On Developing Fair and Orderly Cloud Markets: QoS- and Security-aware Optimization of Cloud Collaboration

Olga Wenge, Dieter Schuller, Christoph Rensing, Ralf Steinmetz
Multimedia Communication Lab (KOM)
Technische Universität Darmstadt, Germany
e-mail: {firstname.lastname}@KOM.tu-darmstadt.de

Abstract. *While cloud markets promise virtually unlimited resource supplies, standardized commodities and proper services, some providers may not be able to offer effectual physical capacity to serve large customers. A solution is cloud collaborations, in which multiple providers unite forces in order to conjointly offer capacities in the cloud markets. Supposably, both the Quality of Service and security properties of such collaborations will be determined by "the weakest link in the chain", therefore resulting in a trade-off between the monetary aggregates, cumulative capacity and the non-functional attributes of a cloud collaboration. Based on our previous research, we examine efficient composition of cloud collaborations from the broker's perspective, considering Quality of Service and information security requirements of multiple cloud providers and users and present an exact approach CCCP-EXA.KOM for building cloud collaborations. Furthermore, we propose a Mixed Integer Programming-based heuristic optimization approach CCCP-PRIOSORT.KOM and provide its quantitative evaluation in comparison with our prior optimal approach.*

Keywords: cloud computing security, cloud collaborations, cloud brokerage, information security governance, cloud markets, quality of service

1. Introduction

Cloud markets promise to supply virtually unlimited capacities and services in a scalable, pay-as-you-go fashion (Buyya, R., Yeo, C., Venugopal, S., Broberg, J., & Brandic, I., 2009). Yet, specifically smaller providers may not be able to satisfy the resource and service demands of large customers on their own due to limited data center capacity and, consequently, limited range of services. A solution lies in cloud collaborations within cloud markets, i.e., the cooperation of multiple providers to aggregate their resources and conjointly satisfy user's demands. Supposably, such cloud collaborations have both Quality of Service (QoS) and information security impacts: as a user may potentially be served by any provider within a collaboration, the aggregated non-functional service attributes - e.g., availability, latency, security protection level, data center location or tiers – will be determined by "the weakest link in the chain", i.e., by the provider with the lowest guarantees.

Take the example of two providers: one provider guarantees 99.5% of availability and another provider guarantees only 99%. If these providers aggregate their capacities and related non-functional guarantees to build a collaboration, the availability guarantee will be determined by the worst one - 99%.

Consideration of country-specific and industry-specific data privacy laws and regulations is another concern by building cloud collaborations within cloud markets. Since providers can reside in different jurisdictions (the European Union, Canada, Singapore, or the United States), where data privacy laws and data classification substantially differ (Wenge, O., Lampe, U., Müller, A., & Schaarschmidt, R., 2014; Wenge, O., Siebenhaar, M., Lampe, U., Schuller, D., & Steinmetz, R., 2012). Also regulatory requirements for banking, medical and health institutions are stricter and harder with respect to confidentiality, integrity and availability of data in comparison with other public enterprises or business areas without confidential data. Therefore, the fulfillment of such requirements may not be achieved once multiple cloud providers enter cloud collaborations.

Based on this scenario and our previous research (Wenge, O., Lampe, U., & Steinmetz, R.,2014) we examine the Cloud Collaboration Composition Problem (CCCP) in the work at hand. Our focus is on a broker within the cloud market, who aims *to maximize his/her profit* through the composition of cloud collaborations from a set of

providers and assignment of users to these collaborations. In that assignment, QoS and security requirements, i.e., non-functional attributes, should also be considered and fulfilled. This work extends the previously introduced CCCP problem and its exact optimization solution approach with a heuristic approach that improves the computational time in the context of cloud markets.

The remainder of this paper is structured as follows: In Section 2, we give an overview of cloud markets and related regulations that must be considered by trading with cloud products. Section 3 describes information security focus within cloud markets and within cloud collaborations. Section 4 explains the role of a cloud broker in cloud markets. In Section 5, we describe the Cloud Collaboration Composition Problem (CCCP) and the formal optimization model. Based on this, the subsequent Section 6 briefly presents an exact optimization approach, called CCCP-EXA.KOM. Section 7 includes evaluation results of this exact approach. Section 8 introduces a heuristic approach, called CCCP-PRIOSORT.KOM which is quantitatively evaluated and compared with the previously results in Section 9. Section 10 gives an overview of related work, and Section 11 concludes the paper with a summary and outlook.

2. Fair and Orderly Cloud Markets

Economic science studies allocation and usage of diverse resources. Some resources allocation approaches are managed by the price systems: low prices are very attractive for consumers, high prices provoke forced savings; high salaries are attractive for employees, low benefits encourage to innovations. But by many instances, the usage of only a price system as a market regulator is not sufficient. There are other aspects that must be considered as well, e.g., legal, ethical, and security requirements.

Furthermore, there are many markets where the price system operates but the traditional assumption of perfect competition is not even approximately satisfied. In particular, many goods are indivisible and heterogeneous, whereby the market for each type of good becomes very thin. How these thin markets allocate resources depends on the institutions that govern trading and transactions.

The current cloud market environments consist of heterogeneous clouds: cloud providers who sell services, cloud users who buy services, and cloud brokers who help to find the perfect match for their clients. In other words, cloud markets present the aggregate of possible buyers and sellers of cloud services and cloud resources and the transactions between them (Garg, S.K., Vecchiola, C., & Buyya, R., 2011). But the current cloud markets are not organized and supervised on the desired level, e.g., in comparison to financial or energy markets (Garg, S.K., Versteeg, S., & Buyyaa, R., 2013; Rahimi, A.F., & Sheffrin, A.Y., 2003).

The financial and energy markets are supervised by exchanges or other organizations that facilitate and oversee the trade, using *physical locations* (e.g., New York Stock Exchange (NYSE), Deutsche Börse (German Stock Exchange in Frankfurt), or European Energy Exchange (EEX) in Leipzig), or *electronic systems* (e.g., NASDAQ (National Association of Securities Dealers Automated Quotations), XETRA (Xchange Electronic Trading).

These are also regulated by different national and international authorities and laws listed in Table 1. These laws demand compliance with data protection requirements and anti-money laundering (AML) policies in all circumstances with respect to trading (Wenge, O., Lampe, U., Müller, A., & Schaarschmidt, R., 2014; Wenge, O., Siebenhaar, M., Lampe, U., Schuller, D., & Steinmetz, R., 2012; Lampe, U., Wenge, O., Müller, A., & Schaarschmidt, R., 2012).

Table 1. *International market regulators and regulations*

Law / Authority /Standard	*Validity Area*
German Federal Data Protection Act (GFDPA)	*effective in Germany*
Data Protection Directive (DPD)	*effective in the European Union (EU)*
the Privacy Act	*effective in the United States of America (USA)*
Conventions of the Organisation for Economic Co-operation and	*effective in 34 countries*

Development (OECD)	
Safe Harbor Principles	*effective for the USA-EU contracts*
the Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism Act (USA Patriot Act)	*effective in the USA*
Sarbanes-Oxley Act (SOX)	*effective for all enterprises that trade in the USA securities markets*
Directive 2006/43/EG (EuroSOX)	*effective for all enterprises that trade in the EU securities markets*
Basel Accords	*effective in 20 countries*
IT Fundamental Right	*effective in Germany*
the Personal Information Protection and Electronic Documents Act (PIPEDA)	*effective in Canada*
the Monetary Authority of Singapore (MAS)	*effective in Singapore and Asian-Pacific region*
Bank secrecy acts	*effective between banks and customers*
United States Code (USC)	*effective in the USA*
ISO Standards	*effective globally*
German Federal Financial Supervisory Authority (GFFSA)	*effective in Germany*
German Banking Act (GBA)	*effective in Germany*
Office of the Comptroller of the Currency (OCC)	*effective in the USA*
Federal Financial Institutions Examination Council – FFIEC	*effective in the USA*
Statement on Auditing Standards No. 70 (SAS-70)	*effective in the USA*
Binding Corporate Rules (BCRs)	*effective in the EU*
Certified Information Systems Security Professional Principles (CISSP)	*recommended globally*
Certified Information Systems Auditor Principles (CISA)	*recommended globally*
Board of Governors of the Federal Reserve System (BGFRS)	*effective in the USA*
Local territorial laws	*effective locally*

Lack of control or supervision is one of main concerns in cloud markets. The development of market supervision techniques and approaches for the current cloud marketplaces, to provide *a fair and orderly cloud market* – a market in which supply and demand for a product are roughly equal, is still in its embryonic stage. The trading of cloud resources within predefined cloud collaborations can be seen as *an interim solution* to provide desired supervision and information security governance (Guitart, J., & Torres, J., 2010; Gomes, E., Vo, Q.B., & Kowalczyk, R., 2012). Two main principles in the market design theory for the establishing of any fair and orderly market are *stability* and *incentive compatibility*. Both principles are derived from *the cooperative* and *non-cooperative game theory* and *the stable marriage problem* and found a very wide application in the world of economics (Shapley, L.S., 1955; Knuth D.E., 1996; Roth, A.E., 2002; Roth, A.E., Sönmez, T., & Ünver, M.U., 2005; Roth, A.E., 2008).

Stability encourages groups of to voluntarily participate in the market. Incentive compatibility discourages strategic manipulation of the market. The main principle of the cooperative game theory is building of coalitions between individuals (or players, or traders, or cloud providers in a market) who are eager to cooperate with each other. A game in coalitional form with *transferable utility* specifies, for each coalition S its worth $v(S)$. This worth is an economic surplus (a sum of money) that coalition S can generate using its own resources. If coalition S forms,

then its members can split the surplus $v(S)$ in any way they want, and each member's utility equals his/her share of the surplus. This feature is called *transferable utility*. The function v is called the characteristic function. Furthermore, the cooperative game theory studies the incentives of individuals to form such coalitions under consideration that any potential conflicts of interest within a coalition can be solved by binding agreements. These agreements induce the coalition members to maximize the surplus (or revenue) of the coalition. In games with transferable utility, it is assumed that the individuals can freely transfer utility among themselves. The idea of stability corresponds to the idea of Nash equilibrium in non-cooperative game theory. In non-cooperative game theory, a Nash equilibrium is a situation such that no individual can deviate and make herself better. In cooperative game theory, a stable allocation is a situation such that no coalition can deviate and make its members better. That is why stability formalizes an important aspect of idealized frictionless marketplaces.

A coalition building between game players in the cooperative game theory with the purpose of increasing their benefit appears to be very similar to our idea of building cloud collaborations. Stability is one of the important drivers for involving new market participants and building collaborations. Incentive compatibility is necessary to prevent manipulations in the market and within cloud collaborations.

In our work we also assume that each cloud provider and each cloud user has his/her requirements (preferences) that must be fulfilled in any trading and transaction. For convenience, we assume that these requirements are strict. To fulfill these requirements a stable matching between market participants is inevitable. The matching is unacceptable to a market participant (provider or user) if it worse than to remain unmatched. We define the matching as stable, if all collaborations between cloud providers and cloud users are composed in the way they bring the most profit and all requirements are matched acceptably.

In (Wenge, O., Siebenhaar, M., Lampe, U., Schuller, D., & Steinmetz, R., 2012) we identified three types of cloud collaborations with respect to the security critical areas: federated collaborations, loosely-coupled collaborations and ad hoc collaborations. Security requirements, relevant for cloud partners within collaboration types can be used as *admission criteria* for cloud recourses trading within those collaboration types.

Federated collaborations assume the usage of a so-called *metapolicy*, which includes all policies of all collaborative clouds. This metapolicy reduces the possibility of the occurrence of security incidents and breaches, as all security configurations and controls are fully pre-agreed between collaborative partners.

Loosely-coupled collaborations are more flexible and cover smaller cloud regional environments, e.g., the EU, the USA, Canada; or industry specific, e.g., banks or medical institutions. In this case, country or industry specific regulations, or *service level agreements* (SLAs) can be used as a basis for their security policies.

Ad hoc collaborations do not presume any kind of pre-agreed security policies or SLAs, the signing of which can be very time-consuming and can hamper the dynamic of data transfer and service delivery. Ad hoc collaborations are the most critical ones and cannot be performed without a proper supervision and information security governance over them in the form of *a trusted security entity* (e.g., cloud broker, identity broker, etc.).

3. Information Security Governance in Cloud Markets

Information security issues are very critical in cloud computing (Kretzschmar, M., & Golling, M., 2011). Many aspects must be examined concerning security risks in the cloud paradigm: legal risks, data privacy and data protection risks, users' and providers' security levels, right to audit and information security governance processes.

Information Security (IS) governance is a significant part of corporate governance in an enterprise and strives towards the understanding of the criticality of information security, endorsing the development and implementation of security programs and their alignment with business strategy. IS governance also takes the responsibility for performance management, reporting and risk management.

In cloud computing the role of IS governance has become enormously important, as enterprises deal with off-premise services with the involvement of sometimes diverse vendors and non-enterprise employees, whose compliance and activity must be monitored and reported (Bernsmed, K., Jaatun, M.G, Meland, P.H., & Undheim, A., 2011; Yang, D., 2011). To the best of our knowledge, there are three security mechanisms to provide security governance over cloud providers: cloud certifications, cloud risk assessments, and trusted security entities.

Cloud certification sounds very promising and gives a certain sense of trust. The most cloud certificates are based on best practices security frameworks and already existing security standards, such as ISO (International Organization for Standardization), NIST (National Institute of Standards and Technology), CSA (Cloud Security Alliance) and FISMA (Federal Information Security Management Act). The main disadvantage of cloud certification is its *generality*. These certificates are not always sufficient for peculiar cases (e.g., for critical data, banking transactions, country laws) and should be adapted or extended with other security governance mechanisms (e.g., risk assessments and audit) (Bernsmed, K., Jaatun, M.G, Meland, P.H., & Undheim, A., 2011; Yang, D., 2011).

Cloud risk assessments are more granular and can be used with respect to different industries (banking, insurance, healthcare). Cloud risk assessments are also based on the existing risk assessments and are extended with specific vendor governance controls for availability, auditing and controlling. These risk assessments are provided by ISO, CSA, BSI (Bundesamt für Sicherheit in der Informationstechnik, Federal Office of Information Security in Germany), ENISA (European Network and Information Security Agency), COBIT (Control Objectives for Information and Related Technology), ISACA (Information Systems Audit and Control Association), Basel Accords, and SOx (Sarbanes-Oxley Act). The risk assessment process is very time-consuming and in case of risk acceptance procedures or a necessary risk remediation, can be followed by numerous complex bilateral agreements (Papish, M., 2012; Bernstein, D., & Vij, D., 2012).

A trusted security entity concept is more dynamic, but currently does not cover all security aspects of cloud computing. It is mostly focused on identity and access management, ignoring infrastructure, network, and application security. The existing solutions, especially for ad hoc collaborations, do not cover the whole aspects of cloud security or need sufficient evidence for their implementation, e.g., monitoring and logging tools in place, or regular auditing (Wood, K., & Anderson, M., 2011; Ates, M., Ravet, S., Ahmat, A., & Fayolle, J., 2011; Goyal, P., 2011; He., Y.H., Bin, W., Xiao, X.L., & Jing, M.X., 2010; Yang, D., 2011).

Therefore, the role of a cloud broker, as a mediator, who is responsible for bringing cloud market actors together with respect to their requirements, is very important and in our research, we aim at the development of an efficient QoS- and security-aware brokerage model.

4. The Role of a Cloud Broker in Cloud Collaborations

Many environments, and especially IT environments, are still far from the perfectly competitive benchmark, as they are still very heterogeneous and without precisely specified rules that can govern trade. In such markets, the participants must be appropriately matched in order to trade with each other, i.e., a role of a broker, who provides this matching is very important and mandatory, if a cloud market must comply with market design principles. Cloud broker is also seen as inevitable market in the NIST (National Institute of Standards and Technology) cloud computing reference architecture (Liu, F., Tong, J., Mao, J., Bohn, R., Messina, J., Badger, L., & Leaf, D., 2011) (Figure1).

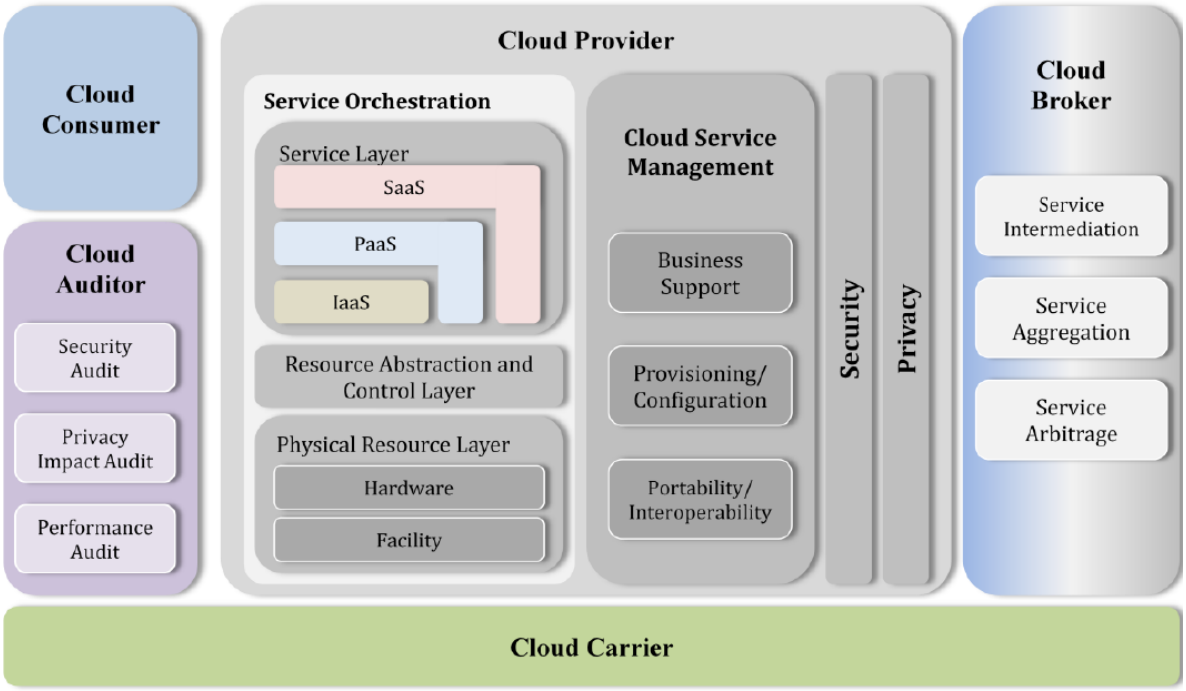

Today's cloud environments are built up of heterogeneous landscapes of independent clouds. The heterogeneity of clouds, as a consequence of still nonexistent technology, security and audit standards, presents a hurdle for a proper collaboration between clouds, necessary for the building of the cloud ecosystem and cloud marketplaces (Kretzschmar, M., & Golling, M., 2011; Corporation Essvale Corporation Limited, 2008).

The reasons for cloud collaborations can be very different: enterprise acquisitions, storage and compute power extensions, disaster recovery plans, sub-contracting and service outsourcing, the necessity for a wider spectrum of services, etc. Such cloud collaborations bring cloud providers further advantages. Besides the eco-efficiency, due to shared usage of data centers and technologies (Guitart, J., & Torres, J., 2010), a better scalability and cost reduction can be achieved by the ad hoc selling of free resources and buying of additional external resources. This exchange of cloud resources forms the basis of the cloud brokerage service model (Uttam Kumar, T., & Wache, H., 2010).

Cloud brokerage enables cloud providers to find an optimally suitable match for each other, i.e., to find a collaborative partner that meets all requirements of intended cloud collaboration. These requirements may include business aspects (pricing, timelines), functional and non-functional technical aspects (compatibility, interoperability, availability), and of course non-functional legal and security aspects (level of data protection, security measures, compliance with different industrial regulations, etc.) (Uttam Kumar, T., & Wache, H., 2010; Lampe, U., Wenge, O., Müller, A., & Schaarschmidt, R., 2012; Siebenhaar, M., Wenge, O., Hans, R., Tercan H., & Steinmetz, R., 2013).

The cloud broker is the leading actor in the cloud brokerage service model, and acts as a mediator between cloud service providers and cloud service consumers, providing matchmaking, monitoring and governance of cloud collaborations (Gomes, E., Vo, Q.B., & Kowalczyk, R., 2012).

The matchmaking of security and legal requirements and especially monitoring of their fulfillment during the cloud collaboration is not trivial. The security risks tend to accelerate by entering into cloud collaborations within cloud marketplaces, because collaborative partners may have different implemented security policies and standards. Therefore, two main requirements must be met to provide secure and compliant cloud collaboration - the cloud broker must perform an optimally reliable security risk assessment prior to the collaboration, or on-demand; and the cloud broker must provide the security governance during the collaboration.

The security risk assessments of cloud providers are widely discussed in the recent research, but, to the best of our knowledge, these assessments are still very time-consuming and cannot be applied to ad hoc cloud collaborations (Schnjakin, M., Alnemr, R., & Meinel, C., 2010).

5. Cloud Collaboration Composition Problem

As mentioned before in our work, we take the perspective of a cloud broker, who acts within a cloud market and unites cloud providers to build cloud collaborations and provides assignment of cloud users to these collaborations. So, the cloud market consists of a set of cloud providers and a set of users, formally denoted as $P = \{1, 2, \dots, P^{\#}\}$ and $U = \{1, 2, \dots, U\}$, respectively.

Each user $u \in U$ exhibits a certain resource demand of $RD_u \in \mathbb{R}^+$ units, for which he/she is willing to pay a total of $M_u^+ \in \mathbb{R}^+$ monetary units. Furthermore, each cloud provider $p \in P$ is able to provide a resource supply of $RS_p \in \mathbb{R}^+$ units at a total cost of $M_p^- \in \mathbb{R}^+$.

QoS and security constraints, which determine requirements by consumption and provision of services, we define by the common term of non-functional constraints. Specifically, we distinguish two sets, $A = \{1, 2, \dots, A^{\#}\}$ and $\hat{A} = \{1, 2, \dots, \hat{A}^{\#}\}$, of quantitative and qualitative non-functional attributes. Quantitative attributes represent numerical properties, e.g., availability or latency.

Qualitative attributes correspond to nominal properties, e.g., applied encryption technology, data center location, and adherence to a certain industry-specific security policy or country-specific data privacy protection controls. The cloud providers make certain guarantees with respect to the non-functional attributes. For each quantitative attribute $a \in A$, the value guaranteed by provider $p \in P$ is denoted as $AG_{a,p} \in \mathbb{R}$ For each qualitative attribute $\hat{a} \in \hat{A}$, the corresponding information is given by $\hat{A}G_{\hat{a},p} \in \{0,1\}$. The cloud users specify also certain requirements concerning their non-functional attributes. With respect to each quantitative attribute $a \in A$, the value required by user $u \in U$ is denoted as $AR_{a,u} \in \mathbb{R}$. Likewise, $\hat{A}R_{\hat{a},u} \in \{0,1\}$ denotes the requirement for each qualitative attribute $\hat{a} \in \hat{A}$, i.e., indicates whether this attribute is mandatory or not.

The objective of the broker is the composition of cloud collaborations, consisting of multiple cloud providers, and subsequently assigning users to them. In that process, all defined constraints must be fulfilled and *the profit maximization*, i.e., the difference between the revenue from the served cloud users and the spending on the incorporated cloud providers, should be achieved. A tangible, simplified example for a CCCP instance is provided in Figure 2. The instance exhibits four users and providers with different resource demands - supplies and non-functional requirements - guarantees, respectively. In the example, providers P_2 and P_4 form a collaboration, which enables them to conjointly serve users U_1 and U_2 under the given constraints. Both providers substantially profit from the collaboration, since their combined resource supply permits to serve larger customers and allows to achieve a higher degree of resource utilization.

Figure 2: Tangible example of a CCCP instance with four users and providers

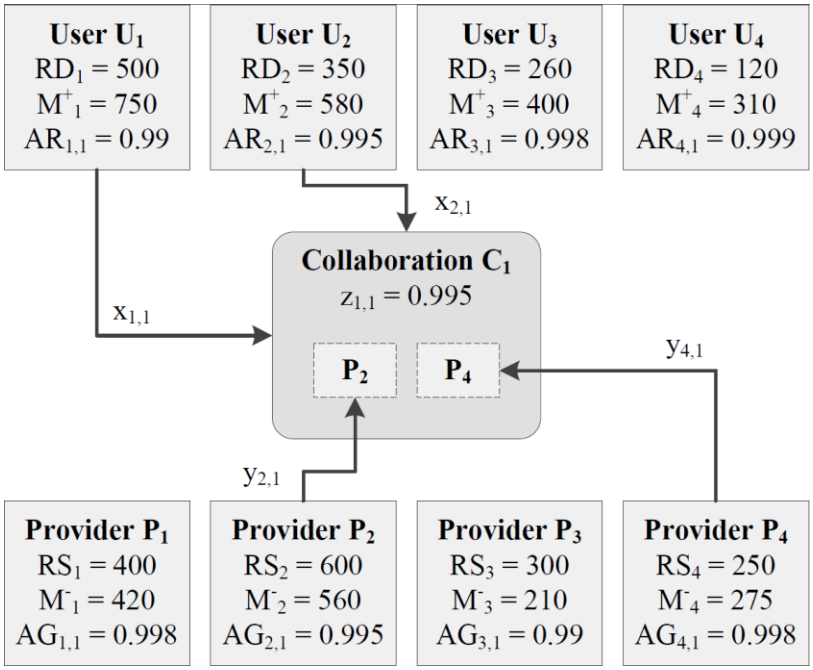

6. Exact Optimization Approach CCCP-EXA.KOM

Based on the notations that were introduced in the previous section, the Cloud Collaboration Composition Problem (CCCP) can be transformed into an optimization model. The result is given in Model 1 and will be explained in the following.

To start with $x_{u,c}$, and $y_{p,c}$ are the main decision variables in the model (cf. Equation 11). They are defined as binary and indicate whether user u or provider p, respectively, has been assigned to collaboration c or not. As additional auxiliary decision variables, we introduce $\acute{y}_{p,c}$, which are also binary and serve as complement to $y_{p,c}$, hence indicating the non-assignment of a provider p to a collaboration c. Furthermore, $z_{a,c}$ and $\hat{z}_{\hat{a},c}$ are specified (cf. Equation 12). They are defined as real and binary, respectively, and represent the cumulative value of the non-functional property a or $\hat{a}$, respectively, for collaboration c. The variables x and y are referred to as main decision variables, since they have a direct impact on the objective function. In contrast, $\acute{y}$, z, and $\hat{z}$ only have an indirect influence. The monetary objective consists in profit maximization (cf. Equation 1).

Model 1 Cloud Collaboration Composition Problem

Objective function

(1) Maximize Profit $(x, y, \acute{y}, z, \hat{z}) =$

$$\sum_{u \in U, c \in C} x_{u,c} \times M_u^+ - \sum_{p \in P, c \in C} y_{p,c} \times M_p^-$$

so that

(2) $\sum_{c \in C} x_{u,c} \leq 1 \quad \forall u \in U$

(3) $\sum_{c \in C} y_{p,c} \leq 1 \quad \forall p \in P$

(4) $y_{p,c} + \acute{y}_{p,c} = 1 \quad \forall p \in P, \forall c \in C$

(5) $\sum_{u \in U} x_{u,c} \times RD_u \leq \sum_{p \in P} y_{p,c} \times RS_p \quad \forall c \in C$

(6) $z_{a,c} \leq y_{p,c} \times AG_{p,a} + \acute{y}_{p,c} \times max_{\,p \in P}(AG_{p,a})$

$\forall p \in P, \forall c \in C, \forall a \in A$

(7) $\hat{z}_{\hat{a},c} \leq y_{p,c} \times \hat{A}G_{p,\hat{a}} + \acute{y}_{p,c}$

$\forall p \in P, \forall c \in C, \forall \hat{a} \in \hat{A}$

(8) $z_{a,c} \geq x_{u,c} \times AR_{u,a} \quad \forall u \in U, \forall c \in C, \forall a \in A$

(9) $\hat{z}_{\hat{a},c} \geq x_{u,c} \times \hat{A}R_{u,\hat{a}} \quad \forall u \in U, \forall c \in C, \forall \hat{a} \in \hat{A}$

(10) $C = \{1, 2, \dots, \min(P^{\#}, U^{\#})\}$

(11) $x_{u,c} \in \{0,1\} \; and \; y_{p,c} \in \{0,1\} \quad \forall u \in U, \forall p \in P, \forall c \in C$

(12) $\acute{y}_{p,c} \in \{0,1\} \; and \; z_{a,c} \in \mathbb{R} \; and \; \hat{z}_{\hat{a},c} \in \{0,1\}$

$\forall p \in P, \forall a \in A, \forall \hat{a} \in \hat{A}, \forall c \in C$

That is, the difference between the revenue from the served cloud users and the spending on the used cloud providers should be maximized, depending on the values of the decision variables. Equations 2 and 3 make sure that each user and provider is assigned to not more than one collaboration. Thus, the broker may opt to not satisfy certain users' demands, but also to not exploit cloud providers as part of a collaboration. Equitation 5 does not allow resource demand to exceed the resource supply. Equation 4 determines the inverse variable $\acute{y}_{p,c}$ for each decision variable $y_{p,c}$. This definition is used in the following two Equations 6 and 7. They determine the cumulative non-functional values for quantitative and qualitative attributes, respectively. Both equations are formulated such that quantitative properties are given by the "worst" value among all providers in a certain collaboration, i.e., the "weakest link in the chain". Equations 8 and 9 make sure that users can only be assigned to such collaborations that make sufficient non-functional guarantees, given the users' specific non-functional requirements.

Lastly, Equation 1 defines a set of potential cloud collaborations. The underlying notion for the given definition is that no user or provider will be assigned to more than one collaboration (recall Equations 2 and 3). Hence, the maximum number of collaborations is given by the number of users or providers, whichever is lower.

We implemented the given model and evaluated the optimal approach in order to obtain an exact (i.e., profit maximal) solution. We used a Mixed Integer Program (MIP), i.e., a special form of Linear Program (LP) that features both integer (in this case, binary) and natural decision variables, and a branch-and-bound off-the-shelf optimization algorithms (Hillier, F., & Lieberman, G., 2005). The evaluation results are presented in the following section.

7. Evaluation of CCCP-EXA.KOM

To assess the practical applicability of our proposed approach CCCP-EXA.KOM, we have prototypically implemented it in Java 7. In order to transfer Model 1 into a programmatic representation, we use the free JavaILP framework1. While this potentially permits for the application of different backend solver framework, we have selected the commercial IBM ILOG CPLEX framework as default due to its favorable performance (Meindl, B., & Templ, M., 2012) and its popularity in related research, e. g., (Hans, R., Lampe, U., & Steinmetz, R., 2013; Mashayekhy, L., & Grosu, D., 2012).

Evaluation Setup and Procedure

The main objective of our evaluation is to assess the required computation time of CCCP-EXA.KOM for different problem sizes. This allows us to judge the applicability of the proposed approach under practical conditions, where time constraints in the decision process play an important role. Thus, formally, we regard computation time as the dependent variable of our evaluation.

As independent variables, we include the number of considered users and providers, i.e., $U^{\#}$ and $P^{\#}$. In contrast, the number of quantitative and qualitative non-functional attributes were fixed ($A^{\#} = 1$ and $\hat{A}^{\#} = 1$); hence, they constitute controlled variables. This is justified by two aspects: First, these variables are likely also predefined in practice. Second, they do not have an impact on the number of decision variables and hence, the size of the solution space. Each specific combination of $U^{\#}$ and $P^{\#}$ results in a test case. For each test case, we created 100 specific CCCP instances with the according dimensions.

The parameter values or distributions that were used in the problem generation process are summarized in Table 2. The specifications of the nonfunctional parameters are based on the notion that the sole quantitative and qualitative attribute represent availability (a QoS aspect) and data center location in the European Union (a security aspect), respectively. Furthermore, monetary parameters were set such that higher availability results in quickly increasing values, based on the observation that each additional "nine" in the availability figure results in doubled cost (Durkee, D., 2010). In contrast, an EU data center location only leads to a moderate increase of 10%, which closely corresponds to the price difference observed for Eastern U.S. and Ireland-located Amazon EC2 VM instances (Amazon Web Services, Inc., 2013).

Following the generation, we computed a solution to each problem instance using our prototypical implementation of CCCP-EXA.KOM. In that process, we imposed a timeout of 300 seconds (i. e., five minutes) per problem instance. Based on the resulting sample of computation times for the successfully solved problems, we computed the mean computation time, as well as the 95% confidence interval.

The evaluation was conducted on a desktop computer, equipped with an Intel Core 2 Duo E7500 processor and 4 GB of memory, operating under the 64-bit edition of Microsoft Windows 7.

Table 2: *Parameter values and distributions used in the problem instance generation. Abbreviations: Uni – Uniform distribution; Ber – Bernoulli distribution.*

Parameter	*Value / Distribution*
$AR_{1,u}$	*Uni (0.99, 0.9995)*
$\hat{A}R_{1,u}$	*Ber (0.5)*
$AG_{1,p}$	*Uni (0.995, 0.9995)*
$\hat{A}G_{1,p}$	*Ber (0.5)*
RD_u	*Uni (1000, 5000)*
RS_p	*Uni (1000, 5000)*
M_u^+	$\alpha_u \times RD_u \times log_{10}(1 - AR_{1,u})^2 \times (1.1^{\hat{A}R_{1,u}})$
M_p^-	$\beta_p \times RS_p \times log_{10}(1 - AG_{1,p})^2 \times (1.1^{\hat{A}G_{1,p}})$
α_u	*Uni (1.5, 1.75)*
β_p	*Uni (1.0, 1.25)*

Evaluation Results and Discussion

The results of our evaluation, i.e., the observed mean computation times per test case, are graphically illustrated in Figure 3. As can be clearly seen, the computation times quickly increase with the problem size, i.e., the considered number of users and providers. The effect is less pronounced for the smallest two problem classes (with $U^{\#} \leq 6$ and $P^{\#} \leq 9$); in fact, for these two test cases, there is no statistically significant difference in mean computation time observable at the 95% confidence level. In absolute terms, we already find absolute computation

times in the order of magnitude of one-hundred seconds and one second respectively for the medium-sized test cases with $U^{\#} \leq 8$. For these test cases, increasing the number of providers increases the computation time by a factor of approximately ten already.

For the four largest test cases (with $U^{\#} \geq 10$ and $P^{\#} \geq 10$), the absolute computation times reach the order of magnitude of seconds and ten seconds. All observed increases are statistically significant at the 95% confidence level. In addition, the ratio of solved problem instances sharply drops with growing problem size. This effect is most notable for the largest problem class that involves 12 users and 18 providers, where only 43% of the 100 problem instances could be solved within the timeout period of five minutes. Given that the considered problem dimensions are still relatively small in the context of a large cloud market, it can be concluded that the practical applicability of the proposed optimization approach CCCP-EXA. KOM is rather limited.

As it has already been explained before, a broker will likely have to decide on the composition of collaborations under rigid time constraints, since users likely require resources at short notice. Hence, an important future challenge consists in the development of appropriate heuristics, which permit to trade reductions in computation time against small degradations in broker profit, and are consequently applicable to practically relevant, large-scale problem instances. In that context – apart from its potential application to small-scale problem instances – CCCP-EXA.KOM can serve as a valuable performance benchmark and we present our heuristic approach CCCP-PRIOSORT.KOM in the next section.

Figure 3: Evaluation results, i.e., observed mean computation times (with 95% confidence intervals) for CCCP-EXA.KOM by test case. Please note the logarithmic scaling of the ordinate.

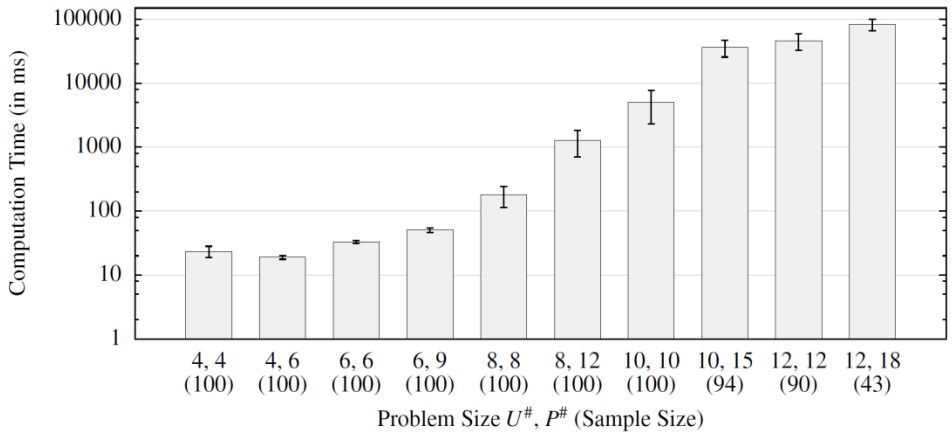

8. Heuristic Solution Approach CCCP-PRIOSORT.KOM

Heuristic optimization algorithms (heuristics for short) seek good feasible solutions to optimization problems in circumstances where the complexity of the problem or the limited time available for its solution does not allow exact solution. The formal intractability, in the sense of NP-hardness (Jonson, D.S., 2012), of many commonly encountered optimization problems and the growing use of real-time control have made the development of heuristics a major area within operations research.

Unlike exact algorithms, where time-efficiency is the main measure of success, there are two burning issues in evaluating heuristics: how fast can solutions be obtained and how close do they come to being optimal.

As discovered in the previous section, the computation time of the proposed CCCP exact solution grows in the dependence on the number of cloud market participants and in the worst case it is exponential. In the following, we propose a heuristic optimization approach CCCP-PRIOSORT.KOM with the improved computation time.

CCCP-PRIOSORT.KOM consists of two phases:

Phase 1: Presorting of cloud providers and cloud users according to the quotients $Q_u = M_u^+ \ / \ RD_u$ and $Q_p = M_p^- \ / \ RS_p$ to provide priority lists. Thereby, the quotients Q_u are sorted in ascending order and the quotients p are sorted in descending order.

Phase 2: Composition of cloud collaborations and assignment of cloud users with the help of the greedy principle for multi-dimensional knapsack problem (Akay, Y., Li, H., & Xu, S.H., 2007).

The building of the priority lists is used as a basis for the subsequent assignment algorithm to determine in which order the cloud provider could be assigned to a collaboration; respectively, to determine which cloud users to which cloud collaboration could be assigned in the cloud market according to their resources demand. The objective of this assignment is again the profit maximization for a broker, i.e., maximization of the difference between revenue and cost.

Phase 1 - Presorting

Priority list for cloud users consists of the quotients $Q_u = M_u^+ \ / \ RD_u$ (willingness to pay / resources demand) that determine willingnesses to pay for a demanded resource unit and are sorted in the ascending order, as due to the objective function, the cloud user with the best willingness to pay will be selected first.

Priority list for cloud providers consists of the quotients $Q_p = M_p^- \ / \ RS_p$ (revenue / resources supply) that determine prices for the bought resource unit and are sorted in the descending order. According to the objective function (namely, profit maximization) the cheapest cloud provider will be selected first to build cloud collaboration.

Phase 2 - Assignment

Further, we apply the assignment algorithm to the priority lists, similar the application of greedy principle to solve a multi-dimensional knapsack problem (Akay, Y., Li, H., & Xu, S.H., 2007), where the units (in our case - users) in knapsacks (in our case - cloud collaborations) are filled according to their values and the maximal weight limits (in our case - aggregated capacities).

Step 1: Search cycles. First, we search in the quotient priority lists (as our solution area) using two cycles - a search cycle in the provider list and a search cycle in the user list. We start with the cheapest cloud provider and insert the second-cheapest cloud provider to build a cloud collaboration. Then we search in the cloud user quotient priority list and insert cloud users with the best willingness to pay first, concerning the aggregated maximal capacity of cloud providers. The search runs forward with the next participants. In this step, the fulfillment and matching of non-functional requirements are not yet considered.

Step 2: Diversification. The next step in the algorithm is the diversification in cloud collaborations. As during the search cycles the cheapest cloud providers were selected firstly, it is in all probability that they also have the worst non-functional guarantees. For this reason, the probability of a successful assignment of cloud users to collaboration is very low. To avoid this side-effect and support diversification, we provide a rotation in the cloud providers' priority list: the first element will be placed at the end of the priority list. Furthermore, the cloud providers with not fulfilled non-functional requirements will be deleted from the list. During the search cycles, the valid compositions of cloud collaborations will be stored (if all cloud users and cloud providers fulfill all requirements and the profit maximization is achieved as well.

Step 3: Composition of Collaborations. Further, the mentioned above greedy principle for the multi-dimensional knapsack problem is used to replace the cloud collaboration with a lower profit by the cloud collaboration with the better profit value. Afterwards, we check whether the cloud collaboration partner can build more than one collaboration. In this case, this partner will be replaced if he can bring more profit; if not, this collaboration will be added to the solution if the objective function > 0. The complete solution - the composition - will be built with the best cloud collaborations.

The asymptotic runtime of CCCP-PRIOSORT.KOM is determined by the search cycles and the rotation step. Thereby, all cloud users $U^{\#}$ and cloud providers $P^{\#}$ will be searched once during the assignment algorithm respectively. The rotation step goes priority lists with the length P through. These steps lead to the asymptotic time $O(P^{\#^2} * U^{\#})$.

9. Evaluation of CCCP-PRIOSORT.KOM

To assess the improvement, we prototypically implemented our proposed heuristic approach CCCP-PRIOSORT.KOM in Java and used the same set up for our evaluation, namely, free JavaILP framework and the commercial IBM ILOG CPLEX framework. The main objective of our evaluation is to assess the required computation time of CCCP-PRIOSORT.KOM for different problem sizes and compare it with the exact optimization approach CCCP-EXA.KOM we provided before. Thus, formally, we regard computation time as the dependent variable of our evaluation. The parameter values or distributions that were used in the problem generation process are the same as summarized in Table 2.

Evaluation Results

The results of our evaluation, i.e., the observed ratio of solved instances and the ratio of the mean computation times in comparison to the CCCP-EXA.KOM approach, are summarized in Table 3. As can be clearly seen, the mean computation times are drastically improved, and even the test case (with $P^{\#} = 12$ and $U^{\#} = 18$) takes only 20.56% of the previously computation time used by the exact approach.

The ratio of the solved instances (from 100 problem instances) goes already down with the test case (with $U^{\#} \geq 8$ and $P^{\#} \geq 8$). Given that the considered problem dimensions are still relatively small in the context of a large cloud market, it can be concluded that the practical applicability of the proposed heuristic optimization approach CCCP-PRIOSORT.KOM is still rather limited. As it has already been explained before, a broker will likely have to decide on the composition of collaborations under rigid time constraints, but also with the best profit. Hence, an important future challenge consists in the improvement of our heuristic with respect to the problem reduction.

Table 3: *Evaluation results of CCCP-PRIOSORT.KOM*

Test case $P^{\#}$; $U^{\#}$	*Ratio of solved instances*	*Ratio of mean computation times*
4 ; 4	*89.79 %*	*3.70 %*
4 ; 6	*88.10 %*	*5.70 %*
6 ; 6	*83.15 %*	*7.90 %*
6 ; 9	*71.79 %*	*11.01 %*
8 ; 8	*66.81 %*	*14.39 %*
8 ; 12	*63.93 %*	*15.06 %*
10 ; 10	*61.19 %*	*17.90 %*
10 ; 15	*54.71 %*	*18.45 %*
12 ; 12	*53.44 %*	*20.03 %*
12 ; 18	*53.79 %*	*20.56 %*

10. Related Work

The work on stable allocations and stable algorithms was recognized as an important theoretical contribution in the 1960s and 1970s, but it was not until the early 1980s that its practical relevance was discovered. The key contribution made Roth, A. (1984) who documents the evolution of the market for new doctors in the U.S. and argues convincingly that a stable algorithm improved the functioning of the market.

The Gale-Shapley allocation mechanisms (1962) rely on a rather abstract idea. If rational people – who know their best interests and behave accordingly – simply engage in unrestricted mutual trade, then the outcome should be efficient. If it is not, some individuals would devise new trades that made them better off. An allocation where no individuals perceive any gains from further trade is called *stable. T*hey examined the case of *pairwise matching*: how individuals can be paired up when they all have different views regarding who would be the best

match. Gale and Shapley analyzed matching at an abstract, general level. They used marriage as one of their illustrative examples. How should ten women and ten men be matched, while respecting their individual preferences? The main challenge involved designing a simple mechanism that would lead to a stable matching, where no couples would break up and form new matches which would make them better off. The solution – the Gale-Shapley "deferred acceptance" algorithm – was a set of simple rules that always led straight to a stable matching.

Niyato, D., Vasilakos, A.V., & Kun, Z. (2011) study the cooperative behavior of multiple cloud providers in order to cooperate and support the establishment of resource pools to offer services to public cloud users. The authors present a stochastic LP game model which takes the random internal demand of cloud providers and a transferable utility into account to define and commit the optimal offer of cooperated cloud providers. In contrast to our work, Niyato et al. do not consider non-functional constraints, i.e., Quality of Service and infsecurity requirements.

In a more recent work, Niyato, D., Wang, P., Hossain, E., Saad, W., & Han, Z. (2012) examine building coalitions between cloud providers as a novel approach to optimize the capacity expansion and maximize the mobile cloud providers' monetary benefits. The authors consider cooperative game theory and the Nash equilibrium principles in their approach and propose admission control and revenue sharing strategies for building cloud provider coalitions and a resource pool for mobile applications. The provided results illustrate improvements in cloud providers' capacity and profit maximization by entering such cloud coalitions. Similar to their previous work, the authors do not consider non-functional constraints, which are an important aspect of our work.

Gohad, A., Ponnalagu, K., Narendra, N.C., & Rao, P.S. (2013) propose a dynamic algorithm for forming self-adaptive cloud collaborations based on the identifying most appropriate healthy set of cloud provider resources (cloud provider capabilities and functional abilities at the SaaS layer), cost modeling and tenancy requirements. The approach is high-lighted with a realistic example. In contrast to us, Gohad et al. focus on ad-hoc resource provisioning, rather than the long-term formation of cloud collaborations, and do not consider security aspects. This specifically includes the cumulative security properties of cloud collaborations that were a focal point of our work.

Song, B., Hassan, M.M., & Huh, E.N. (2010) examine the problem of task selection and allocation to physical machines in the context of dynamic cloud collaborations. Their objective consists in the balancing of resource demands under consideration of different resource types, such as CPU and memory. For that purpose, the authors propose three heuristic optimization approaches, and demonstrate that a cooperative heuristic has benefits with respect to the objective of balanced resource utilization. In contrast to us, Song et al. focus on individual cloud providers and do not regard security requirements.

Mashayekhy, L., & Grosu, D. (2012) model a cloud federation formation problem based on the game theory and formulate a corresponding IP-based optimization approach. In their model, the authors consider the cooperative provisioning of VM instances and storage by federated cloud providers. Their objective consists of profit maximization combined with the formation of stable coalitions, i.e., coalitions in which cloud providers do not have a monetary incentive to switch to different coalitions. In contrast to our work, the authors only consider resource constraints, but do not regard non-functional requirements. Their work also aims at low-level VM provisioning, rather than strategic composition of collaboration.

Kołodziej, J., Khanb, S.U, Wang, L., Kisiel-Dorohinicki, M., Madanie, S.A., Niewiadomska-Szynkiewicz, E., Zomayag, A.Y., Xuh, C-Z., (2012) examine in their research the problem of minimization of the energy consumed in the processes of scheduling and executionof batch of independent tasks submitted in the grid environment. They monitor energy consumption in different grid scenarios based on the security requirements specified by the grid users. The authors define the scheduling issue as a multi-objective Independent Batch Job Scheduling problem in computer grids. Furthermore, they develop genetic-based single- and multi-population meta-heuristics for solving the considered optimization problem. The effectiveness of these algorithms has been empirically justified in two different grid architectural scenarios in static and dynamic modes. In contrast to our work, the authors consider security parameters as a sum of probabilities (a real number between 0 and 1) and use this value to calculate trustworthiness and correctness.

Lampe, U. (2013) introduces the Cloud-oriented Workload Distribution Problem. This problem concerns the distribution of a workload, which comprises multiple computational jobs, across leased infrastructure. This work assumes the position of a cloud user, who aims at cost-minimal deployment under consideration of resource constraints. On the basis of a mathematical optimization model, the author proposes the exact solution approach and

the heuristic optimization approach with the improved computational time. The practical applicability and performance of these optimization approaches is demonstrated using a quantitative evaluation, based on realistic data from the cloud computing market. Furthermore, the author examines the Equilibrium Price Auction Allocation Problem. This problem refers to the allocation of Virtual Machine instances based on an equilibrium price auction scheme. Here, the research is focused on the role of a cloud provider, who pursues the aim of profit maximization. The author formalizes the problem as an optimization model, which permits to deduce the exact optimization approach. In contrast to the work at hand, the author does not consider Quality of Service and security attributes, hence, the usage of optimization models and constraints and quantitative evaluation of results is very similar.

Schuller, D. (2013) examines in his work service marketplaces and the corresponding Service Selection Problem. The author provides optimal as well as heuristic solutions to this problem under consideration of fulfillment of QoS requirements. To achieve this, the author develops an optimization framework specifying and formulating the Service Selection Problem as an optimization problem. In addition to providing the mentioned solution approach for computing optimal solutions to the Service Selection Problem, a heuristic solution method has been developed coping for scalability issues. Both approaches are thereby based on deterministic values for considered non-functional service attributes. In order to assess and reduce potentially negative consequences of differing Quality of Service, the author provides a simulation-based adaptation framework which focuses on reducing the risk of uncertainty and therewith of a potential negative impact of stochastic Quality of Service behavior. Evaluation results show that reductions in total cost up to 30% can be achieved - depending on the considered scenario - by reducing penalty costs that accrue due to the violation of Quality of Service constraints.

Lastly, Hans, R., Lampe, U., Steinmetz, R. (2013) have examined the cost-efficient selection of cloud data centers for the delivery of multimedia services. In that context, the authors propose an exact optimization approach based on IP. While their work is similar with respect to the consideration of resource and Quality of Service constraints, it focuses on a single cloud provider and does neither regard the composition of collaborations nor qualitative non-functional aspects.

In conclusion, to the best of our knowledge, we are the first to examine the profit-maximal, strategic composition of cloud collaborations under consideration of cumulative non-functional properties that result from the very formation of these collaborations, i.e., are determined by the “weakest link in the chain”. Apart from the identification of that specific problem, our main contribution consists in the proposal of both an exact optimization approach, as benchmark, and our heuristic approaches. We consider security requirements and guarantees as quantitative attributes (as binary), i.e., not technically measurable. We also do not consider overall rating approaches to calculate the security level of cloud market participants and use it as a basic or admission criteria for collaboration building.

11. Conclusion and Outlook

While cloud computing promises access to virtually unlimited IT resources, the physical infrastructure of cloud providers is actually limited. Hence, smaller providers may not be able to serve the demands of larger customers. A possible solution is cloud collaborations, where multiple providers join forces to conjointly serve customers. Unfortunately, in such scenario, non-functional Quality of Service and information security properties are determined by the “weakest link in the chain”, rendering the process of composing collaborations cumbersome.

In this work, we introduced the corresponding Cloud Collaboration Composition Problem with our new heuristic optimization approach CCCP-PRIOSORT.KOM, as a complement to our primary work, where we discussed the exact optimization approach for CCCP. Our evaluation results indicated drastic improvement in the computation time, but shows also that it is still applicable to small-scale problem instances, thus indicating the need for further improvements.

In our future work, we aim at the development of heuristic approaches with problem reduction and dynamic changes. In addition, we plan to extend the proposed model to cater for more complex non-functional constraints, such as conditional requirements (e.g., strong data encryption is only required if data is placed outside the European Union). Furthermore, we aim at working on cloud market design. When a market is successfully designed, many cloud market actors are persuaded to participate, thereby creating a fair and orderly market with many trading opportunities. The empirical evidences and quantitative results (e.g., results of case studies or interviews with cloud

providers and cloud auditors) will extend our research scope in order to proper access and understand the functions that markets perform, the conditions required for them to be performed successfully, and what can go wrong if these conditions fail to hold. The cooperation with existing cloud markets is also on our research roadmap that can give us a possibility to gather real market data and use it to evaluate our approaches in the real-world of cloud markets.

Acknowledgment. This work is supported in part by E-Finance Lab e. V., Frankfurt am Main, Germany (http://www.efinancelab.com).

References

Akay, Y., Li, H., & Xu, S.H. (2007). Greedy Algorithm for the General Multidimensional Knapsack Problem. In Annals of Operations Research.

Ates, M., Ravet, S., Ahmat, A., & Fayolle, J. (2011). An Identity-Centric Internet: Identity in the Cloud, Identity as a Service and other delights. In International Conference on Availability, Reliability and Security (pp. 555 – 560).

Bernsmed, K., Jaatun, M.G, Meland, P.H., & Undheim, A. (2011). Security SLAs for Federated Cloud Services. In International Conference on Availability, Reliability and Security (pp. 202 – 209).

Bernstein, D., & Vij, D. (2012). Intercloud Security Considerations. In IEEE International Conference on Cloud Computing Technology and Services (pp. 537 – 544).

Buyya, R., Yeo, C., Venugopal, S., Broberg, & J., Brandic, I. (2009). Cloud Computing and Emerging IT Platforms: Vision, Hype, and Reality for Delivering Computing as the 5th Utility. In Future Generation Computer Systems (pp. 599-616).

Corporation Essvale Corporation Limited. (2008) "Business Knowledge for It in Prime Brokerage".

Durkee, D.(2010). Why Cloud Computing Will Never Be Free. Queue 8(4) (pp. 20-29).

Garg, S.K., Vecchiola, C., & Buyya, R. (2011). Mandi: a market exchange for trading utility and cloud computing services. In Springer Science+Business Media.

Garg, S.K., Versteeg, S., & Buyyaa, R. (2013). A framework for ranking of cloud computing services. In Future Generation Computer Systems, Vol. 29, Issue 4 (pp. 1012–1023).

Gohad, A., Ponnalagu, K., Narendra, N.C., & Rao, P.S. (2013). Towards Self-Adaptive Cloud Collaborations. In Int. Conf. on Cloud Engineering.

Gomes, E., Vo, Q.B., & Kowalczyk, R. (2012). Pure exchange markets for resource sharing in federated clouds. In Concurrency and Computation: Practice and Experience, Vol. 24, Issue 9 (pp. 977–991).

Goyal, P. (2011). Application of a Distributed Security Method to End-2-End Services Security in Independent Heterogeneous Cloud Computing Environments. In IEEE World Congress on Services (SERVICES) (pp. 379 – 384).

Guitart, J., & Torres, J. (2010). Characterizing Cloud Federation for Enhancing Providers' Profit. In IEEE 3rd International Conference on Cloud Computing (CLOUD) (pp. 123-130).

Hans, R., Lampe, U., & Steinmetz, R. (2013). QoS-Aware, Cost-Efficient Selection of CloudData Centers. In 6th Int. Conf. on Cloud Computing.

He., Y.H., Bin, W., Xiao, X.L., & Jing, M.X. (2010). Identity Federation Broker for Service Cloud. In International Conference on Service Sciences (pp. 115 – 120).

Hillier, F., & Lieberman, G. (2005). Introduction to Operations Research. 8th edn. McGraw-Hill .

Jonson, D.S. (2012). A Brief History of NP-completeness. In Documenta Mathematica.

Knuth D.E. (1996). Stable Marriage and Its Relation to Other Combinatorial Problems: An Introduction to the Mathematical Analysis of Algorithms, American Mathematical Society.

Kretzschmar, M., & Golling, M. (2011). Security management spectrum in future multi-provider Inter-Cloud environments - Method to highlight necessary further development. In 5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management (SVM) (pp. 1-8).

Lampe, U. (2013) Monetary Efficiency in Infrastructure Clouds - Solution Strategies for Workload Distribution and Auction-based Capacity Allocation.

Lampe, U., Wenge, O., Müller, A., & Schaarschmidt, R. (2012). Cloud Computing in the Financial Industry - A Road Paved with Security Pitfalls? In 18th Americas Conference on Information Systems (AMCIS), Association for Information Systems (AIS).

Mashayekhy, L., & Grosu, D. (2012). A Coalitional Game-Based Mechanism for Forming Cloud Federations. In 5th Int. Conf. on Utility and Cloud Computing.

Meindl, B., & Templ, M. (2012). Analysis of Commercial and Free and Open Source Solvers for Linear Optimization Problems. Technical report, Technische Universit□at Wien.

Niyato, D., Vasilakos, A.V., & Kun, Z. (2011). Resource and Revenue Sharing with Coalition Formation of Cloud Providers: Game Theoretic Approach. In 11th Int. Symp. on Cluster, Cloud and Grid Computing.

Niyato, D., Wang, P., Hossain, E., Saad, W., & Han, Z. (2012). Game Theoretic Modeling of Cooperation Among Service Providers in Mobile Cloud Computing Environments. In 2012 Wireless Communications and Networking Conf.

Papish, M. (2012). A method for implementing dynamic, cloud-based metadata services based on a unified content ID space across a fragmented CE ecosystem. In IEEE International Conference on Consumer Electronics (ICCE) (pp. 57 – 60).

Rahimi, A.F., & Sheffrin, A.Y. (2003). Effective market monitoring in deregulated electricity markets. In IEEE Transactions on Power Systems, Vol. 18, Issue 2 (pp. 486 – 493).

Roth, A.E. (2002). The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics, In Econometrica, Vol. 70, Issue 4 (pp. 1341–1378).

Roth, A.E. (2008). The Shapley Value: Essays in Honor of Lloyd S. Shapley, Cambridge University Press.

Roth, A.E., Sönmez, T., & Ünver, M.U. (2005). A kidney exchange clearinghouse in New England. In American Economic Review.

Schuller, D. (2013) QoS-aware Service Selection - Optimization Mechanisms and Decision Support for Complex Service-based Workflows.

Shapley, L.S. (1955). Markets as Cooperative Games, RAND Corporation.

Schnjakin, M., Alnemr, R., & Meinel, C. (2010). Contract-based cloud architecture. In Iternational workshop on Cloud data management (CloudDB) (pp. 33-40).

Siebenhaar, M., Wenge, O., Hans, R., Tercan H., & Steinmetz, R. (2013). Verifying the Availability of Cloud Applications. In 3rd International Conference on Cloud Computing and Services Science (CLOSER 2013).

Song, B., Hassan, M.M., & Huh, E.N. (2010). A Novel Heuristic-Based Task Selection and Allocation Framework in Dynamic Collaborative Cloud Service Platform. In 2nd Int. Conf. on Cloud Computing Technology and Science.

Uttam Kumar, T., & Wache, H. (2010). Cloud Broker: Bringing Intelligence into the Cloud. In IEEE 3rd International Conference on Cloud Computing (CLOUD) (pp. 544 – 545).

Wenge, O., Lampe, U., Müller, A., & Schaarschmidt, R. (2014). Data Privacy in Cloud Computing an Empirical Study in the Financial Industry. In 20th Americas Conference on Information Systems.

Wenge, O., Lampe, U., & Steinmetz, R. (2014). QoS- and Security-Aware Composition of Cloud Collaborations. In 4th International Conference on Cloud Computing and Services Science.

Wenge, O., Siebenhaar, M., Lampe, U., Schuller, D., & Steinmetz, R. (2012). Much Ado about Security Appeal: Cloud Provider Collaborations and their Risks. In 1st European Conference on Service-Oriented and Cloud Computing (ESOCC), Springer (pp. 80-90).

Wood, K., & Anderson, M. (2011). Understanding the complexity surrounding multitenancy in cloud computing. In IEEE International Conference on e-Business Engineering (pp. 119 – 124).

Yang, D. (2011). Ad Hoc Aggregation Query Processing Algorithms Based on Bit-Store. In International Conference on in Data Intensive Cloud Cyber-Enabled Distributed Computing and Knowledge Discovery (CyberC) (pp. 313 – 320).