

**SERVICE-ORIENTED RE-ENGINEERING OF
LEGACY JEE APPLICATIONS: ISSUES AND
RESEARCH DIRECTIONS**

*Hafedh Mili and Ghizlane El-Boussaidi and Anas
Shatnawi and Yann-Gaël Guéhéneuc and Naouel Moha
and Jean Privat and Petko Valtchev*

Septembre 2017

Département d'informatique

Université du Québec à Montréal

Rapport de recherche Latece 2017-6



Laboratoire de recherche sur les technologies du commerce électronique

**SERVICE-ORIENTED RE-ENGINEERING OF LEGACY
JEE APPLICATIONS: ISSUES AND RESEARCH
DIRECTIONS**

*Hafedh Mili, Anas
Shatnami, Naouel
Moha, Jean Privat and
Petko Valtchev
Département
d'informatique
Université du Québec à
Montréal
Montréal, Qc, Canada*

*Ghizlane El-Boussaidi
Département de génie
logiciel et des TI
École de Technologie
Supérieure
Montréal, Qc, Canada*

*Yann-Gaël Guéhéneuc
Département de génie
informatique et génie
logiciel
École Polytechnique
Montréal, Qc, Canada*

Laboratoire de recherche sur les technologies du commerce électronique

Département d'informatique
Université du Québec à Montréal
C.P. 8888, Succ. Centre-Ville
Montréal, QC, Canada
H3C 3P8

<http://www.latece.uqam.ca>

Septembre 2017

Rapport de recherche Latece 2017-6

Summary

Service-orientation views applications as orchestrations of independent software services that (1) implement functions that are reusable across many applications, (2) can be invoked remotely, and (3) are packaged to decouple potential callers from their implementation technology. As such, it enables organizations to develop quality applications faster than without services.

Legacy applications are *not* service-oriented. Yet, they implement many reusable functions that could be exposed as *services*. Organizations face three main issues when re-engineering legacy application to (re)use services: (1) to mine their existing applications for reusable functions that can become services, (2) to package those functions into services, and (3) to refactor legacy applications to invoke those services to ease future maintenance.

In this paper, we explore these three issues and propose research directions to address them. We choose to focus on the service-oriented re-engineering of *recent legacy* object-oriented applications, and more specifically, on JEE applications, for several reasons. First, we wanted to focus on architectural challenges, and thus we choose to *not* have to deal with programming language difference between source and target system. We chose JEE applications, in particular, because they embody the range of complexities that one can encounter in recent legacy applications, namely, multi-language systems, multi-tier applications, the reliance on external configuration files, and the reliance on frameworks and container services during runtime. These characteristics pose unique challenges for the three issues mentioned above.

Contents

1	Introduction	1
2	Principles	2
2.1	Reusability = Usefulness + Usability	2
2.2	It is architectural style migration	3
2.3	Focus on source code	4
2.4	KDM as a repository for architectural level knowledge	5
3	Issues in Service Identification	5
3.1	What is a service	5
3.2	A service taxonomy	8
3.3	Service identification: a (brief) survey	11
3.4	Service identification: our strategy	13
3.5	Issues in static code analysis of legacy JEE Application	16
4	Issues in Service Packaging	26
4.1	Computing service interfaces	26
4.2	Designing and implementing service interfaces	27
5	Issues in Refactoring	31
6	Conclusion	34

List of Figures

1	Services can be fairly coarse-grained; components can be fairly fine-grained	8
2	Java code analyses have to be complemented with the analysis of, 1) other program artefacts, and 2) services offered by the run-time infrastructure	19
3	An overview of the components of the J2EE family of technologies, circa 2002 (EJB2.x)	20
4	We need to codify the dependencies that are inherent in frameworks .	22
5	We need to codify the dependencies that are inherent in RMI	23
6	The lifecycle of EJBs. Red methods belong to the <i>home</i> interface, but they end up invoking the blue methods on the EJB class (e.g. a call to <code>create()</code> on the home object lead to a call to <code>ejbCreate()</code> and <code>ejbPostCreate()</code> on the EJB object). The bold methods correspond to method invoked following explicit user actions. Non-bold methods are invoked by the server to manage its resources	25

List of Tables

1	A multidimensional service taxonomy	10
2	Service type identification criteria: a preliminary proposal	15

1 Introduction

IT departments strain under the needs to automate business processes within, and across, organizational boundaries, and to develop and deploy new applications. They also contend with heterogeneous applications that duplicate some of the same business functions but cannot share most of these functions because of the way that these functions are packaged. Thus, they face the problem of encapsulating functions into reusable services.

Service-orientation purports to solve this problem by viewing applications as orchestrations of independent software services that (1) implement functions that are reusable across many applications or application *domains*, (2) can be distributed and invoked remotely, and (3) are exposed such that their interfaces (a) abstract the specificities of their implementation technology and (b) allow *any client*, regardless of its implementation technology, to invoke them.

An organization that applies service-oriented development can develop new applications (1) by performing a high-level decomposition of the functions of the applications into coarse functional blocks, (2) by mapping those functional blocks to available services, (3) by implementing the functional blocks that have no corresponding services, and (4) by implementing the process that orchestrates all the services. The more services we have available, the more productive and effective is service-oriented development.

An organization may also want to re-engineer its legacy applications to use service-oriented development (1) by mining the existing applications for reusable functions that could qualify as services, (2) by packaging these functions as *services* to enable their (re)use, and (3) by rewriting some existing applications to (re)use the newly-identified services. Indeed, while service identification and packaging is useful for *future* application development, they do little for existing applications. Therefore, *refactoring* existing applications into service-oriented applications is also important to modernize legacy applications and ease future maintenance.

In this paper, we gather, define, and discuss the most pressing issues and research directions on (1) service identification, (2) service packaging, and (3) the refactoring of legacy applications into service-oriented applications. We survey the literature related to service identification, in particular the many works on identifying services in legacy applications [11, 18, 4].

We choose legacy JEE applications as the input of the re-engineering effort for many reasons. First, we chose not to have to deal with *language* or *language paradigm* issues, and focus on *architectural* issues. Thus, we needed a legacy whose programming language(s) can support the SOA style. Second, we chose JEE *in particular*, as opposed to plain Java applications, because JEE applications embody the range of complexities that we are likely to encounter in enterprise applications. Indeed, JEE applications are multi-tier, multi-language, rely on external configuration files to link the pieces together, and rely on frameworks and containers during runtime to connect the pieces together. As such, they illustrate the range of complexities one

is likely to encounter in recent legacy enterprise applications. Finally, we chose JEE because the technology has been in use since 1999, and is quite popular for enterprise applications.

This report aims at providing researchers and practitioners with a view on the state of the art on service identification, service packaging, and service-oriented refactoring as well as highlight gaps in the literature. We conclude with a research agenda to discuss with the scientific community and industrial partners.

Section 2 introduces the principles of service-oriented development and our focus on source code of legacy applications. Section 3 discusses issues related to service identification. We first present a taxonomy of services and then discuss methods for identifying services of each category in the taxonomy. Section 4 discusses issues related to the packaging of services. Finally, Section 5 describes approaches to refactor the code of legacy applications to (re)use recently identified and packaged services. Section ?? discusses the issues and current state of the art while Section 6 concludes with a research agenda.

2 Principles

2.1 Reusability = Usefulness + Usability

The idea of building new applications from reusable artifacts is not new [14, 15]. Such artifacts can be analysis-level artifacts (e.g. software models, analysis patterns), design-level artifacts (e.g. architectural styles, design patterns, reference architectures), source-level artifacts (libraries, frameworks), or executables (compiled code, components, services). Reusability has many advertised advantages, including:

- *Enhanced productivity*: by reusing existing artifacts of level i , we save on development tasks *up to* development stage i . By reusing analysis models (or patterns), we save on analysis; by reusing executable components/services, we save on analysis, design, coding, and testing of the functionality implemented by the components/services.
- *Enhanced quality*: not only do *domain engineering* methodology suggests that reusable components be thoroughly tested, but *reused* artifacts, whether by design or by accident, will have been tested in *many* different contexts.
- *Enforcement of patterns* inherent in the reusable artifacts, be they analysis level patterns, architectural patterns, design patterns, coding styles, etc.

We argued in [15] that *reusability* is a combination of two qualities:

- *Usefulness*, which represents the extent to which a reusable artifact can be used in different contexts. A *sorting* function/procedure or a *collection* data structure are examples of very *useful* utility (domain-independent) artifacts.

- *Usability*, which represents the extent to which the functionality is packaged in a way that facilitates its use in those contexts where it is *useful*. If I am programming in Java, a C sort routine is of little use to me.

Our work on service-oriented reengineering will deal with both issues:

- Service identification will deal with issues of *usefulness*: identify those clusters of functionality that are used/invoked in *many* places within the same application, or across *many* different applications.
- Service packaging deals with *usability*, as it attempts to wrap/package that reusable functionality behind *service interfaces*, which will facilitate their reuse (remoteness, technology independence, etc.).

2.2 It is architectural style migration

The problem that we are trying to solve, i.e. the *service-oriented reengineering of legacy JEE applications*, is an instance of the more general problem of *architectural style migration*: how to migrate an application from an architectural style *source* to an architectural style *destination*, *while keeping other things equal*. While migrating legacy COBOL applications towards a Java-based service-oriented applications may be of utmost practical relevance, such an endeavor requires dealing with *many* differences between the source and target technologies:

- Dealing with different programming *languages* that are generations apart, at the programming language *construct* level (typing, procedural invocation, parameter passing, control structures, data reference).
- Notwithstanding statement/construct-level difference, we would have to deal with different programming *paradigms* (as in going from C to C++).
- Dealing with different virtual–and physical–machine models (e.g., how processes are run, how concurrent access to shared resources is done, etc.).
- And, *also*, dealing with different *architectural styles*.

Because we want to focus on the issues dealing with *architectural style migration*, apart from the other issues, we chose a recent legacy—in this case a JEE web application—so that we don't have to deal with language or *programming paradigm* issues. Thus both the source and target application will be written in Java¹. Similar work can be explored with recent legacy in other languages (e.g. C#). While the *tools* that we develop may not be portable across languages, the issues—and solution approaches—should be similar.

¹*Mostly*, in addition to all the other scripting and 'data specification' languages, such as XML-based configuration files

2.3 Focus on source code

Service identification is a difficult problem, in part because there are different kinds of services, which have different characteristics (see section 3.2), and in part because there are domain-specific considerations that may determine what makes a 'good service'. The literature on service identification uses a number of techniques that often combine *many* knowledge sources, *including* source code, but also including domain knowledge, as provided by subject matter/domain experts, or *domain-/analysis-level* models, that provide a more abstract representation of the functionality of the software [references]. Other work combines *static* code analysis with *dynamic* code analysis, that looks at execution traces of the legacy software at hand, which gives a better idea about the *effective/actual usefulness* of a piece of code (how often it *is used*), as opposed to the *potential usefulness* of a piece of code, as established by program dependency graphs, for example².

For this work, we choose to focus on *static analysis of source code*, for several reasons:

- Dynamic code analysis (run-time tracing) is *usually* fairly complicated to set-up, and the techniques that we develop can only be applied to relatively 'small' and 'simple' legacy applications. For example, in a client-server application, both the client and the server need to be traced. When the two do not run on the same machine, it becomes difficult to collect– and correlate– the traces.
- Domain or analysis models, which are sometimes used to complement static code analysis (reference) may not always be available. And when they are, they are *most likely* out of sync with the source code.
- We are not so much interested in advancing the state of the art in *service identification*, which has received *considerable attention* in the past decade or so, as we are interested in issues dealing with *service packaging* and *service refactoring*, which, in our view, did not get the attention they deserve. The technical challenges involved in service packaging and refactoring are, as far as we can tell, mostly independent of the *intrinsic* quality (the *usefulness*) of the services. By foregoing those additional knowledge sources (execution traces or domain/analysis knowledge), we may get lower quality services, but that does not prevent us from, or affect our ability to, explore the issues related to packaging and refactoring.

We do not exclude input from human experts to annotate/qualify intermediate or final results of service identification to enhance the quality of the services, but we will refrain from making such input a *necessary* step in service identification.

²The *presence* of a call relationship between two methods indicates a *possibility*, depending on whether the control path along which the call is made is actually taken during execution

2.4 KDM as a repository for architectural level knowledge

Service identification, repackaging, and refactoring require that we manipulate representations of the legacy application at various levels of abstraction, going from instruction-level representation to detect method/procedure invocations or perform data flow analysis, to representations of *software artifacts* including files of different types (HTML, JSP, JSF, Java, various dialects of XML, property files, etc.), to representations of architectural-level concepts such as the notion of *package*, *layer*, *interface*, and eventually, *service*.

Because of the interdependencies between the different representations— and reengineering tasks— we need a representation that:

- *integrates* the different views/representations of the legacy application, and
- we can share with researchers working on similar problems, or use to compare our work/algorithms to those of other researchers.

To this end, we propose to use OMG's *Knowledge Discovery Metamodel*, which provides a '... common *intermediate representation* for existing software systems and their operating environments' (<http://www.omg.org/technology/kdm/>). The main page of OMG's KDM standard further defines KDM as:

KDM is a metamodel for knowledge discovery in software. It defines a common vocabulary of knowledge related to software engineering artifacts, regardless of the implementation programming language and runtime platform - a checklist of items that a software mining tool should discover and a software analysis tool can use.

Specifically, we will use the MODISCO (partial) implementation of KDM, which is an open-source Eclipse project (<http://www.eclipse.org/MoDisco>). MODISCO, available as an Eclipse plug-in, includes a number of tools that can parse files of different types and build KDM representations of the contents of those files. Within the context of our work, we expect to write our own parsers/builders to analyze the different file types that we are likely to find in a JEE legacy application, and build the corresponding KDM representations. Such parsers/builders may include Eclipse's own JDT (Java Development Toolkit), which knows how to analyse/parse Java projects.

3 Issues in Service Identification

3.1 What is a service

There has been much debate in the community about what constitutes a service. Thomas Erl (Reference), an SOA pioneer, identified eight characteristics of services :

- *Standardized [service] contracts*. As software components, services define their capabilities using a standard, implementation-neutral language.

- *Loose coupling.* The services are loosely coupled, and any dependencies are explicitly stated in their service contracts.
- *Abstraction.* Whereas loose coupling refers to dependencies between services, abstraction refers to dependencies between a service provider and a service consumer. The consumer should not depend on the implementation details of the service.
- *Reusability.* Services embody reusable functionality that can service many consumers. In other work, we defined reusability as usefulness and usability (Reference). Usefulness refers to how often the provided functionality is needed while usability refers to how easy it is to use. Usability embodies many aspects, including the existence of (standardized) service contracts (see above), as well as discoverability, composability, and interoperability, discussed below.
- *Autonomy.* From the perspective of the consumer, services should be perceived as self-contained components with total control over their resources and environment. The consumer should be able to assume that the service needs no more than the parameters specified in its service contract to do its job. Naturally, behind the scenes, a service may in turn depend on other services. For example, business services can depend on a layer of shared technical services.
- *Statelessness.* We can understand statelessness of services in two complementary ways. To be able to 'service' many consumers, a service should not have to rely on implicit state information about its consumers; all of the data needed to service a particular consumer's request should be explicitly passed as parameter. The second aspect of statelessness is related to multiple interactions with the same consumer. This means that a consumer can invoke the operations of the service as many times as they want, in any order they wish, and always get the same result. In practice, of course, these two conditions are seldom attainable-and not necessarily desirable. If I am using a flight booking service, I sure hope that my interactions with the service have a lasting effect on the state of the world: the creation of a booking in the booking database. Erl writes "Applying the principle of service statelessness requires that measures of realistically attainable statelessness be assessed, based on the adequacy of the surrounding technology architecture to provide state management delegation and deferral options" (see footnote 17).
- *Discoverability.* This refers to the ability of services to document and advertise their capabilities so that service consumers can find them. The documentation of the capabilities of a service needs to be expressed in a domain language that is distinct from the language used to express the service contract.
- *Composability.* This refers to dual capability of services to, a) be composed at arbitrary levels of aggregation to form more complex services, and b) address

many needs. This, in turn, influences two design aspects of services, a) the modalities for interacting with the service, and b) the way the capabilities of the service are distributed among its operations.

All of the characteristics, but for *discoverability*³ apply, to a large extent, to *software components*. There has been much debate in the literature about the difference between services and components. We see three differences:

- **Packaging.** With *services*, there are existing *standards* that prescribe, a) how services expose their functionalities (e.g. using WSDL files, for the case of web services, or SCA interfaces, for SCA), and b) how their functionalities are invoked (e.g. using the HTTP/SOAP or REST protocols). However, no such standardized protocols exist for *components*⁴.
- **Lifecycle.** *Typically*, components are 'dormant' when not needed. They come to life when solicited, service the incoming request, and die out upon termination; in other words, their lifecycle is managed, directly or indirectly, by their client applications ('consumers')⁵. By contrast, services *tend to be mostly* independently running processes/services, with their lifecycle managed independently from their client applications.
- **Granularity.** Components can be as small as an ActiveX control (a button, or a selection-in-list widget), while services can be as big as a 'business process' (see Section 3.2). This is perhaps a corollary of the previous two characteristics: the overhead involved in the packaging and lifecycle management of services makes it ineffective for small-grain functionality— although, *microservices* tend to blur the lines.

Figure 1 illustrates the difference in granularity between services and components.

We suspect that, with the exception of granularity, the differences between services and components will not matter for *service identification*, as they deal mostly with packaging and deployment of the identified services, but not with their functional content, or patterns of usage, per se. More on the criteria for service identification in section 3.4.

³One might want to distinguish between *definitional properties/characteristics* of services, which *all* services must have, from *characteristic properties/characteristics*, which *most* services will have but which are not necessary. We could also distinguish between content-related characteristics, and packaging-related characteristics. Within the context of *web services*, discoverability relates to the extent to which a description of the service functionality is available and published in a service registry à la UDDI.

⁴There are, of course, *component models*, such as Microsoft's COM/DCOM models, or Java-based EJBs, however, they are language/technology/vendor-specific.

⁵The lifecycle of components, within a particular technology, may be managed by *containers* for that technology, as is the case for the Windows operating system, for COM components, and for JEE servers, for the case of EJB entities




Figure 1: Services can be fairly coarse-grained; components can be fairly fine-grained

3.2 A service taxonomy

Many authors/standards pointed out that it is important to define and use a service taxonomy when implementing an SOA [7, 5, 1]. Many service taxonomies were proposed in the literature (e.g., [7, 5, 8, 2]). These taxonomies tend to classify services according to a hierarchical layered scheme that mainly supports communication between stakeholders during the implementation of SOA initiatives. Most of these taxonomies classify services according to their granularity (e.g., [7, 5, 2]) and/or their reuse (e.g., [5, 8]).

In the Opengroup SOA reference architecture (RA), services are classified according to the capabilities they are providing [1]. The Opengroup taxonomy describes a broad functional categorization scheme which includes: strategy and planning services, management services, development services, business services, infrastructure services, asset and registry services, interaction services, process services, information services, access services, partner services, business application services and life-cycle services. Most of the other taxonomies (e.g., [7, 5]) concentrate only on the core services implementing an SOA solution, i.e., services implementing the business processes, domain-specific functions, and infrastructure/utility functions.

All proposed taxonomies distinguish between domain-specific services and domain-neutral (or generic) services. Domain-specific services, commonly referred to as busi-

ness/application services represent/implement business processes and domain-specific functions. Depending on their granularity and scope of reuse, most of existing taxonomies identify the following types of domain-specific services:

- **Business services:** Named process or task services in some taxonomies (e.g., [7, 5]), they correspond to business processes or use cases. These are the services that are generally consumed by the end-users; for instance a service that enables an on-line booking. To implement business processes, these services compose the capabilities/functions provided by the Enterprise, Application and Entity services described below.
- **Enterprise services:** Named capabilities in [5], they are of finer granularity than Business services. They implement generic business functions that are generally reused across different applications; for instance a service that computes taxes.
- **Application services:** These services implement business functions that are specific to an application. They may be created to support reuse within the application scope or to enable the decomposition of a complex business process [5]. Similar to Enterprise services, these services are of finer granularity than Business services.
- **Entity services:** They are also called information or data services depending on the taxonomy. They provide access to and management of the persistent data of the business. They generally support actions on data (create, read, update and delete), and may have side-effects (i.e., they modify shared data).

In their broadest definition, domain-neutral services provide common facilities that enable using, managing, developing and integrating domain-specific services. We distinguish two categories of these technical services:

- **Utility services:** They may be seen as services that do not support directly the business processes of the company. They generally embody some cross-cutting capabilities required by domain-specific services. Logging and security services are examples of such services.
- **Infrastructure services:** These services represent the facilities that are required for deploying and running an SOA application. Examples of such services are integration services which include services for routing, protocol conversion, message processing and transformation. These services are typically supported by an Enterprise Service Bus (ESB). Compared to utility services, infrastructure services have a broader scope of reuse.

Considering the different categories of services identified above, we propose a multidimensional taxonomy that aims at supporting the identification of services from the legacy source code. The dimensions of this taxonomy are:

- **Domain:** domain-specific (business) versus domain-neutral (technical).

- **Granularity:** fine-grained versus coarse-grained. The granularity of a service depends on the complexity of the capability/function the service provides.
- **Scope of reuse:** Enterprise versus Application. This dimension is dependent on the two previous ones: domain and granularity. Fine-grained services are likely to be more shareable as they may be composed to provide different coarse-grained services that may not belong to the same application. Likewise domain-neutral services, in particular infrastructure services, may be reused across different applications.
- **Side-effects:** computation-only services versus services with side-effects. Services with side-effects are those that manage the application state; i.e. they modify the persistent data of the application. These are mainly Entity services. To maintain data consistency, these services require the implementation of some transactional and/or compensation mechanisms.
- **Visibility:** This refers to the fact that end-users/external services/partner services interact directly with the service or not. Coarse-grained domain-specific services (i.e., Business services) are generally those that are visible.

Table 1 summarizes the properties of the different categories of services according to the dimensions of our taxonomy. In the context of a legacy to SOA migration process, service categories and their properties must be taken into consideration while building a service identification approach (see discussion in Section 3.4).

Service category	Domain	Granularity	Scope of reuse	Side-effect	Visibility
Business	Specific	Coarse	Application	No	Yes
Enterprise	Specific	Variable	Enterprise	No	No
Application	Specific	Variable	Application	No	No
Entity	Specific	Fine	Variable	Yes	No
Utility	Neutral	Variable	Variable	No	No
Infrastructure	Neutral	Variable	Enterprise	No	No

Table 1: A multidimensional service taxonomy

Beyond the fact that different service types require different identification strategies, the most appropriate SOA deployment technology (e.g., Web services versus

REST), may depend on the service type. While REST seems to be a very good fit for implementing Entity services with its uniform interface, it's not granted that it is a good fit for the other types of services. Many factors, including the needs for more advanced caching features and security requirements than those provided by REST, makes it worth investigating the features supported by SOAP⁶.

3.3 Service identification: a (brief) survey

Generally existing service identification methods (SIMs) use different techniques to identify services. These techniques include model-driven approaches, clustering and architecture recovery techniques and ontology mappings.

Alahmari et al. [2] propose a model-based SIM. They propose a service classification that is based on the granularity of the services where the granularity of a service is determined by the number of messages and the complexity of data exchanged. To identify services, they propose a hybrid model-driven approach. They manually build a knowledge portfolio that contains analysis models describing main functions and components of the system. This is done using questionnaires, interviews and available documents. They also derive UML class diagrams from legacy code. Relying on the business functions defined in the knowledge portfolio, they manually build activity diagrams from these class diagrams. These activity diagrams are transformed automatically into BPMN models which are used to identify atomic processes and related business entities. The identified processes are stored in a process portfolio. Candidate services are identified from the process portfolio using metrics and rules that gather/partition atomic processes and their related entities. These metrics and rules are not detailed. The candidate services are then classified and evaluated according to an SOA meta-model and a set of rules to generate services with an optimal level of granularity.

Zhang et al. [20] propose an architecture-based SIM. The proposed approach starts by recovering design and architecture information using static and dynamic analysis. The extracted information is documented using an architecture description language. Service identification is carried out in two steps. The first step consists of analyzing the domain and building a domain model. The model is then used to identify the business functionalities that need to be provided as services, called logical services. In the second step, the recovered architecture information is analyzed to identify legacy components. This is done using hierarchical clustering techniques and modularization criteria (e.g., coupling, cohesion). A legacy component is then matched to the logical service that must offer the functionalities embedded in the component. Legacy components are integrated with newly built components and packaged into Web services.

⁶REST relies on the HTTP caching and security features. In specific contexts (e.g., Stock exchange), dynamic resources will require a more appropriate caching than that provided by HTTP. Also, HTTP provides a very limited support to security when compared to WS-Security the SOAP specification for applying security to Web services.

Greiger et al. [9] proposed an approach whose focus is to identify composite services which are the entry points for end users. The approach does not really identify other finer services. In particular, an initial service design is built by mapping each single module to a business service as defined in the classification in [2]; i.e. one of enterprise, application or entity service, according to our taxonomy. A composite service is introduced to aggregate and orchestrate these services. This initial service design is improved through a two-step process. In the first step, hierarchical clustering is applied to the business services depending on the navigation flow between them. Composite services are then introduced to aggregate the resulting clusters of business services. Hierarchical clustering is performed in an iterative way: at each iteration, new clusters are computed for a hierarchical layer and composite services are introduced to realize these clusters. In the second step, the functionalities of several services are partitioned into parts and common parts are moved into a single service. This partitioning is based on detection and removal of software clones.

Service identification methods were studied in a number of surveys (e.g., [10, 4, 11]). We briefly present some of these surveys in the following subsection.

3.3.1 Overview of existing SIM surveys

With the goal of supporting practitioners in selecting a service identification method (SIM), Gu and Lago [10] carried out a systematic review of 30 SIMs. The review focused on three critical aspects of SIMs: the inputs to these methods; the type and the format of the services they produce; and the strategies and techniques they use. The authors identified seven types of inputs including business processes, application domain (i.e., models or documents that describe different aspects of the domain) and legacy system. According to the results of this survey, most of existing SIMs adopt a top-down approach when implementing SOA and take as input business processes and domain information. Bottom-up SIMs rely on source code and its architecture to identify services. Very few methods adopt a hybrid approach that combines legacy systems and other types of inputs. Regarding the service type, most of existing SIMs identify business services and few target technical services. Most of the SIMs describe the identified services in an informal way using a list of terms (e.g., description, input, output). To identify services, the primary strategies used by the studied SIMs involve the decomposition of business processes or business function models. These strategies are combined with different techniques including formal rules encoded in the form of algorithms and less formal techniques provided as guidelines.

Cai et al. [4] surveyed a number of SIMs with the goal of highlighting their shared high-value activities and practices. In particular, they study the relation between service identification and the SOA engineering process (i.e., forward engineering vs re-engineering). In a forward engineering process, service identification is part of the requirement specification and analysis processes and it includes defining, analyzing, refining and modeling functional and QoS requirements. In a re-engineering process, service identification is part of the reverse-engineering stage and it involves analyzing both code and documents of the existing application, wrapping existing modules and

consolidating dispersed functions. Though SIMs adopt different processes, quite a few activities, called high-value activities, are shared among these methods. Cai et al. argue that selecting and composing these activities may be an effective way to identify services. They identify model-driven and decomposition activities as high-value activities in top-down SIMs while identification of reusable legacy assets is the most frequently used activity in bottom-up SIMs.

Most of existing surveys focus on the service type identified by existing approaches. Some of these surveys consider the techniques used by these approaches, the input they require and the output they produce. Very few studies considered how these approaches manage the granularity of the identified services. None of the surveys analyzed the impact of the SOA targeted technology on the identification method.

3.4 Service identification: our strategy

Most of the work on service (and component) identification methods described above rely on the 'modularity' aspect of a service or component—in addition to domain semantics—as reflected in its *functional cohesion*, i.e. the extent to which the different 'parts' of a component or service are related to the same functional objective, and its *coupling* with the outside world, i.e. the extent to which it depends on the 'outside world' to perform its functions. Accordingly, several of the approaches rely on *code coupling metrics* that measure code-level dependencies between lower-level constructs such as *classes*, which are then used by *clustering algorithms* where dependency clusters are proposed as candidate components/services.

In light of the different kinds of services described in section 3.2, we believe that *not only* cohesion and coupling are not *sufficient* to characterize services, among all potential clusters of functionality, but they may not even be *necessary*. Take the example of a *security service* that stands between a program and its clients, checking whether a particular client has the right to invoke a particular function of the program. We can have two programming models:

- Clients request the execution of functions by talking to the security manager, who checks their credentials against a stored security policy, and then decides whether to forward the request or not, or
- Clients talk directly to servers' public API, who first double-check with the security manager, before proceeding with the requested call.

In the first case, the security manager will have a strong coupling with both clients and servers: it will be called by *all* client programs, and it will call *all* server functions. In the second case, the security manager will be called by *all* server functions. Either way, the functionality offered by the security manager will likely not show-up *alone* within its own dependency cluster, and will not be proposed as a candidate service.

Thus, we argue that any service identification approach needs to take into account the specific type of service that we wish to identify, and use the metrics that are appropriate for that type. More generally, we believe that *component identification*

algorithm for a target architectural style needs to take into account the specifics of the components types for that style. As a first step, we need to identify, for each service type, the kind of 'signature' that such a service type will have on the code, *and then*, develop algorithms to detect such signatures. We believe that not all service types have distinct *code smells*, and two different service types may leave similar or indistinguishable traces in the code.

Further, to go back to our security manager example, we may be able to identify such 'interceptor-type services' by looking for classes that have both high fan-in ('everybody calls them') and high fan-outs (they call everybody). However, this relies on the fact that the functionality of the security manager has been packaged within the same class. For example, the security manager might have distinct methods for checking whether a caller can, a) create an object of a particular type, b) read the value of objects of a certain type (or certain objects), c) modify objects of a particular type, or d) delete objects of a certain type. However, if those methods are not part of the same class, we will have two additional complications:

- The code signature (high fan-in, high fan-out) may not be as strong
- It would not be easy to guess that the four independent functions are part of the same service.

Thus, our service identification algorithms need to rely on code smells of *service functionality that is not yet packaged as a service*.

Table 2 shows a preliminary mapping between service types and service identification criteria. Let us elaborate on some of the cells of the table:

- With regard to *entity services*, the way to detect their occurrences depends on the kind of operation: C, R, U, or D. For C and U, concentrated use of constructors, setters, and more generally, methods returning void, for R, use of getters on individual objects or configurations of objects (e.g. navigating associations).
- With regard to *enterprise services*:
 - regarding their characteristics, we said that are supposed to have no side-effects. However, when the computations are complex, we may need complex objects to return them, thus they may have side-effects in the *programming sense*, but do not result into modifying persistent data
 - regarding their code signatures, they may follow similar patterns to an R pattern for an entity service, with significant computation (perhaps more than simple aggregation) that involves external data/constants.
 - regarding ways to uncover them: the best way to test their (re)use across applications *is to have several applications*: if we had a portfolio of enterprise applications, then we can look for reusable clusters.

Service type	Examples	Characteristics	Code smell	How to recognize	related work
Utility service	A security service, a logging service	Implements a virtual machine-like function	Methods that fit a particular pattern call the service	High fan-in from call sites that share some characteristics	work on aspect mining such as [19, 16, 13]
Infrastructure service	a persistence service, a message broker, etc	pervasive services	high fan-in and low fan-out, a layer in a layered architecture	fan-in from higher layer, fan-out to lower layers	extracting layers in layered architectures (e.g. [6, 3])
Entity service	Creating an order, retrieving payments made on an order	CRUD operations on <i>configurations</i> of <i>domain objects</i>	different patterns, depending on C, R, U or D	see below	criteria for REST services might be useful
Enterprise service	Computing taxes, scoring applications, rendering decisions	Supposed to have no side effects. See below.	similar to entity services, except for side effects	same as with entity services, except side effects	Not sure
Business services	Booking a flight, making an on-line purchase	they implement a business process/use case	complex orchestrations of smaller domain services	there is not a single pattern to look for	Existing work relies on external knowledge such as BP models
Application services	Etc	Smaller granularity than business services; see below	Complexity, combination of more elementary services	May not be distinguishable from business services	Unknown

Table 2: Service type identification criteria: a preliminary proposal

- With regard to *business services*:
 - concerning their code smells, how do we detect/guess that a piece of code implements a business process? it is a combination of factors. The functionalities may exhibit a significant error handling or compensation component; they are transactional, possibly involving different databases or servers; they typically result into many changes to disjoint databases (multi-party transactions), multi-threading, asynchronous execution, etc.
 - regarding ways to detect them, there is no single pattern: it is a preponderance of evidence from the presence of each one of the previous characteristics (transaction frontiers, use of messaging, complex error handling, etc.)
- With regard to application services:
 - Regarding their characteristics, although they represent orchestrations of other domain services (e.g. entity or enterprise services), they tend to be application-specific, reusable within application, but not across applications.
 - regarding ways to uncover them, they will exhibit similar levels of complexity to business services in terms of error handling, multi-threading/asynchrony, compensation, and so forth, and thus, will not be distinguishable from business services. However, we probably do not care about the distinction, which is mostly granularity and reuse scope

This table provided some fairly preliminary hypotheses. More thought needs to go into the different categories. Also, we need to decide whether the distinction between different types actually matters. In the next section, we discuss the issues raised by the static analysis of legacy JEE applications.

3.5 Issues in static code analysis of legacy JEE Application

Having decided to perform service identification using source code as our only input, we now look at some of the issues raised by analyzing the source code of JEE applications. We first present an overview of the issues, then discuss some of the issue in some detail.

Static code analysis relies on a set of graph-like structures that represent the relationships between program elements. Such structures provide the basis for computing the various metrics that we will need for our service identification algorithms. As a preliminary step to service identification, we need to compute such structures in a way that identifies *all* of the relationships that exist between elements, *despite* the different language, technology, and development environment mechanisms that combine to 'abstract away'—and thus obfuscate—such dependencies, that we are likely to encounter in modern, or recent legacy applications. To take a small example, in order

to make an application localized (i.e. have its textual output adjusted to the *locale* of the operating system running the application), we have to:

- Replace literal output strings in program statements by named variables
- Specify language-specific values for those variables in external property files, with language specific file extensions
- Have the application read the locale of the underlying operating system during run-time, to figure out the appropriate file extension
- Have the application load the property file that has the appropriate file extension, and initialize the named variables with the proper language-specific values

This simple example illustrates some of the levels of indirection that we are likely to find in modern or recent legacy applications simply to find the value of a string constant that is used within the program.

Generally speaking, the code of modern or recent legacy applications is difficult to analyze for a variety of reasons:

- They tend to be multi-language. For example, a typical JEE application will combine several languages, including: 1) Java, for both the server side and the client side (embedded within HTML, or JSP, or JSF tags), 2) Javascript, on the client side, embedded within HTML or JSP or JSF tags, 3) various property files, and 4) various configuration files (web.xml, ra.xml, etc.). For example, to figure out which Java method is called on the server side when the user presses a button on their browser, we typically need to scripts or files written in different languages, each with its own semantics
- They make heavy use of late/dynamic binding, wherever possible, through a combination of mechanisms, including:
 - Java reflexion/introspection, which enables a 'client program' to invoke the functionality of a class, for example, *intensionally*, as opposed to *literally/nominatively*, as shown in the following code excerpts

```
...
Customer myCustomer = new Customer(...);
...
Class itsClass = myCustomer.getClass();
String methodName = "setName";
...
Class[] argumentTypes = {java.lang.String.class};
...
Method nameSetter = itsClass.getMethod(methodName, argumentTypes);
```

```
Object[] arguments = {"John Smith"};

nameSetter.invoke(myCustomer, arguments)
...
}
```

In this case, it is difficult to guess that the method `void setName(String s)` of class `Customer` is called by this code fragment since the method is not invoked by name as in `myCustomer.setName("John Smith")`. Practically, we need to find the *value* of the variable `methodName` to find which method is ultimately called⁷.

- Data-driven control. For example, the *observer* pattern introduces a level of indirection which makes it harder to guess which *observer* method is called when a specific *change* happens to the *observable*⁸. Again, data flow analysis would be needed in this case to figure out the *values* of the data control variables.
- Runtime input. Since JEE applications tend to be interactive, some of the routing control will depend on user input, and thus, we won't know until run-time which path the control flow will take at a given point in the execution
- Reliance on frameworks and container services. Such applications are typically run within *containers*, which provide a number of *services* (persistence, security, transaction) which imply, among other things, that specific methods from user code (so-called *callback methods* will be called by the container/server, at specific points in the application lifecycle. Such call relationships will not be visible in the user code. We need to find a way of closing the gap.

All of these considerations mean that to get a complete representation of relationships between program elements, we need to augment the traditional *unilingual* program static analysis techniques with other kinds of analyses, involving other kinds of artifacts, but also, possibly, involving the codification of services offered by containers/application servers. Figure 2 illustrates this.

In the remainder of this section, we will discuss some of these issues in more detail.

⁷In this example, the variable `methodName` is initialized in the same scope, and the value is easy enough to find, but the value could have been returned by a function, in which case, a full dataflow analysis is needed

⁸Recall that methods in the *observable* that make changes that are worth notifying the *observer* about, need to notify the observer *before exiting* by calling a generic notification method as in `this.changed(changeType)`; where `changeType` is a string or a value from an enumeration that is agreed upon between the *observable* and the *observer* that would tell the observer what to do. In one variant, the observer use a `switch` statement on the `changeType` to select the appropriate handle of change.

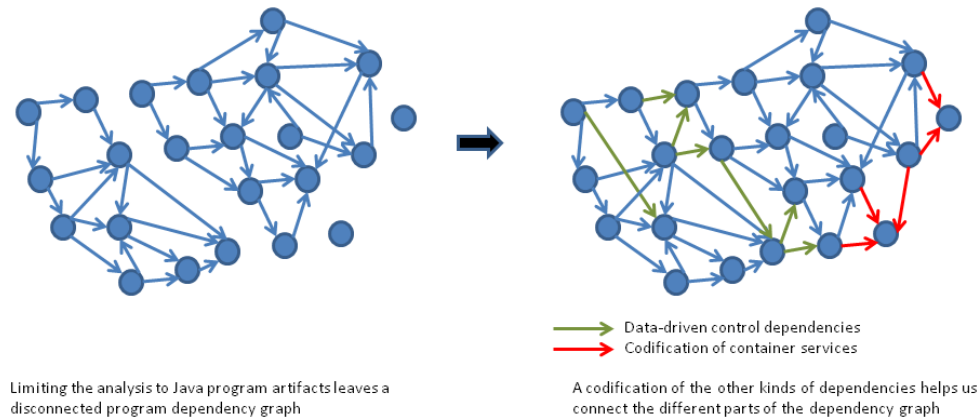


Figure 2: Java code analyses have to be complemented with the analysis of, 1) other program artefacts, and 2) services offered by the run-time infrastructure

3.5.1 Characterizing and codifying invocation patterns

Figure 3 shows an overview of the J2EE technology, circa 2002 (EJB 2.x). The figure shows a number of *technologies* (frameworks, tools, API) that are built on top of basic Java APIs, that abstract some of the services offered by the J2EE technology. Ultimately, in a web application, when all is said and done, when the user interacts with any widget on their client interface, some (typically server-side) user java code (as opposed to infrastructure code) will be called. The different technologies that are provided, such as servlets, Java Server Pages, or Java Server Faces, provide abstraction layers enabling client-side developers to focus on the visual aspects, and wrote as little application code as possible, to link up with the server side logic.

Thus, we need to examine the different technologies and figure out how that invocation works. In [17], we examined, among other things, invocation patterns for JSP, and identified *half a dozen ways* that a JSP page can be invoked⁹, including:

- Through the web.xml file, which shows the correspondence between JSP page names, URL's, and servlet classes

⁹JSP pages are mapped, by JSP containers, into *servlets*. The mapping can be done at deployment time or even during runtime. We will refer interchangeably/abusively to JSP pages, and the corresponding servlets

- Through annotations within Java classes (since Java 3)
- Through different kinds of explicit calls between servlets, can be found in servlet code, or in Java code fragments embedded within JSP tags within JSP pages, or Java code fragments embedded within HTML tags within HTML pages
- Etc. (see [17]).

When analyzing a JEE application, we need to look for the different invocation patterns that can be found in the different kinds of artifacts.

Similar analyses need to be made for other JEE technologies, including Java Server Faces (see [17]). For the purpose of this project, we will not perform an *exhaustive* analysis of *all* the technologies. We will limit ourselves to the most commonly used technologies in the early JEE applications ('recent legacy'), and/or to the technologies needed for our experimentation.

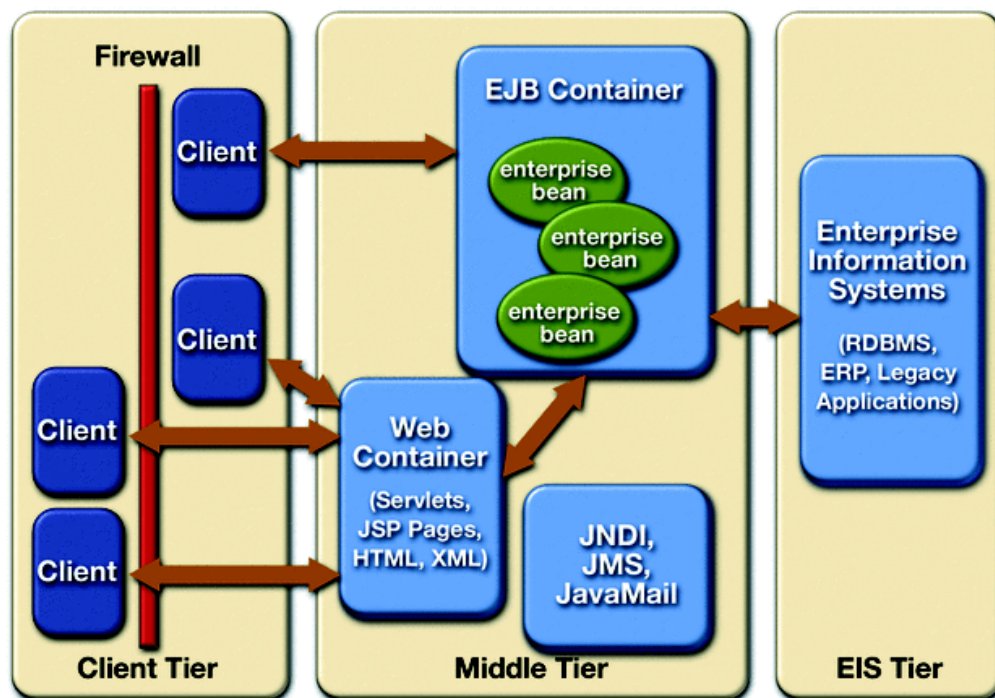


Figure 3: An overview of the components of the J2EE family of technologies, circa 2002 (EJB2.x)

3.5.2 Don't call us we will call you

Unlike code libraries, which developers can (re)use by *explicitly* calling the functionality that they need, *applications frameworks* embody reusable functionality through a

combination of, 1) services that are supplied by the framework to user code, *provided that the user code implements some predefined functionality*, and 2) reusable artifacts that developers can reuse, in a way similar to libraries. The 'service contract' between the framework developer and the framework user relies on *inversion of control*, or, as early OO pioneers called it, 'the Hollywood principle: don't call us we will call you'.

Web applications, be they JEE or otherwise, rely *heavily* on frameworks to ease the development, and manage the dependencies between the various layers of the application, including MVC frameworks, and persistence frameworks. To take advantage of the functionality provided by such frameworks (e.g. linking up and synchronizing models to views and controllers), developers have to implement call back methods which will be called by the 'framework infrastructure' at specific application execution times/events. However, if we want to figure out which server-side method, say, is called when the user presses a button, I need to be able to figure out how a particular framework associates application events with user code. Thus, for a given JEE application, we need to explicitly codify/represent the dependencies that are managed by the frameworks it uses. Figure 4 illustrates the idea.

Note that having the source code of the framework available does not necessarily make the problem any easier. To be able to offer a set of services to 'arbitrary' user code, frameworks tend to use, themselves, to have all of the characteristics of JEE applications as a whole: 1) multi-lingual, 2) heavy use of late binding techniques, such as the ones mentioned in Section 3.5, and the use of code generators. Not to mention that the source code is not always available. Thus, we are better off codifying the 'framework contracts' *explicitly*, as opposed to detecting them through code analysis. More on this when we talk about container services.

3.5.3 Codifying container service contracts

JEE application servers offer a number of *services* to contained/hosted applications, that *do not require* end user programming, including remote method invocation, lifecycle management, persistence, security, persistence, and transactions. In fact, this was one of the selling points of the JEE technology, as a technology for enterprise-class, distributed applications, as compared to the CORBA standard, for example, which defined many of the same services, *in the form of APIs that enterprise application developers needed to invoke*. Thus, it was argued, let application developers focus on the business logic, while we (JEE application servers) worry about the infrastructure services needed to support your applications.

A number of these services come out-of-the-box, and are offered to all the hosted applications, including remote method invocation and lifecycle management. Let us start with remote method invocation. Figure 5 shows the example of a two-tier client-server JEE application that uses EJBs— an *entity bean* in this case. We assume in this case that we are developing a client application (class `MyClient`) that manipulates customer objects which are managed by a remote server. The blue boxes represent developer-supplied code, which includes, 1) the client application (`MyClient` class), and 2) the customer EJB. In turn, to develop an EJB to represent *shared, remote*




Figure 4: We need to codify the dependencies that are inherent in frameworks

instances of customer, we need to supply, 1) a Java interface representing the remotely available public methods of customers (`Customer`), 2) a Java interface to represent some sort of a remote 'customer factory' (`CustomerHome`), and 3) the actual server-side class that implements customers (`CustomerEJB`).

Typically, such an application would be coded in (at least) two separate projects, one for the client side, and one for the server side. To compile the client side project, all I need is the Java interfaces (`Customer` and `CustomerHome`), and all I know is that the method `MyClient.foo(...)` is able to create an object *of a class that implements Customer*, but I don't know which class that is.¹⁰

¹⁰This is a simplification of reality, which is actually more complicated. In J2EE and J3EE, when we *deploy* an EJB, the deployment tool generates, among other things, a JAR file to be included in client projects. Such a JAR file will contain the Java interfaces *and the client side proxies*, represented in Figure 5 by the shaded Java classes `Customer.Proxy` and `CustomerHome.Proxy`. If I decided to complement my analysis of the client application with those two proxies, I might be able to figure out that the variable `newCust` of method `foo(...)` is *necessarily* of type `Customer.Proxy` (for example, by looking at which class is instantiated in the `create(...)` method of `CustomerHome.Proxy`), I would have no way of finding that the method `getName()` of `Customer.Proxy` calls the method `getName()` of the server side class `CustomerEJB`. The reason for that is that prior to 'traveling down the wire', method calls are packed into generic request invocation objects (that include identity of remote object, method name, and method parameters), which are *unpacked at the other/server end by the ORB (Object Request Broker)*, which *locates the appropriate instance of Customer_Impl*,

Similarly, on the server side, I may have a bunch of classes and methods that seem to be waiting to be instantiated and/or called, but I have no code to instantiate those classes/call those methods. The missing link *can't even be identified if I had the source code of the server*, because the server makes heavy use of the Java reflection API (see footnote). Thus, the only way to identify the missing links is *by understanding the underlying technology and by codifying its services explicitly*. In this particular case, having understood how remote method invocation works in JEE, I can *safely and surely* add a call link between the method `MyClient.foo(...)` and the method `CustomerEJB.setName(...)`. In fact, I don't even need to worry about polymorphism or the runtime type of (the server side) `newCust`: it can *only be* an instance `CustomerEJB`. This example illustrates the kind of work that we need to do for JEE services.

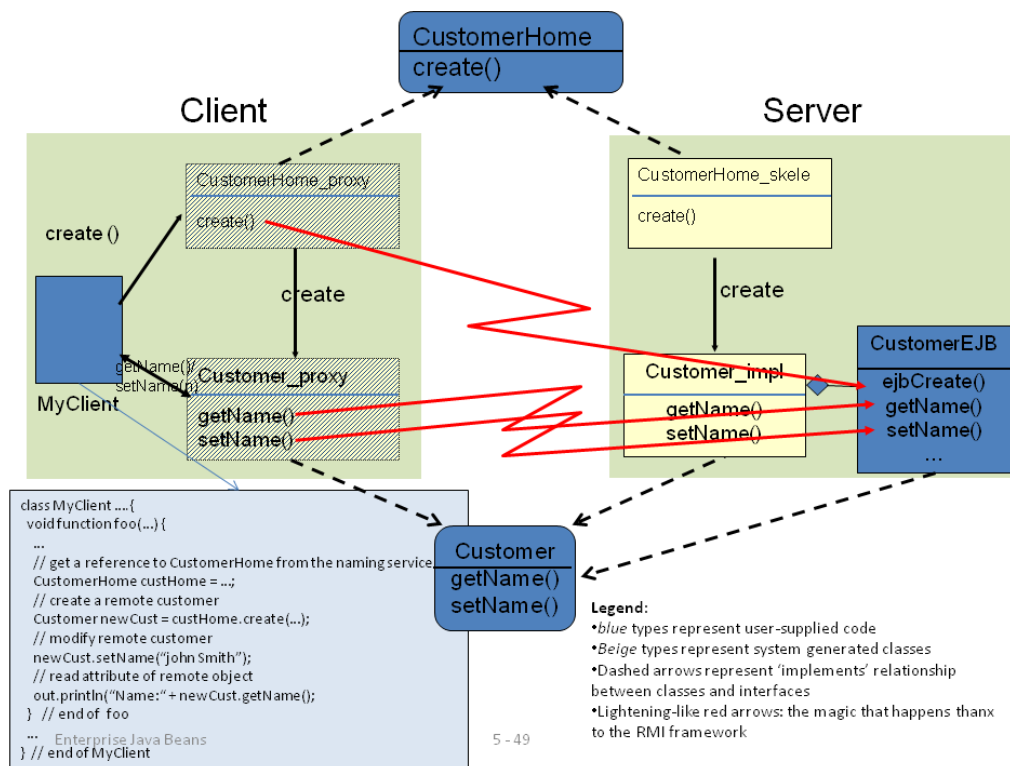


Figure 5: We need to codify the dependencies that are inherent in RMI

Another example of JEE container services is the *lifecycle management* service, which will manage the lifecycle of EJBs in a way that optimizes resource usage within the server. Figure 6 shows the lifecycle of *entity* beans, and it provides examples of the funny things servers do. First, we have to distinguish between three states for an entity bean: 1) null-state, which is the initial and final state, corresponding to the instance 'not-existing', 2) the *pooled state*, where an instance is created and is which then delegates the call the `CustomerEJB`.

made available to service client applications, and 3) the *ready state*, where the object is assigned to a specific business object, and is ready to service client application requests. To understand the meaning of *pooled state*, we should mention that JEE servers maintain *pools of EJB objects* during runtime, that they may assign to specific client programs. This is done, in part to save on object instantiation time, which is typically time-consuming, especially for large objects. Thus, when an object is no longer needed (for example, the shopping cart from your amazon session), instead of garbage-collecting it, the server "empties its fields", and returns it to the pool: that is the purpose of the `remove()` method. It is similar to the Java `release()` method, in the sense that you are supposed to reset the field of the object (to remove hanging references), and to release any resources it may hold (database connections, files, etc.). The `remove()` method on the home object actually calls the `ejbRemove()` on the EJB object. Similarly, when a client application ask a 'home object' to create an entity bean, the server pulls one from the pool, initializes its fields based on the parameters of `create(...)` (for example, a customer ID). That call will provoke a call to `ejbCreate(...)` with the same parameters (the same customer ID), which, in *bean managed persistence*, can be used to query the database and load the other fields. The method `ejbPostCreate()` acquires any additional resources that are needed, beyond object loading/initialization. It is at this point that a `CustomerEJB` object is 'loaded' with the data for an actual customer from the database, and is able to respond to queries.

When the server runs low on memory, it will call the `ejbPassivate(...)` method to typically *serialize* the EJB object on hard disk, so that it may be activated later. Serialization and deserialization to/from files (e.g. XML) are typically faster than executing a SQL query, and loading the object from the result set, especially because we can serialize and deserialize computed attributes, which are *not* stored in the database. This can lead to interesting design trade-offs between memory and CPU usage, as we weigh how much to save/serialize during passivation, against how much we need (and can) recompute during activation. But it illustrates a 'difficulty' with containers:

- They will call some methods without 'telling us' when. We cannot precisely characterize 'when the server runs low on memory'.
- Some of these methods may contain (or invoke) quite a bit of business logic. For example, if we decide to not serialize computed/recoverable attributes during `ejbPassivate(...)`, we need to recompute those attributes in `ejbActivate()`.

Generally speaking, container services fall into two general categories:

- out of the box services that apply to all deployed applications, including RMI and lifecycle management
- configurable services such as persistence, transactions, and security, where the parameters of the service can be found in many places

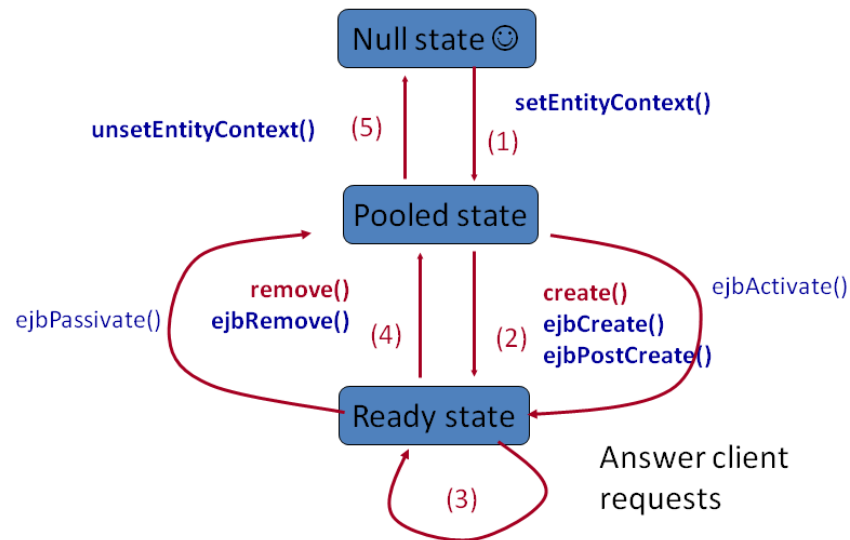


Figure 6: The lifecycle of EJBs. Red methods belong to the *home* interface, but they end up invoking the blue methods on the EJB class (e.g. a call to `create()` on the home object lead to a call to `ejbCreate()` and `ejbPostCreate()` on the EJB object). The bold methods correspond to method invoked following explicit user actions. Non-bold methods are invoked by the server to manage its resources

For example, with transaction services, the developer can specify which of the methods of an EJB need to be executed within transaction boundaries. This can be done in one of two ways: 1) in the EJB deployment descriptors ('old style'), or 2) as code annotations of the 'transaction methods'. Similarly, the *security service* enables developers to specify restrictions on the credentials of the client applications that call a particular method. Such restrictions can be specified in a security policy file, for coarse-grained security policies, or in deployment descriptors.

We make an additional distinction between container services:

- Services that are configured/specified *extensionally*, i.e. by listing the *explicitly* the entities/software components that benefit from the service. The security and transaction service are examples of such *extensionally specified* services, and
- Services that are specified *intensionally*, for example in terms of properties that the hosted applications/software components must satisfy to take advantage of the service. For example, all classes that implement a particular interface, or

that extend a particular class, will take advantage of the service. Examples of such services are RMI or lifecycle management.¹¹

The difference between the last two categories is important, for our purposes, because we will need to look in two different places to figure out whether we need to add a dependency or not. In the first case, we look at annotations and deployment descriptors. In the second case, we submit the software artifacts to a number of checks to figure out if they qualify for the service or not, and if they do, we add the dependency.

4 Issues in Service Packaging

Once we have identified a cluster of functionality as a candidate service, the next step consists of packaging that cluster 'behind a service' interface. We can think of the service identification step as recognizing those clusters of functionality that are *useful*, whereas service packaging as the step of making those clusters *usable*. Broadly speaking, there are two sets of issues:

- specifying/computing the service interfaces, based on the 'contents' of the service, and how it is used. For example, if we determine that two classes A and B, each with its set of methods, make up a potential service, we need to determine which methods of classes A and B need to be exposed as service functions.
- packaging them in appropriate service interfaces. Having determined which functionality needs to be exposed (previous step), now the issue is to figure out *how* to expose it. This involves a number of *design* and *implementation*-level issues.

We will discuss the two issues in turn in the next two subsections.

4.1 Computing service interfaces

At a basic level, given a bunch of Java classes that offers a cluster of functionality that was deemed worth (re)using as a whole, we need to figure out:

- the set of functions that are *offered* by the cluster. These will be presented as *provided interfaces*
- the set of functions that are *needed* by the cluster. These will be presented as *required interfaces*

¹¹For example, the lifecycle of an EJB depends on whether it is an *entity bean* or a *session bean*. The nature of the EJB is specified in the deployment descriptor, but is also implicit in the interfaces/classes that are implemented/extended by the EJB components.

A naive approach could compute the *provided interface* as the set of all the methods within the functional cluster that are called by methods that are *outside* the cluster. Similarly, we could compute the *required interface* as the set of all methods called by the code of the classes of the cluster, that do *not* belong to the cluster. However, the resulting interfaces would not be very accurate, or useful, we explained below.

With regard to the *provided interface*, there are two problems with the naive approach. First, some classes may provide functionality that is *not* related to the service at hand. Including all the methods called from the outside would clutter the provided interface with methods that have little relation with the functionality provided by the service. Thus, perhaps instead of including *all* of the methods of the classes that are part of the candidate service, we will only include the ones that contributed to the specific service pattern that was used to identify the candidate service. For example, if we are dealing with an interceptor-type technical service (see Section 3.2), the provided interface will be limited to those methods that have high fan-in, ignoring the others.

The second problem is more of a usability problem: a service will typically provide *several* interfaces, each one of which consisting of a subset of methods, exhibiting a stronger cohesion than the whole. Thus, we have to identify 'subclusters', within the provided interface of the service, and expose those as separate provided services. This can be framed as an optimization problem, where we need to balance the size of individual interfaces (not too big), with the number of interfaces (not too big either).

4.2 Designing and implementing service interfaces

Once we have identified the functions that are part of the interfaces, we need to find a way of exposing them. Broadly speaking, we need to hide the functionality of the service behind one or several faades. A naive approach would, for example, map each method f of class C with parameters p_1, \dots, p_i ($f(T_1 p_1, T_2 p_2, \dots, T_i p_i)$) to a method $g(C p_0, T_1 p_1, \dots, T_i p_i)$ of the faade, defined as follows:

```
public class MyServiceFacade {  
  
    ...  
    public T g(C p0, T1 p1, ..., Ti pi) {  
        return p0.f(p1, ..., pi);  
    }  
    ...  
}
```

However, we need to take into account a number of issues, some related to API design, in general, while others are specific to the implementation of the faade pattern.

For example, in general API (or library) design, it is usually recommended to have functions pass simple data types, as opposed to complex data structures. This insures, among other things, portability of the functionality. For example, this was one of the

design principles behind the X-Windows interfaces. To guarantee (or simplify) its portability across operating systems, its designers chose to pass the representation of various operating system objects (e.g. files or processes), as simple/individual attributes, instead of complex structures. For example, instead of passing *file* objects as C structures (`struct`), the functions that manipulated files took several `int` and `char*` parameters to represent the properties of a file (e.g. name, path, file type, etc.). This made it possible to have each operating system use its own representation that is accessed internally by the code of the functions.

If we apply the first pattern, that means that we should replace object parameters by a list of the object attributes. However, not all attributes may be needed by the specific functions. Thus, if we wanted only to pass the data that is needed by the function ... we need to know what data that is. This becomes a slicing issues: checking which slice of an object is read/written by a specific method, and passing only the attributes of the slice as parameters.

The issue of passing the object, versus its attributes, may depend on the service type. Could it be that with *business services*, we can pass the attributes—typically an object ID—where as with technical services we pass the entire object? Let us take the example of a business service for processing insurance claims, which checks the claimed expenses against the coverages included in the policy. Figure ?? shows excerpts of the relevant object model. It is a common design practice for such remote services to *not* pass the entire `InsurancePolicy` object, along with the `Claim` object, but simply pass the insurance policy *identifier*, and let the claim processing service fetch the policy object from persistent storage, if it needs to. This way, the policy object does not 'travel down the wire' Large objects consume time and space to `serialize, send, and deserialize at the receiving end`. The next code excerpts show what the insurance claim class might look like to make sure that insurance policy objects do not travel.

```
public class InsurancePolicy {
    ...
    private String id;
    ...
}

public class InsuranceClaim {
    ...
    private String insurancePolicyId;

    /**
     * transient fields are neither serialized nor
     * persisted. They may be 'populated' from
     * persistent storage on demand/at some point,
     * using insurancePolicyId, and accessed in a
```

```

        * read mode (unless invalidated by some other
        * mechanism, if modified)
        */
private transient InsurancePolicy policy;
...
}

```

One could also argue that technical services *typically* need the entire object, including, for example, serialization, and persistence.

In addition to the above, the facade pattern raises its own set of issues. One of the issues is the relationship between the service client programs, and the classes that are behind the service facade. To go back to our methods $f(T_1, T_2, \dots, T_i)$ and $g(C, T_1, T_2, \dots, T_i)$, if we make the object of type C as a parameter of $g(\dots)$, does that mean that clients of the service now need to 'create their own' instances of C to pass them as parameters of $g(\dots)$? if so, the service facade does not sound very helpful; it is actually more cumbersome. The answer depends, in part, on whether a *single* instance of C is needed to serve/service all of the clients of method $g(\dots)$, or whether each client needs its own instance of C :

- if a single instance of C can serve all the clients of the service, actually we do not even *need* to pass the C instance as a parameter to $g(\dots)$: we could have the faade manage a singleton, and then delegate the calls to that singleton, as shown below:

```

public class MyServiceFacade {

    private static C singleton = C.getInstance();
    ...
    public T g(T1 p1, ..., Ti pi) {
        return singleton.f(p1, ..., pi);
    }
    ...
}

```

- each client has its own instance of C . Strictly speaking, even in this case, we do not need to have the clients manage the C objects that serve them, or pass the corresponding object as a parameter. Clients may ask the faade to create and manage the C objects for them, and the faade ensures that they will be served/serviced by their assigned C objects. The following code excerpts illustrate this:

```

public class MyServiceFacade {
    private HashMap<Object, C>
        clientsServers = new HashMap<Object, C>();
    ...
}

```

```
public void assignMeAC(Object client) {
    C assignedC = clientsServers.get(client);
    if (assignedC == null)
        // the client has no C instance assigned to
        // it yet
    {
        // I create an instance of C specifically
        // for the client
        assignedC = new C(client);
        // register it for future reference
        clientsServers.put(client,assignedC);
    }
}

/**
 * here g(...) requires that we pass the
 * client/caller, as a parameter
 */
public T g(Object client, T1 p1, T2 p2, ..., Ti pi)
        throws NoCServerAssignedException {
    // first get the assigned C server
    C assignedC = clientsServers.get(client);

    // if the client has not yet been assigned
    // a C instance, throw an exception
    if (assignedC == null) throw
        new NoCServerAssignedException(client);

    // else, execute f(...)
    return assignedC.f(p1,p2,...,pi);
}

/**
 * @deprecated
 */
public C whoIsMyC(Object client) {
    return clientsServers.get(client);
}
}
```

Strictly speaking, the method `whoIsMyC()` may not be needed. But if we worry about forgetting to hide a needed `C` method behind the facade interface, we can provide this method for convenience, and mark it as deprecated to discourage its use.

A somewhat contrived example of the 'one object serves all' case of a remote security manager, where one object serves all clients. An example of the one-to-one case is how a JEE container assigns a different remote *session bean* to each active client.

There may be other problems related to the use of the faade pattern, since hiding classes behind faade may restrict the possibilities of their reuse. Kiczales and Lamping have shown in [12] that designers and users of class libraries may, inadvertently/abusively, rely on implementation details *inherent in class/subclass relationships* within class libraries, which tightly couples 'client programs' (class library users) to implementation details that are meant to be hidden, and that library developers are not bound to respect when they evolve the library, which may break existing code. A similar situation may occur with faades.

5 Issues in Refactoring

Now that we have identified functional clusters as candidate services (see Section 3), and that we have wrapped that functionality behind services interfaces (see Section 4), we need to refactor existing applications that *use* that cluster of functionality, so that they invoke the functionality through the newly minted service interfaces. This raises a number of issues, which we will touch upon below.

The first issue has to do with the scope of the refactoring. Assume that the method $T f(T_1 p_1, \dots, T_i p_i)$ of class C has been identified as part of a service S that was packaged under a faade, as explained in Section 4.2. Should I consider *all* invocations of $f(\dots)$ as being part of a service call, and reroute that call through the service interface? That is a good question. For one thing, the class C could be part of several services, or could be used both within and outside the service. If we hide C behind the service interface, could clients manage? if we decide *not* to include all of the methods of C in the service interface of S , then a first-cut decision could be to route calls to service-included methods to the service, and leave the others untouched. Even so, there may be (definitely *is*) an overhead in calling a service interface, which could be remote, as opposed to manipulating the C class directly. If the class C is part of some library, then we are in trouble. For example, a business service that computes the yearly repayment of a loan of based on some capital amount, a yearly interest rate, and an amortization period, might expose a 'utility' method that computes the compound interest rate over a number of periods. Should have every call to that method go through the service interface, or should I make it possible to invoke the method independently of the service?

Part of the answer *may lie* in the invocation pattern of the caller: what other methods of the service, if any, does it call, in addition to method $f(\dots)$. Is there a typical invocation sequence that uniquely characterizes the functionality of the service? For example, a flight booking service may offer methods to, 1) search for flights (call it `searchFlights(...)`), 2) to get the details of a flight (`getFlightDetails(...)`), and 3) to book a seat at a particular flight (`bookFlight(...)`). I can imagine the `searchFlights(...)` method as being part of an aggregation service that does price

comparison. When I find a call to `searchFlights(...)`, should I consider it as an invocation of the booking service, or of the comparative shopping service. Part of the answer may lie in which other methods were called before or after. Thus, to be able to characterize a method call as being part of the invocation of a service, I may need to check whether other functions of the service (the faade) are invoked, within a certain scope from the first call. Dynamic traces would make this pattern detection easier—provided that we can characterize the typical, or expected, or permissible invocation sequences.

Once I have determined that a client call to a method that was included in a service is a call to that service, I need to focus on the mechanics of the code transformation itself. Depending on what *flavor* of service-orientation we use, but the refactoring may involve a *paradigm shift*. For example, in going from Java EE to SOA, I might chose an event-based implementation of SOA, in which case I needed to migrate some functional code from a *call-and-return* style, to an asynchronous *event-based style*. This presents its own challenges:

- Going from structural/flow-based functional composition to data-level composition:
 - with structural/flow based composition, function calls are hardcoded in the program—reflection notwithstanding
 - with data-level composition, composition happens through shared messages.

Anyone who has coded variants of the *observer* pattern knows the challenges of implementing it in such a way as to communicate *complex* changes that may involve *complex* data structures/objects, and doing so through a simple interface¹².

- Going from a synchronous call style to an asynchronous call style. This involves two challenges, 1) at the algorithmic level, 2) the architectural level.

At the algorithmic level, we illustrate the kind of changes that we will need to make. Consider the function `f(...)` that invokes functions `g(...)`, `h(...)`, `i(...)`, which are now part of service interfaces. We will ignore object-orientation because the issues are the same, whether we use procedural code or object-oriented code.

```
public T1 f(T2 x) {
    T3 y = g(x);
```

¹²Coding the `changed(...)` method of the `Observable` requires some ingenuity to communicate complex changes and complex data through two parameters, as in `this.changed(changeAspect, changeData)`. A realistic implementation of the `Observer` either requires different variants of the `changed(...)` method, or a fairly complex representation of the change data, which tightly couples the `Observer` which needs to know how to unpack the change data for specific change aspects, defeating the purpose of the *observer* pattern

```
    T4 z = h(y);
    T1 u = i(z);
    return u;
}
```

Here, we have a simple composition where the output produced by $g(\cdot)$ is used as an input to $h(\cdot)$, whose output is used as input to $i(\cdot)$, whose output is returned by $f(\cdot)$. Thus $f(x) = i(h(g(x)))$. If we were to transform $f(\cdot)$ into a message driven style (e.g. *la* Java Messaging Service), we would have each of $g(\cdot)$, $h(\cdot)$, and $i(\cdot)$ receive their inputs through named input queues, and put their outputs on named output queues. To handle the coordination between them, we would then need an *message/event listener*, then listens for messages arriving on queues, and takes the appropriate action. Thus, the code for performing this composition would look something like the following:

```
public void onMessage(aMessage) {
    // first get the message type to figure
    // out where to send it next
    int messageType = aMessage.getType();

    // then switch on the message type
    switch (messageType) {
        case F_REQUEST:
            // unpack aMessage and get x value

            x = aMessage.get();

            // create an G message to send x to G next
            anG_Message = createG_Message(x);

            // put anG_Message on G in_queue;
            ...
            return;
        case G_RESPONSE:
            // unpack aMessage and get y value

            y = aMessage.get();

            // create an H message to send y to H next
            anH_Message = createH_Message(y);

            // put anH_Message on H in_queue;
            ...
            return;
    }
}
```

```
case H_RESPONSE:
    // unpack aMessage to get the z value
    ...
    z = aMessage.get(...)

    // create an I message to send z to I next
    anI_Message = createI_Message(z);

    // put anI_Message on I in_queue
    ...
return
```

This example illustrates how cumbersome it is to implement a simple sequence in the JMS style. Typically, the functions that are accessible through message queues tend to be of relatively coarse granularity. From an architectural point of view, if you use an asynchronous/message-driven style, at a particular level of functional aggregation, it pretty much imposes that style at the *levels above it*. In this case, it means that if $f(\dots)$ is implemented in a message-oriented way, then its callers will *likely* need to be, themselves, message oriented¹³.

Of course, you are not going to code a sorting algorithm this way. But this raises interesting issues in terms of execution sequence/trace equivalence, and error handling, among other things.

6 Conclusion

In this paper, we discussed issues and the state of the research related to the re-engineering of legacy JEE applications into service oriented ones. We argued that this migration requires that we perform three steps:

1. Identify clusters of functionality within the legacy application(s) that qualify as candidate services
2. Package these functional clusters into a service-like packaging so that enterprise applications can invoke these functionalities using the service interfaces, and
3. Refactor existing applications so that they invoke these 'newly minted' services using their service interfaces, thereby increasing their future maintainability.

¹³If the caller of $f(\dots)$ does not care about the output produced by it, but just needs to make sure that it is called, then it need not be asynchronous/message oriented. I can have a blocking function call launch an independent thread and return, considering its job is done. That thread ($f(\dots)$ in this case) will eventually terminate and do whatever needed to be done. But if the caller needs the output of $f(\dots)$ to continue, then it *too* needs to be message-oriented

It is fair to say that these steps received attention in the literature by decreasing order.

A lot of work has been performed on component and service identification. But like we mentioned in section 3, most of the existing work looks for services based on their functional cohesion and low coupling with other parts of business applications, regardless of service types. We argued that different service types have different code smells/signatures in legacy code, and hence, service identification methods should take into account service types. Thus, we proposed a service taxonomy in section 3.2, and proposed a preliminary set of service-type specific identification criteria in section 3.4. As our approach relies on source code, it becomes important to identify all of the static dependencies within legacy applications. We identified a number of issues related to the static code analysis of JEE applications, including the fact that they are multi-tiers, multi-language, their reliance on a number of 'extra-lingual' mechanisms for linkage and bindings, and their reliance on frameworks and containers that hide some call dependencies. Thus, prior to looking for functional clusters that can qualify as candidate services, we need to map out all the dependencies that are inherent in a JEE application, beyond ones that static analysis of Java code can find.

With regard to service packaging, discussed in section 4, we identified a number of issues that need to be resolved, which can think of as *interface specification* (section 4.1, and *interface design* (see section 4.2). Roughly speaking, specification deals with ways to partition the set of functions offered or consumed by a functional cluster into a set of relatively independent and more strongly cohesive interfaces so that client programs may be exposed to—and use—those interfaces, as opposed to the full gamut of functions available. We discussed some potential approaches to do this. *Interface design* deals with the detailed design of individual service function *signatures*, and encompasses things such as granularity of function parameters, and lifecycle management of objects that live within the service.

With regard to the refactoring of existing applications, we barely started exploring the relevant issues. Basically, there are two sets of issues:

1. The client side consequence of service interface design. Indeed, as we hide functional parameters, service-side object creation and lifecycle management, client programs need to manipulate the service functionalities through the newly defined (and constraining) service interfaces
2. The general issue of changing architectural style in the migration process. In particular, we looked at the specific example of going from a call-and-return invocation style—prevalent in JEE applications—to a message-oriented style¹⁴, to identify some of the relevant issues.

In this report, we raised more questions than we answered, and that was the intent. Our priority is on identifying the problems. In those cases where our thinking has

¹⁴Technically speaking, JMS and Message-Driven Beans are also part of JEE

progressed, we presented the first elements of an approach, but for the most part, all of the research problems are still fairly open.

More reports—and publications—will follow this report, by expanding on some of the issues, and proposing a solution approach, or elaborating on some of the elements presented here. We will not evolve this report: progress that we make on any topic will be described in separate reports.

References

- [1] “The open group soa reference architecture,” in http://www.opengroup.org/soa/source-book/soa_refarch/. [Online; accessed 19-July-2017].
- [2] S. Alahmari, E. Zaluska, and D. De Roure, “A service identification framework for legacy system migration into soa,” in *2010 IEEE International Conference on Services Computing*, pp. 614–617, 2010.
- [3] A. Belle, G. El Boussaidi, and H. Mili, “Recovering software layers from object oriented systems,” in *ENASE 2014 - Proceedings of the 9th International Conference on Evaluation of Novel Approaches to Software Engineering*, pp. 1–12, 2014.
- [4] S. Cai, Y. Liu, and X. Wang, “A survey of service identification strategies,” in *IEEE Asia-Pacific Services Computing Conference*, pp. 464–470, 2011.
- [5] S. Cohen, “Ontology and taxonomy of services in a service-oriented architecture,” *The Architecture Journal*, vol. 11, no. 11, pp. 30–35, 2007.
- [6] G. El Boussaidi, A. B. Belle, S. Vaucher, and H. Mili, “Reconstructing architectural views from legacy systems,” in *Proceedings - Working Conference on Reverse Engineering, WCRE*, pp. 345–354, 2012.
- [7] T. Erl, *SOA Principles of Service Design*. Upper Saddle River, NJ, USA: Prentice Hall PTR, 2007.
- [8] A. Fuhr, T. Horn, and V. Riediger, “Using dynamic analysis and clustering for implementing services by reusing legacy code,” in *18th Working Conference on Reverse Engineering*, pp. 275–279, 2011.
- [9] M. Grieger, S. Sauer, and M. Klenke, “Architectural restructuring by semi-automatic clustering to facilitate migration towards a service-oriented architecture,” *Softwaretechnik-Trends*, vol. 34, no. 2, 2014.
- [10] Q. Gu and P. Lago, “Service identification methods: A systematic literature review,” in *Third European Conference ServiceWave Proceedings*, pp. 37–50, Springer Berlin Heidelberg, 2010.

-
- [11] R. S. Huergo, P. F. Pires, F. C. Delicato, B. Costa, E. Cavalcante, and T. Batista, “A systematic survey of service identification methods,” *Serv. Oriented Comput. Appl.*, vol. 8, pp. 199–219, Sept. 2014.
 - [12] G. Kiczales and J. Lamping, “Issues in the design and specification of class libraries,” in *Conference Proceedings on Object-oriented Programming Systems, Languages, and Applications*, OOPSLA '92, (New York, NY, USA), pp. 435–451, ACM, 1992.
 - [13] M. Marin, L. Moonen, and A. Van Deursen, “A common framework for aspect mining based on crosscutting concern sorts,” in *Proceedings - Working Conference on Reverse Engineering, WCRE*, pp. 29–38, 2006.
 - [14] H. Mili, F. Mili, and A. Mili, “Reusing software: Issues and research directions,” *IEEE Trans. Softw. Eng.*, vol. 21, pp. 528–562, June 1995.
 - [15] H. Mili, A. Mili, S. Yacoub, and E. Addy, *Reuse-based Software Engineering: Techniques, Organization, and Controls*. New York, NY, USA: Wiley-Interscience, 2001.
 - [16] M. Robillard and G. Murphy, “Concern graphs: finding and describing concerns using structural program dependencies,” *Proceedings of the 24th International Conference on Software Engineering. ICSE 2002*, pp. 406–416, 2002.
 - [17] A. Shatnawi and al, “Towards end-to-end static call graph of complex j2ee web-applications,” june 2016.
 - [18] S. Suwisuthikasem and M. H. Samadzadeh, *Migration from Legacy Systems to SOA Applications: A Survey and an Evaluation*, pp. 609–614. Cham: Springer International Publishing, 2015.
 - [19] P. Tonella and M. Ceccato, “Aspect mining through the formal concept analysis of execution traces,” in *Proceedings - Working Conference on Reverse Engineering, WCRE*, 2004.
 - [20] Z. Zhang, R. Liu, and H. Yang, “Service identification and packaging in service oriented reengineering,” in *SEKE*, vol. 5, pp. 620–625, 2005.