This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications

Juan Manuel Rodriguez, Marco Crasso, Cristian Mateos, Alejandro Zunino, Marcelo Campo
ISISTAN Research Institute - UNICEN University
Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET)
Tandil, Buenos Aires, Argentina.
Corresponding Author's Email: mcrasso@conicet.gov.ar

***Abstract*—The Service-Oriented Computing (SOC) paradigm has gained a lot of attention in the software industry since represents a novel way of architecting distributed applications. SOC is mostly materialized via Web Services, which allows developers to structure applications as building blocks exposing a clear, public interface to their capabilities. Although conceptually and technologically mature, SOC still lacks adequate development support from a methodological point of view. We present the EasySOC project, a catalog of guidelines to build service-oriented applications and services. This catalog synthesizes best SOC development practices that arise as a result of several years of research in fundamental SOC-related topics, namely WSDL-based technical specification, Web Service discovery and Web Service outsourcing. In addition, we describe a plug-in for the Eclipse IDE that has been implemented to simplify the utilization of the guidelines. We believe that both the practical nature of the guidelines, the empirical evidence that supports them, and the availability of IDE support that enforces them may help software practitioners to rapidly exploit our ideas for building real SOC applications.**

***Keywords*-Service-Oriented Computing, Service-oriented development guidelines, Web Services, WSDL anti-patterns, Web Service discovery, Web Service consumption**

I. Introduction

Service-Oriented Computing (SOC) [1] is a relatively new computing paradigm that radically changed the way applications are architected, designed and implemented. SOC has mainly evolved from component-based software engineering by introducing a new kind of building block called *service*, which represents functionality that is described, discovered and remotely consumed by using standard protocols. Service-oriented software systems started as a more flexible and cost-effective alternative for developing Web-based applications, but their usage eventually spread to gave birth to a wave of contemporary infrastructures and notions including Service-Oriented Grids and Software-As-A-Service [2].

The common technological choice for materializing the SOC paradigm is Web Services, i.e. programs with well-defined interfaces that can be published, discovered and consumed by means of ubiquitous Web protocols [1] (e.g. SOAP [3]). The canonical model underpinning Web Services encompasses three basic elements: service providers, service consumers and service registries (see Figure 1). A service provider, such as a business or an organization, provides meta-data for each service, including a description of its technical contract in WSDL [4]. WSDL is an XML-based language that allows providers to describe the functionality of their services as a set of abstract operations with inputs and outputs, and to specify the associated binding information so that consumers can actually consume the offered operations.

To make their WSDL documents publicly available, providers employ a specification for service registries, called UDDI, whose central purpose is the representation of meta-data about Web Services. Apart from the data model, UDDI defines an inquiry API, in terms of WSDL, for discovering services. Consumers use the inquiry API to find services that match their functional needs, select one, and then consume its operations by interpreting the corresponding WSDL description. Both the model and the API are built on Web Service technologies, as the aim of WSDL and UDDI is to offer standards to enable interoperability among applications and services across the Web. As a consequence, for example, an application implemented in a programming language can talk to a Web Service developed in another language. Ideally, such interoperability levels would allow consumers to switch among different providers of the same functionality, according to non functional requirements such as cost per service consumption, response time or availability, without modifying the applications involved.

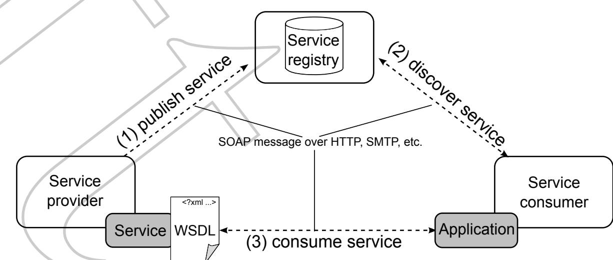

Figure 1: The Web Services model.

Unfortunately, the promises of Web Services of guaranteeing loose coupling among applications and services, providing agility to respond to changes in requirements, offering transparent distributed computing and lowering ongoing investments are still eclipsed by the high costs of outsourcing Web Services of current approaches for service-enabling applications as well as the ineffectiveness of Web Service publication systems. On one hand, unless appropriately specified by

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

providers, service meta-data can be counterproductive and instead obscure the purpose of a service, thus hindering its adoption. For example, a WSDL document without much comment of its operations can make the associated Web Service difficult to discover and understand. On the other hand, service consumers often have to invest much effort into providing code to invoke dicovered Web Services afterward. Moreover, the outcome of the second task is software containing service-aware code. Therefore, the software is more difficult to test and to modify during its maintenance phase.

In this paper, we describe EasySOC, a set of provider and consumer guidelines for avoiding these problems. Roughly, these guidelines represent a compilation of best practices for simplifying the activities illustrated by the arcs of Figure 1, while improving the quality of the artifacts implementing services and consumers' applications. EasySOC is based on previous research carried out by the authors in the subareas of WSDL-based technical contract design and specification [5], Web Service discovery [6] and service-oriented development and programming [7, 8].

Complementary, the contribution of this paper is to provide a uniform, conceptualized and synthesized view of these findings to provide, on one hand, clear and precise hints of how to adequately exploit the SOC paradigm and its related technologies regardless their context of usage, i.e. when implementing services or applications. At the same time, another contribution is to delineate potential concrete materializations of these hints into a software tool so as to enforce the promoted best practices. With respect to the latter, we have built a plug-in for the popular Eclipse IDE and the Java language, thus we believe our ideas can be readily employed in the software industry [9]. The software can be downloaded from http://sites.google.com/site/easysoc.

The rest of the paper is structured as follows. The next section focuses on discussing the aforementioned guidelines, emphasizing on clarifying their scope and the usage scenarios in which they are applicable. Later, Section III presents the EasySOC Eclipse plug-in and its modules. Then, Section IV surveys relevant related efforts. Finally, Section V concludes the paper.

II. The EasySOC project

Even when Web Service technologies are far more mature and reliable than they were years ago, the definition of guidelines for developing service-oriented software is still an incipient research topic. Thus, the following paragraphs present a catalog of identified best practices for SOC development, which are related to the roles and activities that are commonly performed by developers of both services and consumers' applications. Schematically, according to the model of Figure 1, two distinctive roles are established: providers and consumers. Providers are responsible for making a piece of software publicly available as a Web Service, while ensuring that such a service can be discovered and understood by third-parties. Consumers are responsible for discovering and incorporating external services into their applications, or from now on *client applications*. Sometimes the same actors can play both roles, as occurs when developing services that need of other services to accomplish the functionality they expose.

Depending on the role(s) played by a SOC developer, there are three possible different development scenarios. Table I lists these scenarios by relating them to the EasySOC guidelines that developers are encouraged to pay attention to.

Table I: SOC usage scenarios and the EasySOC guidelines.

Scenario/Guidelines	Guidelines for service publication	Guidelines for service discovery	Guidelines for service consumption
Developer only exposes a functionality as a service	Yes	No	No
Developer only consumes services	No	Yes	Yes
Developer exposes as a service a functionality that consumes other services	Yes	Yes	Yes

A. Guidelines for improving service descriptions

Many of the problems related to the efficiency of standard-compliant approaches to service discovery stem from the fact that the WSDL specification is incorrectly or partially exploited by providers. Although the intuitive importance of properly describing services, some practices that attempt against the discoverability of services, such as poorly commenting offered operations or using unintelligible naming conventions, are frequently found in publicly available WSDL documents.

There is no silver bullet to guarantee that potential consumers of a Web Service will effectively discover, understand and access it. However, we have empirically shown that a WSDL document can be improved to simultaneously address these issues by following six steps [10] (we assume the reader is familiar to the WSDL basics):

1) Separating the schema –i.e. XSD code– from the definitions of the offered operations.
2) Removing repeated WSDL and XSD code.
3) Putting error information within Fault messages and only conveying operation results within Output ones.
4) Replacing WSDL element names with self-explanatory names if they are cryptic.
5) Moving non-cohesive operations from their port-types to a new port-type.
6) Properly commenting the operations.

The first step means moving complex data-type definitions into a separated XSD document, and adding the corresponding import sentence into the WSDL document. However, when data-types are not going to be reused or are very simple, they can be part of the WSDL document to make this latter self-contained.

The second step deals with redundant code in both the WSDL document and the schema. Repeated WSDL code

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

usually stem from port-types tied to a specific invocation protocol, whereas redundant XSD is commonly a result from data definitions bounded to a particular operation. Therefore, repeated WSDL code can be removed by defining a protocol-independent port-type. Similarly, to eliminate redundant XSD code, repeated data-types should be abstracted into a single one. This change must be consequently made visible in operation messages by updating their data-types so as to reference the newly derived types.

The third step encourages to separate error information from output information or service invocation results. To do this, error information should be removed from Output messages and placed on Fault ones, a special construct provided by WSDL to specify errors and exceptions. Moreover, as many Fault messages as kinds of errors exist should be defined for the operations of the Web Service.

The fourth step aims to improve the representativeness of WSDL element names by renaming non-explanatory ones. Grammatically, the name of an operation should be in the form <verb> "+" <noun>, because an operation is essentially an action. Furthermore, message, message part and data-type names should be a noun or a noun phrase because they represent the objects on which the operation executes. Additionally, the names should be written according to common notations, and their length should be between 3-15 characters because this facilitates both automatic analysis and human reading, respectively. With respect to the former hint, the name "theelementname" should be rewritten for example as "theElementName" (camel casing).

The fifth step is to place operations in different port-types based in their cohesion. To do this, the original port-type should be divided into smaller and more cohesive port-types. This step should be repeated while the new port-types are not cohesive enough.

Finally, all operations must be well commented. An operation is said to be well commented when it has a concise and explanatory comment, which describes the semantics of the offered functionality. Moreover, as WSDL allows developers to comment each part of a service description separately, then a very good practice is to place every <documentation> tag in the most restrictive ambit. For instance, if the comment refers to a specific operation, it should be placed in that operation.

It is worth noting that, except for steps 4 and 6, the other steps may require to modify service implementations. Moreover, as a result of applying these guidelines, there will be two versions of a revised service description. Despite being out of the scope of this article, some version support technique is necessary to allow service consumers that use the old service version to continue using the service until they migrate to the new version.

B. *Guidelines for making effective queries*

Queries play an important part in the process of service discovery since service consumers may greatly benefit from generating clear and explanatory descriptions of their needs. This is because the underpinnings of UDDI-based registries rely upon the descriptiveness of the keywords conveyed in both publicly available service interfaces and queries.

We have empirically proved that the source code artifacts of client applications may carry relevant information about the functional descriptions of the potential services that can be discovered and, in turn, consumed from within applications [6]. The idea is that developers should be focused on building the logic of their applications, while using automatic heuristics to pull out keywords standing for queries of the required services from the code implementing such applications. In this line, best practices for building an application that contains useful information about the services the application needs comprise [6]:

1) Defining the expected interface of every application component that is planned not to be implemented but outsourced to a Web Service.
2) Revising the functional cohesion between the implemented (i.e. internal) components that directly invoke, and hence depend on the interfaces of, the components defined in step 1.
3) Naming and commenting each defined interface and internal component by using self-explanatory names and comments, respectively.

The first step encourages developers to think of a third-party service as any other regular component providing a clear interface to its operations. The idea of defining a functional interface before knowing the actual exposed interface of a service that fulfills an expected functionality aligns with the *Query-by-Example* approach to create queries. This approach allows a discoverer to search for an entire piece of information based on an example in the form of a selected part of that information. This concept suggests that because of the structure inherent to client applications and Web Service descriptions in WSDL, the expected interface can be seen as an example of what a consumer is looking for. This is built on the fact that, via WSDL, publishers can describe their services as object-oriented interfaces with methods and arguments. Therefore, in the context of client applications, the defined interfaces stand for *examples*.

The second step bases on an approach for automatically augmenting the quantity of relevant information within queries called *Query Expansion*, which relies on the expansion of extracted examples by gathering information from the source code representing internal components that directly interact with the interfaces representing external services. The reasoning that supports this mechanism is that expanding queries based upon components with strongly-related and highly-cohesive operations should not only preserve, but also improve, the meaning of the original query. Therefore, the second step deals with ensuring that defined interfaces are strongly-related and highly-cohesive with those components that depend on them.

The above two steps deal with identifying the source code parts of an application that may contain relevant information for generating queries and discovering Web Services afterward. Also, a third step exists for checking that the

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

identified code parts actually have relevant information for that purpose. Since automatic query generation heuristics gather keywords from operation names and comments present in source codes, developers should follow conventional best practices for naming and commenting their code.

C. Guidelines for shielding applications from service specifics

Maintaining client applications can be a cumbersome task when they are tied to specific providers and WSDL documents. The common approach to call a Web Service from within an application is by interpreting its associated WSDL document with the help of invocation frameworks such as WSIF, CXF [11], or the .NET Web Services Description Language Tool (WSDL.exe)[1]. These frameworks succeed in hiding the details for invoking services, but they still fail at isolating internal components from the interfaces of the services. Consequently, applications result in a mix of pure logic and sentences for consuming Web Services that depend on their operation signatures and data-types. This approach leads to client applications that are subordinated to third-party service interfaces and must be modified and/or re-tested every time a provider introduces changes. In addition, this also hinders service replaceability, which means how easily a service could be replaced with another functionality equivalent service.

In this sense, we have shown that the maintenance of service-oriented applications can be facilitated by following certain programming practices when outsourcing services [8]:

1) Defining the expected interface of every component that is planned to be outsourced.
2) Adapting the actual interface of a selected service to the interface that was originally expected, i.e. the one defined in the previous step.
3) Seamlessly injecting adaptation code into each internal component that depends on the expected interface.

Step 1 provides a mean of shielding the internal components of an application from details related to invocating third-party services. To do this, a functionality that is planned to be implemented by a third-party service should be programmatically described as an abstract interface. Note that this is the same requirement as the first step of Section II-B. Accordingly, internal application components depending on such an abstractly-described functionality consume the methods exposed by its associated interface, while adhering to operation names and input/out data-types declared in it.

The second step takes place after a service has been selected. During this step, developers should provide the logic to transform the operation signatures of the actual interface of the selected service to expected the interface defined previously. For instance, if a service operation returns a list of integers, but the interface defined at step 1 returns an array of floats, the developer should code a service adapter that performs the type conversion. By properly accomplishing steps 1 and 2, client components depend on neither specific service implementations nor interfaces. Therefore, from the perspective of the application logic, services that provide equivalent functionality can be transparently interchangeable at the expense of building specific adapters.

Finally, the third step is for separating the functional code of an application from configuration aspects related to binding a client component that depends on an interface with the adapter component in charge of adapting it into a selected service. A suitable form of doing this, in terms of source code quality, involves delegating to a software layer or container the administrative task of assembling interfaces, internal components and services together.

In the following section we describe a software tool, implemented as a plug-in for the Eclipse IDE, which enforces the aforementioned guidelines for developing SOC applications and Web Services written in Java.

[1].NET WSDL.exe, http://msdn.microsoft.com/en-us/library/7h3ystb6(v=VS.71).aspx

III. THE EASYSOC ECLIPSE PLUG-IN

The EasySOC Eclipse plug-in comprises three modules, each one materializing the set of guidelines explained before. Sections III-A through III-C discuss the design and implementation of these modules.

A. WSDL discoverability Anti-Patterns Detector

The Anti-patterns Detector is the EasySOC module for automatically checking whether the WSDL document describing the technical contract of a Web Service conforms to the guidelines of Section II-A. The module receives this name since its construction was driven by the catalog of WSDL document discoverability anti-patterns that we introduced in the study published in [10]. Besides measuring the impact of each anti-pattern on service discovery, the study assessed the implications of anti-patterns on developers' ability to make sense of WSDL documents. The catalog consists of eight anti-patterns and provides a name, a problem description, and a sound refactoring procedure for each anti-pattern. Although the results of the study motivate anti-patterns refactoring, manually looking for an anti-pattern in WSDL documents might be a time consuming and complex task. Thus, the Anti-patterns Detector comprises heuristics to automatically detect the anti-patterns in the aforementioned catalog.

Since those heuristics are based on the anti-pattern definition, they can be classified according to the analysis required to detect the anti-patterns. In this sense, a taxonomy comprising types of anti-patterns was derived [10]. Basically, antipatterns can be divided into two categories: those that can be detected by analyzing only the structure of a WSDL document, and anti-patterns whose detection requires a semantic analysis of the names and comments in the WSDL document.

The heuristics to detect the first kind of anti-patterns are simple rules based on the commonest anti-pattern occurrence form. Examples of such anti-patterns are redundant XML code for defining both data-types and port-types, data-types embedded in a WSDL document, and data-types that allow transferring data of any type (*Whatever types* in EasySOC

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

terminology). Moreover, the second problem difficults data-type reuse, and the third one hinders understandability. The rule that detects redundant port-types verifies that any pair of port-types has the same number of operations and that they are equally named. In this case, the heuristic does not verify the similarity between the messages of the port-types because they are likely to change in accord with the underlying binding protocol.

As it was previously mentioned, detecting the remaining anti-patterns requires analyzing the semantics of names and comments. Basically, there are three problems that are detected by the associated heuristics: two naming issues, namely operations of different domains in the same port-type, and fault information within standard output messages. Firstly, our heuristic deals with names being too short or too long, using a rule to check that each name has a length between 3-30 characters is provided. Moreover, our heuristic concerns name structure, i.e. message part names should be nouns or noun phrases, while operations should be named with a verb plus a noun. This is verified by using a probabilistic context free grammar parser. For example, Figure 2 depicts the parsing trees of different message part names generated by the parser. The first and second names do not present problems, whereas the third name does because it starts with a verb.

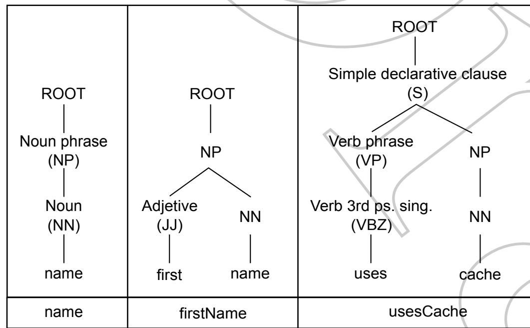

Figure 2: Parsing trees of message part names.

Secondly, our heuristic for determining whether two operations belong to the same domain is based on a text classification technique, because the only "semantic" information an operation of a WSDL document provides consists of the names and comments of the operation. In particular, the Rocchio's TF-IDF classifier has been selected because empirical studies show that it outperforms other classifiers in the Web Services area [8]. Rocchio's TF-IDF represents textual information as vectors, in which each dimension stands for a term and its magnitude is the weight of the term related with the text. Having represented all the textual information of a domain as vectors, the average vector, called centroid, is built for representing the domain. Then, the domain of an operation is deduced by representing it as a vector and comparing it to each domain centroid. Finally, the domain associated with the most similar average vector is returned as the domain of that operation.

Finally, our heuristic for detecting error information within output messages checks whether an operation has no fault

Table II: Anti-pattern detection: Confusion matrixes.

Automatic detection results per anti-pattern		Manual detection results	
		Negative	Positive
Enclosed data model	Negative	116	6
	Positive	0	270
Redundant port-types	Negative	161	4
	Positive	0	227
Redundant data models	Negative	221	2
	Positive	3	166
Whatever types	Negative	339	0
	Positive	3	50
Lack of comments	Negative	135	0
	Positive	0	257
Low cohesive operations in the same port-type	Negative	272	10
	Positive	78	32
Ambiguous names	Negative	67	0
	Positive	9	316
Undercover fault information within standard messages	Negative	351	3
	Positive	4	34

message defined, and the comment or some name related with the output contains one of the following words, which are commonly related with error conditions while executing a service: *error, fault, fail, exception, overflow, mistake* and *misplay*.

These heuristics have been experimentally validated with a real-world data-set, showing an averaged accuracy of 98.5%. The methodology followed in the evaluation first involved manually analyzing each WSDL document to identify the anti-patterns it has, peer-reviewing manual results afterwards (at least three different people reviewed each WSDL document), automatically analyzing WSDL documents based on the proposed heuristics, and finally comparing both manual and automatic results. Results were organized per anti-pattern, and if a WSDL document has an anti-pattern it is classified as "Positive", otherwise it is classified as "Negative". When the manual classification for a WSDL document is equal to the automatic one, it means that the heuristic accurately operates for that WSDL document. Achieved results are shown in Table II by using a confusion matrix. Each row of the matrix represents the number of WSDL documents that were automatically classified using the heuristic associated with a particular anti-pattern. In addition, the columns of the matrix show the results obtained manually, i.e. the number of WSDL documents that actually had each anti-pattern.

In the experiments, we used a data-set of 392 WSDL documents [10]. Once each heuristic was applied on this data-set, we built the confusion matrixes. Then, we assessed the accuracy and false positive/negative rates for each matrix. The accuracy of each heuristic was computed as the number of classifications matching over the total of analyzed WSDL documents. For instance, the accuracy of the *Redundant data model* heuristic was $\frac{221+166}{221+2+3+166} = 0.987$. The heuristic for

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

detecting *Low cohesive operations within the same port-type* anti-pattern achieved the lowest accuracy: 0.775. One hypothesis that could explain this value relates to potential errors introduced by the classifier, thus more experiments are being conducted. Nevertheless, the average accuracy for all the heuristics was 0.958.

The false positive rate is the proportion of WSDL documents that an heuristic has wrongly labeled as having the corresponding anti-pattern. In opposition, the false negative rate is the percentage of WSDL documents that an heuristic has wrongly labeled as not having the corresponding anti-pattern. A false negative rate equals to 1.0 means that a detection heuristic has missed all anti-pattern occurrences. Therefore, the lower the achieved values the better the detection effectiveness. The average false positive rate was 0.036, and the average false negative rate was 0.052, which we believe are encouraging.

B. Query Builder

The EasySOC Query Builder module gathers information to build service queries from the source code of client applications. This module provides a graphical tool that guides consumers through the generation of the queries. When a consumer selects "Find services for..." by clicking on an interface that stands for an external service to be outsourced (see Figure 3 step 1), a wizard dialog starts. The wizard uses the Eclipse JDT Search Engine[2] for automatically discovering the client-side classes that depend on the interface and presents them to the user. Then, the user may select or discard the resulting classes (Figure 3 step 2). Similarly, the wizard presents a list of argument classes. This list is automatically built by analyzing the interface to retrieve the class names associated with each argument. If an argument is neither primitive (e.g., int, long, double, etc.) nor provided within a built-in Java library package (e.g., Vector, ArrayList, String, etc.) it is included in the list of argument classes. Finally, those manually selected classes along with the Java interface are used as input for the text-mining process depicted in the center of Figure 3. The module allows users to customize queries and test the retrieval effectiveness when using different classes as input, making query building interactive or semi-automatic. Alternatively, by clicking on the "Finish" button, the wizard selects all target classes on behalf of the consumer, making query expansion fully automatic.

We evaluated the retrieval effectiveness of the Query Builder by using the previous collection of 392 WSDL documents to feed a service registry [6]. Moreover, undergraduate students played the role of service consumers in the context of the "Service-Oriented Computing"[3] course of the Systems Engineering BSc. program at the Faculty of Exact Sciences (Department of Computer Science) of the UNICEN. The students were assigned an exercise consisting on deriving 30 queries, in which each query comprised a Java interface describing the functional capabilities of a potential service. The header and the operations of each interface were commented. For those operations with non-primitive data-types as arguments, their corresponding classes were also commented. Then, for each query, the students implemented and commented the internal components that depended on the interface. This methodology allowed us to separately evaluate five combinations of different sources of terms associated with an *example*, namely its "Interface", "Documentation", "Arguments" and "Dependants". Finally, a fifth alternative was used by combining all these four sources. In this context, documentation does not refer to extra software artifacts but to textual comments embedded within the classes.

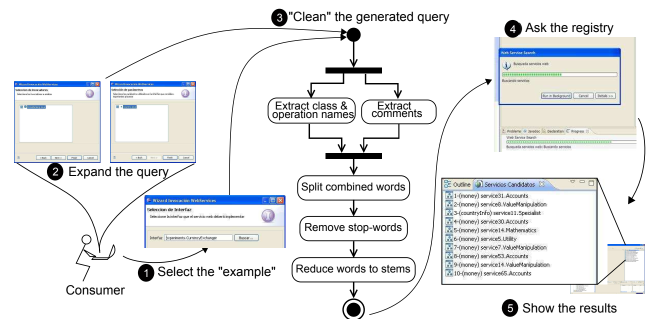

Figure 3: Wizzard for generating service queries.

To evaluate the discovery performance resulted from employing the different sources of terms, we used the Precision-at-*n*, Recall-at-*n*, *R*-precision and Normalized Recall (NR) information retrieval metrics. In this sense, the goal was to evaluate our Query Builder in terms of the proportion of relevant services in the retrieved list and their positions relative to non-relevant ones. We applied each metric for the 30 queries by individually using each one of the combination of sources (a total of 150 experiments per measure), and then we averaged the results over the 30 queries. As some of these metrics require to know the set of all services in the collection that are relevant to a given query, we exhaustively analyzed the data-set to determine the relevant services for each query. An important characteristic regarding the evaluation is the definition of "hit", i.e. when a returned WSDL document is actually relevant to the user. We judged a WSDL document as being a hit or not depending on whether its operations fulfilled the expectations previously specified in the Java code. For example, if the consumer required a Web Service for converting from Euros to Dollars, then a retrieved Web Service for converting from Yens to Dollars was not considered relevant, even though these services were strongly related. In this particular case, only Web Services for converting from Euros to Dollars were relevant. Note that this definition of hit makes the validation of our discovery mechanism very strict. Additionally, it is worth noting that for any query there are, at most, 8 relevant services within the data-set. Besides, there are 10 queries that have associated only one relevant service.

Each bar in Figure 4 stands for the averaged metric results that were achieved using a particular query expansion

[2] Eclipse Java Development Tools (JDT), http://www.eclipse.org/jdt
[3] http://www.exa.unicen.edu.ar/~cmateos/cos

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

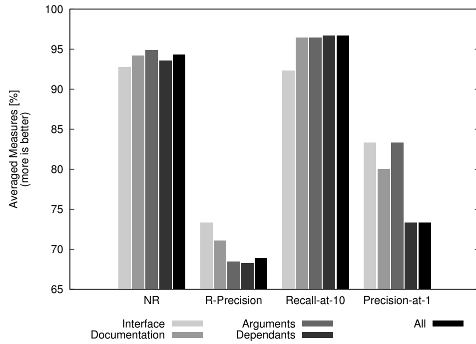

Figure 4: EasySOC Query Builder: Retrieval effectiveness.

alternative. Achieved results pointed out that by following the conventional Query-By-Example approach to build queries (the alternative named “Interface”) query-specific results were ranked first. When using more general, elaborated queries via the Query Expansion approach (e.g. the “All” alternative), the chance of including a relevant service at the top of the list decreased as the possibilities of including it before the 11^{th} positions increased. All in all, for this experiment our Query Builder alleviated discovery by concentrating relevant services within a window of 10 candidates.

C. Service Adapter

The Service Adapter module has been created to automatically perform the steps 2 and 3 for the guidelines of Section II-C. Once a consumer has selected a candidate service, this module performs three different tasks to adapt service interfaces and assemble internal components to it. The first task builds a proxy for the service. Second, the module builds an *adapter* to map the interface of the proxy onto the abstract interface internal components expect. Third, the module indicates a container how to assemble internal components and adapters, which is done through Dependency Injection (DI), a popular pattern for seamlessly wiring software components together that is employed by many development frameworks. Figure 5 summarizes the steps that are needed to proxy, adapt, and inject services into applications or another service implementation.

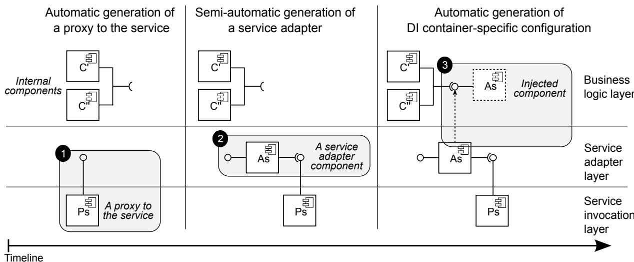

Figure 5: EasySOC steps for outsourcing services.

The current implementation of the Service Adapter module uses the Axis2 Web Service library for building service proxies, and Spring as the container supporting DI. Building a proxy with Axis2 involves giving as input the interface description of the target service (a WSDL document) to a command line tool. To setup the DI container, the names of dependant components and services must be written in an XML file. For adapting external service interfaces to the expected ones, we have designed an algorithm based on the work published in [12].

Our algorithm takes two Java interfaces as input and returns the Java code of a service adapter. To do this, it starts by detecting to which operations of one interface should be mapped the operations offered by the other. The algorithm assesses operation similarity by comparing operation names, comments, data-types and argument names. Data-type similarity is based on a pre-defined similarity table that assigns similarity values to pairs of simple data-types. The similarity between two complex data-types is calculated in a recursive way. Once a pair of operations has been determined, service adapter code is generated. The algorithm adapts simple data-types by taking advantage of type hierarchies and performing explicit conversions, i.e. castings. Complex data-types are resolved recursively as well. Clearly, not all available mismatches can be covered by the algorithm. Therefore, developers should revise the generated code.

In order to quantify the source code quality resulting from employing our plug-in, we conducted a comparison with the more traditional way of consuming Web Services, in which coding the application logic comes after discovering and knowing the description of the external services to be consumed. Basically, we used these two alternatives for developing a simple, personal agenda by outsourcing services from a given data-set comprising several services offering similar functionality but exposed by different providers.

After implementing the two variants, we randomly picked one service already included in the applications and we changed its provider. Then, we took metrics on the resulting source codes to have an assessment of the benefits of the EasySOC guidelines for software maintenance with respect to the traditional approach. To this end, we employed the well-known SLOC (Source Lines Of Code), Ce (Efferent Coupling), CBO (Coupling Between Objects) and RFC (Response For Class) software engineering metrics.

Table III: Personal agenda: Source code metrics.

Variant		Id	SLOC	Ce	CBO	RFC
Initial Web Service providers	Traditional	T_1	242	7	4.50	30.00
	EasySOC	E_1	309	7	1.70	7.20
Alternative Web Service providers	Traditional	T_2	246	10	4.67	22.67
	EasySOC	E_2	327	10	2.00	7.45

Table III shows the resulting metrics values for the four implementations of the personal agenda: traditional, EasySOC, and two additional variants in which another provider for a service was chosen from the Web Service data-set. For convenience, we labeled each implementation with an identifier (*id* column), which will be used through the rest of the paragraphs

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

of this section. To perform a fair comparison, a uniform formatting standard for all source codes was employed, Java import statements within compilation units were optimized, and the same tool to generate the underlying Web Service proxies was used.

From Table III, it can be seen that the variants using the same set of service providers resulted in equivalent Ce values: 7 for T_1 and E_1, and 10 for T_2 and E_2. This means that the variants generated via our plug-in (E_x), did not incur in extra efferent couplings with respect to the traditional variants (T_x). Moreover, if we do not consider the corresponding service adapters, Ce for the EasySOC variants drops down to zero, because relying on EasySOC effectively push the code that depends on service descriptions out of the application logic. Interestingly, the lower the Ce value is, the less the dependency between the application code and the Web Service descriptions is, which simplifies service replacement.

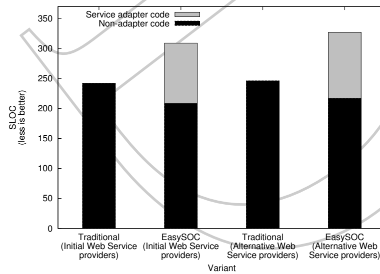

Figure 6: Source Lines of Code (SLOC) of the different applications.

Figure 6 shows the resulting SLOC. Changing the provider for a random service caused the modified versions of the application to incur in a little code overhead with respect to the original versions. The non-adapter classes implemented by E_1 were not altered by E_2 at all, whereas in the case of the traditional approach, the incorporation of the new service provider caused the modification of 17 lines from T_1 (more than 7% of its code).

The variants coded under EasySOC had an SLOC greater than that of the traditional variants. However, this difference was caused by the code implementing service adapters. In fact, the non-adapter code was smaller, cleaner and more compact because, unlike its traditional counterpart, it did not include statements for importing/instantiating proxy classes and handling Web Service-specific exceptions. Additionally, there are positive aspects concerning service adapters and SLOC. A large percentage of the service adapter code was generated automatically, which means programming effort was not required. Besides, changing the provider for the target service triggered the automatic generation of a new adapter skeleton, kept the application logic unmodified, and more importantly, allowed the programmer to focus on supporting the alternative service description only in the newly generated adapter class. Conversely, replacing the same service in T_1 involved the modification of the classes from which the service was accessed (i.e. statements calling methods or data-types defined in the service interface), thus forcing the programmer to modify more code. In addition, this practice might have introduced more bugs into the already built and tested application.

CBO and RFC metrics were also computed (Figure 7). Particularly, high CBO is undesirable, because it negatively affects modularity and prevents reuse. The larger the coupling between classes, the higher the sensitivity of a single change in other parts of the application, and therefore maintenance is more difficult. Hence, inter-class coupling, and specially couplings to classes representing (change-prone) service descriptions, should be kept to a minimum. Similarly, low RFC implies better testability and debuggability. In concordance with Ce, which resulted in greater values for the modified variants of the application, CBO for both the traditional approach and EasySOC exhibited increased values when changing the provider for a service. RFC, on the other hand, presented a less uniform behavior.

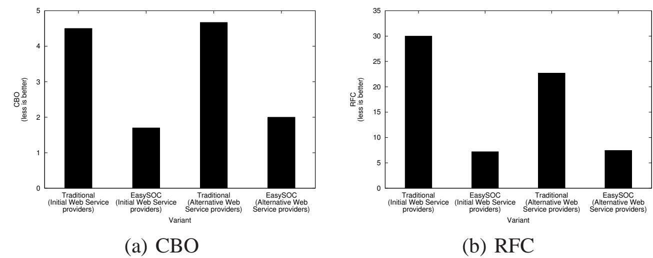

(a) CBO (b) RFC

Figure 7: Coupling Between Objects (CBO) and Response For Class (RFC) of the different applications.

As quantified by Ce, EasySOC did not reduce the amount of efferent couplings from the package implementing the application logic. Naturally, the reason of this is that the service descriptions to which E_x adhere are exactly the same as T_x. However, the EasySOC applications reduced the CBO with respect to the traditional implementations, because the access to the various services utilized by the application, and therefore their associated data-types, is performed within several cohesive compilation units (i.e. adapters) rather than within few, more generic classes. This in turn improves reusability and testability since application logic classes do not directly depend on services.

As depicted in Figure 7 (b), this separation also helped in achieving better average RFC. Moreover, although the plain sum of the RFC values of the E_x were greater compared to T_x, the total RFC of the classes implementing application logic (i.e. without taking into account adapter classes) were both smaller. This suggests that the pure application logic of E_1 and E_2 is easier to understand. In large projects, we reasonably may expect that much of the source code of EasySOC applications will be application logic instead of service adapters. Therefore, preserving the understandability of this kind of code is crucial.

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

IV. Related work

This work is somehow related to a number of preliminary methodologies that have emerged to address the demand for process guidance for SOC. These methodologies build upon existing techniques, such as EA and BPM, but also agile processes, such as XP and RUP [13]. However, it is out of the scope of this work to provide a SOC methodology. Instead, we aim at cataloging a set of proven best practices to design and implement both client applications and Web Services. It is worth to remark that the proposed guidelines have been followed to produce SOC-based software by playing both consumer and provider roles, thereby the collected empirical evidence supports that the proposed guidelines are indeed best practices. Additionally, another point of difference between our work and efforts like [14] is that, at least to the best of our knowledge, EasySOC is the first attempt towards a tool-aided *step-by-step* guide for materializing both Web Services and client applications.

V. Conclusions

The software industry is embracing the Service-Oriented Computing (SOC) paradigm as the premier approach for achieving integration as well as interoperability in heterogeneous, distributed computing environments. However, SOC presents many intrinsic challenges that both Web Service providers and consumers must unavoidable face.

Historically, catalogs of best practices have been widely recognized as a very valuable and helpful mean for software practitioners to deal with common problems in many different contexts. In this sense, this paper presented a set of concrete, proven guidelines for avoiding recurrent problems when developing Web Services and client applications. The proposed catalog is a set of 3 guidelines comprising 12 steps. One guideline covers 6 steps that service providers should take into consideration when exposing their services, in order to allow potential consumers to effectively discover, understand and access them. The other two guidelines cover aspects that service consumers should consider when discovering and consuming services. Regarding discovery, the proposed guideline consists of 3 steps that should be followed to easily build effective queries, which alleviates consumers' task by narrowing down the number of candidate services. With respect to service consumption, following the 3 steps of the associated guideline results in more maintainable code within client applications, but also in those services that invoke other services to accomplish their tasks.

The practical implications of each guideline have been corroborated experimentally, which suggests that the guidelines can be conceived as being best practices and can be readily employed in the software industry. In particular, we have assessed the impact of improving Web Service descriptions according to the corresponding guideline, by employing three registries simultaneously supporting service discovery and human consumers, who had the final word on which service is more appropriate. Results showed that improved descriptions are easier to understand than their "raw" counterpart [10]. Similarly, the positive effect on service discovery of the guideline for generating and expanding queries has been also measured [6]. Also, the implications on clients' maintenance of the corresponding guidelines have been formally and experimentally shown in [7] and [8] respectively.

Clearly, building truly loose coupled client applications using the corresponding guidelines imposes a radical shift in the way such applications are developed. This means that a company willing to employ EasySOC to start producing service-oriented applications would have to invest much time in training its development team, which results in a costly start-up curve. The impact of EasySOC on the software development process itself from an engineering point of view has been empirically assessed in [9]. Concretely, we performed further experiments to test the following hypothesis: understanding pervasive design patterns (i.e. Adapter and DI) and the "first build your application and then servify it" philosophy are the only required intellectual activities to start developing service-oriented applications with EasySOC, which should sharpen the associated learning curve. The hypothesis has been confirmed with 45 postgraduate and undergraduate students of the Systems Engineering program at the UNICEN during 2009. Results showed that they perceived that the proposed approach is convenient and thus may be easily adopted.

In the near future, we will conduct experiments with other students and real development teams to further validate our claims.

Acknowledgments

We thank the SCCC'10 chair Dr. Sergio F. Ochoa and the anonymous referees for their helpful comments to improve the paper. We also acknowledge the financial support provided by ANPCyT through grant PAE-PICT 2007-02311.

References

[1] J. Erickson and K. Siau, "Web Service, Service-Oriented Computing, and Service-Oriented Architecture: Separating hype from reality," *Journal of Database Management*, vol. 19, no. 3, pp. 42–54, 2008.

[2] M. Campbell-Kelly, "The rise, fall, and resurrection of software as a service," *Communications of the ACM*, vol. 52, no. 5, pp. 28–30, 2009.

[3] W3C Consortium, "SOAP version 1.2 part 1: Messaging framework." W3C Recommendation, http://www.w3.org/TR/soap12-part1, June 2007.

[4] T. Erl, *SOA Principles of Service Design*. Prentice Hall, 2007.

[5] M. Crasso, J. M. Rodriguez, A. Zunino, and M. Campo, "Revising WSDL documents: Why and how," *IEEE Internet Computing*, vol. 14, no. 5, pp. 30–38, 2010.

[6] M. Crasso, A. Zunino, and M. Campo, "Combining query-by-example and query expansion for simplifying Web Service discovery," *Information Systems Frontiers*, 2009. To appear.

This article is a pre-print of the article "J. M. Rodriguez, M. Crasso, C. Mateos, A. Zunino and M. Campo: "The EasySOC Project: A Rich Catalog of Best Practices for Developing Web Service Applications". JCC - INFONOR 2010. pp. 33-42. IEEE. ISBN 978-1-4577-0073-6."
The published version is available at http://dx.doi.org/10.1109/SCCC.2010.12

[7] C. Mateos, M. Crasso, A. Zunino, and M. Campo, "Separation of concerns in service-oriented applications based on pervasive design patterns," in *Web Technology Track (WT) - 25th ACM Symposium on Applied Computing (SAC '10)*, pp. 2509–2513, ACM Press, 2010.

[8] M. Crasso, C. Mateos, A. Zunino, and M. Campo, "EasySOC: Making Web Service outsourcing easier," *Information Sciences*, 2010. To appear.

[9] C. Mateos, M. Crasso, A. Zunino, and M. Campo, "An evaluation on developer's acceptance of EasySOC: A development model for Service-Oriented Computing," in *11th Argentine Symposium on Software Engineering (ASSE2010) - 39th JAIIO*, SADIO, 2010.

[10] J. M. Rodriguez, M. Crasso, A. Zunino, and M. Campo, "Improving Web Service descriptions for effective service discovery," *Science of Computer Programming*, vol. 75, no. 11, pp. 1001–1021, 2010.

[11] Apache Software Foundation, "Apache CXF: An Open Source Service Framework." http://cxf.apache.org, 2009.

[12] E. Stroulia and Y. Wang, "Structural and semantic matching for assessing Web Service similarity," *International Journal of Cooperative Information Systems*, vol. 14, no. 4, pp. 407–438, 2005.

[13] E. Ramollari, D. Dranidis, and A. Simons, "A survey of service oriented development methodologies," in *2nd European Young Researchers Workshop on Service Oriented Computing (YR-SOC 2007)*, 2007.

[14] M. Papazoglou and W.-J. van den Heuvel, "Service-oriented design and development methodology," *International Journal of Web Engineering and Technology*, vol. 2, no. 4, pp. 412–442, 2006.

Draft