Challenges of Quality-Driven Resource Discovery

Bernd Amann[1], Laure Berti-Equille[2], Zoé Lacroix[3], and María-Esther Vidal[4]

[1] Université Pierre et Marie Curie, Paris, France
[2] University of Rennes 1, Rennes, France
[3] Arizona State University, Tempe, Arizona, USA
[4] Universidad Simón Bolívar, Caracas, Venezuela

Abstract. We report on the concluding panel of the Third International Workshop on Resource Discovery. The panel followed two invited presentations that addressed the problem of quality in the context of resource discovery. They are *Assuring Quality of Service and Quality of Data: New Challenges for Service and Resource Discovery* by Laure Berti-Equille and *Optimization Techniques for QoS-Aware Workflow Realization in Web Services Context* by Joyce El Haddad. The questions discussed by the panelists covered modeling issues, formats, languages, semantics, applications, and benchmarks.

1 Introduction

The growth of Internet technologies has unleashed a wave of innovations that have a significant impact on the way people and organizations interact with each other as well as in the ability for anyone to publish, share, and discover resources. In particular, the recent adoption of Web services and resource discovery technologies demonstrates the effective automation of business-to-business and interpersonal collaborations with new models for automated interactions among distributed and heterogeneous applications [1]. A resource may be a data repository, a database management system, a linked data endpoint, a link between resources, a semantic wiki, or a Web service. Resources are characterized by core information including a name, a description of its functionality, its URLs. Many Web services and resources provide overlapping, similar or conflicting functionalities or contents, albeit with different levels of Quality of Service (QoS) and degrees of Quality of Data (QoD). Resource discovery is the process of identifying, locating and selecting existing resources that satisfy specific functional and non-functional requirements. Current research includes crawling, indexing, ranking, clustering, and rewriting techniques, for collecting and consuming the resources for a specific request. State-of-the-art approaches support the identification of services to participate in a given service composition or resources to satisfy the user's requirements both in terms of content and service quality.

This paper reports on the concluding panel of the Third International Workshop on Resource Discovery that addressed the challenges that emerged through the shift from the traditional Web to Web 2.0 in the scope of QoS and QoD management. Panelists discussed past work in the areas of quality-aware service modeling and composition, resource discovery, and data quality management, introduced

Z. Lacroix and M.E. Vidal (Eds.): RED 2010, LNCS 6799, pp. 181–189, 2012.

recent methods and techniques, and highlighted new directions for assessing both Quality of Service and Quality of Content in data-centric service discovery.

2 Quality for Resource Discovery

First the panelists shared their views on resources, what they are, what quality measures can be assigned to them, and the problem of resource discovery. As discussed in our previous workshops, a resource refers to different things depending on the user or expert; it includes an agent, an application, a service, a data source, etc [16]. The problem of resource discovery itself can be approached in various ways each with its own quality measures.

3 Traditional QoS Approaches

Although resource discovery raises specific issues, many existing approaches designed for Quality of Service may be exploited and revised for that purpose. Analysis of models, techniques, and tools for QoS negotiation and Quality Assurance requires a preliminary understanding on how quality of a service or a service-based application can be described and modeled. Various models for describing QoS have been proposed by practitioners; standardization consortia and frameworks for creating agreement on QoS aspects between companies that endorse the service oriented paradigm can be adapted. In the presence of multiple Web services with overlapping or identical functionality, users will discriminate these alternatives based on their QoS or preferences. QoS is a broad concept that encompasses a number of non-functional properties such as price, availability, reliability, and reputation [13, 3, 9]. In addition, services are distributed across the Internet and some of their QoS properties (e.g., availability and successful execution rate) are affected by the communication link and should be measured from the perspective of the requestor rather than the provider [14, 15]. Preferences represent opinions or feedback of the users about the behavior of the available services; simple logic-based languages can be used to express preferences such as hard constraints, QoS preferences, and conditional preferences [12].

4 Workflow-Driven Quality

Recent approaches look at a resource in the context it is used, often composed with other resources in a workflow. The panelists discussed issues specific to service composition, workflow optimization, continuous information discovery (RSS), resource discovery for data integration, etc. The problem of selecting the services that best fit the QoS-based criteria, is known as the QoS-aware service selection or composition problem. This is a combinatorial optimization problem that has been shown to be NP-hard. Thus, the main challenge is to define heuristics able to identity a relatively good solution in a reasonably short period of time. Another important issue to solve, is to provide ranking techniques able to highly rank the resources that in addition to satisfy the functional user request, best meet the user QoS restrictions.

First, to develop effective and efficient heuristics, resources need to be semantically described. Traditionally, service functionality has been defined in terms of ontology annotated input and output arguments and pre- and post-conditions. However, because of the lack of expressiveness, this approach does not allow the possibility to represent properties of the objects produced by the resource. For example, a service that given a US city retrieves the cities reached by two-leg one-stop flights that pass through Los Angeles, would not be precisely described using this commonly used formalism. Recently, Ambite et al. [2] have proposed to describe resource functionality as local views of the concepts that comprise a domain ontology. This description approach is taken from the Local as View paradigm traditionally successfully used in integration systems, and they have developed the tool DEIMOS that is able to semi-automatically generate these mappings from existing Web data sources. Furthermore, the transactional information and user preferences must be specified either in user requests and in resource descriptions. For instance, consider an abstract workflow which has to be atomic, so all the services selected to implement the workflow, have to be compensable; thus, the resource description formalism must ensure the description of this property. Additionally, suppose a user searching for services retrieving round trip flights between two American cities, prefers services that retrieve flights operated by the same airline. The framework to describe the user request must be able to assign preferences to services that combine flights, and according to these scores, identify the services that best meet the user request [6, 12].

Second, the design of heuristics that usually exploit background information explicitly expressed in the resource descriptions can be used to solve the QoS-aware Service Selection problem, guide the search, and prune the space of possibilities. Thus, existing approaches focus on two main issues, scale up to a large number of resources, while the selected resources best meet the user request. For example, Rahmani et. al. [21] present a distant metric-based heuristic that guides a backward algorithm; this metric induces an order of the services in a way that services that do not satisfy the user request, are unlikely to be visited. Ko et al. [13] propose a constraint-based solution that encodes the non-functional permissible values as a set of constraints whose violation needs to be minimized; a hybrid heuristic that combines the Tabu search and Simulated Annealing meta-heuristic is proposed to efficiently traverse the space of possibly service compositions. Others have considered planning-based approaches to efficiently solve the QoS-aware the service selection problem [20, 19, 22]. Although existing approaches achieve the scalability issue, because of the lack of expressiveness of the description formalisms used to describe the resources, the selected resources may be far from the resources that best meet the user request. Recently Izquierdo et al. [12] devise the problem of selecting the best services as the problem of identifying the best query rewritings of a conjunctive query in terms of a set of views. Queries represent user requests, views semantically describe services, and user preferences are specified with a simple logic-based language; preferences are used to rank the solutions of the problem that best meet the user request.

Finally, once the resources that meet a user request are identified, ranking metrics are required to discriminate among the selected resources, those that best meet the QoS and user preference requirements where all the QoS parameters and preferences are equally important. This condition implies that some services may be incomparable because there is no other one that has the best values in all the QoS conditions or user preferences. Traditionally, Skyline [7] and Top-k Skyline [10] approaches have been used to efficiently identify the set of incomparable objects or the top-k objects among a set of incomparable objects. Although state-of-the-art techniques effectively achieve the problem of identifying the best elements, techniques that exploit semantics about QoS parameters are required to effectively prune the space of possibilities and efficiently identify the best services [11].

5 Quality Dimensions

Existing models and formats to represent resources often lack the mechanisms needed to capture and publish the metadata required to express quality. The panelists discussed the metadata needed to support quality measures and identified the limitations of existing formats. Quality dimensions provide a shared knowledge about how a quality can be defined in terms of names, metrics, relationships and dependencies among quality dimensions. The set of quality dimensions that are considered relevant for a resource will be used by all the actors involved to express the requirements and capabilities, to compare them, and to select the best resource. In the literature, quality dimensions are classified in different ways. Focusing on the product quality specification, we need to consider two main issues: the technical quality dimensions and the domain dependent quality dimensions [22]. Technical quality dimensions include all the quality dimensions that characterize the service provisioning and that are relevant regardless of the kind of service. In the Web service community, a set of technical quality dimensions is grouped in four main categories (runtime, transaction, configuration management, and security), that are considered relevant for describing a Web service: i.e, scalability, capacity, performance, reliability, availability, flexibility, exception handling, integrity, regulatory, supported standard, stability, cost, completeness.

6 Data, Metadata, and Quality

Many quality factors have been proposed for modeling data quality. Typically, data quality can be defined and measured along four dimensions: relevance, granularity, and level of detail for data *views*, accuracy, consistency, currency and completeness for data *values*, format and ease of interpretation for data *presentation*, and privacy, security, trust, and ownership. Quality factors can be classified according to semantic criteria, process-oriented criteria or goal-oriented criteria. Other works propose formal frameworks for describing quality factors and deal with the management and storage of the appropriate metadata. However, there is no consensus in the definition of quality factors. Each application domain has its specific vision of data quality as well as a battery of (generally ad hoc) solutions to solve quality problems [5]. Furthermore, although quality factors are frequently treated as being independent, there exist many relationships among them. The great number of quality factors and their inter-relationships cause quality evaluation to be a complex problem of many variables. Improving the quality of a system corresponds to optimizing a problem of N variables, which is of high complexity if done in a general context. As a consequence, it is difficult to consider many quality factors at a time. Regarding the assessment of the quality of a data resource, many works have been proposed to classify existing assessment techniques, types of metrics, units and aggregation functions, stressing the need of an algebra for combining resource quality metadata. Other works have analyzed the combination of resource quality values for some specific quality factors of accuracy and completeness, but the dynamic resource selection based on quality of content is an open problem that requires continuous resource profiling. Despite the existence of specific solutions for database profiling and cleaning [5], the land of QoD improvement and QoD-aware design of resource/service is almost unexplored.

7 Discovery Language and Customization

Users expect to be able to express queries that support resource discovery and specify the measures that best capture the quality measure they expect to optimize. Resource selection and quality have a significant impact on the provenance of data. Providing efficient access to resources received a sustained interest in the recent years but relatively few approaches have been proposed to ensure both quality of service and quality of content in distributed and Web syndicated applications. This raises several particularly challenging issues due the nature and characteristics of the resources, their structure and data quality, and the domain ontology (e.g., resource volatility, data volume, large heterogeneity spectrum, complex data type: multimedia, XML, relational records, RSS, etc.). In this context, many challenges are raised in the design of service oriented architectures integrating QoS-driven service composition and QoD-driven resource discovery. Building such architectures requires addressing several interesting research issues.

Description formalisms. One of the first challenges is to define new paradigms to semantically describe the functional and non-functional properties of the resources. And capture the user requirements and preferences in terms of QoS and QoD.

QoD- and QoS-aware query languages. The idea is to devise declarative query languages that operate on data, resources and services with quality of data, user preferences, and quality of service (QoS) constraints.

Computation models. The resolution of any quality-aware query may involve an iterative process between different distributed systems within the infrastructure. We need to devise a computation model for the interaction of the different (sub-) systems (e.g., Web portals/Web services/resource and providers, etc.).

Optimization models. Performance has a prime importance in successfully deploying a QoS- and QoD-aware service oriented architecture. It mainly relates to query optimization. One challenge is to define appropriate metrics to characterize and measure QoD and QoS dimensions depending on the application domain, the systems capabilities, and the limited resources/performances. The different query planning strategies focus generally on finding feasible and optimal sub-goal orderings based on available bindings and supported conditions. Proposed techniques assume a full knowledge of the query capabilities of every resource that can be potentially discovered. They rely heavily on the way that information sources are described and the objective function of the optimizer (e.g., number of sources, response time, etc.). Using the same source description and data quality description models may not always be possible across a large spectrum of information sources.

Quality-aware adaptive services. Another interesting trend is the use of adaptive or dynamic approaches in dealing with quality-aware services. This is motivated by the intrinsic dynamics of the distributed and autonomous resources where unpredictable events may occur during the execution of a query affecting the QoS and/or QoD. The types of actions that are proposed in these approaches fall into one of the following cases: i) change the query execution plans in order to privilege data quality of query results, ii) change the scheduling of operations in the same query execution plan or in different concurrent query plans, iii) introduce new operators to cater for the unpredictable events (e.g., in case of QoD or QoS degradation. Adaptive techniques have demonstrated their applicability to various Web service applications [4]. There is a need to show how they react under heavy quality of data and quality of service fluctuations. To the best of our knowledge, the issues of data quality-awareness adaptation in on-line service composition or query processing have not been much investigated and constitute a very challenging perspective of research. Taking into consideration the user feedback is also a critical issue that needs to be addressed by e-Service community, in particular to fix as quickly as possible certain errors in the resource description and resource mismatch with respect to the user query.

8 Benchmarks

The design of benchmarks is critical to evaluate and compare different solutions. These benchmarks raise similar questions from a different point of view (performance, data sets, queries) independently of a particular solution.

9 Applications

The problem of resource discovery covers a large spectrum of applications. Some domains such as the biomedical domain [17] seem to express a particular interest in the development of solutions. Several approaches exploiting metadata have been developed in the context of biological resource discovery most of which can be applied to any domain. The challenge of resource discovery is to extend the optimization of quality measures of local criteria that capture the identification of the resource that best implement a task to the optimization of global criteria where several resources are identified so that connected into a workflow they fulfill a complex network of tasks. Discovery queries expressed as sequential connected tasks were asked against data sources mapped to the concepts that their structural metadata could be mapped to in [8]. More complex discovery queries were expressed against a semantic graph of resources were each data sources is mapped to the concepts its structurally correspond to and each application or service is mapped to the conceptual relationship it implements in [18].

10 Conclusion and Perspectives

Quality related aspects relevant for service- and resource-based applications cover a broad field of research, including work on quality modeling and specification, QoS and Services Level Agreement (SLA) negotiation as well as constructive and analytical quality assurance (like testing, monitoring and static analysis). Designing data-centric services raises several interesting challenges related to i) the accuracy of data from data resources that are of different quality either intrinsically or from one resource to another, and over time, ii) the completeness and coverage of data resources to provide effective and fine-grained content-based resource discovery with up-to-date metadata and efficient maintenance, iii) the usability of the services in terms of personalized resource discovery with respect to data and content quality requirements and user feed-back gathering to improve the services and quickly fix the errors. Ensuring both quality of service and quality of content is indeed a very rich and challenging research direction that needs further investigations.

References

1. Alonso, G., Casati, F., Kuno, H., Machiraju, V.: Web Services: Concepts, Architectures, and Applications. Springer, Heidelberg (2004)
2. Ambite, J., Darbha, S., Goel, A., Knoblock, C., Lerman, K., Parundekar, R., Russ, T.: Automatically Constructing Semantic Web Services from Online Sources. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) ISWC 2009. LNCS, vol. 5823, pp. 17–32. Springer, Heidelberg (2009)
3. Artz, D., Gil, Y.: A survey of trust in computer science and the semantic Web. Web Semant. 5(2), 58–71 (2007)

4. Bai, X., Chen, Y., Shao, Z.: Adaptive web services testing. In: Proceedings of the 31st Annual International Computer Software and Applications Conference, COMPSAC 2007, Washington, DC, USA, vol. 02, pp. 233–236. IEEE Computer Society, Los Alamitos (2007)
5. Berti-Equille, L.: Quality Awareness for Data Managing and Mining. Habilitation à diriger des recherches, Université de Rennes 1 (June 2007)
6. Bistarelli, S., Montanari, U., Rossi, F.: Semiring-based constraint satisfaction and optimization. Journal of the ACM 44(2), 201–236 (1997)
7. Börzsönyi, S., Kossmann, D., Stocker, K.: The skyline operator. In: Proceedings of the 17th International Conference on Data Engineering, Washington, DC, USA, pp. 421–430. IEEE Computer Society, Los Alamitos (2001)
8. Cohen-Boulakia, S., Davidson, S., Froidevaux, C., Lacroix, Z., Vidal, M.-E.: Path-Based Systems to Guide Scientists in the Maze of Biological Resources. Journal of Bioinformatics and Computational Biology 4(5), 1069–1095 (2006)
9. Deora, V., Shao, J., Gray, W.A., Fiddian, N.J.: A quality of service management framework based on user expectations. In: Orlowska, M.E., Weerawarana, S., Papazoglou, M.P., Yang, J. (eds.) ICSOC 2003. LNCS, vol. 2910, pp. 104–114. Springer, Heidelberg (2003)
10. Goncalves, M., Vidal, M.-E.: Reaching the top of the skyline: An efficient indexed algorithm for top-k skyline queries. In: Bhowmick, S.S., Küng, J., Wagner, R. (eds.) DEXA 2009. LNCS, vol. 5690, pp. 471–485. Springer, Heidelberg (2009)
11. Goncalves, M., Vidal, M.-E., Regalado, A., Ayadi, N.Y.: Efficiently selecting the best web services. In: Lacroix, Z. (ed.) RED 2009. LNCS, vol. 6162, pp. 120–139. Springer, Heidelberg (2010)
12. Izquierdo, D., Vidal, M.-E., Bonet, B.: An expressive and efficient solution to the service selection problem. In: Patel-Schneider, P.F., Pan, Y., Hitzler, P., Mika, P., Zhang, L., Pan, J.Z., Horrocks, I., Glimm, B. (eds.) ISWC 2010, Part I. LNCS, vol. 6496, pp. 386–401. Springer, Heidelberg (2010)
13. Ko, J.M., Kim, C.O., Kwon, I.-H.: Quality-of-Service Oriented Web Service Composition Algorithm and Planning Architecture. Journal of Systems and Software 81(11), 2079–2090 (2008)
14. Kritikos, K., Plexousakis, D.: Semantic QoS metric matching. In: European Conference on Web Services, pp. 265–274. IEEE (2006)
15. Kritikos, K., Plexousakis, D.: Requirements for QoS-based Web service description and discovery. In: 31st Annual International Computer Software and Applications Conference, Washington, DC, USA, vol. 2, pp. 467–472. IEEE Computer Society, Los Alamitos (2007)
16. Lacroix, Z.: Editorial. International Journal on Metadata, Semantics, and Ontologies 5(3), 167–169 (2010); Special Issue on Resource Discovery
17. Lacroix, Z., Kothari, C.R., Mork, P., Rifaieh, R., Wilkinson, M., Cohen-Boulakia, S., Freire, J.: Biological Resource Discovery. In: Liu, L., Özsu, M.T. (eds.) Encyclopedia of Database Systems, pp. 220–223. Springer, Heidelberg (2009)
18. Lacroix, Z., Vidal, M.-E., Legendre, C.: Customized and Optimized Service Selection with ProtocolDB. In: Hameurlain, A., Tjoa, A.M. (eds.) Globe 2009. LNCS, vol. 5697, pp. 112–123. Springer, Heidelberg (2009)
19. Lécué, F.: Optimizing qoS-aware semantic web service composition. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) ISWC 2009. LNCS, vol. 5823, pp. 375–391. Springer, Heidelberg (2009)

20. Lohmann, M., Mariani, L., Heckel, R.: A model-driven approach to discovery, testing and monitoring of web services. In: Baresi, L., Nitto, E.D. (eds.) Test and Analysis of Web Services, pp. 173–204. Springer, Heidelberg (2007)
21. Rahmani, H., GhasemSani, G., Abolhassani, H.: Automatic web service composition considering user non-functional preferences. In: Proceedings of the 2008 4th International Conference on Next Generation Web Services Practices, Washington, DC, USA, pp. 33–38. IEEE Computer Society, Los Alamitos (2008)
22. Sohrabi, S., McIlraith, S.A.: Optimizing web service composition while enforcing regulations. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) ISWC 2009. LNCS, vol. 5823, pp. 601–617. Springer, Heidelberg (2009)