

A Survey of SOA Migration in Industry

Maryam Razavian and Patricia Lago

Department of Computer Science, VU University Amsterdam, The Netherlands*
{m.razavian,p.lago}@vu.nl

Abstract. Migration of legacy software to service-based systems is an increasingly important problem area. So far, many SOA migration approaches have been proposed in both industry and academia. There is, however, considerable difference between SOA migration approaches defined in academia and those emerged in industry. This difference pinpoints a potential gap between theory and practice. To bridge this gap, we conducted an industrial interview survey in seven leading SOA solution provider companies. Results have been analyzed with respect to migration activities, the available knowledge assets and the migration process. In addition, industrial approaches have been contrasted with academic ones, hence discussing differences and promising directions for industry-relevant research. As a result we found that, in fact, all companies converge to the same, one, common SOA migration approach. This suggests that, with experience, enterprises mature toward a similar approach to SOA migration.

1 Introduction

Migration of legacy systems to service-based systems enables enterprises to achieve advantages offered by SOA, while reusing the business functions embedded in the legacy systems. Enterprises nowadays have many software systems that are needed to be modernized because they are difficult to change and they cannot cope with everlasting requirements changes. Service-enabling the legacy systems allows enterprises to modernize their pre-existing business functions as added-value services, and therefore achieve SOA promises such as agility and flexibility. Hence, identification of migration strategies for service engineering is critical for migration of legacies, and SOA adoption in industrial setting.

So far, many SOA migration approaches have been proposed in both industry and academia with the ultimate goal of adoption in practice. There is, however, considerable difference between SOA migration approaches defined in academia and those emerged in industry. For example, while scientific approaches mainly take a reverse engineering perspective, industrial practitioners developed best practices in forward engineering from requirements to SOA technologies, where legacy code is not transformed but used as a reference. This difference pinpoints

* This research has been received funding from Jacquard (contract 638.001.206 SAPI-ENSA: Service-enAbling PreexIsting ENterprISe Assets); and FP7 contract 215483 (S-Cube). We would like to thank all architects that participated in this study.

a potential gap between theory and practice. One of the key causes of such a gap is that the approaches proposed in academia do not fully fit the main goals and needs of practice. To bridge this gap, it is necessary to understand the properties of migration approaches that are both feasible and beneficial for practice.

This paper provides deeper understanding of the types of migration approaches in industrial practice. To this end, we conducted an industrial interview survey in seven leading SOA solution provider companies. To the best of our knowledge, this is the first survey of this kind. With the objective of understanding the industrial migration approaches, we designed and executed the interviews. Each interview was analyzed considering the constituent conceptual elements of a migration process as proposed in [1], including the activities carried out, the available knowledge assets, and the overall organization of migration process. Furthermore, we looked for the best practices that companies have developed out of experience for successful legacy migration.

As a result we found that, in fact, all companies converge to the same, one, common SOA migration approach. This suggests that industrial migration approaches converge to a similar set of activities, process organization, and best practices, in other words, with experience enterprises mature toward a similar approach to SOA migration. In addition, we contrasted the industrial approaches with academic ones, which we identified from a previous Systematic Literature Review (SLR) on SOA migration [2]. Here we use the results of the SLR to discuss the differences and draw promising directions for industry-relevant research.

2 Results

To gain an understanding on industrial migration approaches, we needed to typify the approaches in a unified manner. For this purpose, we used the SOA Migration Framework (SOA-MF) introduced in our earlier work [1] (see Fig. 1.I). The analysis of the approaches revealed patterns common among various companies¹. These are listed in four key findings presented in this section. Each finding is summarized in a Reflection Box, which is followed by detailed discussion of the finding. Furthermore, each finding is compared with the results of our previous study on academic SOA migration approaches (the SLR mentioned in Section 1). Major differences between industrial approaches and academic ones can *reflect* gaps between theory and practice.

2.1 Migration Activities

Reflection Box.1.

- **F1.1.** Different companies share the same set of activities for migration.
- **F1.2.** Industrial migration approaches converge to one, common, type of migration.

¹ Due to space constraints, our research methodology including the research questions, the study design, and data analysis method are made available at <http://www.few.vu.nl/~mrazavi/IndustrialSurveyAppendix.pdf>




Fig. 1. Industrial Approaches Mappings

To answer *what is done* in industrial approaches, we identified the constituent activities of various approaches and mapped them on SOA-MF. Fig. 1.III, represents the schematic forms of those mappings. Mappings revealed two main findings: a) industrial approaches share the same set of activities for migration and b) industrial approaches are convergent to a subset of those activities. The two findings are further discussed in the following.

Finding F1.1. Various companies, independent from the company type (i.e., consultancy vs. in-house) and migration application domain, share the same set of activities for migration. This is evident from Fig. 1.III, where the activities correspond to three graphically similar coverage patterns. It should be noted that the similarity among coverage patterns, thanks to expressiveness of SOA-MF, indicates the conceptual similarity of constituent activities and artifacts of the migration approaches. According to [2], SOA migration approaches with similar set of activities constitute a migration family. Similarly, the three similar approaches identified in the interviewed companies belong to the same family.

Contrast with theory. While the industrial approaches are all members of one family, the SLR revealed that the academic approaches belong to eight very different families. By covering different sets of activities each of these eight families provide a very different view on what SOA migration entails. For instance, one family reverse engineers the legacy code and transforms the extracted code segments to services, another family only covers the forward engineering subprocess. Considering the industrial approaches, all the approaches are categorized into (only) one of the eight families. Interestingly, the size of that family,

called industrial family, is the smallest as compared to the others (i.e. 3% of academic approaches). Thus, one could conclude that 97% of the academic approaches do not fit in industrial family. This may indicate that academic research might be digging into aspects (like sub-processes and techniques) that are less relevant for industry. On the contrary, by looking at the characteristics of the industrial family research could better focus on the open research questions pertaining such family and hence have a better chance to close the gap between academic research results and industry needs.

Finding F1.2. By further analyzing the activities of the industrial approaches, we found that those common among all approaches, called *core activities*, are the ones shown in Fig. 1.II with bold boxes. The *variable* activities, i.e., those not common to all approaches, pertain to the coverage of the two transformation activities shown in Fig. 1.II by dashed line boxes. Furthermore, we observed that the core activities are those performed more frequently and systematically, while the variable activities are carried out less frequently and in an ad-hoc manner. More precisely, the limitations posed by legacy systems makes the variable activities less frequent. Several of the interview participants mentioned that, transformations that require decomposing the legacy systems are rarely carried out because they are not feasible as legacy systems are mainly monolithic. Furthermore, we observed that core activities are mainly supported by the state-of-the-practice methodologies and techniques such as SOMA [3]. The variable activities, however, are mainly carried out using local best practices. Consequently, we argue that, due to higher feasibility of the core activities and support of well-established methodologies and techniques, the industrial migration approaches are characterized by core activities.

Contrast with theory. None of the migration approaches in the SLR fully covers the core activities. I.e., none of the academic approaches comprehensively supports the type of migration that is both feasible and beneficial in industrial setting. This indicates an important gap between the migration activities emerged from practice and the ones researched in academia.

2.2 Sequencing of Migration Activities

Reflection Box.2.

- **F2.** In the industrial migration approaches the To-Be situation initiates and drives the migration.

By providing the mappings on SOA-MF, previous section addressed what activities are covered in the industrial migration approaches. Here we focus on what is the sequencing of those activities. There are two main types of sequencing of activities in the migration approaches, namely arc-shaped and bowl-shaped [4]. In summary, in arc-shaped approaches migration is driven by As-Is situation, while it is the To-Be situation that drives the bowl-shaped ones. All the industrial approaches elicited by our study were bowl-shaped.

This categorization of approaches is based on the graphical representation resulted from mapping their sequencing of activities on SOA-MF (e.g. Fig. 1.IV). The sequencing of activities in an arc-shaped approach starts from the reverse engineering sub-process. In this category, the As-Is situation initiates and drives the migration. Unlike the arc-shaped category, the bowl-shaped one starts from forward engineering and the To-Be situation is the main driver of migration.

Finding F2. The bowl-shaped sequencing of activities in industrial approaches implies the following: in all of the migration approaches the To-Be situation, characterized by requirements or properties of the target service-based system, drives and shapes the migration. To shape the migration process, first the To-Be situation is defined within the forward engineering sub-process; further, the To-Be situation is compared with the As-Is and as such, the legacy elements are selected and re-shaped to services. A question that arises is why industries perform migration in a bowl-shaped manner. Some of the participants, in one way or another, stated that in order to reach the migration goals they need to have the To-Be situation as the primary shaping force behind migration. As such, we conclude that to ensure achieving the migration goals, companies shape their migration decisions primarily by the To-Be situation.

Contrast with theory. Unlike the industrial migration approaches, the academic ones are mainly arc-shaped. In the SLR only 30 % of the primary studies are categorized as bowl-shaped approaches and the rest are arc-shaped. As such, 70% of the approaches do not support To-Be driven migration, which is considered as the best practice among the practitioners. This highlights promising opportunities for research to focus on how to support To-Be driven migration.

2.3 Legacy Understanding through Personalization

Reflection Box.3.

- **F3.1.** The industrial migration approaches do not use reverse-engineering techniques to understand the legacy systems.
- **F3.2.** The required knowledge is elicited from the stakeholders who own the knowledge.

Understanding the legacy systems plays an important role in SOA migration as it enhances extracting the best candidates among existing legacies for migration to SOA. In traditional software engineering, this understanding is gathered by extracting the representation of the legacy systems using reverse engineering techniques. As shown in Fig. 1.III, we observed that in the industry-defined approaches none covers the reverse-engineering subprocess. This observation resulted in two key findings discussed in the following.

Finding F3.1. To gain the required understanding of the legacy system, the industrial approaches do not use reverse engineering techniques. This is due to the following two reasons: *a)* the knowledge about the pre-existing system mainly resides in the stakeholders' minds (e.g. maintainer, developer, and architect). As such, the stakeholders know what functionalities are supported, and where they

are located in the legacy system. As a result, reverse engineering of the pre-existing system is not favorable considering the little Return On Investment (ROI) it brings.

b) the legacy systems are usually comprised of a set of heterogeneous systems that are implemented in different programming languages ranging from COBOL to Java. As a result, for reverse engineering of the code different tools are needed and this implies a considerable amount of costs.

Contrast with theory. To understand the legacy systems, more than 60% of the approaches in the SLR use reverse engineering techniques. Those approaches extract the representations of the legacy systems using techniques such as code analysis and architectural recovery. Only one of the academic approaches (out of 39), supports the legacy understanding without reverse engineering techniques (i.e. using structured interviews)[5]. This indicates an important gap between theory and practice since reverse engineering is not favorable in practice.

Finding F3.2. We further observed that the industrial migration approaches elicit the relevant knowledge by directly asking the stakeholders, who own, developed, or maintained that system. More precisely, knowledge about the legacy system mainly remains tacit in stakeholders minds. As such, understanding is achieved by person-to-person knowledge elicitation. We argue that, this type of knowledge elicitation is in-line with *personalization* knowledge management strategy [6]. Personalization deals with exchanging tacit type of knowledge. Using personalization, the legacy understanding is gained by knowing ‘who knows what’ and consequently sharing the tacit knowledge about the legacy systems in that regard.

Contrast with theory. In the SLR, all the approaches focus on capturing the knowledge by documenting it. As such, they are in-line with codification strategy addressing explicit documentation of the knowledge. The results of this study, however, suggests the importance of personalization. As such, research is needed to improve elicitation techniques, especially targeted for SOA migration, supporting personalization strategy.

2.4 Service Extraction by Defining the Ideal Services

Reflection Box.4.

- **F4.1.** The main driver in extraction of the legacy assets for migration is the portrait of ideal service.
- **F4.2.** Approaches emerged out of more experience portray the ideal services in more detail.

Finding F4.1. The migration approaches, inherently, embrace trade-off analysis between the level of reuse of legacy elements and characteristics of the ideal services. We observed that, in this trade-off analysis, the industrial approaches assign considerably higher weight to the later rather than the former. To do so, first they determine the ideal services during the forward engineering sub-process. Later, those ideal services are re-shaped in a way that the reuse of

pre-existing assets are realized. This way, the portrait of the ideal service is the main driver of service extraction. That is, the services identified from the pre-existing capabilities would likely be substantially different in the absence of that portrait of the ideal service. This is in-line with our other finding that all the migration approaches are bowl-shaped meaning that the To-Be candidate services guide the analysis and transformation of the As-Is legacy elements.

Contrast with theory. A characteristic of the bowl-shaped approaches is having the ideal services (To-Be situation) as the main driver in service extraction. As such, this finding points out the same gap between theory and practice as discussed in finding F.2, namely inadequate support of To-Be driven migration.

Finding F4.2. We further observed that, industrial approaches vary in the level of detail in which they portray their ideal services. Some of the approaches only define the capability of the desired services at conceptual level (e.g. order business service), while some others also provide the design of such services along with its associated service contract (e.g. order software service design). Some of the approaches externalize the constraints which each service should meet, while some others do not explicitly consider any constraints. Interestingly, we observed that the companies with more experience in providing service-based solutions tend to define the ideal services more detailed compared to the ones with less experience. Hence, we argue that the extent to which the ideal service is codified is an indicator of the maturity of the migration approach.

Contrast with theory. Detailed description of the ideal services is a best practice that companies have developed with experience. Interestingly, we could not trace back this best practice to the academic approaches.

3 Discussion

In software engineering as an applied science, research in principle should serve the final purpose of being applied in practice. The extent to which this principle is supported by research, however, has been subject of debate for decades, and remains a still unsolved problem. The premier conference on software engineering featured in 2011 a panel on “What Industry Wants from Research” discussing the current gaps between theory and practice, and how to address them. All panel members in one way or another hinted the following cause of such gap: what research proposes does not fit the fundamental problems, goals, strategies and weaknesses of practice. We argue that, this paper is a step towards filling the theory and practice gap as it sheds light on how migration is performed in practice and further contrasting it with how academic research addressed the migration problem. By identifying the characteristics which make these approaches favorable for practice, we could identify directions for future research that have better chance of adoption by practitioners.

I) Migration approaches fitting core activities. Getting back to finding F1, we argue that core activities can act as a frame of mind confining the migration approaches that are more aligned with practice. From that perspective,

one would see that, for instance, the approaches addressing wrapping the applications as a whole are more in-line with practitioners concerns, compared to the ones addressing the automatic recovery of the legacy architecture. Hence, this frame of mind pinpoints the types of industry-relevant research in SOA migration methodologies and techniques.

II) To-Be driven migration approaches. As noted in finding F2, inadequate support for the bowl-shaped approaches in academia highlights promising opportunities for research to focus on how to support To-Be driven migration. For instance, future research can focus on addressing the following challenge of the practitioners: how to systematically elicit and capture the migration drivers and how to shape the migration process using those drivers.

III) Legacy understanding without reverse-engineering. Although reverse engineering is not covered in industrial migration approaches (see finding F3), elicitation of the knowledge about the legacy system is crucial for a successful migration. In this regard, research can benefit practice by providing methods, techniques, or guidelines that facilitate elicitation of migration-relevant knowledge from different sources of such knowledge.

IV) Legacy evaluation from multiple perspectives. As noted, companies evaluate and extract the legacy assets for migration to SOA by depicting their ideal services. This is, however, done in an ad-hoc manner, which may hinder successful service extraction. An immediate concern calling for further research is how to systematically evaluate pre-existing legacy assets based on different aspects of the ideal services.

4 Conclusions

This paper explored the types of migration approaches employed by leading SOA solution providers in practice. Results show that by supporting similar set of activities, process organization, and best practices, industrial migration approaches do converge to one, common, type of migration. As such, this paper suggests that the industrial approaches mature towards a similar approach to SOA migration. Further findings (removed for sake of space) show that industrial approaches, strictly follow incremental migration.

In spite of what academics think, practitioners still face difficulties in consolidating to a successful yet cost-effective migration approach. The many available methods often prove to be abstract or commercial to be applicable. By contrasting the industrial migration approaches and the academic ones, this paper emphasizes important gaps between theory and practice and consequently sketches the promising industry-relevant research directions. Those research directions enable finding solutions to problems that industrial practice confronts in real-world migration cases and is tailored to individual needs.

When a company wants to devise or select a specific approach for migration of its pre-existing assets to services there are many issues that need to be resolved. In this study we identified the type of industrial migration approaches that

is feasible in practice. What issues, goals, assumptions, and decisions explicitly make that specific type of migration favorable in practice, though, is yet unclear. We are carrying out follow-up studies to identify the goals, assumptions and issues that shape the migration decision making process.

References

1. Razavian, M., Lago, P.: Towards a Conceptual Framework for Legacy to SOA Migration. In: Fifth International Workshop on Engineering Service-Oriented Applications (WESOA 2009), pp. 445–455 (2010)
2. Razavian, M., Lago, P.: A Frame of Reference for SOA Migration. In: Di Nitto, E., Yahyapour, R. (eds.) ServiceWave 2010. LNCS, vol. 6481, pp. 150–162. Springer, Heidelberg (2010)
3. Arsanjani, A., Ghosh, S., Allam, A., Abdollah, T., Gariapathy, S., Holley, K.: SOMA: a method for developing service-oriented solutions. *IBM Syst. J.* 47, 377–396 (2008)
4. Razavian, M., Lago, P.: A Dashboard for SOA Migration (under submission, 2011)
5. Lewis, G., Smith, D.B.: Developing Realistic Approaches for the Migration of Legacy Components to Service-Oriented Architecture Environments. In: Draheim, D., Weber, G. (eds.) TEAA 2006. LNCS, vol. 4473, pp. 226–240. Springer, Heidelberg (2007)
6. Hansen, M.T., Nohria, N., Tierney, T.: What's your strategy for managing knowledge? *Harvard Business Review* 77(2) (1999)