tau_i^{a1} * f\right)} \tag{5}$$

With the first own successful improvement proposal, a negotiating party enters the *Improvement Stage*. If this happens without any own unsuccessful improvement proposal before (i.e., $f = 0$), the risk factor ρ^a is duplicated, since higher gains in utility are achievable (cf. Equation 6a).

$$\rho^a = \begin{cases} \rho^a * 2 & \text{if f=0} \quad (6a) \\ \rho_{min}^a + (\rho_{max}^a - \rho_{min}^a) * e^{\left(-\tau_i^{a2} * f\right)} & \text{if f} \neq 0 \quad (6b) \end{cases}$$

If the new, duplicated risk factor ρ^a is also successful with $f = 0$, it is duplicated again. This process repeats, until ρ^a leads to an unsuccessful proposal. In this case, the current value of ρ^a is stored as ρ_{max}^a (cf. Equation 7a) representing the lowest unsuccessful value of ρ^a so far. If a party enters the *Improvement Stage* with $f \neq 0$, ρ_{max}^a is set to the last unsuccessful value of ρ^a stored in ρ_{last}^a (cf. Equation 7b). In any case, the last own successful ρ^a is stored in ρ_{min}^a.

$$\rho_{max}^a = \begin{cases} \rho^a & \text{if f=0} \quad (7a) \\ \rho_{last}^a & \text{if f} \neq 0 \quad (7b) \end{cases}$$

The interval $[\rho_{min}^a, \rho_{max}^a]$ determined when entering the improvement stage represents a range where the opponent's MRS lies. Hence, during the improvement stage, ρ^a is again (further) reduced according to Equation 6b based on the number f of own unsuccessful improvement proposals so far using cooling factor τ_i^{a2}. Whenever ρ^a leads to a successful improvement proposal again, the lower and upper bounds of the interval are adapted accordingly. In doing so, both negotiating parties iteratively approach the opponent's MRS and a Pareto-efficient outcome.

IV. Evaluation of Negotiation Performance

Our negotiation mechanism has been implemented using the agent-based simulation platform Repast Simphony[1]. This section presents selected results of our evaluation.

A. Simulation Environment and Experimental Setup

In order to assess the performance of our negotiation mechanism, we examine the Pareto-efficiency of our approach as dependent variable. For this purpose, we compute a representation of the Pareto frontier as optimal solution space for a given negotiation result and express the distance to that frontier as the maximum amount in utility that a consumer or provider still could achieve without making the other party worse off (cf. [9]). Furthermore, we consider the three independent variables greed factor, risk factor, and patience and only vary one independent variable at each point in time while we keep the other independent variables fixed (The values assumed for the fixed variables are underlined in Table Ib). Furthermore, we keep all cooling factors fixed and make use of a value suggested by Di Nitto et al. [7]. The resulting setup of the GRiP-based strategies applied in the *Trade-Off Phase* is summarized in Table Ib. All in all, our evaluation comprises 15 test cases in total. Concerning the time-dependent strategies used in the *Agreement Phase*, the configuration can be obtained from Table Ic. Basically, in each negotiation, a certain type of strategy is applied by the broker as specified by the consumer and a certain type of strategy is applied by all the providers, but their concession amounts differ. This setup permits to analyze the impact of nine different strategy settings. Each of the 15 test cases mentioned above is evaluated in all nine strategy settings. In each setting, we randomly generate 1.000 negotiation instances. For the generation of each instance, the number of providers, the negotiation deadline and the constraint intervals for each issue under negotiation are drawn from a uniform distribution. The corresponding ranges for the generation of the constraint intervals are listed in Table Ia. Finally, we also randomly generate the weights of each negotiating party for the different issues based on a uniform distribution in the open interval $(0, 1)$ and normalize them, so that they add up to 1.

B. Results and Discussion

Figure 2 shows selected results of our evaluation. The influence of each independent variable is depicted for the strategy setting *provider: conceder - consumer: boulware*. In each figure, the y-axis displays the remaining distance in utility to the Pareto frontier for the consumer or provider and the x-axis shows the values of a particular independent variable. When using our negotiation mechanism, the distance in utility to the Pareto frontier grows with an increasing risk factor up to 1.1% on average in the worst case on provider-side and decreases to less than 0.75% on average in the best case on consumer-side. All in all, a lower risk factor between 0.01% and 0.05% yields the best results. Considering the greed factor, the distance in utility to the Pareto frontier also grows with an increasing greed factor up to around 1.6% on average in the worst case on provider-side and decreases to less than 0.75% on average in the best case on consumer-side. Concerning the independent variable patience, it can be obtained from the results that the greatest reduction in the distance from the Pareto frontier is already achieved when switching from a patience level of 1 up to a level of 3. In doing so, a reduction of 0.75% on average is obtained on both sides in the best case. When increasing the patience level up to 5, the remaining distance in utility to the Pareto frontier is further reduced to less than 0.9% on average. From that point on, only marginal improvements are made when further increasing the patience level. Hence, our negotiation approach is appropriate to improve an initial agreement to a near Pareto-efficient solution already at a low level of patience.

V. Summary and Outlook

In the work at hand, we have explored the design and realization of an automated negotiation mechanism in a setting, where a broker acting on behalf of a consumer concurrently

[1] http://repast.sourceforge.net/

TABLE I: Parameters used in the evaluation

(a) Ranges for consumer and provider proposals

Issue	**Consumer**		**Provider**	
Price	best: $U\,[5, 10]$	worst: $U\,[15, 20]$	best: $U\,[17, 22]$	worst: $U\,[7, 12]$
Exec. Time	best: $U\,[100, 400]$	worst: $U\,[600, 900]$	best: $U\,[700, 1000]$	worst: $U\,[200, 500]$
Availability	best: $U\,[98.95, 99.95]$	worst: $U\,[97.55, 98.55]$	best: $U\,[97.5, 98.5]$	worst: $U\,[98.9, 99.9]$

(b) Parameters of the GRiP-based strategies

Independent Variable	**Symbol and Values**
Risk factor:	$\rho_i^a \in \{\underline{0.01}, 0.05, 0.1, 0.25, 0.5\}$
Greed factor:	$\gamma_i^a \in \{0.05, \underline{0.1}, 0.15, 0.2, 0.25\}$
Patience:	$\pi_i^a \in \{1, \underline{3}, 5, 10, 20\}$
Cooling factors:	$\tau_j^{a1}, \tau_j^{a2} = 0.0025$

(c) Parameters of the time-dependent strategies

Variable	**Values**
Consumer concession:	boulware: 0.5| linear: 1| conceder: 2.5
Provider concession:	boulware: $U\,(0.0, 1.0)$ | linear: 1.0| conceder: $(1.0, 5.0]$
Number of providers:	$U\,[2, 100]$
Deadline:	$U\,[1, 50]$

Fig. 2: Selected results of the evaluation

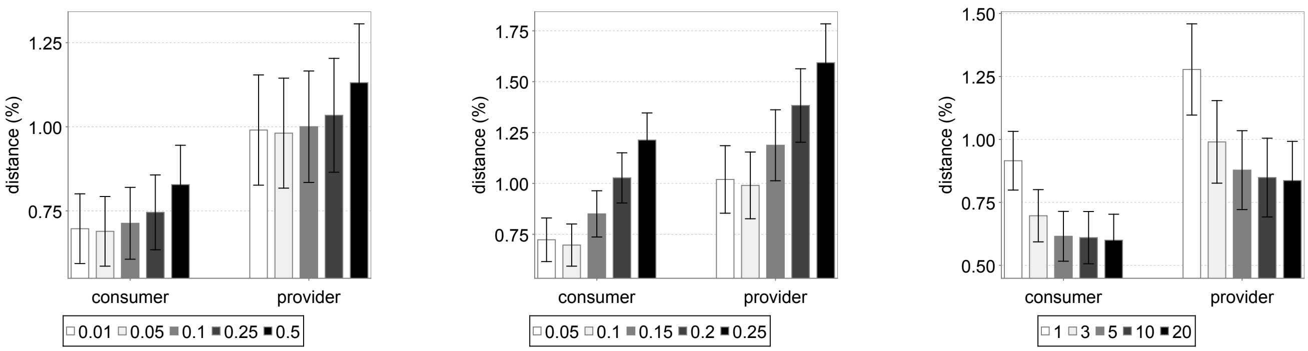

(a) Risk factor (prov:conc-con:boul)

(b) Greed factor (prov:conc-con:boul)

(c) Patience (prov:conc-con:boul)

negotiates with multiple providers over multiple issues of a desired service. Since rational negotiating parties want to obtain the best result that is achievable while not disclosing any private information to the opponent(s), we have proposed a negotiation mechanism that allows to approach Pareto-efficient outcomes despite incomplete information. In doing so, we have proposed a new two-phase protocol and new strategies based on risk, greed and patience. Our approach is also applicable to further improve preliminary, inefficient agreements achieved using a different negotiation mechanism. Our evaluation revealed that by applying our negotiation mechanism, the distance in utility to the Pareto frontier can be decreased to less than 0.75% and to around 0.9% on average in the best case on consumer- and provider-side, respectively.

ACKNOWLEDGMENT

This work was supported in part by the E-Finance Lab Frankfurt am Main e.V. (http://www.efinancelab.com) and the German Research Foundation (DFG) in the Collaborative Research Center (SFB) 1053 – MAKI.

REFERENCES

[1] 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, no. 6, pp. 599–616, 2009.

[2] M. Siebenhaar, T. A. B. Nguyen, U. Lampe, D. Schuller, and R. Steinmetz, "Concurrent Negotiations in Cloud-based Systems," in *Proceedings of the 8th International Workshop on Economics of Grids, Clouds, Systems, and Services (GECON2011), in conjunction with ICSOC 2011*. Springer, 2012, pp. 17–31.

[3] G. Lai, K. Sycara, and C. Li, "A Decentralized Model for Multi-Attribute Negotiations with Incomplete Information and General Utility Functions," in *Rational, Robust, and Secure Negotiations in Multi-Agent Systems*, ser. Studies in Computational Intelligence, T. Ito, H. Hattori, M. Zhang, and T. Matsuo, Eds. Springer, 2008, vol. 89, pp. 39–57.

[4] S. Wied-Nebbeling and H. Schott, *Grundlagen der Mikroökonomik*, 3rd ed. Springer, 2005.

[5] C. Sierra, P. Faratin, and N. R. Jennings, "A Service-Oriented Negotiation Model between Autonomous Agents," in *Proceedings of the 8th European Workshop on Modelling Autonomous Agents in a Multi-Agent World: Multi-Agent Rationality*, 1997, pp. 17–35.

[6] P. Faratin, C. Sierra, and N. Jennings, "Using Similarity Criteria to make Issue Trade-Offs in Automated Negotiations," *Artificial Intelligence*, vol. 142, no. 2, pp. 205 – 237, 2002.

[7] E. Di Nitto, M. Di Penta, A. Gambi, G. Ripa, and M. Villani, "Negotiation of Service Level Agreements: An Architecture and a Search-based Approach," in *Proceedings of the 5th International Conference on Service-Oriented Computing (ICSOC 2007)*, 2007, pp. 295–306.

[8] M. Siebenhaar, D. Schuller, O. Wenge, and R. Steinmetz, "Heuristic Approaches for Robust Cloud Monitor Placement," in *Service-Oriented Computing*. Springer, 2014, pp. 321–335.

[9] D.J.A. Somefun, E.H. Gerding, and J.A. La Poutré, "Efficient Methods for Automated Multi-Issue Negotiation: Negotiating over a Two-Part Tariff," *International Journal of Intelligent Systems*, vol. 21, no. 1, pp. 99–119, 2006.