Normal Distributions and Multi-issue Negotiation for Service Composition*

Silvia Rossi[1], Dario Di Nocera[2,**], and Claudia Di Napoli[3]

[1] Dipartimento di Ingegneria Elettrica e Tecnologie dell'Informazione
University of Naples "Federico II", Napoli, Italy
silvia.rossi@unina.it

[2] Dipartimento di Matematica
University of Naples "Federico II", Napoli, Italy
dario.dinocera@unina.it

[3] Istituto di Calcolo e Reti ad Alte Prestazioni
C.N.R., Napoli, Italy
claudia.dinapoli@cnr.it

Abstract. Software negotiation is gaining an increased popularity as a viable approach to establish agreements between service providers and consumers of QoS-aware Service-Based Applications (SBA) composed of services provided by different agents. In most cases, QoS preferences are expressed as end-to-end quality requirements on the whole application, and different service agents have to provide services with QoS values that, once aggregated, have to meet them. In the present work we analyze the properties of a hybrid iterative negotiation mechanism occurring among a composer agent and service provider agents on the QoS attributes of the required SBA. The proposed negotiation relies on normal probability distributions to model service provider agents, and it allows to model single-issue and multi-issue negotiation within the same negotiation framework in terms of adopted concession strategy, utility and protocol.

1 Introduction

The increased popularity of Service Oriented Computing [1] is enhancing the development of Service Based Applications (SBAs), i.e. distributed applications composed of services provided by independent and autonomous providers, in a loosely coupled manner that collectively fulfill a requested task. Usually, users requesting an SBA specify also non-functional requirements, referring to Quality

* The research leading to these results has received funding from the EU FP7-ICT-2012-8 under the MIDAS Project (Model and Inference Driven - Automated testing of Services architectures), Grant Agreement no. 318786, and the Italian Ministry of University and Research and EU under the PON OR.C.HE.S.T.R.A. project (ORganization of Cultural HEritage for Smart Tourism and Real-time Accessibility).

** Ph.D. scholarship funded by Media Motive S.r.l., Napoli - Italy, POR Campania FSE 2007-2013.

J.B. Perez et al. (eds.), *Trends in Practical Applications of Heterogeneous Multi-agent Systems. The PAAMS Collection,* Advances in Intelligent Systems and Computing 293,
DOI: 10.1007/978-3-319-07476-4_1,

of Service (QoS) attributes of the application, that need to be fulfilled by the providers of the component services. Typical examples of QoS attributes are price, response time, reliability, reputation, and so on. Users may not need to be part of the composition process as long as the functional and non-functional requirements they specify are satisfied [2]. Negotiation mechanisms were shown to be a suitable approach to deal with QoS-aware SBAs [2,3,4], allowing to create composition of services that meet the users QoS requirements. Negotiation occurs on the QoS attributes of the services composing the application that represent the issues of negotiation.

However, practical negotiation mechanisms for SBA applications must be computationally efficient [4], and negotiation strategies should be developed based on the assumption of bounded rather than perfect rationality [5]. Moreover, negotiation protocols for these applications have to be more complex than traditional bilateral negotiation. In most cases, when dealing with QoS-aware SBAs, there are multiple QoS attributes representing non-functional characteristics of the SBA component services, so negotiation has to be modeled as a multi-issue one. While one-issue negotiation is widely studied in literature, negotiation on multiple-issue is less mature [6]. Moreover, the tractability requirement is fundamental in order to apply negotiation mechanisms when services are made available on the market to end users with QoS values depending on market trends. In such settings, service applications providers and consumers have to engage in interactions easy to model and that quickly converge to an agreement.

In this paper we show that negotiation in SBAs is inherently multi-issue even in the case of a single issue negotiation. In fact, also in the case of one QoS attribute for each service composing the application, the QoS value of the complete application, that have to meet the end-to-end constraint required by the user, is obtained by composing the single QoS values provided by the different component services. In particular, we show that when dealing with composition of services, the same utility functions and strategies used to model a negotiation on r issues, representing the QoS attributes of a single service, can be used for the case of one-issue negotiation for r services composing the application. This means that the complexity of the negotiation for SBAs depends on the number of issues and on the number of services composing the SBA in an uniform way. Finally, we show that non linear utility functions, as well as concession strategies, can be modeled through the use of normal distributions with an uniform strategy approach. Normal distributions convolution properties when scaling up in dimensions, allow to deal with computational tractability requirements necessary in real market of services. These properties allow to use the same negotiation mechanism in terms of protocols, strategies and utilities for both single issue and multi-issue negotiation, when dealing with composition of services. Furthermore, the use of normal distributions allows to simulate the stochastic behaviour of service providers with zero-intelligence that can be used to approximate the trends of a volatile and open market of services [7]. Stochastic behavior can be often observed in practical multi-agent negotiation applications [8].

2 One-to-Many-to-Many Asymmetric Negotiation

In many real market situations, provider agents may adopt negotiation strategies to formulate offers, while composer may only have range of acceptance (in some case flexible) on the complete set of offers for the SBA QoS attributes they require. They are not able to provide single counteroffers for each service, but they can only evaluate acceptable and unacceptable offers for the complete package. So, symmetric protocols requiring a strong symmetry between composers and providers [9] are not appropriate.

In [3], we proposed a *one-to-many-to-many* negotiation mechanism allowing only the provider agents to formulate new offers for the issue to be negotiated upon, and only the composer agent (acting on behalf of the user) to evaluate them both individually and globally. The rationale of this choice is that offers for a single functionality cannot be evaluated independently from the ones received for the other functionalities when the value of the issue to be negotiated upon results from the composition of the values provided by the component services. In fact, in such a case, when the value of the issue provided by one services changes, also the values of the same issue provided by the other component services have to change accordingly in order to meet the user's constraint. The protocol of the negotiation is based on an Iterative Contract Net Protocol. It allows a composer agent to negotiate separately with all the agents available for each service, and it may be iterated for a variable number of times (*rounds*) until a *deadline* is reached or the negotiation is successful. A successful negotiation occurs if a complete set of offers, one for each service composing the requested SBA, is accepted.

3 Single and Multi-issue Strategy for the Provider

In order to prepare an offer, an agent uses a set of tactics to generate new values for each negotiated issue [10]. In fact, agents must be provided with strategies to formulate offers and they must be equipped with algorithms to evaluate their utilities for the offers. This is done by evaluating an offer in terms of agent utility with respect to the offer. For each provider, a negotiation strategy on a single parameter (q_1) can be modeled by a Gaussian function [3], as shown in Figure 1a. Such distribution is used both to map values of the single QoS into an utility value for the provider, but also as a strategy to select concession values of the utility. In particular, the Gaussian function represents the probability distribution of the offers in terms of the provider's utility as follows: $U_x(q_1) = \exp\left(-\frac{1}{2}\left(\frac{q_1-\mu_1}{\sigma_1}\right)^2\right)$.

As shown in Figure 1a, the mean value of the Gaussian μ represents the best offer the provider agent may propose in terms of its own utility ($U(\mu) = 1$), but, at the same time, the QoS value with the highest probability to be selected. The standard deviation σ represents the attitude of the provider to concede during negotiation (i.e. greater deviation corresponds to higher concession rate), and the

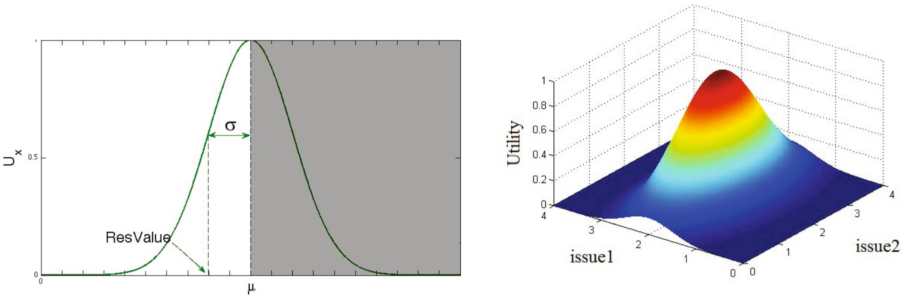

(a) Probability distributions for one-issue.

(b) Probability distribution for multi-issues.

Fig. 1. Utilities functions and Gaussian distributions

reservation value ($ResValue$) is set equal to $\mu - \sigma$. In this way the reservation value for a provider is related to its concession strategy, and it is also a value with a known probability to be selected according to its probability distribution. The negotiation set (i.e. the negotiation space for the provider) to be considered is only $[\mu - \sigma; \mu]$ (or $[\mu; \mu + \sigma]$), so only the white (or grey) section of Figure 1a is considered. Values of the utility function are in the domain [0,1], i.e. it is normalized so the values of the QoS attributes are evaluated according to the same scale in order to avoid inaccurate evaluation due to different measurement metrics used for different QoS attributes.

At each negotiation round, a provider generates, following its probability distribution, a new utility value corresponding to a new offer. In order to follow a monotonic concession protocols, if the utility value of this new offer is lower than the one offered in the previous round and within the negotiation set, then the provider proposes the new value. If this value is greater than the one offered in the previous round, or it is outside the negotiation set, the provider proposes the same value offered in the previous round. This strategy allows to simulate different and plausible behaviours of providers that prefers not having a constant loss in utility, even though by increasing the number of negotiation rounds the probability for the provider to move towards its reservation value increases.

The chosen function was shown to be both time and resource dependent. In fact, it takes into account both the *computational load* of an agent, driving its attitude to concede, and the *computational cost* of the provided service corresponding to its best utility value. The computational load of the provider accounts for its workload in terms of the amount of resources it has to provide the service implementations it committed to deliver; while the time dependency is intrinsic in the use of a probabilistic function.

In this work, we show that the negotiation strategy for one-issue negotiation with multiple providers of different services, can be easily extended to multi-issues cases, and that such strategy has relevant properties in the negotiation

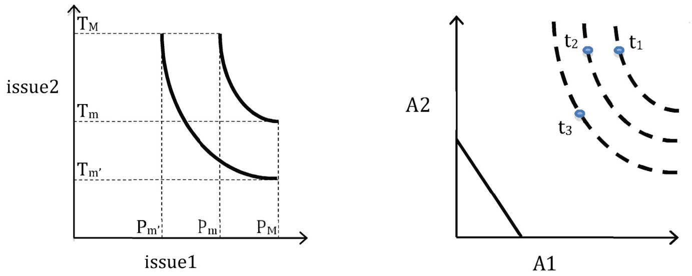

(a) Two-issues indifference curves for different utility values.

(b) One-issue negotiation spaces for a two service composition.

Fig. 2. Negotiation spaces

process. In the multi-issue case, instead of the mono-dimensional Gaussian function used for one-issue negotiation, a multi-dimensional one is adopted. It models both provider's utility, and its attitude to concede by providing offers varying on the Gaussian function. The provider's utility is modeled as follows:

$$U_x(q_1,\ldots,q_r) = \prod_{i=1}^{r} \left[\exp\left(-\frac{1}{2}\left(\frac{q_i-\mu_i}{\sigma_i}\right)^2\right)\right]$$

where, for each issue q_i, σ_i models the concession attitude when all the other $r-1$ issues are kept fixed, and μ_is are the values for the issues corresponding to the best utility ($U(\mu_1,\ldots,\mu_r)=1$) for the providers. Values of the utility function are still in the domain [0,1], that is one dimensional (see Figure 1b).

This general representation allows on one hand to model an utility function with non linear dependencies among different issues, but at the same time allows to model a "probabilistic" concession strategy that takes into account different concessions attitudes (σ_i) on different issues. Starting from this multidimensional Gaussian function, an utility level corresponds to an indifference curve, that includes a combination of values, one for each issue, having the same utility value for the provider. In Figure 2a we show different negotiation spaces (section of ellipses) generated for different values of U_x. Such spaces correspond to a negotiation domain that is rational and strictly convex (properties widely applied in economics) [6]. Differently from the one-issue case, here the agent can do tradeoffs between values with the same utility. Tradeoffs in a continuous space may become intractable. Efficient heuristics to find pareto or quasi-pareto optimal solutions exist, but such approaches rely on the availability of counteroffers from the other agents [6].

Moreover, the agent can concede in utility selecting a new negotiation space. When conceding in utility the agent fixes $r-1$ issues and makes a concession on a single issue. Convexity of the utility function ensures that the agent preference on each issue is monotone when fixing the others. So, if the value increases (or decreases), the utility always decreases (or increases). In the multi-issue case, as in the one-issue case, the provider agent generates a new value of utility corresponding to a new offer following its normal distribution. For example, in Figure 2a the rightmost curve represents points for a starting constant value of utility: $U_x(p,t) = const$, with $p \in [P_m, P_M]$ and $t \in [T_m, T_M]$. Fixing $t = T_M$ (for the $issue_2$ in Figure 2a) the agent selects a new value for p (for the $issue_1$ in Figure 2a) from a single issue normal distribution $p = P_{m'}$. To this new set corresponds a new marginal utility value $(U_x(P_{m'}, T_M))$, where $U_x(P_{m'}, T_M) < U_x(P_m, T_M)$.

4 Single Issue Negotiation for Service Composition

In this section we show that the utility functions and strategies used to model a negotiation on r issues for each service of an SBA, are the same as for the case of a single issue negotiation for r services when the value of the single issue of the required SBA is given by the aggregation of the r component values provided by each service of the r services in the SBA. The aggregation function depends on the considered issue. When the issue is additive, as considered in this paper, and each component value is modeled as a normal distribution, such aggregation function is a convolution of their probability distributions.

As an example, here we consider the case of an SBA composed of 2 services $S1$ and $S2$, where the issue under negotiation is the price given by the sum of p_1 and p_2, that are respectively the QoS price values for $S1$ and $S2$, and the end-to-end QoS user's requirement is the global price ($globalPrice$) for the complete application ($p_1 + p_2$). Since the values of the two variables p_1 and p_2 vary according to normal distributions $f_{S1}(p_1)$, $f_{S2}(p_2)$, then the distribution of the variable $z(p_1, p_2) = p_1 + p_2$ is still a normal distribution obtained as the convolution of $f_{S1}(p_1)$ and $f_{S2}(p_2)$:

$$(f_{S1} * f_{S2})(z) = \int f_{S2}(z - p_1) f_{S1}(p_1) dp_1 = \exp\left[-\frac{(z-(\mu_{S1}+\mu_{S2}))^2}{2\left(\sigma_{S1}^2+\sigma_{S2}^2\right)}\right]$$

Hence, this convolution can be used to evaluate the distance of the end-to-end QoS requirements from the aggregated QoS values received at a given round as in the case of a multi-issue negotiation with one single service.

In Figure 2b, the corresponding negotiation space is depicted.The dotted curves represent the projection, in a bi-dimensional space, of the section of the Gaussian resulting from the convolution of the two Gaussian functions $f_{S1}(p_1)$ and $f_{S2}(p_2)$ obtained by intersecting the Gaussian with a plane representing a constant utility; the points on the dotted curves represent the aggregated price obtained by a combination of offers received at round t; the line represents the end-to-end QoS constraint, i.e. $p_1 + p_2 = globalPrice$.

The compositor agent adopts the same strategy for both one-issue and multi-issue negotiation evaluating an Euclidean distance between the aggregated value of the received offers and the QoS end-to-end constraint. For one-issue negotiation in a composition of r services, such distance is calculated in a r-dimensional space. The same is for a multi-issue negotiation on r issues with one service. In both cases, the compositor accepts the offers when the distance is equal to 0.

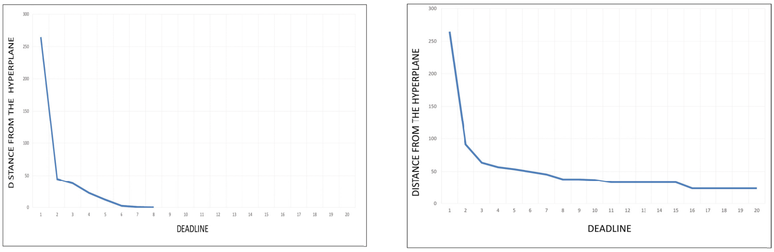

(a) Success case. **(b)** Failure case.

Fig. 3. Global distances trends in a case of successful and failed negotiations

4.1 Numerical Evaluation

We report a simple numerical simulation on the trends of the negotiation for the scenario reported in the previous section, considering 5 services in the SBA, and 4 provider agents for each service. In particular, we evaluate the trends of the offers received for each service by calculating, for all the offers, the utility of the offer provided by the jth provider for the ith service, with respect to the offers received by different providers for the same service (local evaluation), normalized with respect to the range of minimum and maximum values of the offers for all services. Such utility is computed, at each negotiation round, using the approach formulated in [11]:

$$U_{local}(price_{i,j}(t_k)) = \frac{max_i(price_{i,j}(t_k)) - price_{i,j}(t_k)}{\sum_{i=1}^{m} max_i(price_{i,j}(t_k)) - \sum_{i=1}^{m} min_i(price_{i,j}(t_k))}$$

where i identifies one of the m services (with $m = 5$) and the j identifies one of the n providers (with $n = 4$). Such utility is normalized to be in [0,1]. For each ith service the compositor agent selects the most promising offer ($price_{i,s}$), i.e. the one with maximum value of U_{local}. For each promising offer at round t_k, the global requirements satisfaction is evaluated by computing the Euclidean global distance of the composition of the values of the selected offers from the hyper-plane representing the end-to-end user's requirement (see Figures 3a and 3b). In the case of successful negotiations such distance (Figure 3a) converges to zero.

5 Conclusions

In the present work the use of software agent negotiation is used as a means to select service implementations required by an SBA by taking into account the Quality of Service that providers offer for their services, and the end-to-end QoS requirements expressed by a user requesting the application. We showed that negotiation in SBAs is inherently multi-dimensional even in the case of a single issue negotiation. Such multi-dimensions are given by the requirements of composing different services with provided QoS values that, once aggregated, meet an end-to-end constraint.

In this paper we showed that the negotiation implementation of a strategy on a single issue split among different provider agents available for the different services composing a requested SBA, can be easily extended for a multi-issue negotiation, so allowing to use the same negotiation mechanism, in terms of protocol and strategies, for both one- and multi-issue negotiation. Non linear utility functions, as well as concession strategies, can be modeled through the use of normal probability distributions, whose properties in scaling up in dimensions, allow to easily scale from one- to multi-issue negotiation by simply scaling the normal distribution dimensions.

References

1. Papazoglou, M., Traverso, P., Dustdar, S., Leymann, F.: Service-oriented computing: State of the art and research challenges. IEEE Computer 40(11), 38–45 (2007)
2. Yan, J., Kowalczyk, R., Lin, J., Chhetri, M.B., Goh, S.K., Zhang, J.: Autonomous service level agreement negotiation for service composition provision. Future Generation Computer Systems 23(6), 748–759 (2007)
3. Di Napoli, C., Pisa, P., Rossi, S.: Towards a dynamic negotiation mechanism for qoS-aware service markets. In: Pérez, J.B., Hermoso, R., Moreno, M.N., Rodríguez, J.M.C., Hirsch, B., Mathieu, P., Campbell, A., Suarez-Figueroa, M.C., Ortega, A., Adam, E., Navarro, E. (eds.) Trends in Pract. Appl. of Agents & Multiagent Syst. AISC, vol. 221, pp. 9–16. Springer, Heidelberg (2013)
4. Lau, R.Y.K.: Towards a web services and intelligent agents-based negotiation system for b2b ecommerce. Electronic Commerce Research and Applications 6(3), 260–273 (2007)
5. Lomuscio, A., Wooldridge, M., Jennings, N.: A classification scheme for negotiation in electronic commerce. Group Decision and Negotiation 12(1), 31–56 (2003)
6. Lai, G., Sycara, K.: A generic framework for automated multi-attribute negotiation. Group Decision and Negotiation 18(2), 169–187 (2009)
7. Gode, D.K., Sunder, S.: Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. Journal of Political Economy 101(1), 119–137 (1993)
8. Baarslag, T., Hindriks, K.V.: Accepting optimally in automated negotiation with incomplete information. In: Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems, AAMAS 2013, pp. 715–722 (2013)
9. Hanson, J.E., Tesauro, G., Kephart, J.O., Snible, E.C.: Multi-agent implementation of asymmetric protocol for bilateral negotiations (extended abstract). In: Proceedings of the 4th ACM Conference on Electronic Commerce, EC 2003, pp. 224–225. ACM, New York (2003)
10. Faratin, P., Sierra, C., Jennings, N.R.: Negotiation Decision Functions for Autonomous Agents. Robotics and Autonomous Systems 24, 3–4 (1998)
11. Alrifai, M., Risse, T.: Combining global optimization with local selection for efficient qos-aware service composition, pp. 881–890 (2009)