



A Trade-Off Negotiation Strategy for Pareto-Optimal Service Composition with Additive QoS-constraints

Claudia Di Napoli¹ · Silvia Rossi²

Accepted: 10 September 2020 / Published online: 22 September 2020
© The Author(s) 2020

Abstract

QoS-based service composition enables the development of complex business applications. They are composed of distributed services characterized by QoS attributes representing non-functional characteristics, such as cost, execution time, and reliability. With the proliferation of services on the Internet, more candidates for each component service may be available with different QoS values. Hence, the ones satisfying QoS global constraints required for the application have to be selected. Automated negotiation is adopted to select component services in a dynamic market of services where provided QoS values may vary according to different market strategies. Nevertheless, when dealing with multiple QoS attributes, multiple service providers, and not shared information, it is difficult to guarantee the formal properties of the negotiation outcomes. In the present work, we propose a trade-off negotiation strategy that allows reaching a Pareto-optimal agreement, if it exists. The agreement consists of the QoS values of component services that are the ones selected to provide the complete application. The strategy exploits both the competition, that is due to multiple services providing the same functionality with different QoS values, and the cooperation among the providers of the different component services, that are necessary to meet the required end-to-end QoS constraints.

Keywords QoS-based service composition · QoS constraints · Multi-attribute negotiation · Trade-off strategies

✉ Silvia Rossi
silvia.rossi@unina.it

Claudia Di Napoli
claudia.dinapoli@cnr.it

¹ Istituto di Calcolo e Reti ad Alte Prestazioni, C.N.R., Naples, Italy

² Department of Electrical Engineering and Information Technologies, University of Naples Federico II, Naples, Italy

1 Introduction

Service composition encompasses a set of techniques enabling the development of complex distributed business applications, known as Service-Based Applications (SBAs), resulting from the composition of elementary business services. Services are atomic, self-contained computational entities, independent of the execution platform, not subject to centralized control, and usually provided by different owners in highly dynamic environments as Internet (Papazoglou et al. 2007). They provide a specific functionality and communicate with each other through well-defined and interoperable protocols. Technological advances of the last years, together with the growing pervasiveness of services available on the Internet are the driving forces for service composition, a complex process including service discovery, service selection, service orchestration, service binding, and execution.

Services are characterized not only by the functionality they provide but also by non-functional attributes, referred to as *Quality of Service* (QoS) attributes, that represent a measure of how much a service is useful to the customer (Shehu et al. 2014). Typical QoS attributes are the price a customer has to pay to use the service, the availability (i.e. the probability that a service is accessible), the execution time (i.e. the time interval between the service invocation and the output reception), and the reliability (i.e. the ratio between the number of service responses and the number of service requests) (Cardoso et al. 2004). Due to the proliferation of services now available on-demand on the Internet, it is very likely that more services with the same functionality but different QoS attributes are available (Bouguettaya et al. 2017).

In realistic business scenarios, users specify QoS preferences at the level of the entire application, known as end-to-end QoS requirements, and they are not aware of how these preferences can be subdivided locally among the component services. Examples of QoS-based Service-Based Applications are Cloud applications. In these cases, component services have to be executed on computing resources belonging to different Cloud providers at a price depending both on the required computational power, and on the adopted market strategies (De Francesco et al. 2015). The response time, availability, pay-per-use, cost, and reputation are critical QoS attributes in cloud computing (Feng et al. 2016).

Selecting component services means to determine the services whose QoS values, once aggregated, meet these requirements. QoS aggregation rules depend both on the QoS type and on the composition patterns of the application. Typical aggregation rules are summation for attributes such as price and response time, average or min or max for reputation and throughput, multiplication for availability, and reliability. The problem of determining the component QoS values to meet global QoS requirements can be reduced to the Multidimensional Multiple-choice Knapsack Problem (MMKP) (Alrifai and Risse 2009) that is NP-complete, making it necessary to adopt approximation approaches that cause loss of generality (Strunk 2010) and that assume fixed QoS values. However, in a market-based service provision of business applications, these values may vary according to market trends, users' requirements, and providers' market strategies.

To deal with the dynamic nature of an Internet-based market of services, and with the conflicting interests of service consumers and providers regarding QoS attribute values, a viable approach is automated negotiation on the QoS attribute values of component services among the consumer and the service providers (Di Napoli 2009). Nevertheless, most approaches in automated negotiation require complex decision models, making it difficult for their adoption in real settings characterized by a dynamic and competitive market of conflicting requirements (Shojaiemehr et al. 2018).

In the present work, we propose a trade-off negotiation strategy applicable to a market of services characterized by multi-dimensional QoS attributes that allows us to obtain a Pareto-optimal agreement, if it exists, on the QoS attribute values. The providers reaching the agreement are the ones selected to deliver the application. Trading-off is a common market strategy when multiple decision variables can be traded-off against each other, such as paying a higher price to obtain an earlier delivery time of an application, or waiting for a longer time to obtain higher reliability, and so on. The proposed trade-off strategy relies on the possibility to exploit competition among agents providing services with the same functionality, and cooperation among agents providing services with different functionality necessary to meet the global QoS constraints of the required application. In fact, in QoS-based service composition, agents providing service with different QoS values compete to be selected, while agents providing services with different functionality cooperate to provide the QoS values that meet the required global QoS constraints of the application.

The remainder of this paper is organized as follows. In Sect. 2 an overview of related works on QoS-based service selection is reported. In Sect. 3 an overview of the problem addressed and the reference negotiation model is given. In Sect. 4, the proposed trade-off strategy to obtain Pareto-optimal agreements is presented, and its analytical properties supported by simulation results are described in Sect. 5. Section 6 reports the numerical experimentation carried out to show the efficiency and scalability of the proposed strategy in competitive contexts. Conclusions are reported in Sect. 7.

2 Related Works

QoS-aware service selection to meet the overall performance of a Service-Based Application is typically formalised as an optimization problem solved through integer programming techniques (Zeng et al. 2004), or mixed integer programming models, as in Ardagna and Pernici (2007). These methods are not applicable when the number of concrete services increases leading to an NP-hard problem. In these cases, heuristic approaches are used. In Chen et al. (2017) a genetic-based algorithm is proposed to select the best component services among the ones generated at each evolution and ranked according to the Compromise Ratio Method comparing the new solution generated with the top-K ranked solutions of the previous evolution. A genetic-based algorithm is proposed also in Mardukhi et al. (2013) to decompose end-to-end QoS constraints into a set of local constraints for the component

services, while the local selection process relies on a search algorithm to select the best service for each local task.

However, these approaches do not consider the possibility for providers to change QoS values, as in a market-based service provision that is a more realistic scenario for business applications. In these settings, QoS-based negotiation approaches are usually adopted, as in the present work. Many negotiation-based approaches for service composition perform negotiation for each individual service, relying on bilateral negotiation mechanisms between a specific provider, selected according to some criteria, and a service requester (Paurobally et al. 2007; Siala and Ghedira 2011), without considering global QoS constraints. When negotiating to find QoS values for each service in a composition that satisfy end-to-end QoS requirements, the agreement for one service depends on the agreements of the others in the composition. An agreement has to be reached on an aggregation of QoS attributes values offered for the individual services. Hence, it is not possible to evaluate the offers for an individual service without considering the offers for the other components, making approaches based on independent bilateral negotiations not viable.

Other approaches use negotiation as a mechanism to dynamically select the appropriate service providers whose provided service best matches QoS requirements (Alrifai and Risse 2009), as in this work. But negotiation is carried out for each required service independently from the others. As before, these approaches are based on one-to-many concurrent and independent negotiations for each required component service. Attempts to propose a coordinated negotiation with all the providers of the different required services in a composition have been proposed, as in Yan et al. (2007), where a Negotiation Coordinator instructs the negotiation on the single component services by decomposing end-to-end QoS requirements in local QoS. Nevertheless, a decomposition approach prevents from deriving formal properties of the negotiation outcome since it is not possible to determine the most promising decomposition.

Other approaches propose automated auctions to select component services according to QoS values. An iterative multi-attribute combinatorial auction is proposed in He et al. (2014), where bidders are allowed to improve their bids round by round to find a satisfactory solution. At each round, the winner determination problem is solved as a constraint satisfaction problem, and then the solutions found are ranked according to an objective function. The solution with the optimal ranking value is the solution to the winner determination problem of the auction. In Wang et al. (2017) a dynamic pricing mechanism is proposed as an incentive to obtain truthful bidding from candidate services leading to a Bayesian Nash equilibrium that maximizes the utility of the composition of services.

With multi-dimensional QoS attributes, trade-off negotiation strategies can be used allowing a negotiator to generate a new proposal that enhances the opponent's utility without reducing its own utility. In Wu et al. (2009) a trade-off strategy is proposed to reach a Pareto-optimal agreement for the allocation of a single resource among n agents. An end-to-end QoS constraint can be considered as a single resource to be allocated among n agents with incomplete information about others. Nevertheless, the proposed trade-off strategy cannot be applied in our negotiation

scenario where there are not exactly n agents among which to distribute the global QoS constraints, but more agents competing for the allocation of the same service.

3 Negotiation Formalization

Negotiation on the QoS attributes values of different component services leads to a cooperative negotiation among the providers of different component services to reach an agreement, due to the interdependence among them to satisfy the end-to-end QoS requirement. On the other hand, in market-based scenarios, it is likely that there are more providers for the same service. Since only one provider for each component service has to be selected, a competitive negotiation occurs among them. Therefore, negotiation for end-to-end QoS-aware service composition is neither purely distributive nor purely integrative, but competition and cooperation coexist. The proposed trade-off negotiation strategy exploits both competition among service providers of the same component service, and also cooperation among service providers of the different component services required to deliver the application.

To simulate a more realistic market of services, service providers' strategies and utility functions are considered private information, due to the market competitiveness, therefore negotiation takes place in an incomplete information setting (An et al. 2007), and without direct interaction among service providers. Since the providers' strategies are not shared, it is not possible to foresee which providers are more promising during negotiation (An et al. 2011), so all service providers available for each requested service are involved throughout the negotiation process.

We model service providers as software agents responsible for the provision of service implementations, called *Service Providers* (SPs), that interact with the agent responsible for managing an SBA user request, called *Service Compositor* (SC) (Di Napoli 2009). A QoS-based Service Based Application request is specified by its abstract representation, named an *Abstract Workflow* (AW), and by the QoS attributes values required for the entire application. An Abstract Workflow is an acyclic oriented graph $AW = (AS, P)$ in which each node, called an *Abstract Service* (AS), is a requested functionality, and each arc, called a *Precedence* (P), is the precedence relation, in terms of control and data dependencies, among two nodes. The precedence relations define their execution order according to basic composition patterns, such as sequential, conditional, loop, and parallel.

The Service Compositor and the Service Providers for each Abstract Service negotiate to find an agreement on the values of the component QoS attributes to obtain an aggregated value that meets the end-to-end QoS constraints. Only if an agreement is reached for each Abstract Service, the SBA can be delivered with the required QoS.

The adopted negotiation protocol for trading-off is based on the alternating offers protocol initiated by the Service Compositor requesting offers by the Service Providers available for each Abstract Service. The protocol is iterated until one offer for each required Abstract Service is accepted, or a *negotiation deadline* (T), representing the maximum number of allowed iterations, is reached. QoS attributes are

assumed to be multidimensional, such as the execution time and the cost of each component service, and they are negotiated according to a package-deal procedure.

At each iteration, the Service Compositor issues an offer request for each Abstract Service in the composition to all Service Providers available, and it receives back the replies consisting of the offered QoS values. It evaluates the aggregated values of the received offers, and if QoS requirements are satisfied, there is an agreement and the negotiation ends successfully. Otherwise, if the deadline is reached without an agreement a failure is declared to all participants. An iteration, denoted with t , is called a *negotiation round*. If no agreement is reached at iteration t , and $t + 1 \leq T$, another negotiation round is performed.

Let's consider an Abstract Workflow composed of $n \geq 2$ Abstract Services, each one characterized by $m \geq 1$ QoS attributes. For each Abstract Service there are $k \geq 1$ distinct Service Providers. Without loss of generality, an acyclic Abstract Workflow with a sequential composition pattern is considered since it is always possible to reduce a generic workflow to a sequential one (Cardoso et al. 2004). The QoS attributes values are assumed to be normalized, i.e. each attribute is a value in the interval $[0, 1]$ in \mathbb{R} . An offer for an Abstract Service with m QoS attributes, is defined as follows:

Definition 1 (*Offer*) At round t , an *offer* of a Service Provider SP_j providing a service implementation for the i th AS (AS_i) with m QoS attributes' values, is a point $\mathbf{x}_{i,j}^t = (x_{i,j,1}^t, \dots, x_{i,j,m}^t)$ within the unit hypercube.

At each round, the offers are collected by the Service Compositor and aggregated in packages of offers to be evaluated. A package of offers is defined as follows.

Definition 2 (*Package of offers*) At round t , a *package of offers* is a vector of n offers: $(\mathbf{x}_1^t, \dots, \mathbf{x}_n^t)$.

The notation \mathbf{x}_i^t (with $1 \leq i \leq n$) denotes an offer for the i th Abstract Service when it is not needed to specify which SP proposed the offer.

For each attribute h (with $1 \leq h \leq m$) a global constraint C_h is required, representing the aggregated QoS value for all Abstract Services of the application. This value is assumed to be a real positive number. For example, considering as QoS attributes the cost and the execution time of each service, the global constraints represent the requirement that the total cost and the total execution time of the composition of services do not exceed given values.

A package of offers represents an agreement if and only if it meets all the end-to-end QoS constraints, as follows:

Definition 3 (*Agreement*) In case of additive QoS attributes, at round t a package of offers $(\mathbf{x}_1^t, \dots, \mathbf{x}_n^t)$ is an *agreement* if and only if

$$\sum_{i=1}^n x_{i,j,h}^t \leq C_h, \quad \forall h \in \{1, \dots, m\} \quad (1)$$

Thus, the vector $C = (C_1, \dots, C_m)$ represents an upper bound for the aggregated QoS attribute values for all ASs.

At each round, the SPs generate offers according to their own utility functions.

Definition 4 (Utility function) Let t be a round, AS_i the i th Abstract Service, and SP_j a Service Provider for AS_i , the *utility function* of SP_j is:

$$u_j(x_{ij}^t) = \begin{cases} 0 & \text{if } t > T \\ 0 & \text{if } t = T \text{ and Eq. 1 does not hold} \\ v_j(x_{ij}^t) & \text{if } t \leq T \text{ and Eq. 1 does hold} \end{cases} \tag{2}$$

where $v_j(x_{ij}^t)$ is an evaluation function.

Offers with the same utility value constitute the indifference curve of the Service Provider, defined as follows.

Definition 5 (Indifference curve) Let t be a round, AS_i the i th AS, and SP_j a provider for AS_i , the *indifference curve* of SP_j during t is the set:

$$U_{ij}^t = \{x_{ij}^t \in \mathbb{R}^m \mid u_j(x_{ij}^t) = \rho\} \tag{3}$$

where ρ is a constant utility value.

The considered indifferent curves are continuous and strictly convex functions typically used in microeconomics (Lai and Sycara 2009; Mas-Colell et al. 1995). An example a of strictly convex function is the Cobb Douglas one, defined as follows:

$$\gamma = \prod_{h=1}^m x_h^{\lambda_h}, \tag{4}$$

where γ is a constant, m is the total number of attributes, x_h is the h th attribute, and λ_h is the exponent of x_h . This means that when changing the value of an attribute x_h , the others have to change accordingly to maintain constant their product. An example of a Cobb–Douglas function for two attributes is reported in Fig. 1 together with circles indicating offers with the same utility value, but different values for each attribute.

In the case of a successful negotiation, the SPs that provided the offers whose values combined with the others meet the QoS constraints, are the winners of the negotiation, and so the ones selected for providing the service at the agreed conditions. The selected SPs are assumed to commit to providing their services when the negotiation ends, so no penalty mechanism is considered.




Fig. 1 Cobb–Douglas utility function with bids

4 The Proposed Trade-Off Strategy

In multi-issue bilateral negotiation, trade-off strategies are adopted to achieve a Pareto-optimal solution, if it exists. The Orthogonal Bidding Strategy (OBS) (Wu et al. 2009) is a trade-off strategy applied when n agents negotiate with each other for the distributed allocation of a resource. The OBS is based on the concept of *reference point*, calculated for each agent as a point that aggregates the proposals of all other $n - 1$ agents. The n th agent uses this reference point to compute a new offer on its indifference curve, as the one with the minimum Euclidean distance to the reference point. Negotiation proceeds by moving reference points as close to the available offers, until a Pareto solution is reached (or the negotiation ends) (Susetyo et al. 2013). The reference point represents the joint preference that the $n - 1$ agents have on the n th agent's offer: the more this offer is close to the reference point, the more the $n - 1$ agents are satisfied. To find component services providing QoS values that aggregated satisfy end-to-end QoS constraints is similar to allocate a resource among n agents, by considering the multi-dimensional vector of QoS global constraints as the resource to be allocated. As n agents negotiate to get the allocation of a portion of the resource, service providers negotiate to be selected to provide a component service if they can provide the right portion of the required end-to-end QoS value. In the scenario of a market of services, the available providers have incomplete information about others (Di Napoli et al. 2015).

We consider only the trade-off phase of negotiation, and we propose a trade-off strategy that guarantees to achieve a Pareto-optimal agreement if it exists. We assume that no concessions on utilities are made since sellers in a market tend to avoid lowering their utility, while they are available to trade-off to increase their chance to win the negotiation. During trade-off, the utility level of each SP_j is private and it is fixed, i.e. $U_{ij}^s = U_{ij}^t, \forall s, t \leq T, \forall i$, and $\forall j$, and each SP_j generates a new offer on its indifference curve.

In the OBS model, agents interact with each other directly, so they know the offers of the other agents. In our model, they negotiate with the Service Compositor in charge of evaluating the resulting package of offers. Since the Service Compositor collects all offers at each round, it computes the reference points, so it is not feasible for each SP to compute its own reference point as in Wu et al. (2009). In addition, since for each Abstract Service there is a set of SPs offering QoS values, a reference point is computed for all SPs of each Abstract Service, instead of a single SP. At this purpose, it is necessary to determine a representative offer for each Abstract Service to be used for computing the reference point.

The proposed trade-off strategy is composed of a two-step process: in the first one, the most promising offer for each Abstract Service is determined to compute the corresponding reference point; in the second one, offers for each Abstract Service are generated by the corresponding SPs engaged in the negotiation, referring to the reference point.

In our approach, reference points for each Abstract Service are calculated by the Service Compositor as a sort of counteroffers, based on the most promising offers selected at each round. The reference point of each Abstract Service is sent to all the corresponding SPs, and not only to the one that provided the most promising offer.

We defined two opposite criteria to determine the most promising offers, one for each Abstract Service, called respectively *less-is-better* and *more-is-better*. The first one is based on the intuition that the selection of offers with the lower possible values could, in principle, make it easier to obtain a composed solution that meets the global constraint (that is indeed an upper-bound), that is true in the case of a single-issue negotiation. The second one, called *more-is-better* is the opposite, and it is proposed to include all the possibilities to reach a Pareto-optimal solution.

Definition 6 (*Less-is-better* criterion) At round t , the best offer $\check{\mathbf{b}}_i^t$ for the i th AS is:

$$\check{\mathbf{b}}_i^t = \arg \max_{\mathbf{x}_{i,j}^t} \left(\sum_{h=1}^m \frac{Q_{max}^t(i, h) - x_{i,j,h}^t}{Q_{max}^t(h) - Q_{min}^t(h)} \right) \tag{5}$$

where $\mathbf{x}_{i,j}^t$ is the offer for the i th AS at round t proposed by an SP_j , with each offer component $x_{i,j,h}^t$ referring to the h th attribute, and:

$$Q_{max}^t(i, h) = \max_{\mathbf{x}_{i,j}^t} (x_{i,j,h}^t)$$

$$Q_{max}^t(h) = \sum_{i=1}^n \max_{\mathbf{x}_{i,j}^t} (x_{i,j,h}^t)$$

$$Q_{min}^t(h) = \sum_{i=1}^n \min_{\mathbf{x}_{i,j}^t} (x_{i,j,h}^t)$$

We use the notation $\check{\mathbf{b}}_i^t$ for the best offer for the i th Abstract Service when it is not needed to specify which SP_j proposed the offer. In general, we omit the j index when it is not necessary to specify the Service Provider.

At round t , the *best package of offers* according to the *less-is-better* criterion is the vector $\check{b}^t = (\check{b}_1^t, \dots, \check{b}_n^t)$.

Definition 7 (*More-is-better* criterion) At round t , the *best offer* \hat{b}_i^t for the i th AS is:

$$\hat{b}_i^t = \arg \max_{x_{i,j}^t} \left(\frac{\sum_{h=1}^m x_{i,j,h}^t - Q_{min}^t(i, h)}{Q_{max}^t(h) - Q_{min}^t(h)} \right) \tag{6}$$

where $x_{i,j}^t$ is the offer for the i th AS at round t proposed by an SP_j , with each offer component $x_{i,j,h}^t$ referring to the h th attribute, and:

$$Q_{min}^t(i, h) = \min_{x_{i,j}^t} (x_{i,j,h}^t)$$

$$Q_{min}^t(h) = \sum_{i=1}^n \min_{x_{i,j}^t} (x_{i,j,h}^t)$$

$$Q_{max}^t(h) = \sum_{i=1}^n \max_{x_{i,j}^t} (x_{i,j,h}^t)$$

At round t , the *best package of offers* according to the *more-is-better* criterion is the vector $\hat{b}^t = (\hat{b}_1^t, \dots, \hat{b}_n^t)$.

The term, $Q_{max}^t(i, h)$ (resp. Q_{min}) denotes the highest offer (resp. lowest) for the attribute h of the i th AS at round t , and $Q_{max}^t(h)$ (resp. Q_{min}) denotes the sum of the highest offers (resp. lowest) for the attribute h of each AS at round t . The Eqs. 5 and 6 compute the best offer for the i th AS by evaluating the QoS attributes of the standing offers both locally, considering the single service (numerator), and globally, considering the whole composition (denominator), with the convention that if $Q_{max}^t(h) = Q_{min}^t(h)$, then the denominator is 1.

More formally, for all ASs, from Eqs. 5 and 6 we have that all the selected offers x_i^t have at least one component (i.e., an attribute $x_{i,h}^t$) that satisfies the following inequalities:

$$\check{b}_{i,h}^t \leq x_{i,h}^t \quad \text{or} \quad x_{i,h}^t \leq \hat{b}_{i,h}^t \tag{7}$$

This derives from the maximisation (respectively minimisation) of a simple additive weighting method for the selection of the best offers. In fact, if $\forall h \quad x_{i,h}^t \geq \hat{b}_{i,h}^t$, then the *more-is-better* criterion would have selected it as the best offer. The same holds for *less-is-better* criterion.

A reference point is defined as follows:

Definition 8 (*Reference point*) At round t , the *reference point* for the set of Service Providers of an AS_i is:

$$r_i^t = C - \sum_{l \in N - \{i\}} b_l^t \tag{8}$$

where C is the vector of global constraints on QoS attributes, where N is the set $\{1, \dots, n\}$, and $b_l^t \in \{\check{b}_l^t, \hat{b}_l^t\}$.

If we use \check{b}_l^t (resp. \hat{b}_l^t) in Eq. 8, we call r_i^t a *less-is-better* (resp. *more-is-better*) reference point. Consequently, we distinguish between *less-is-better* and *more-is-better* negotiations depending on the chosen criterion. The notation b^t is used when no criterion is specified.

Given a reference point for an AS_i , the corresponding Service Providers generate their trade-off offers as follows.

Definition 9 (*Trade-off offer*) Let $t > 1$ be a round, i an AS, and r_i^t the reference point for i , the offer of every SP_j at round t is the point on the indifference curve of SP_j with the minimum Euclidean distance from r_i^t , i.e.,

$$x_{i,j}^t = P_{U_{i,j}^t}[r_i^t] = \arg \min_{q \in U_{i,j}^t} \|q - r_i^t\| \tag{9}$$

where $P_{U_{i,j}^t}[r_i^t]$ is the Euclidean projection of r_i^t on $U_{i,j}^t$, and $\|\cdot\|$ is the Euclidean norm.

The previous definition implies that if $r_i^t \in U_{i,j}^t$, then $P_{U_{i,j}^t}[r_i^t] = r_i^t$. In other words, the projection of a point on the set $U_{i,j}^t$ it belongs to, is the point itself. If $r_i^t \notin U_{i,j}^t$, and $U_{i,j}^t$ is a compact set with its boundary defined by a differentiable function, then $P_{U_{i,j}^t}[r_i^t]$ lies on the boundary of the set $U_{i,j}^t$, and the line joining r_i^t and $P_{U_{i,j}^t}[r_i^t]$ is normal to $U_{i,j}^t$. In other words, $P_{U_{i,j}^t}[r_i^t]$ is the foot of the perpendicular from r_i^t to the boundary of $U_{i,j}^t$ (Zheng et al. 2016).

5 Analytical Properties of the Proposed Strategy

We will now investigate the two obtained criteria from an analytic point of view, to determine which one allows to reach a Pareto-optimal agreement.

To describe the process to reach an agreement, we first give the following definition.

Definition 10 (*Remaining values*) At round t , the vector of *remaining attributes* is:

$$R^t = C - \sum_{l=1}^n b_l^t \tag{10}$$

where C is the vector of global constraints on the QoS attributes, and $b_l^t \in \{\check{b}_l^t, \hat{b}_l^t\}$.

If R^t has at least one negative component, then the agreement condition of Eq. 1 is not satisfied at round t . Hence, $\|R^t\|$ is a measure of how close the negotiation is to an agreement. In fact, when reference points and best offers join, or equivalently

when $\|\mathbf{R}^t\|$ tends to 0 because \mathbf{R}^t converges to the origin of \mathbb{R}^m , then \mathbf{b}^t is an agreement. Based on these considerations, we demonstrate the following theorems.

Theorem 1 *At round t , if for each AS_i the less-is-better reference point Pareto-dominates every offer $\Rightarrow \check{\mathbf{b}}^t$ is an agreement.*

Proof By hypothesis, $\forall j, \forall h, r_{i,h}^t \geq x_{i,j,h}^t$. Let's suppose by contradiction that $\check{\mathbf{b}}^t$ is not an agreement. Then, there is at least one attribute for which the following inequality holds:

$$\sum_{l \in N} \check{b}_{l,h}^t = \check{b}_{i,h}^t + \sum_{l \in N - \{i\}} \check{b}_{l,h}^t > C_h \rightarrow \check{b}_{i,h}^t > C_h - \sum_{l \in N - \{i\}} \check{b}_{l,h}^t \stackrel{\text{def. 8}}{=} r_{i,h}^t$$

From Definition 8, $\check{b}_{i,h}^t > r_{i,h}^t$, so there is a combination of offers $\check{b}_{i,h}^t$ that dominates the reference point. This contradicts the hypothesis ($r_{i,h}^t \geq x_{i,j,h}^t$) and so the assumption that $\check{\mathbf{b}}^t$ is not an agreement is wrong. \square

It should be noted that if $\check{\mathbf{b}}^t$ is an agreement, it is possible that the reference point computed using the *less-is-better* criterion does not Pareto-dominate each offer. This is shown in Fig. 2 where the offers evolution for a negotiation with 2 ASs, each one provided by 3 SPs using Cobb–Douglas indifference curves, and 2 QoS attributes with global constraints $\mathbf{C} = (0.8, 0.8)$ is reported. In particular, at every round, the reference point for AS_1 does not Pareto-dominate the offer of SP_2 since all reference points calculated for AS_1 lie below SP_2 's curve. On the contrary, in Fig. 3, reporting the offers evolution with the same negotiation settings but using the *more-is-better* criterion, the last reference point lies on SP_2 's curve in the final round. In fact, with the *more-is-better* criterion, the following theorem can be proved.

Theorem 2 *At round t , $\hat{\mathbf{b}}^t$ is an agreement \Leftrightarrow for each AS_i , the more-is-better reference point Pareto-dominates every offer $\mathbf{x}_{i,j}^t$, that is $\forall j$ and $\forall h, r_{i,h}^t \geq x_{i,j,h}^t$.*

Proof (\Rightarrow) Let's suppose that there is an offer \mathbf{x}_j^t for an AS_i that is not Pareto-dominated by the i th *more-is-better* reference point, that is $r_{i,h}^t < x_{i,j,h}^t, \forall h$. Then:

$$r_{i,h}^t \stackrel{\text{def.}}{=} C_h - \sum_{l \in N - \{i\}} \hat{b}_{l,h}^t < x_{i,j,h}^t \rightarrow C_h < x_{i,j,h}^t + \sum_{l \in N - \{i\}} \hat{b}_{l,h}^t \tag{11}$$

Since $\hat{\mathbf{b}}^t$ is an agreement, then:

$$\sum_{l \in N} \hat{b}_{l,h}^t = \hat{b}_{i,h}^t + \sum_{l \in N - \{i\}} \hat{b}_{l,h}^t \leq C_h \tag{12}$$

From Eqs. 11 and 12 it follows that:

$$\hat{b}_{i,h}^t + \sum_{l \in N - \{i\}} \hat{b}_{l,h}^t < x_{i,j,h}^t + \sum_{l \in N - \{i\}} \hat{b}_{l,h}^t \rightarrow \hat{b}_{i,h}^t < x_{i,j,h}^t \quad \forall h \tag{13}$$




Fig. 2 Offers evolution for *less-is-better* negotiation

which contradicts Eq. 7.

(\Leftarrow) The proof is the same as Theorem 1 since it does not rely on the different formal definitions of \check{b}^t and \hat{b}^t . □

Thus, the double implication of Theorem 2 holds only when the *more-is-better* criterion is adopted. Since the theorem is a requirement to demonstrate the Pareto-optimality of an agreement, only the *more-is-better* criterion is considered.

5.1 Pareto-Optimality

Any convergence of R^t to the origin of \mathbb{R}^m does not necessarily imply the achievement of a Pareto-optimal agreement, because it may still be possible for some




Fig. 3 Offers evolution for more-is-better negotiation

SP to trade-off and to get a Pareto improvement. In fact, as shown in the example reported in Fig. 4, the reference point converges to the best offer, but unlike Figs. 2 and 3, the convergence trajectory is above every indifference curve. Since the curves are strictly convex, any point (i.e. an offer) above SP_3 's curve of Fig. 4 has greater utility for SP_3 than the one on its curve.

Let's define the conditions for which Pareto improvements are not possible from a geometrical point of view.

Given a round t , and an AS_i , the notation $p[b_i^t]$ indicates the service provider offering b_i^t . $\mu[b_i^t]$ is the utility value of its indifference curve, and $U[b_i^t]$ the set of points composing it.

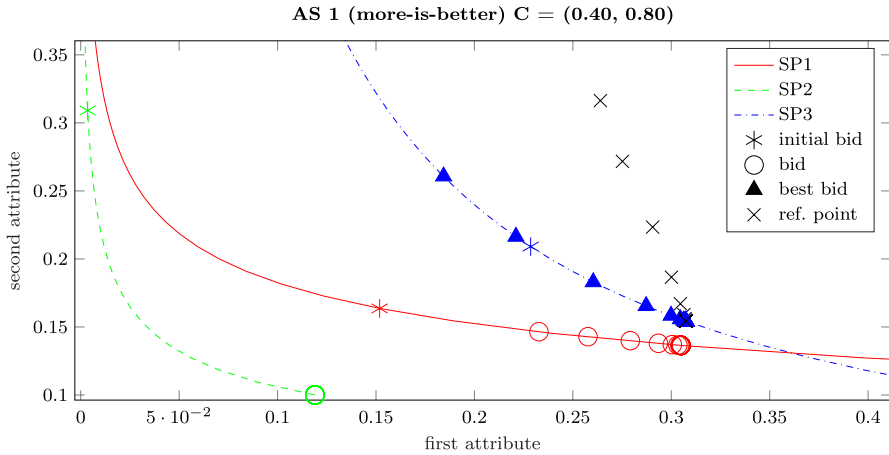


Fig. 4 Offers evolution for *more-is-better* negotiation with low utility levels w.r.t. global constraints

Definition 11 (Agreement area) Given a round t and an AS_i , the *agreement area* for AS_i is the following set:

$$X_i^t = \left\{ x_i^t \mid \forall l \neq i \exists x_l^t : v_{p[b_l^t]}(x_l^t) \geq \mu[b_l^t] \wedge \forall h \quad x_{i,h}^t + \sum_{l \neq i} x_{l,h}^t \leq C_h \right\} \quad (14)$$

This means that any offer for AS_i in the agreement area X_i^t allows to reach an agreement with a utility value that is at least as good as the utility of the most promising offer.

From the Theorem 2, if \hat{b}^t is an agreement, then each reference point Pareto-dominates all the offers for the AS_i . It follows that X_i^t is a set whose Pareto-frontier contains the *more-is-better* reference points for AS_i . From Definitions 4 (utility function) and 8 (reference point), r_i^t is the complement of $\sum_{l \in N - \{i\}} b_l^t$ with respect to C , and every b_l^t belongs to strictly convex sets. Hence, the points in X_i^t constitute a strictly concave set since the sum of convex functions is still a convex function.

In the following theorems, we assume for simplicity that the number m of QoS attributes is 2.

Theorem 3 Given a round $t \leq T$ and an AS_i , $|X_i^t \cap U[b_i^t]| > 1 \Leftrightarrow$ an agreement $x^t = (x_1^t, \dots, x_n^t)$ exists such that $u_{p[b_i^t]}(x_i^t) > \mu[b_i^t]$.

Proof (\Rightarrow) According to the hypothesis, we have that $|X_i^t \cap U[b_i^t]| > 1$. Let $y_i \neq z_i$ be two different points in $X_i^t \cap U[b_i^t]$. Let x_i be a convex combination of y_i and z_i . Since the indifference curve $U[b_i^t]$ is strictly convex, x_i is on its right side, and so $v_{p[b_i^t]}(x_i) > \mu[b_i^t]$. Moreover, since $X_i^t \cap U[b_i^t]$ is a convex subset of X_i^t , then $x_i \in X_i^t$, and so an agreement x^t exists such that $x_i \in x^t$ and $u_{p[b_i^t]}(x_i, t) = v_{p[b_i^t]}(x_i^t) > \mu[b_i^t]$.

(\Leftarrow) According to the hypothesis, an agreement x^t exists, with $x_i^t \in x^t$ and $u_{p[b_i^t]}(x_i^t) = v_{p[b_i^t]}(x_i^t) > \mu[b_i^t]$. Let $y_i = (y_{i,1}, y_{i,2})$ be a point with utility value $\mu[b_i^t]$

such that $y_{i,1} = x_{i,1}$, and $z_i = (z_{i,1}, z_{i,2})$ the point with utility value $\mu[b_i^t]$ such that $z_{i,2} = x_{i,2}$. Both points exist because the utility is a continuous function. Moreover, they belong to X_i^t since $y_{i,2} \leq x_{i,2}$, $z_{i,1} \leq x_{i,1}$ and the utility function is increasing monotonic with respect to each attribute because it is strictly convex. Hence, $X_i^t \cap U[b_i^t]$ includes these two points. \square

In Fig. 5 the two implications of Theorem 3 are described. On the left, it is shown that if the intersection of the agreement area with the indifference curve is greater than 1 (i.e., $y_i, z_i \in X_i^t$), then there is a point x_i leading to an agreement with $u_{p[b_i^t]}(x_i, t) > \mu[b_i^t]$. On the right, it is shown that if there is a point x_i of an agreement such that $u_{p[b_i^t]}(x_i, t) > \mu[b_i^t]$, then the agreement area intersects the indifference curve in more than one point.

If $X_i^t \cap U[b_i^t] = \emptyset$, it is not possible to reach an agreement with the considered utility indifference curves. Hence, if the indifference curve and the agreement area, for each AS_i , intersect only in one point, the agreement including the intersection point is Pareto-optimal.

Lemma 1 Given a round $t \leq T$, and an AS_i , $|X_i^t \cap U[b_i^t]| = 1 \Leftrightarrow$ a Pareto-optimal agreement exists $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$ such that $u_{p[b_i^t]}(\mathbf{x}^t) = \mu[b_i^t]$.

Proof (\Rightarrow) According to the hypothesis, $|X_i^t \cap U[b_i^t]| = 1$, hence only one agreement exists, and it is Pareto-optimal.

(\Leftarrow) Let's suppose that a Pareto-optimal $\mathbf{x}^t = (x_1^t, \dots, x_n^t)$ agreement exists such that $u_{p[b_i^t]}(x_i, t) = \mu[b_i^t]$. Let's suppose, by contradiction, that for an AS_i $|X_i^t \cap U[b_i^t]| \neq 1$. If $|X_i^t \cap U[b_i^t]| = 0$, it is not possible to obtain an agreement with the current utility indifference curves. If $|X_i^t \cap U[b_i^t]| > 1$, from Theorem 3 derives that other solutions with a better utility may exist. Hence, \mathbf{x}^t is not a Pareto-optimal agreement. \square

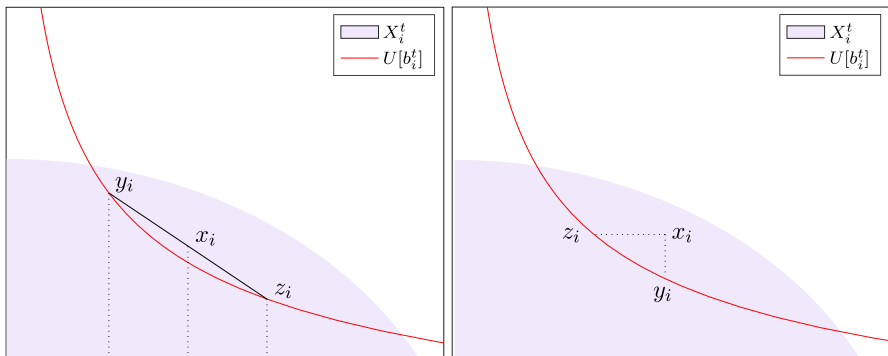


Fig. 5 Examples of agreements area for an AS_i with 2 QoS attributes

Concluding, in the trade-off phase, a set of offers is a Pareto-optimal agreement if and only if once an offer is selected for the agreement, the intersection of its utility indifference curve with respect to the agreement area is unique. Obviously, this does not mean that there is only one Pareto-optimal agreement since different Service Providers proposing offers for the same AS_i could have the utility indifference curves with this property, so different Pareto-optimal agreements may exist. From Theorem 2 and Lemma 1, if at round t a Pareto-optimal agreement exists, the agreement is $\mathbf{r}^t = \hat{\mathbf{b}}^t$.

6 Experimental Results

Simulation experiments aimed at validating the assumptions of our model are carried out. First, we analyze if the market competitiveness due to negotiating with all available service providers allows to obtain Pareto-optimal solutions. In addition, we analyze the scalability of the strategy in terms of the number of rounds to reach an agreement when increasing the number of Service Providers.

In all tests the following settings are used:

- number of QoS attributes $m = 2$;
- maximum number of rounds $T = 250$;
- number of Abstract Services n between 2 and 9;
- number of Service Providers $k \in \{2, 4, 6, 8\}$;
- each global constraint is between 0.5 and 1;
- each SP has a Cobb–Douglas indifference curve ($\gamma = x_1^\alpha x_2^\beta$) with the following parameters: $\alpha, \beta \in (0.2, 1)$, $\gamma \in (5, 7.5)$, $u \in (0.6, 1)$;
- each SP always proposes as first offer a random point on its indifference curve.

The values of the indifference curves' parameters are chosen so that their points are in $[0, 1]^2$ for normalization. For each experimental setting, 100 distinct tests are executed.

6.1 Competitiveness

For each considered setting, the following 3 test cases were executed:

1. **Global**: consisting in negotiating with all available Service Providers;
2. **First round best**: consisting in negotiating with only the most promising Service Providers for each AS selected at the first round, that is $p[\mathbf{b}^1] = (p[\mathbf{b}_1^1], \dots, p[\mathbf{b}_n^1])$;
3. **Last round best**: consisting in negotiating with only the Service Providers that resulted to be the winners of the **Global** case, that is $\mathbf{w} = (p[\mathbf{b}_1^T], \dots, p[\mathbf{b}_n^T])$.




Fig. 6 Number of rounds to reach agreement in the considered test cases

Table 1 Number of rounds for reaching NPA and PA in the considered test cases

	$(NPA, PA)^1$	$(NPA, PA)^2$	$(NPA, PA)^3$
Global	(2, 17)	(2, 13)	(3, 16)
First-round best	(3, 21)	(2, 10)	(41, 233)
Last round best	(3, 19)	(2, 11)	(2, 14)

Obviously, in the last two cases, no competition among Service Providers of the same Abstract Service occurs. The **Global** negotiation is used as baseline since it is always successful (a Pareto-optimal agreement is found).

For each considered test case, the numbers of rounds required for reaching the first agreement (also not Pareto), the first near Pareto-optimal agreement (NPA), and the Pareto-optimal agreement (PA) are reported. An NPA is an agreement with the following conditions: $R^s \leq R^t, \forall s \leq t, e R^t = -0.01$.

In Fig. 6 each column accounts for a specific configuration setting while the rows represent the three considered test cases. Diagonal cases represent the best results in terms of the obtained Pareto-optimal agreement with respect to the number of rounds, for which the norm of the vector of remaining attributes $\|R^t\| = \|C - \sum_{i \in N} b_i^t\|$ and the distance between reference points and best offers $\sum_{i \in N} \|r_i^t - b_i^t\|$ converge to 0 in fewer rounds.

In Table 1, we show the number of rounds necessary to reach Near Pareto Agreement (NPA) and Pareto Agreement (PA) for each case of Fig. 6. The results show that there are cases in which the First-round best is faster than Global (e.g., the second column where First-round best reach a PA agreement in 10 rounds while

Global in 13), and cases in which it is significantly slower (e.g., the third column where **First-round best** reach a PA agreement in 233 rounds while **Global** in 16). Indeed, in the **First-round best** there are cases not convergent/oscillating behavior of the reference points. Figure 7 shows that for a configuration with 3 Abstract Services and 4 Service Providers, **Global** converges to the NPA in 3 rounds and to the PA in 23 rounds, while both **First round best** and **Last round best** do not converge to 0. Figure 8 reports the negotiation trends of the previously considered case. The columns represent the test cases, while the rows represent the 3 Abstract Services. Note that, in the case of **First-round best** and **Last round best** the Service Providers selected for the negotiation are different. In addition, as previously shown, their offers do not converge to a Pareto-optimal agreement.




Fig. 7 Example of not convergence to zero in non-competitive cases

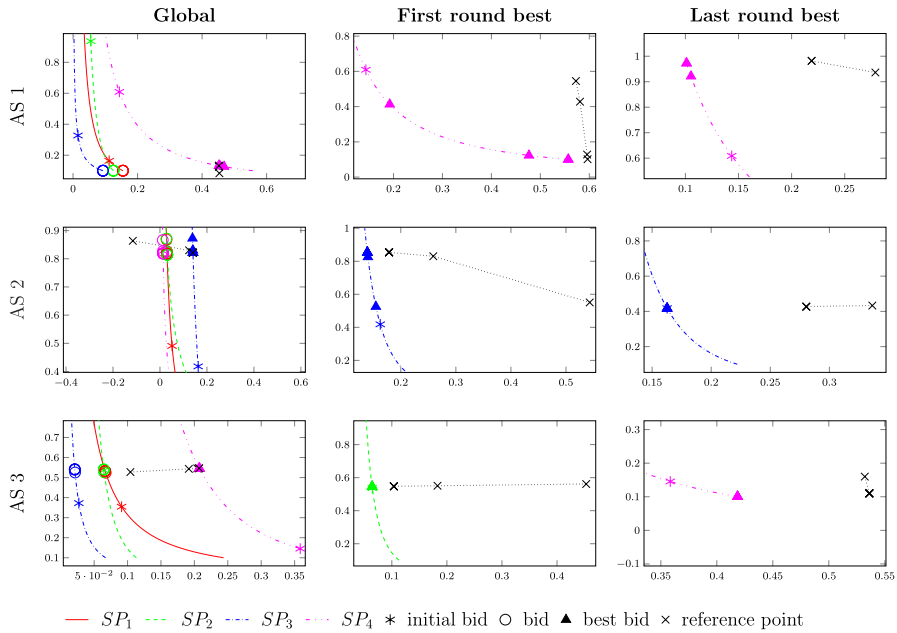


Fig. 8 Negotiations for the example of Fig. 7

In conclusion, while choosing to negotiate only with 1 Service Provider for each AS reduces the negotiation space, the convergence to the optimal agreement is not guaranteed even when the selected SP is more promising. In particular cases, negotiating with only 1 SP could lead to better performances, but, in general, some regions in the negotiation space where there is a Pareto-optimal agreement could be excluded. In other words, eliminating SPs which could become more competitive during the negotiation might invalidate the Pareto-optimality conditions.

6.2 Scalability

We run some tests with an increasing number of AS and SP to evaluate their impact on the number of rounds to reach an agreement. For every (n, k) combination, we calculated the mean and standard deviation of the rounds necessary for reaching NPAs and PAs. Figure 9 shows the obtained results.

Results show that, for the considered cases, the average number of rounds to reach both NPAs and PAs seems to be influenced more by the number of SPs for each AS, than by the number of ASs. In any case, no negotiation required more than 7 rounds for an NPA, and 80 for a PA. This is an encouraging result since the proposed strategy may be adopted in the case of more complex (i.e., with more ASs) abstract workflows. Finally, experiments show that for $n \in [6, 8]$, PA and NPA are obtained in the same number of rounds. This will be further investigated in future works to evaluate results with more SPs.

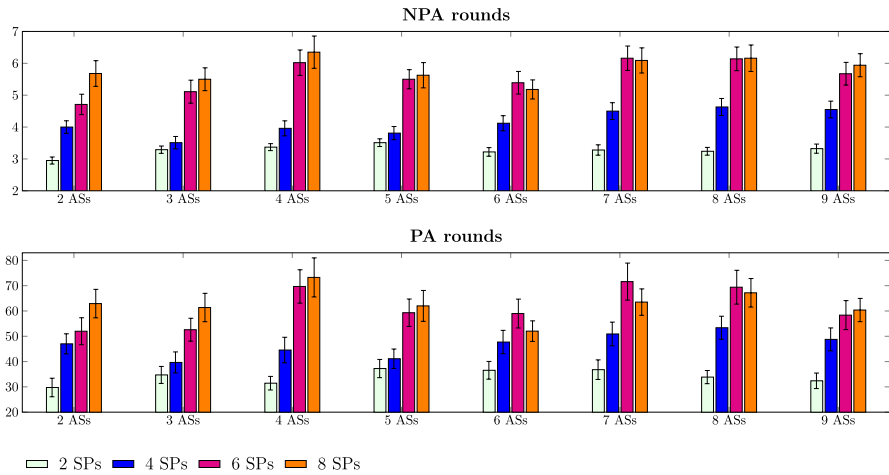


Fig. 9 Scalability of NPAs and PAs

7 Conclusions and Future Work

Automated negotiation proved to be a viable approach allowing a service consumer to dynamically negotiate with service providers to reach agreements on end-to-end QoS constraints that have to be met by the aggregation of the QoS values of component services. The providers reaching the agreement are selected for the composition of services with required QoS values. The negotiation involves two aspects: one aspect is the negotiation between the service consumer and one or more service providers for a single service in the composition; the other aspect is the coordination of the negotiation for the different services in the composition to ensure end-to-end QoS values. Services are characterized by multiple QoS attributes, so the negotiation is multi-agent and multi-issue. For this type of negotiation, it is more difficult to derive formal properties, and to determine the Pareto-optimality of its solutions.

In this work, we propose a trade-off strategy that extends the orthogonal bidding strategy presented in Wu et al. (2009) to include the possibility to negotiate QoS values with all providers for the same service required in the composition and not with only one provider for each component service. By doing so, the proposed strategy exploits the competition among the providers of the same service. At each negotiation round, the most promising offers for each component service are selected among the available ones, and they are used to compute a reference point for each set of providers for the same service. This point represents the most satisfying offer for that service, and it is used as a reference to generate offers by all providers of that service.

In multi-issue negotiation between two rational agents, there is a trade-off between the quality of the agreement reached and the amount of information revealed. In this work, we showed that the possibility to exploit both competition and cooperation allows us to obtain Pareto-optimal agreements also when no information is shared among providers, as in the case of a market oriented QoS-based service composition. Cooperation between providers of different services is exploited by pushing

the providers of each service of the composition to offer QoS values that are more beneficial for the providers of the remaining component services, while maintaining the same utility. Competition between providers of the same service is exploited by including all of them throughout the negotiation process, so the providers which might become more satisfying for the customer, as the negotiation proceeds, are not discarded.

The two criteria proposed to determine the set of most promising offers differ from each other because the *less-is-better* privileges the lowest aggregation of QoS attributes, while the *more-is-better* privileges the highest one. We showed that only by using the *more-is-better* criterion, so privileging the offer that outperforms the others, it is possible to formally characterize the conditions of agreement and Pareto-optimality. With the *less-is-better* criterion, it is possible to find solutions more quickly than with the *more-is-better* one, and they may also be Pareto-optimal in some cases. However, in some cases, an oscillating behaviour of the offers is possible so the Pareto-optimality is not guaranteed.

Lastly, we provided experiments that indicate the robustness and scalability of the proposed strategy by varying the number of Service Providers, and of Abstract Services.

We plan to extend this work by considering also not additive QoS attributes. This means that the agreement conditions have to be redefined and, consequently, a new criterion for defining the most promising offers needs to be identified. So, the agreement and the Pareto-optimality conditions will have to be demonstrated again.

Funding Open access funding provided by Università degli Studi di Napoli Federico II within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alrifai M, Risse T (2009) Combining global optimization with local selection for efficient qos-aware service composition. In: Proceedings of the 18th international conference on world wide web, ACM, pp 881–890
- An B, Miao C, Shen Z (2007) Market based resource allocation with incomplete information. In: Proceedings of the 20th international joint conference on artificial intelligence, Morgan Kaufmann Publishers Inc., IJCAI'07, pp 1193–1198
- An B, Lesser V, Sim KM (2011) Strategic agents for multi-resource negotiation. *Auton Agent Multi-Agent Syst* 23(1):114–153
- Ardagna D, Pernici B (2007) Adaptive service composition in flexible processes. *IEEE Trans Softw Eng* 33(6):369–384

- Bouguetaya A, Singh M, Huhns M, Sheng QZ, Dong H, Yu Q, Neiat AG, Mistry S, Benatallah B, Medjahed B, Ouzzani M, Casati F, Liu X, Wang H, Georgakopoulos D, Chen L, Nepal S, Malik Z, Erradi A, Wang Y, Blake B, Dustdar S, Leymann F, Papazoglou M (2017) A service computing manifesto: the next 10 years. *Commun ACM* 60(4):64–72
- Cardoso J, Sheth A, Miller J, Arnold J, Kochut K (2004) Modeling quality of service for workflows and web service processes. *Web Semant Sci Serv Agents World Wide Web* 1(3):281–308
- Chen F, Li M, Wu H (2017) Gacrm: a dynamic multi-attribute decision making approach to large-scale web service composition. *Appl Soft Comput* 61:947–958
- De Francesco A, Di Napoli C, Giordano M, Ottaviano G, Perego R, Tonello N (2015) Midas: a cloud platform for soa testing as a service. *Int J High Perform Comput Netw* 8(3):285–300
- Di Napoli C (2009) Software agents to enable service composition through negotiation. *Stud Comput Intell* 170:275–296
- Di Napoli C, Di Nocera D, Rossi S (2015) Computing pareto optimal agreements in multi-issue negotiation for service composition. In: *Proceedings of the 2015 international conference on autonomous agents and multiagent systems, International Foundation for Autonomous Agents and Multiagent Systems, AAMAS '15*, pp 1779–1780
- Feng Y, Zhijian W, Qian H (2016) A novel qos-aware mechanism for provisioning of virtual machine resource in cloud. *J Algorithms Comput Technol* 10(3):169–175
- He Q, Yan J, Jin H, Yang Y (2014) Quality-aware service selection for service-based systems based on iterative multi-attribute combinatorial auction. *IEEE Trans Softw Eng* 40(2):192–215
- Lai G, Sycara K (2009) A generic framework for automated multi-attribute negotiation. *Group Decis Negot* 18(2):169–187
- Mardukhi F, NematBakhsh N, Zamanifar K, Barati A (2013) Qos decomposition for service composition using genetic algorithm. *Appl Soft Comput* 13(7):3409–3421
- Mas-Colell A, Whinston MD, Green JR (1995) *Microeconomic theory*, vol 1. Oxford University Press, New York
- Papazoglou M, Traverso P, Dustdar S, Leymann F (2007) Service-oriented computing: state of the art and research challenges. *IEEE Comput* 40(11):38–45
- Paurobally S, Tamma V, Wooldridge M (2007) A framework for web service negotiation. *ACM Trans Auton Adapt Syst* 2(4):14
- Shehu U, Epiphaniou G, Safdar GA (2014) A survey of qos-aware web service composition techniques. *Int J Comput Appl*
- Shojaimehr B, Rahmani AM, Qader NN (2018) Cloud computing service negotiation: systematic review. *Comput Stand Interfaces* 55:196–206
- Siala F, Ghedira K (2011) A multi-agent selection of web service providers driven by composite qos. In: *Proc. of 2011 IEEE symposium on computers and communications (ISCC)*, IEEE, pp 55–60
- Strunk A (2010) Qos-aware service composition: a survey. In: *8th European conference on web services, IEEE*, pp 67–74
- Susetyo C, Timmermans H, de Vries B (2013) Computer-mediated e-negotiation support system. In: *Proceedings of the 13th international conference on computers in urban planning and urban management*
- Wang P, Liu T, Zhan Y, Du X (2017) A Bayesian Nash equilibrium of qos-aware web service composition. In: *2017 IEEE international conference on web services, ICWS 2017, Honolulu, HI, USA, June 25–30, 2017*, pp 676–683
- Wu M, de Weerd M, La Poutré H (2009) Efficient methods for multi-agent multi-issue negotiation: allocating resources. In: *International conference on principles and practice of multi-agent systems*, Springer, pp 97–112
- Yan J, Kowalczyk R, Lin J, Chhetri MB, Goh SK, Zhang J (2007) Autonomous service level agreement negotiation for service composition provision. *Future Gener Comput Syst* 23(6):748–759
- Zeng L, Benatallah B, Ngu AHH, Dumas M, Kalagnanam J, Chang H (2004) Qos-aware middleware for web services composition. *IEEE Trans Softw Eng* 30(5):311–327
- Zheng R, Chakraborty N, Dai T, Sycara K (2016) Automated multilateral negotiation on multiple issues with private information. *INFORMS J Comput* 28(4):