

All the Services Large and Micro: Revisiting Industrial Practice in Services Computing

Gerald Schermann^(✉), Jürgen Cito, and Philipp Leitner

University of Zurich, Zurich, Switzerland
{schermann,cito,leitner}@ifi.uzh.ch

Abstract. Services computing is both an academic field of study looking back at close to 15 years of fundamental research and a vibrant area of industrial software engineering. Industrial practice in this area is notorious for its ever-changing nature, with the state of the art changing almost on a yearly basis based on the ebb and flow of various hypes and trends (e.g., microservices). In this paper, we provide a look “across the wall” into industrial services computing. We conducted an empirical study based on the service ecosystem of 42 companies, and report, among other aspects, how service-to-service communication is implemented, how service discovery works in practice, what Quality-of-Service metrics practitioners are most interested in, and how services are deployed and hosted. We argue that not all assumptions that are typical in academic papers in the field are justified based on industrial practice, and conclude the paper with recommendations for future research that is more aligned with the services industry.

1 Introduction

Since the inception of standardized XML-based service definition, description and discovery languages and approaches [6] (i.e., the WS-* stack) around the year 2002, academic research has zealously embraced the ideas of service-oriented computing and Service-Oriented Architecture (SOA) to build and organize large-scale distributed applications. However, services computing is not a static field. Over the years, various new industry-driven technological trends (e.g., REST [16], enterprise service buses or ESBs [17], cloud computing [4], or most recently, microservices [12, 14]) have appeared, and became integrated into how academic researchers think about services. The disadvantage of this integrative approach is that, by now, the term “service-based application” (SBA) can mean any number of things, ranging from dynamic SOAP- and WSDL-based applications built using the traditional triangle of publish-find-bind [11], WS-BPEL-based compositions of public Web services, large-scale, heterogenous, enterprise services connected via an ESB, all the way to microservices-based cloud applications (or any combination thereof).

Orthogonally, but relatedly, a sometimes voiced criticism of current academic services research is that it is too removed from industrial practice. To give just one example, Prof. Anthony Finkelstein (University College London) has in a

blog entry remarked that research in service discovery deals with a problem that very few practitioners actually have¹. However, non-anecdotal data about the practical impact and relevance of services research is hard to come by.

In this paper, we aim to provide the academic community with a glance “across the wall” into industrial services computing. We conducted a small-scale survey of the state of practice in services computing, with the primary goal of understanding what practitioners mean when they talk about services, how they technically implement and host services, and what issues they struggle with. Our study has been set up with a specific focus on the recent trend of microservices. We hope to contribute to services research by painting a clearer picture of how service-based applications actually look like in practice, which issues require better approaches, and which traditional research areas in the field are simply not all that relevant in practice. Note that we focus specifically on *technical* issues of services computing in this paper. To keep the size of the research manageable, we excluded economic and cultural topics in this study. Further, we limit our research to technical services, and exclude questions on human-provided services, workflows and business processes.

It should be noted that the goal of this workshop paper is to provide a starting point for fruitful discussion, not to criticize individual researchers or the community at large. The third author of this paper has himself published on all individual research ideas that are going to be put into question in the following. Further, given that our sample size is not overly large, we do not claim to have all the answers. There is certainly potential for more large-scale and more rigorous follow-up research. We primarily follow an empirical approach. Using a Web-based survey, we questioned 42 companies with one or more service-based products. Our results show that most service ecosystems are of quite manageable size. While public Web services are in use, most services are actually internally developed and operated. REST and HTTP are almost ubiquitous, while SOAP is falling out of favor fast. Most companies do not make use of a centralized composition engine or service bus. Instead, most service interactions follow what researchers would call a service choreography style. Cloud computing and QoS monitoring is indeed of large industrial relevance today.

2 Study Setup and Method

We conducted our research as a quantitative, Web-based survey. We targeted developers and companies that self-identify as building a service-based product or making use of a service-oriented architecture. To acquire participants, we advertised our study on multiple programming-related Web sites, as well as through personal contacts and social media. We were able to acquire 42 participants, which were close to equally distributed over companies of all sizes, ranging from 1–20 employees up to global enterprises with more than 1000 employees. About 50 % of our study participants are working as software developers. The bulk of the remaining participants were either team leads, DevOps engineers,

¹ <http://blog.prof.so/2012/06/bottom-10-software-engineering.html>.

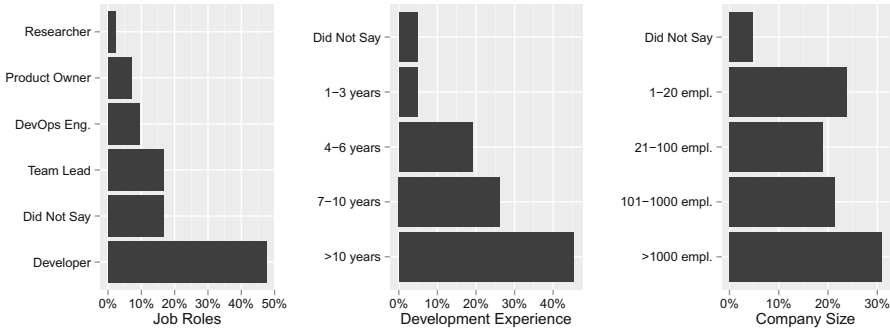


Fig. 1. Demography of study participants. Most participants are experienced and work as software developers in large enterprises.

or product owners. Most of our participants are experienced software developers, with close to three quarters reporting seven years of experience or more. This data is summarized in Fig. 1.

Our study consisted of 25 questions, designed as either multiple choice, single choice, or open-ended free text questions. When designing our study we strived for a good compromise between keeping the study short for the participants and collecting enough material that is related to a wide range of currently “hot” topics in academic services research. Finally, and after discussions with our industrial partners and internal testing of the survey, we decided to ask questions about technical fundamentals, middleware, service discovery, QoS monitoring, and cloud deployment, leading to a Web-based survey that took our participants less than 10 min in the median to complete. A complete list of questions, as well as all resulting data, is available as part of the online appendix².

3 The State of Practice in Service-Based Applications

We now discuss the outcomes of our research. For reasons of brevity, we only summarize the most important outcomes.

3.1 Fundamentals

First, we discuss the technical fundamentals of services. What is the typical size and complexity of services, to what extent are external services used, and what are the common programming languages for developing services?

Services Vary in Size, but Few are Truly “Micro”. An evergreen question in services computing is the “optimal” size of individual services. Erl refers to services as “coarse-grained entities” [7], implying that each service should carry substantial business logic. The current microservices trend emphasizes tiny services, with “10 to 100 lines of code” (LOC) each³. We asked our participants

² <http://wp.ifi.uzh.ch/leitner/?p=743>.

³ <http://guidesmiths.com/blog/the-granularity-of-a-micro-service/>.

about the typical size of services within their organization in LOC. We observed that services in the range of 1'000 to 10'000 LOC are dominant, as stated by 51 % of our participants, followed by 100 to 1'000 LOC chosen by 43 %. Services following the microservices rule of thumb (less than 100 LOC) are rare (3%), as are very large services with more than 10'000 LOC.

Services are Dedicated. Another interesting question is how many separate concerns an individual service covers. Following most literature, services are supposed to be dedicated to a single task. As a metric to measure this, we asked our participants how many public operations a service typically provides. The resulting data shows that services seem to indeed typically be dedicated to relatively narrow tasks, exposing between 1 to 9 (46 %) or 10 to 20 (45 %) public operations. The remaining 19 % of our participants operate services with relatively large public interfaces (between 20 and 50 operations).

Most Service Ecosystems are Actually Not Very Large. Much academic research in services computing is motivated by a presumed large number of services to compose applications from. To this end, we have asked our participants how many services they actually have access to within their organization, including in-house and usable external services. Our respondents stated to, in the median, only have access to 30 services. However, the individual answers to this question varied enormously in a range of 7 to 20'000 services. This is because the background of the study participants also varied. Clearly, developers in globally operating enterprises typically have access to substantially more services than startup employees. However, only 25 % of all participants actually deal with service ecosystems of substantial size (more than 100 services) on a regular basis.

Most Companies use External Services. A similar common assumption in academic research is that companies often make use of external services to implement their business goals. According to our responses, almost two thirds (64 %) of our participants indeed make use of external services. However, 68 % of all services that they use are actually internally developed. That is, most companies use external services, but the majority of services in use are still developed and operated internally.

Java is Still the Most Common Way to Implement Services. To conclude technical fundamentals, we were interested in how services are actually developed. The microservices trend often emphasizes a heterogeneity of programming languages within an organization (“the right tool for the job”). This was not confirmed in our research. Indeed, we found that 45 % of respondents use at most two programming languages, followed by 25 % using three or four and 20 % using five or six programming languages. The remaining 10 % seem to have a strongly heterogeneous service architecture, as they implement their services in more than 6 different programming languages.

Further, we asked our participants what concrete programming languages they use for developing their services. As illustrated by Fig. 2, Java was selected by 67 % of our participants. Besides Java, we identified a strong focus on scripting

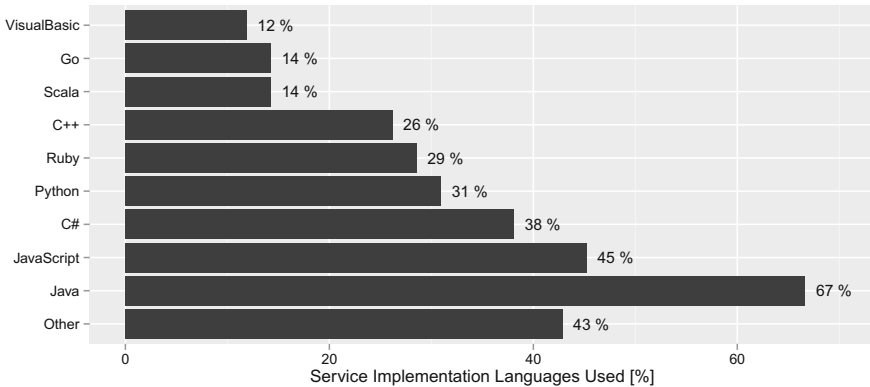


Fig. 2. Java is used to implement most services, followed by JavaScript and C#. Multiple selections were possible.

languages, such as JavaScript, Python, and Ruby, excelling prominent compiled languages such as C++, C#, and Visual Basic. Finally, it should be noted that WS-BPEL or any other service composition language has not been mentioned as a typical service implementation language by any participant.

Key Points. Services vary substantially in size, but true “micro” services are rare. Many companies use some external services, but most services in use are internal. Java and JavaScript are the most common service implementation languages today.

3.2 Communication Between Services

In this section, we focus on how services communicate with each other. Are services more commonly implemented using SOAP/WSDL, or is REST by now more relevant?

HTTP and REST are Ubiquitous. The selection of a communication protocol or technology strongly depends on whether synchronous or asynchronous service-to-service communication is preferred. Almost all participants (95%) operate synchronous services that are based on HTTP and REST. The dominance of HTTP and REST fits well with the characteristic of microservices being built on top of lightweight communication mechanisms. However, services that communicate via message queues (e.g., AMQP), an asynchronous, event-based communication style, are also used by 57% of our participants. Interestingly, every participant that stated to use message queues for service-to-service communication also operates at least one REST service. Only 21% still rely on RPC-based communication technologies, such as RMI or XML-RPC.

JSON is More Common than Plain XML, SOAP is Less Common than Either.

In terms of data exchange formats, we have seen that JSON has widely superseded XML as the primary service data exchange format. As summarized in Fig. 3, 90% of our participants stated to use JSON as data exchange format, while plain XML was the choice of 57%. This can partially be

explained with the increasing importance of JavaScript-based frameworks, such as Node.js. SOAP is not overly wide-spread in our study (40%). Google's Protobuf is on the rise, but still relatively rare with 17% usage across participants.

Dedicated Service Middleware is Not Often used. Finally, our study has shown that centralized, heavy-weight middleware (e.g., ESBs or composition engines) are not overly common. Rather, 45% are not using any middleware at all (excluding simple messaging middleware), which is an interesting fact and highlights the tendency to a more decentralized, choreographed approach rather than a central orchestration point. 31% of our participants use ESB technology for communication between services. The vast majority of those participants are employed at companies with 100 employees or more. API gateways (e.g., Swagger, Tyk, or Strongloop) are only in use at 19% of our respondents.

Key Points. HTTP and REST are used by 95% of our participants' companies. JSON has replaced XML as the most common data exchange format. Half of our of participants does not use any middleware at all, there is a trend towards a more choreography-style of managing service coordination.




Fig. 3. JSON has replaced XML as primary data exchange format. SOAP is not overly common. Multiple selections were possible.

3.3 Monitoring and Quality-of-Service

Measuring and monitoring Quality-of-Service (QoS) [13] has historically been considered an important and valuable field of research, but to what extent do practitioners care about QoS?

Standard QoS Attributes are Indeed Widely Monitored and used. As indicated in Fig. 4, our study participants indeed monitor and use a broad spectrum of infrastructure (e.g., CPU utilization, network traffic) and application metrics (e.g., response times, failure rates). However, those metrics are rarely used to select services. Rather, metrics are used at runtime to monitor the health and performance of services. Our study participants use a wide range of monitoring tools, the most common of which are Logstash, Nagios, and NewRelic.

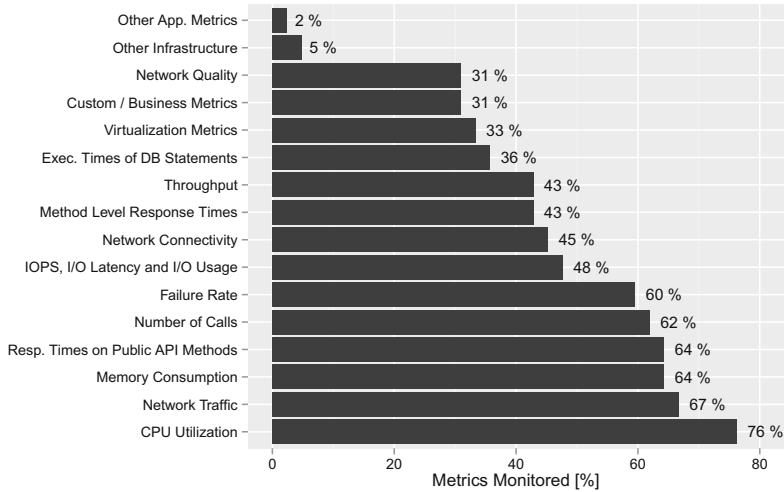


Fig. 4. Participants reported a wide range of QoS metrics being monitored, including system-level and application-level metrics. Business metrics are only used by 31 % of all participants. Multiple selections were possible.

Business Metrics are Rare. A more interesting result is that only 31 % of participants monitor any custom or business metrics⁴. Given that various microservices proponents regularly emphasize the importance of business metrics as basis for development and business decisions, this number was behind our expectations. It seems that for most companies, standard performance metrics are sufficient today.

Key Points. Companies monitor a wide range of standard metrics on application and infrastructure level. These metrics are used to observe the health state of the application rather than to select services. Business metrics are used less than expected.

3.4 Service Discovery

Assuming a service ecosystem with hundreds or thousands of services, discovering the right service is challenging. How do companies handle service discovery in practice, do they use service registries and how do they find out how to invoke services?

Registries are Not Commonly used in Practice. Even though actively researched in the previous 15 years, service registry and discovery concepts such as UDDI have never gained much attention in industrial practice. This is also reflected in our participants' responses when we asked them how they know if certain functionality is available as a service. Only 18 % stated that they have a middleware for registering and querying services. 28 % of participants are

⁴ <http://www.klipfolio.com/resources/articles/what-are-business-metrics>.

manually maintaining a list, website, or WIKI page of available services. 25 % stated that service discovery is a minor issue as they do not have that many services and just know what is available. Similarly, 18 % mentioned that there is a contact person within the organization to ask about what services are available.

Client-Side Dynamic Binding of Services is Not Typically used. We have not seen a strong indication that practitioners actually follow the “SOA triangle” of publish-find-bind in any real way. 70 % of our participants rely on a documented fixed configuration which does not change, or use server-side approaches (e.g., DNS) to manage service binding. Our participants generally do not make use of client-side dynamic binding approaches, such as QoS-aware service selection.

Key Points. 28 % of our participants manually maintain a list of available service functionality. Documented, fixed configurations are the most common way to bind clients to services.

3.5 Service Hosting and Deployment

Finally, we were interested in how services in the wild are actually hosted and deployed. The academic literature has widely embraced the notion of cloud computing as means to house and provision services, but to what extent does this reflect industrial practice?

Cloud Computing is Mainstream. Our study results, which are also depicted in Fig. 5, indicate that by now cloud computing has indeed found its way into the mainstream of industrial services computing. Two thirds of all participants use either public or private cloud systems to host their applications. Only 50 % of all participants even still have services that are hosted in-house on non-virtualized infrastructure. A third of our respondents are still using long-term external hosting providers. Despite cloud computing being sometimes branded as primarily interesting to startup companies [9], there is no significant difference between the data of small companies and large enterprises in our study. Interestingly, even though recent study results indicate that elasticity and automated scaling are primary drivers in cloud adoption [5], most services (70 %) are currently actually *not* scaled automatically. However three quarters of participants use the cloud to redundantly deploy their services, mostly to improve fault tolerance. Half of our respondents even deploy their cloud services redundantly over three or more nodes.




Fig. 5. 66 % of all participants use public or private cloud services. Only 50 % still use in-house hosting. Multiple selections were possible.

Application Packages are Still Widely used. Another interesting outcome of our research is that building and copying application archive packages (e.g., Java JAR or WAR files) is still the most common way (64 %) to provision the implementation code of cloud services, presumably as part of a Continuous Integration toolchain (e.g., Jenkins). Container technologies (e.g., Docker or LXC) are on the rise, but currently only in use at 34 % of all participants. 21 % use virtual machine formats (e.g., Amazon Machine Images or VMWare Images). Small minorities of respondents used UNIX packaging mechanisms or provisioned code directly out of the version control system onto cloud instances (e.g., they clone the code from Git as part of provisioning).

Key Points. Cloud computing is already widely adopted in practice. However, most cloud services are not scaled automatically. Provisioning is still mostly done by building and deploying application packages as part of a Continuous Integration toolchain.

4 Recommendations for Research

The main goal of our study was to survey the state of practice to guide future services computing research. Hence, we now (somewhat provocatively) discuss some implications of our results for a number of common research themes in the field.

Do Not Assume that Service Ecosystems are Huge. Many academic works on service selection are motivated by a presumed extensive number of services to choose from. This is typically not the case, except within a few international corporations. Most service ecosystems are quite easy to track even manually (e.g., via WIKIs). One aspect of this is also that public, external Web services are not quite as prevalent as some research works seem to assume.

Do Not Assume that there are Many Alternative Services to Choose From. In our study we have not seen any particular indication that practitioners indeed commonly need to choose from a list of functionally comparable services. As most services are internal, there is typically exactly one (in some cases two, including a legacy system) service that implements any particular business need. In light of this, academics should reflect whether more attention to approaches for client-side dynamic binding and dynamic service selection is warranted.

Do Not Assume that Web Services Always use SOAP. Our results have shown that SOAP is certainly not a de-facto standard in the Web services field anymore. If a research work needs to assume a specific service style, it should probably be REST and HTTP rather than SOAP. Consequently, the importance of WSDL should also be reconsidered.

Do Research on Choreography Rather than Orchestration. Our participants largely do not make use of centralized composition engines or service buses. Particularly, in smaller service ecosystems, services are composed in an ad-hoc, decentralized, choreography style. We argue that these kinds of service

compositions deserve more research attention, particularly in the light of the current microservices trend.

Do Research on QoS, But for Monitoring Rather than Service Selection. Our results show that QoS is indeed a “hot topic” in practice. Even though the state of practice in this area is quite mature, it is our impression that there are still interesting research questions to be addressed. However, academics should not assume that QoS is primarily used as a distinguishing factor between functionally comparable services.

Do Research on Cloud Computing, but Do Not Assume that Every Cloud-Deployed Service is Elastic. Cloud computing is indeed often used in practice, and we argue that the current research attention is warranted. However, there seems to be a trend among current research works to equate cloud computing with elasticity. Our results have shown that there are many, heterogeneous reasons why practitioners use the cloud. Academics should not assume that every service deployed to the cloud is necessarily elastic.

5 Related Work

Quantitative empirical research methodologies, such as the one used in our study, are not overly common in the services field. A small number of empirical studies are available, but those are typically focused on a single product (e.g., IBM Jazz [2], SAP [1]) or domain (e.g., telecommunications [8], the financial industry [10]). While many publications present (more or less sophisticated) case studies (e.g., [18–20]), we are not aware of any recent academic publication that systematically validated some of the long-standing assumptions of the research area on a larger and more heterogeneous sample of practitioners. Consequently, the research roadmaps of the field (e.g., [15]), as well as reference architectures (e.g., [3]), have historically been driven primarily by academic interests and opinions rather than quantified industrial needs. We argue that this has led to a positive feedback loop for some topics, where many published papers on the topic signified relevance to academics, leading to even more papers on the topic being published, despite little actual industry uptake. It is our hope that our research can serve as a reality-check for researchers that allows them to evaluate whether their assumptions are plausible for industrial practice. However, ultimately, more and more rigorous empirical data will be necessary to move services computing forward.

6 Conclusion

The goal of this study was to provide a peek into the current state of practice in services computing. We surveyed 42 practitioners working in companies of widely varying size. Our results indicate that most service ecosystems are small and consist mostly of internal services. The REST paradigm is very popular. Service choreography is more commonly used than central orchestrators. Cloud

computing is of large industrial relevance, but not everybody who uses the cloud does so because of elasticity.

Our goal with this paper was primarily to motivate researchers working on services computing to reflect on the practical relevance of their work, and to occasionally revisit long-standing and often-repeated assumptions. We argue that the services computing field would benefit from more empirical studies being conducted to ground the basic research. Due to the small sample size and large scope, our work can only serve as a first step into this direction.

Acknowledgment. The authors would like to thank all participants of our survey. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 610802 (CloudWave), and from the Swiss National Science Foundation (SNF) under project Whiteboard (no. 149450).

References

1. Akkiraju, R., Ivan, A.: Discovering business process similarities: an empirical study with SAP best practice business processes. In: Maglio, P.P., Weske, M., Yang, J., Fantinato, M. (eds.) ICSOC 2010. LNCS, vol. 6470, pp. 515–526. Springer, Heidelberg (2010)
2. Anderson, L., et al.: Enhancing collaboration with IBM's rational jazztm. In: Maglio, P.P., Weske, M., Yang, J., Fantinato, M. (eds.) ICSOC 2010. LNCS, vol. 6470, pp. 501–514. Springer, Heidelberg (2010)
3. Arsanjani, A., Zhang, L.-J., Ellis, M., Allam, A., Channabasavaiah, K.: S3: a service-oriented reference architecture. *IT Prof.* **9**(3), 10–17 (2007)
4. Buyya, R., Yeo, C.S., Venugopal, S., Broberg, J., Brandic, I.: Cloud computing and emerging it platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Gener. Comput. Syst.* **25**(6), 599–616 (2009)
5. Cito, J., Leitner, P., Fritz, T., Gall, H.C.: The making of cloud applications an empirical study on software development for the cloud. In: Proceedings of the 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT International Symposium on Foundations of Software Engineering (ESEC/FSE), ACM, New York (2015)
6. Curbera, F., Duftler, M., Khalaf, R., Nagy, W., Mukhi, N., Weerawarana, S.: Unraveling the web services web: An introduction to SOAP, WSDL, and UDDI. *IEEE Internet Comput.* **6**(2), 86–93 (2002)
7. Erl, T.: *Service-Oriented Architecture (SOA): Concepts, Technology, and Design*. Prentice Hall, Upper Saddle River (2005)
8. Griffin, D., Pesch, D.: A survey on web services in telecommunications. *IEEE Commun. Mag.* **45**(7), 28–35 (2007)
9. Gupta, P., Seetharaman, A., Raj, J.R.: The usage and adoption of cloud computing by small and medium businesses. *Int. J. Inf. Manag.* **33**(5), 861–874 (2013)
10. Lawler, J., Li, Z., Javed, N., Anderson, D., Hill, J., Howell-Barber, H.: A study of web services projects in the financial services industry. *IS Manage.* **22**(1), 66–76 (2005)
11. Leitner, P., Rosenberg, F., Dustdar, S.: DaioS: Efficient dynamic web service invocation. *IEEE Internet Comput.* **13**(3), 72–80 (2009)

12. Lewis, J., Fowler, M.: Microservices, March 2014. <http://martinfowler.com/articles/microservices.html>
13. Daniel, A.: Menascé.: Qos issues in web services. *IEEE Internet Comput.* **6**(6), 72–75 (2002)
14. Newman, S.: *Building Microservices*. O’Reilly, Sebastopol (2015)
15. Papazoglou, M.P., Traverso, P., Dustdar, S., Leymann, F.: Service-oriented computing: State of the art and research challenges. *Computer* **40**(11), 38–45 (2007)
16. Pautasso, C., Zimmermann, O., Leymann, F.: Restful web services vs. “big” web services: Making the right architectural decision. In: *Proceedings of the 17th International Conference on World Wide Web, WWW 2008*, pp. 805–814. ACM, New York (2008)
17. Schmidt, M.-T., Hutchison, B., Lambros, P., Phippen, R.: The enterprise service bus: Making service-oriented architecture real. *IBM Syst. J.* **44**(4), 781–797 (2005)
18. Stein, S., Kühne, S., Drawehn, J., Feja, S., Rotzoll, W.: Evaluation of OrViA framework for model-driven SOA implementations: an industrial case study. In: Dumas, M., Reichert, M., Shan, M.-C. (eds.) *BPM 2008*. LNCS, vol. 5240, pp. 310–325. Springer, Heidelberg (2008)
19. Tsai, W.T., Wei, X., Chen, Y., Xiao, B., Paul, R., Huang, H.: Developing and assuring trustworthy web services. In: *Autonomous Decentralized Systems, ISADS 2005, Proceedings*, pp. 43–50, April 2005
20. van der Aalst, W.M.P., Reijers, H.A., Weijters, A.J.M.M., van Dongen, B.F., Alves de Medeiros, A.K., Song, M., Verbeek, H.M.W.: Business process mining: An industrial application. *Inf. Syst.* **32**(5), 713–732 (2007)