Cohesion-Driven Decomposition of Service Interfaces Without Access to Source Code

Dionysis Athanasopoulos, Apostolos Zarras, George Miskos, Valérie Issarny, Panos Vassiliadis

▶ To cite this version:

Cohesion-Driven Decomposition of Service Interfaces Without Access to Source Code

Dionysis Athanasopoulos, Apostolos V. Zarras, *Member, IEEE,* George Miskos, Valerie Issarny, and Panos Vassiliadis *Member, IEEE*

Abstract—Software cohesion concerns the degree to which the elements of a module belong together. Cohesive software is easier to understand, test and maintain. Improving cohesion is the target of several refactoring methods that have been proposed until now. These methods are tailored to operate by taking the source code into consideration. In the context of service-oriented development, cohesion refers to the degree to which the operations of a service interface belong together. In this context, *we propose an approach for the cohesion-driven decomposition of service interfaces*. The very philosophy of services dictates that all that is exported by a service is the service specification. Hence, our approach for the cohesion-driven decomposition of service interfaces is not based on how the services are implemented. Instead, it relies only on information provided in the specification of the service interfaces. We validate the approach in 22 real-world services provided by Amazon and Yahoo. We show the effectiveness of the proposed approach, concerning the cohesion improvement and the size of the produced decompositions. Moreover, we show that the proposed approach is useful, by conducting a user study, where developers assessed the quality of the produced decompositions.

Index Terms—Cohesion, Decomposition, Service interface

✦

1 INTRODUCTION

Alice in the Web services world: Alice is an ordinary Java developer. Some time ago, she discovered the benefits of using Web services for developing software. Alice finds them very handy. As it is typically done, the applications that she develops access services via JAX-WS[1] proxies, generated from the WSDL specifications of the services interfaces. A JAX-WS proxy looks much like an ordinary Java class, but its methods delegate calls to service operations and bring the results back to the application.

However, using services also has its drawbacks. Often, new versions of the Web service interfaces are released and Alice spends quite some time to test and maintain her software, when this happens. For instance, one of the projects that Alice is involved relies on the Amazon SQS service[2] (Figure 1). This service provides functionalities for message-based communication. Messages are stored/retrieved to/from message queues, which are allocated on the Amazon Cloud. Apart from the messaging operations, the service provides operations for the management of queue attributes, the management of queue access grants and the management of message visibility timeouts. Since 2007, the main interface of the service

Fig. 1. A client application that relies on the `MessageQueue` interface.

has been changed more than 4 times [3]. Whenever the `MessageQueue` interface changes, the continuous integration development platform (CIDP) that is used in the project traces that the `MessageQueue` proxy has changed. Following, the CIDP rebuilds the whole application and retests all the classes since they depend on the changed proxy. This overall process takes too long. Worst, Alice spends much time on checking the built logs and the test results to find out which tasks went right, or wrong.

On the back of her head, Alice has an idea that could save her from this burden. The idea is to split the `MessageQueue` interface into a set of new interfaces and develop a corresponding set of surrogate classes that implement these interfaces (Figure 2). The methods of the surrogate classes would then call the actual `MessageQueue` operations, via the `MessageQueue` proxy. Making the application use the surrogate classes, instead of directly using the

- *D. Athanasopoulos, A. Zarras, G. Miskos and P. Vassiliadis are with the Department of Computer Science, University of Ioannina, Greece. E-mail: {dathanas, zarras, pvassil}@cs.uoi.gr, gmiskos@gmail.com*
- *V. Issarny is with the Inria research center of Paris-Rocquencourt, France E-mail: Valerie.Issarny@inria.fr*

1. download.oracle.com/otndocs/jcp/jaxws-2_0-fr-eval-oth-JSpec/
2. aws.amazon.com/sqs/

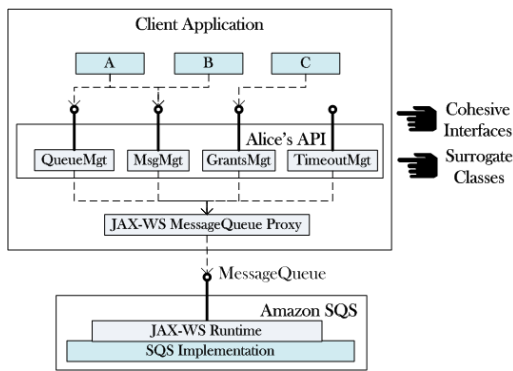3. aws.amazon.com/articles/Amazon-SQS/1148

Fig. 2. A client application that relies on a cohesive decomposition of the `MessageQueue` interface.

`MessageQueue` proxy, will decouple the constituent parts of the application from service operations that are not actually used in these parts. In this setting, changes to the `MessageQueue` interface shall affect certain surrogate classes. Then, only the parts of the application that use the affected surrogate classes will have to be re-built and re-tested. In fact, this idea would be useful for many others that use Amazon SQS. So, Alice plans to make her new interfaces and the surrogate classes that implement them available as an open source Java API. Alice thinks that the same idea could also be useful in the case of services that provide a large number of operations. Amazon EC2[4], for instance, provides 87 operations grouped in a single interface. Similarly, the Yahoo KeywordService[5] provides 34 operations, grouped in a single interface. The decomposition of such large interfaces could be used to develop APIs that provide the developers with higher-level views of what the services do.

Having in mind a larger community of developers that could benefit from her idea, puts Alice into deeper thoughts about the proper splitting of service interfaces. The decomposition of service interfaces should be done in a principled way. Alice recalls the the fundamental notion of *cohesion*. In general, software cohesion refers to *the degree to which the elements of a module belong together* [1]. Cohesive software is easier to understand, test and maintain. In the case of Web service interfaces, the splitting should rely on a certain notion of cohesion that reflects the relatedness of the operations which are grouped in the same interface [2], [3], [4], [5], [6].

Technical challenge: Unfortunately, Alice cannot obtain her desideratum of splitting a service interface into a set of cohesive interfaces via the state of the art cohesion-driven refactoring methods [7], [8], [9], [10], [11], [12], [13], [14], [15]. On the one hand, like all Web services, the ones that Alice uses *do not expose their internals, i.e., their source code*; on the contrary,

the very philosophy of Web services dictates that all that is exported by a Web service is *the Web service specification*. On the other hand, the cohesion-driven refactoring methods are tailored to operate by taking the source code into consideration. To overcome this problem, in this paper we propose an approach that facilitates *the cohesion-driven decomposition of service interfaces, without information of how the services are implemented*.

Contribution: The contributions of this paper are as follows.

First, we define a suite of *cohesion metrics* that quantify the degree to which the operations of a service interface are related based on *interface-level relations*, extracted from the service interface specification. The first two metrics, namely the Lack of Message-Level Cohesion (LoC_{msg}) and the Lack of Conversation-Level Cohesion (LoC_{conv}), rely on relations between operations that have similar types of inputs/outputs. The third metric, called Lack of Domain-Level Cohesion (LoC_{dom}), considers relations between operations that are characterized by similar domain-level terms, which are extracted from the names of the operations. The proposed metrics extend our early ideas introduced in [6].

Second, we propose *a cohesion-driven decomposition method* that accepts as input a cohesion metric and a service interface. The given interface is progressively split into interfaces that are more cohesive than the given interface. If it is no longer possible to produce more cohesive interfaces, the decomposition stops.

Third, we have validated the proposed approach based on 22 case studies that concern services provided by Amazon and Yahoo. We have evaluated the *effectiveness* of our approach from a quantitative perspective by assessing the cohesion improvement and the size of the decompositions produced by the proposed approach. Moreover, we have assessed the *usefulness* of our approach, by conducting a user study, where developers assessed the quality of our method's outcomes, as well as the success of each of the proposed cohesion metrics.

The rest of this paper is structured as follows. In Section 2, we discuss our contribution with respect to the state of the art. In Section 3, we present the cohesion metrics that we employ for the decomposition of service interfaces. In Section 4, we detail the *modus operandi* of the proposed decomposition method. In Section 5, we discuss the results that we obtained. Finally, in Section 6 we summarize our contribution and discuss the future perspectives of this work.

2 RELATED WORK

In this section we discuss in further detail the contribution of our approach with respect to the state of the art. More specifically, in Section 2.1 we discuss the relation of our approach with previous efforts on

---

4. aws.amazon.com/ec2/
5. developer.searchmarketing.yahoo.com/docs/V6/reference/services/

software refactoring. Then, in Section 2.2 we focus on the cohesion metrics that we employ and their relation with cohesion metrics that have been proposed in the object-oriented and the service-oriented paradigms.

2.1 Refactoring

Refactoring is a behavior preserving changing process that improves the quality of a software system [16]. To this end, several interesting approaches have been proposed, concerning various development paradigms (e.g., object-oriented refactoring [17], [18], [19], component-based refactoring [20], aspect-oriented refactoring [21], [22], etc.). For an excellent survey on refactoring the interested reader may refer to [23].

Our approach, is more closely related to metrics-driven refactoring methods, which employ metrics to discover and repair design problems. To achieve this goal, the state of the art methods rely on *implementation-level relations*, derived from source code (Appendix A, Table 5). Specifically, in [24], Harman & Tratt focus on the improvement of coupling. To this end, they rely on dependencies between classes, quantified based on the well-known CBO metric [25]. In [8], [9] and [7], the goal is to improve the cohesion of classes, by taking into account relations between class methods and attributes (or other methods), used by the methods. Certain approaches consider the improvement of multiple quality properties. In particular, the methods proposed in [10] and [11] focus on both coupling and cohesion. The methods proposed in [12], [13] account for coupling, cohesion and code complexity. In [14] the proposed method considers the improvement of coupling, cohesion and code size. Finally, the method proposed in [15] accounts for coupling, cohesion, code complexity and size.

Concerning their *modus operandi*, the metrics-driven refactoring methods can be divided in two categories. In the first category, the methods require more involvement from the developer [8], [13], [7]. In particular, based on the values of the metrics that are considered, the methods identify possible refactorings that can improve the values of the metrics. Following, the developer is supposed to select and apply the refactoring that suits his/her preferences. In the second category, the methods do more work on behalf of the developer. These methods consider the refactoring as an iterative process. As long as the design of the classes can be improved the refactoring process keeps going. The algorithms that are used to realize the refactoring process vary. We have methods that rely on meta-heuristic optimization algorithms [26], [27] (e.g., hill-climbing [24], [14], simulated annealing [14], genetic algorithms [12], [15]). Moreover, we have methods that are based on clustering [9].

Our approach differs from the state of the art methods in two main points:

- It deals with a novel problem, i.e., the cohesion-driven decomposition of service interfaces, without knowledge of how the services are implemented.
- To address this problem, our approach is based on *interface-level relations*, instead of being based on implementation-level relations.

2.2 Cohesion Metrics

In the early 90's Chidamber and Kemerer proposed the well-known LCOM metric (Lack of Cohesion of Methods) for measuring the cohesion of object-oriented software [25]. The interested reader may refer to [28] and [29] for two interesting surveys of the cohesion metrics that have been proposed since the seminal work of Chidamber and Kemerer. In the service-oriented paradigm, cohesion was recognized as an important principle of service design in several interesting approaches that concern the overall service-oriented development methodology [2], [3], [4]. The first efforts for measuring cohesion have been made in the work of Perepletchikov et al. [5]. The first study that investigated the issue of cohesion in the case of real-world services is reported in Athanasopoulos & Zarras [6]. Finally, another interesting work that concerns the cohesion of services is presented in [30].

In the object-oriented paradigm, the majority of the cohesion metrics measure the degree to which the methods of a class are related based on implementation-level relations (Appendix A, Table 6). Two class methods are considered as being related if they use common class attributes (or methods). In the object-oriented paradigm, we also have cohesion metrics that assess the relatedness of methods based on interface-level relations. In these metrics, two methods are considered as being related if they have parameters of the same type.

In the service-oriented paradigm, the SIIC metric [5] measures the relatedness of service operations, with respect to implementation-level relations. According to SIIC, two operations are related if the use common implementation elements. Similarly, in the SCV metric [30] two operations are considered as being related if they access related business entities, where the term business entity refers to an information entity involved in a particular business process. The SIDC and the SISC metrics [5] rely on interface-level relations. In particular, in SIDC two operations are considered as being related if they have input (or output) parameters of the same type. In SISC, if the type of an output parameter of one operation satisfies the type of an input parameter of another operation, the two operations are considered as being related. The SIUC metric [5], also operates at the interface-level; two operations are related if they are used by the same service clients.

Concerning the state of the art metrics that we previously discussed and our prior work on cohesion

Fig. 3. The meaning of LoC_*.

metrics for service interfaces [6], the added value of this paper is the following:

- LoC_{msg} and LoC_{conv} take into account the structure of input/output data types, instead of considering only relations between operations that have inputs/outputs of the same type.
- LoC_{dom} is an additional metric that extends the suite of metrics that we introduced in [6] and follows a completely different direction for measuring the cohesion of service interfaces, as it relies on relations between operations that are characterized by similar domain-level terms, which are extracted from the names of the operations.

3 INTERFACE-LEVEL COHESION METRICS

In this section, we focus on the cohesion metrics that we employ for the decomposition of service interfaces. In Section 3.1 we discuss the basic concepts that are shared among the metrics. Then, in Section 3.2 we define the metrics. Finally, in Appendix B we provide an analytic validation of the metrics, with respect to the theoretical framework of Briand et al. [31].

3.1 Basic Concepts

Our overall approach for measuring cohesion is based on a generic view of the notion of service interface, which is given in the following definition.

Definition 1: **(Service interface)** A service interface, si, is characterized by a name and a set of operations, $si.O$ (Table 1(1)). An *operation* is characterized by a *name*, an *input message* and an *output message* (Table 1(2)). A *message* is a set of parameters (Table 1(3)). Each *parameter* has a *name* and a *type*, which may be either an XML build-in type, or an XML complex type (Table 1(4)).

To measure cohesion, we further employ the concept of *interface-level graph*, which represents the interface-level relations between the operations of a given service interface. In general, two operations are related if their properties (e.g., names, parameters) are similar to some extent, according to a particular similarity function. More formally, the definition of the interface-level graph is given below.

Definition 2: **(Interface-level graph)** An interface-level graph, $G_{si}^* = (V_{si}, E_{si}, OpS_*)$, for a service interface, si, and a similarity function, $OpS_* : si.O \times si.O \to [0, 1]$ that reflects the degree to which the operations of si are related, is a weighted graph with the following properties (Table 1(5)):

- The nodes, V_{si}, of the graph represent the operations of si.
- The edges, E_{si}, of the graph represent interface-level relations between pairs of operations.
- An edge, (op_i, op_j), belongs to E_{si}, iff $OpS_*(op_i, op_j) > 0$; the weight that characterizes the edge is $OpS_*(op_i, op_j)$.

TABLE 1
Basic concepts.

$$
\begin{aligned}
si &= (name : string, O) & (1)\\
O &= \{op : operation\} \\
operation &= (name : string, & (2)\\
& \quad in : message, out : message) \\
message &= \{p : parameter\} & (3)\\
parameter &= (name : string, type : anyType) & (4)\\
G_{si}^* &= (V_{si}, E_{si}, OpS_*) & (5)\\
LoC_*(si, OpS_*) &= 1 - \frac{\sum_{\forall(op_i, op_j) \in E_{si}} OpS_*(op_i, op_j)}{\frac{|V_{si}| * (|V_{si}| - 1)}{2}} & (6)
\end{aligned}
$$

Based on the concept of interface-level graph, the lack of interface-level cohesion is defined as follows.

Definition 3: **(Lack of interface-level cohesion)** Ideally, a service interface, si, would be fully cohesive if every operation of si is related with all the others and the similarity between every pair of operations is maximum (Figure 3, left). Hence, the ideal interface-level graph $G_{ideal}^* = (V_{si}, E_{ideal}, OpS_*)$ for si has two properties: (1) G_{ideal}^* is complete; (2) for all, $(op_i, op_j) \in E_{ideal}$, $OpS_*(op_i, op_j) = 1$. Intuitively, the lack of cohesion for the service interface si measures the amount of transformation that the actual interface-level graph $G_{si}^* = (V_{si}, E_{si}, OpS_*)$ of si must withstand to become identical to the ideal graph, G_{ideal}^*. This practically amounts to adding the missing edges and complementing the weights of the existing edges to become equal to 1. Formally speaking, the lack of cohesion of si measures the relative difference between G_{si}^* and G_{ideal}^*, i.e.,

$$LoC_*(si, OpS_*) = \frac{|E_{ideal}| - \sum_{\forall(op_i, op_j) \in E_{si}} (OpS_*(op_i, op_j))}{|E_{ideal}|}.$$

Given that $|E_{ideal}| = \frac{|V_{si}| * (|V_{si}| - 1)}{2}$, with simple algebraic calculations we get the formula that is given in Table 1(6).

3.2 Metrics Definitions

The proposed cohesion metrics refine the generic definition of LoC_* that was given in Section 3.1. In

particular, the definitions of the metrics that we provide in the following paragraphs employ the notion of interface-level graph; the interface-level graph that is used for each metric relies on a different similarity function between operations.

Message-Level Cohesion

The notion of *message-level cohesion* assumes that two operations are related if their input (respectively, output) messages are similar. To measure the similarity between two messages we employ the notion of *message-level graph* that is defined below.

Definition 4: **(Message-level graph)** A message-level graph, $G_m = (V_m, E_m)$, for a message, m, is a generic representation of the structure of m. Specifically, G_m consists of 3 layers:

- The first layer, comprises a node, $v_m \in V_m$, that represents the message itself.
- The second layer, comprises a set of nodes, which represent the parameters that constitute m. The edges between the first layer and the second layer nodes represent whole-part relations.
- The third layer, details the structure of the types of the parameters that constitute m. Hence, the third layer may contain nodes that represent primitive XML elements, or complex XML elements that consist of further (primitive or complex) XML elements. The edges between the second layer and the third layer nodes represent type-of relations.

TABLE 2
Similarity functions.

$$MsgS(m_i, m_j) = \frac{|V_{m_i \cap m_j}|}{|V_{m_i \cup m_j}|} \quad (1)$$

$$OpS_{msg}(op_i, op_j) = \frac{MsgS(op_i.in, op_j.in)}{2} + \quad (2)$$
$$\frac{MsgS(op_i.out, op_j.out)}{2}$$

$$OpS_{conv}(op_i, op_j) = \frac{MsgS(op_i.in, op_j.out)}{2} + \quad (3)$$
$$\frac{MsgS(op_i.out, op_j.in)}{2}$$

$$OpS_{dom}(op_i, op_j) = \frac{|T_{op_i} \cap T_{op_j}|}{|T_{op_i} \cup T_{op_j}|} \quad (4)$$

Definition 5: **(Message similarity)** The similarity between the two messages, m_i, m_j, (Table 2(1)) is measured with respect to the message-level graphs, G_{m_i}, G_{m_j}, of m_i, m_j. Specifically, let $G_{m_i \cap m_j} = (V_{m_i \cap m_j}, E_{m_i \cap m_j})$ denote the maximum common subgraph of G_{m_i}, G_{m_j} that represents a syntactically correct XML schema. Moreover, let $G_{m_i \cup m_j} = (V_{m_i \cup m_j}, E_{m_i \cup m_j})$ be the union of G_{m_i}, G_{m_j} (i.e., $V_{m_i \cup m_j} = V_{m_i} \cup V_{m_j}$ and $E_{m_i \cup m_j} = E_{m_i} \cup E_{m_j}$). Then, the similarity, $MsgS(m_i, m_j)$, between m_i and m_j is the number of nodes of $G_{m_i \cap m_j}$, divided by the number of nodes of $G_{m_i \cup m_j}$.
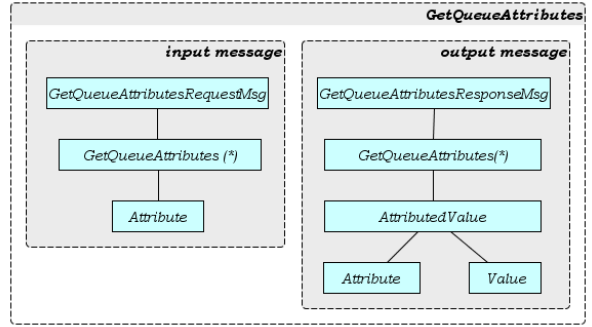
(a) Messages of the `GetQueueAttributes` operation.
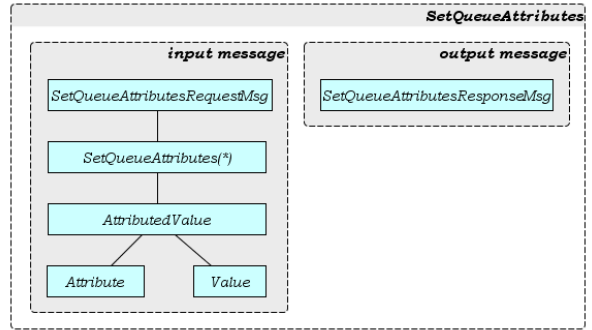
(b) Messages of the `SetQueueAttributes` operation.

Fig. 4. Examples of message-level graphs for `MessageQueue`.

In general, finding the maximum common subgraph for two graphs is a hard problem to solve. However, solving the problem for two message-level graphs is much easier because the search space is limited to a set of subgraphs that represent the XML types of the parameters that constitute the messages. Typically, the definitions of these XML types are provided in the service descriptions, separately from the specification of the service interfaces (see WSDL 1.1, WSDL 2.0 [32]).

Taking a step further, we define the message-level similarity between two operations as follows.

Definition 6: **(Message-level operation similarity)** The message-level similarity, OpS_{msg}, between two operations, $op_i, op_j \in si.O$ of a service interface, si, is the average of (Table 2(2)):

1) the similarity between the input messages of op_i and op_j and
2) the similarity between the output messages of op_i and op_j.

Taking the example of Amazon SQS, Figure 4(a) shows the message-level graph for the input message of the `GetQueueAttributes` operation. The `GetQueueAttributesRequestMsg` node represents the message (first layer). The `GetQueueAttributes` node is a parameter that comprises of sequence of attributes (second layer). The `Attribute` node represents a primitive XML

string element (third layer).

Figure 4(a), further gives the message-level graph for the output message of the `GetQueueAttributes` operation. The `GetQueueAttributesResponseMsg` node represents the message (first layer). The `GetQueueAttributes` node represents a parameter that comprises of sequence of attribute value pairs (second layer). The `AttributedValue` node represents a complex XML element, which consists of two primitive XML string elements, represented by the `Attribute` and the `Value` nodes.

Similarly, Figure 4(b) gives the message-level graphs for the input and the output messages of the `SetQueueAttributes` operation.

The maximum common subgraph for the message-level graphs of the two input messages comprises only the `Attribute` node. The union of the two graphs consists of 7 nodes. Hence, the similarity between the two input messages is $\frac{1}{7}$. On the other hand, the message-level graphs of the two output messages have nothing in common. Thus, the similarity between the two output messages is 0. Overall, the message-level similarity between the two operations is $\frac{\frac{1}{7}+0}{2}$.

Based on the message-level similarity between operations, we refine the LoC_* metric.

Definition 7: **(Lack of message-level cohesion)** For a service interface, si, the lack of message-level cohesion, $LoC_{msg}(si)$, is an alias for $LoC_*(si, OpS_{msg})$. Specifically, $LoC_{msg}(si)$ measures the relative difference between the interface-level graph, $G_{si}^{msg} = (V_{si}, E_{si}, OpS_{msg})$, defined based on the message-level similarity function, OpS_{msg}, and the ideal interface-level graph, G_{ideal}^{msg}.


Fig. 5. $G_{MessageQueue}^{msg}$ for `MessageQueue`.

In our example, Figure 5, gives the interface-level graph for the `MessageQueue` interface that is derived with respect to OpS_{msg}. For presentation purposes the edges' width is proportional to the similarity between the operations. We see that the graph is not complete. Moreover, the message-level relations between

the operations are weak; the similarities between the operations range from 0.07 to 0.21. Overall, the lack of message-level cohesion is $LoC_{msg}(MessageQueue) = 0.98$.

Conversation-Level Cohesion

The notion of *conversation-level cohesion* assumes that an operation is related with another if the former's input (respectively output) message is similar with the latter's output (respectively input) message. More formally, we define the conversation-level similarity between two operations as follows.

Definition 8: **(Conversation-level operation similarity)** The conversation-level similarity between two operations, $op_i, op_j \in si.O$, of a service interface, si, is the average of (Table 2(3)):

1) the similarity between the input message of op_i and the output message of op_j and
2) the similarity between the output message of op_i and the input message of op_j.

Returning to our example, the input message of `GetQueueAttributes` (Figure 4(a)) and the output message of `SetQueueAttributes` (Figure 4(b)), have nothing in common. On the other hand, the maximum common subgraph for the output message of `GetQueueAttributes` and the input message of `SetQueueAttributes` includes three nodes (`AttributedValue`, `Attribute` and `Value`). Hence, the conversation-level similarity between the two operations is $\frac{\frac{3}{7}+0}{2}$.

Given the conversation-level similarity between operations, we introduce the following refinement of the LoC_* metric.

Definition 9: **(Lack of conversation-level cohesion)** For a service interface, si, the lack of conversation-level cohesion, $LoC_{conv}(si)$, is an alias for $LoC_*(si, OpS_{conv})$. In particular, $LoC_{conv}(si)$ measures the relative difference between the interface-level graph, $G_{si}^{conv} = (V_{si}, E_{si}, OpS_{conv})$, defined with respect to the conversation-level similarity function, OpS_{conv}, and the ideal interface-level graph, G_{ideal}^{conv}.

Regarding our example, the interface-level graph that shows the conversation-level relations for the `MessageQueue` interface is given in Figure 6. As in the case of message-level cohesion, the graph is not ideal. The overall lack of conversation-level cohesion is $LoC_{conv}(MessageQueue) = 0.98$.

Domain-Level Cohesion

The basic intuition behind the notion of *domain-level cohesion* is that the names of the operations that are provided by a service reflect what these operations do. More specifically, the names of the operations comprise *terms that correspond to certain actions* (e.g., set, get) and *terms that correspond to concepts of the domain that is targeted by the service* (e.g., queue, attribute, message). Based on this intuition, two operations
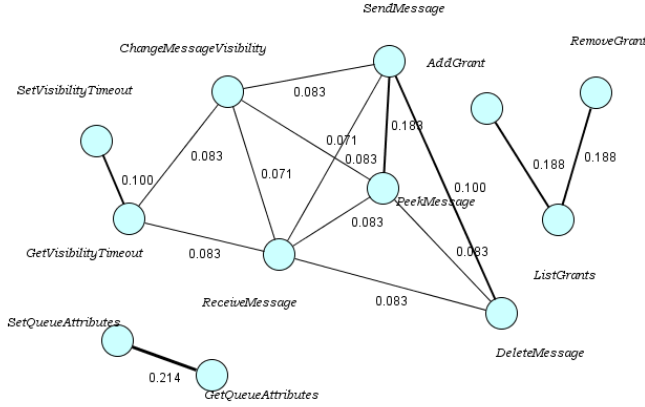
Fig. 6. $G^{conv}_{MessageQueue}$ for `MessageQueue`.

Fig. 7. $G^{dom}_{MessageQueue}$ for `MessageQueue`.

are considered as being related if their names share domain-level terms.

In our approach, we extract the domain-level terms from the names of the operations, based on the naming conventions of the particular coding style that is followed by the service provider. For instance, the Amazon services follow the PascalCase coding style[6] (the names of operations are sequences of terms with the first letter of each term being capitalized). On the other hand, the Yahoo services follow the Java coding style[7]. Then, we measure the domain-level similarity between two operations with the following similarity function.

Definition 10: **(Domain-level operation similarity)** Let T_{op_i} and T_{op_j} denote the sets of the domain-level terms that are extracted from the names of two operations, $op_i, op_j \in si.O$, of a service interface, si. The domain-level similarity between the two operations (Table 2(4)) is the Jaccard similarity for T_{op_i} and T_{op_j} (i.e., the size of the intersection divided by the size of the union of T_{op_i} and T_{op_j}).

Getting back to our example, the name of `GetQueueAttributes` consists of the action term `Get`, which is related with two domain-level terms, `Queue` and `Attributes`. The name of `SetQueueAttributes` comprises the action term `Set`, which is also related with `Queue` and `Attributes`. Therefore, the domain-level similarity between the two operations is $\frac{2}{2}$.

The refinement of the LoC_* metric, with respect to the domain-level similarity between two operations, is given below.

Definition 11: **(Lack of domain-level cohesion)** The lack of domain-level cohesion, $LoC_{dom}(si)$, for a service interface, si, is an alias for $LoC_*(si, OpS_{dom})$. $LoC_{dom}(si)$ measures the relative difference between the interface-level graph, $G^{dom}_{si} = (V_{si}, E_{si}, OpS_{dom})$, defined with respect to the domain-level similarity

function, OpS_{dom}, and the ideal interface-level graph, G^{dom}_{ideal}.

Concerning our example, in Figure 7, we have the interface-level graph that shows the domain-level relations for the `MessageQueue` interface. As in Figures 5 and 7 the graph is not complete. However, the domain-level relations are generally strong; the similarities between operations range from 0.3 to 1. Overall, the lack of domain-level cohesion is $LoC_{dom}(MessageQueue) = 0.81$.

4 COHESION-DRIVEN DECOMPOSITION

In this section, we detail the method that exploits the metrics defined in Section 3 for the cohesion-driven decomposition of service interfaces. Moreover, in Appendix C we provide a complementary analysis that concerns the stopping criteria and the complexity of the method.

At a glance, the cohesion-driven decomposition of service interfaces accepts as input a service interface, si, which is progressively split in more cohesive interfaces (Algorithm 1). Note that although we assume as input a single service interface, the method can be easily applied in the case of a service that provides multiple interfaces. In such case, the interfaces of the service can be merged into a single interface. Then, this interface can be given as input to the cohesion-driven decomposition method. As the decomposition proceeds, si is split in several interfaces, all of which are candidates to be further divided. To this end, we employ a queue, Q, which contains the interfaces that are candidates for decomposition (Algorithm 1, line 1). Initially, Q contains only the given interface, si (Algorithm 1, line 3). During each step (Algorithm 1,

6. msdn.microsoft.com/en-us/library/x2dbyw72(v=vs.71).aspx
7. www.oracle.com/technetwork/java/codeconventions-135099.html

Algorithm 1: $decomposeInterface()$

Input: $si : Interface$ /* An interface that is decomposed in more cohesive interfaces */
Output: $R_I = \{r : Interface\}$ /* The produced set of interfaces */
1. **var** $Q : Queue$ /* A queue that stores intermediate interfaces */
2. $R_I \leftarrow \emptyset$
3. $Q.enqueue(si)$
4. **repeat**
5.   **var** $r_i : Interface$ /* holds an intermediate interface picked from the queue Q */
6.   **var** $r_s : Interface$ /* the splinter interface that comprises operations removed from r_i */
7.   **var** $r_r : Interface$ /* the interface that comprises the remaining operations of r_i */
8. $r_i \leftarrow Q.dequeue()$
9. $r_s \leftarrow null$
10. $r_r \leftarrow null$
11. $(r_s, r_r) \leftarrow createSplinter(r_i)$ /* Phase 1: Returns the splinter r_s and r_r. */
12.   **if** $r_s = null$ **then**
13. $R_I \leftarrow R_I \cup \{r_i\}$ /* $LoC_*(r_i)$ can not be further improved. */
14.   **else**
15. $(r_s, r_r) \leftarrow populateSplinter(r_s, r_r, r_i)$ /* Phase 2: Populate r_s. */
16. $Q.enqueue(r_r)$
17. $Q.enqueue(r_s)$
18.   **end if**
19. **until** $Q.size() = 0$
20. **return** R_I

Algorithm 2: $createSplinter()$

Input: $r_i : Interface$ /* An intermediate interface picked from Q */
Output: $r_s : Interface$ /* The splinter interface that is created */
    $r_r : Interface$ /* The interface that contains the remaining operations of r_i */
1. **var** $op_s : Operation$ /* The operation, whose removal maximizes the cohesion improvement of r_i */
2. **var** $\delta_{max} : Float$ /* The max cohesion improvement that can be achieved by removing an operation from r_i */
3. **var** r_{tmp} /* A temporary interface used to simulate the interface that results from the removal of an operation from r_i */
4. $r_s \leftarrow null$
5. $r_r \leftarrow null$
6. $\delta_{max} \leftarrow 0$
7. $op_s \leftarrow null$
8. $r_{tmp} \leftarrow null$
9. **for all** $op_i \in r_i.O$ **do**
10. $r_{tmp} \leftarrow$ **new** $Interface$
11. $r_{tmp}.O \leftarrow r_i.O - \{op_i\}$
12.   **if** $LoC_*(r_i) - LoC_*(r_{tmp}) > \delta_{max}$ **then**
13. $\delta_{max} \leftarrow LoC_*(r_i) - LoC_*(r_{tmp})$
14. $op_s \leftarrow op_i$
15.   **end if**
16. **end for**
17. **if** $\delta_{max} > 0$ **then**
18. $r_s \leftarrow$ **new** $Interface$
19. $r_s.O \leftarrow \{op_s\}$
20. $r_r \leftarrow$ **new** $Interface$
21. $r_r.O \leftarrow r_i.O - \{op_s\}$
22. **end if**
23. **return** (r_s, r_r)

lines 5-20), the method dequeues from Q an intermediate interface r_i and checks whether it is possible to improve the cohesion of r_i, by removing a set of operations, which form a new interface r_s. Hereafter, we use the term *splinter interface* to refer to r_s, while r_r denotes the interface that comprises the rest of the operations of r_i.

The construction of the splinter interface takes place in two phases. The first phase, called $createSplinter$, checks if it is possible to improve the cohesion of r_i, by removing an operation (Algorithm 1, line 11). If this phase fails to find such an operation, ri is inserted in the results set, R_I (Algorithm 1, lines 12-13). Otherwise, the splinter interface, r_s, that contains the operation is returned as a result of $createSplinter$, along with the interface r_r that contains the remaining operations of r_i. The second phase, called $populateSplinter$, further improves the cohesion of r_s and r_r, by moving operations from r_r to r_s (Algorithm 1, line 15). Finally, the two interfaces, r_s, r_r are inserted in Q (Algorithm 1, lines 16-17).

In further detail, the two phases of the decomposition are discussed below.

The $createSplinter$ phase accepts as input the intermediate interface, r_i, that is picked from Q (Algorithm 2). Following, it iterates over the operations of r_i (lines 9-16). Each iteration checks whether the removal of a single operation, op_i, from r_i improves the cohesion of the interface (line 12). To this end, the removal of op_i is simulated with the help of a temporary interface, r_{tmp}. Moreover, each iteration keeps track of the maximum cohesion improvement, δ_{max}, that can be achieved, and of the operation, op_s,

that should be removed to achieve this improvement (lines 12-15). After this iterative process, if $\delta_{max} > 0$, the splinter interface, r_s, that contains op_s is created, along with the interface, r_r that contains the remaining operations of r_i (lines 17-22). The two new interfaces are returned as the results of $createSplinter$ (line 23). On the other hand, if is not possible to improve the cohesion of r_i, by removing an operation (i.e., $\delta_{max} = 0$), the results of $createSplinter$ are null.

The $populateSplinter$ phase accepts as input the intermediate interface, r_i, and the newly created interfaces, r_s, r_r (Algorithm 3). Then, it repeatedly moves operations from r_r to r_s (lines 9-33) as follows:

- The $populateSplinter$ iterates over the operations of r_r (lines 15-28). Each iteration checks if an operation, op_i, can be moved from r_r to r_s. To perform this check, the movement of the operation is simulated with the help of two temporary interfaces, $r_{r_{tmp}}$, $r_{s_{tmp}}$. In particular, $r_{r_{tmp}}$ is employed to calculate the cohesion improvement, δ_{r_r}, that can be achieved for r_r, if the operation is moved (lines 16-18). Similarly, $r_{s_{tmp}}$ is employed to calculate the cohesion improvement, δ_{r_s}, that can be achieved for r_s (lines 19-21). The operation, op_i, is considered as a candidate to be moved if the following conditions hold (line 22): (a) the cohesion of r_r, after the move, is improved, i.e., $\delta_{r_r} > 0$, (b) the cohesion of r_s, after the move, is also improved, i.e., $\delta_{r_s} > 0$, and (c) the lack of cohesion of r_s, after the move, is smaller than the lack of cohesion of the intermediate interface r_i that was picked from Q. Each iteration further keeps track of the total cohesion improvement that can be achieved, by moving op_i, from r_r to

Algorithm 3: $populateSplinter()$

Input: $ri : Interface$ /* An intermediate interface picked from Q */
  $r_s : Interface$ /* The splinter interface that was created in Phase 1 */
  $r_r : Interface$ /* The interface that contains the remaining operations of r_i */
Output: $r_s : Interface$ /* The populated splinter interface that comprises operations removed from r_r */
  $r_r : Interface$ /* The interface that contains the remaining operations of r_i */

1. **var** $\delta_{r_r} : Float$ /* The cohesion improvement that can be achieved by removing an operation from r_r */
2. **var** $\delta_{r_s} : Float$ /* The cohesion improvement that can be achieved by adding an operation to r_s */
3. **var** $\delta_{total} : Float$ /* The total cohesion improvement ($\delta_{r_r} + \delta_{r_s}$) that can be achieved by moving an operation from r_r to r_s */
4. **var** $op_s : Operation$ /* The operation that is moved from r_r to r_s */
5. **var** $r_{r_{tmp}}, r_{s_{tmp}} : Interface$ /* Temporary interfaces used to simulate the interfaces that result after moving an operation from r_r to r_s */
6. $r_{r_{tmp}} \leftarrow null$
7. $r_{s_{tmp}} \leftarrow null$
8. /* Move operations from r_r to r_s */
9. **repeat**
10. $op_s \leftarrow null$
11. $\delta_{r_r} \leftarrow 0$
12. $\delta_{r_s} \leftarrow 0$
13. $\delta_{total} \leftarrow 0$
14.   /* Find the operation op_s that improves the cohesion of r_r and r_s, and maximizes δ_{total} */
15.   **for all** $op_i \in r_r.O$ **do**
16. $r_{r_{tmp}} \leftarrow$ **new** $Interface$
17. $r_{r_{tmp}}.O \leftarrow r_r.O - \{op_i\}$
18. $\delta_{r_r} \leftarrow LoC_*(r_r) - LoC_*(r_{r_{tmp}})$
19. $r_{s_{tmp}} \leftarrow$ **new** $Interface$
20. $r_{s_{tmp}}.O \leftarrow r_s.O \cup \{op_i\}$
21. $\delta_{r_s} \leftarrow LoC_*(r_s) - LoC_*(r_{s_{tmp}})$
22.   **if** $((\delta_{r_r} > 0) \wedge (\delta_{r_s} > 0) \wedge (LoC_*(r_{s_{tmp}}) < LoC_*(r_i)))$ **then**
23.   **if** $\delta_{r_s} + \delta_{r_r} > \delta_{total}$ **then**
24. $\delta_{total} \leftarrow \delta_{r_s} + \delta_{r_r}$
25. $op_s \leftarrow op_i$
26.   **end if**
27.   **end if**
28.   **end for**
       /* Move the operation op_s */
29.   **if** $op_s \neq null$ **then**
30. $r_r.O \leftarrow r_r.O - \{op_s\}$
31. $r_s.O \leftarrow r_s.O \cup \{op_s\}$
32.   **end if**
33. **until** $op_s = null$
34. **return** (r_s, r_r)
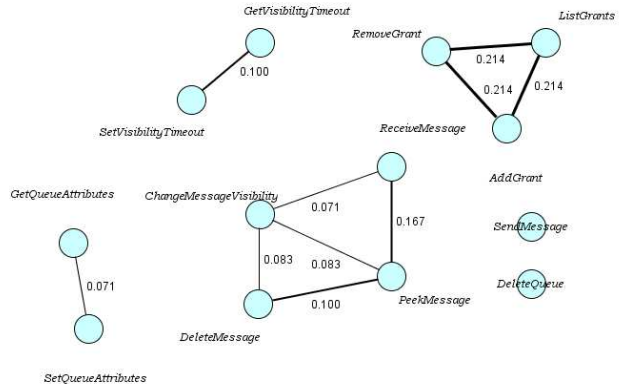

Fig. 8. Decomposition of `MessageQueue`, based on LoC_{msg}.
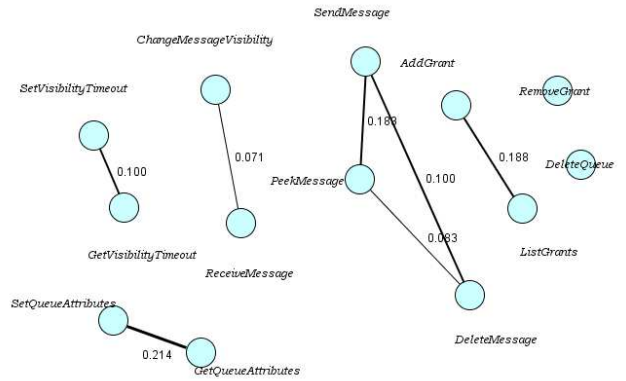
Fig. 9. Decomposition of `MessageQueue`, based on LoC_{conv}.

r_s. Moreover, it keeps track of the operation op_s that maximizes the total cohesion improvement that can be achieved (lines 23-26).

- The operation, op_s, that maximizes the total cohesion improvement, δ_{total}, is moved from r_r to r_s (lines 29-32).
- The whole process stops when $op_s = null$ (line 33) and the updated r_s, r_r are returned (line 34).

Back to our example, Figures 8, 9 and 10, give the three different decompositions of `MessageQueue` that result based on LoC_{msg}, LoC_{conv} and LoC_{dom}, respectively [8]. In particular, the message-level decomposition of `MessageQueue` consists of 6 interfaces. The average lack of message-level cohesion of the interfaces is 0,92. Hence, an improvement has been made compared to the initial interface (Figure 5), but the improvement is small. This result is anticipated

8. The input to the method was the 2007 version of the interface, aws.amazon.com/articles/Amazon-SQS/1148

because the message-level relations between the operations of `MessageQueue` are not strong (Figure 5). The conversation-level decomposition consists of 7 interfaces. The average lack of conversation-level cohesion in this case is 0.88. Again, the improvement compared to the initial interface is small, because the conversation-level relations between the operations of `MessageQueue` are not strong (Figure 6). The domain-level decomposition of `MessageQueue` consists of 4 interfaces and the average lack of domain-level cohesion is 0.13. The improvement in this case is high, since the domain-level relations between the operations of `MessageQueue` are quite strong (Figure 7).

Figure 11, provides an detailed view of the execution of Algorithm 1 in the case of LoC_{dom}. Specifically, in the first step, the general queue management operations (`DeleteQueue`, `SetQueueAttributes` and `GetQueueAttributes`) are removed from `MessageQueue`. These operations constitute the splinter interface, r_{s_1} (in Figure 2, this interface appears with the name `QueueMgt`). The remaining operations form r_{r_1}. Overall, the lack of cohesion of r_{s_1} is 0.33, while the lack of co-

Fig. 10. Decomposition of `MessageQueue`, based on LoC_{dom}.

Fig. 11. Decomposing the `MessageQueue` interface.

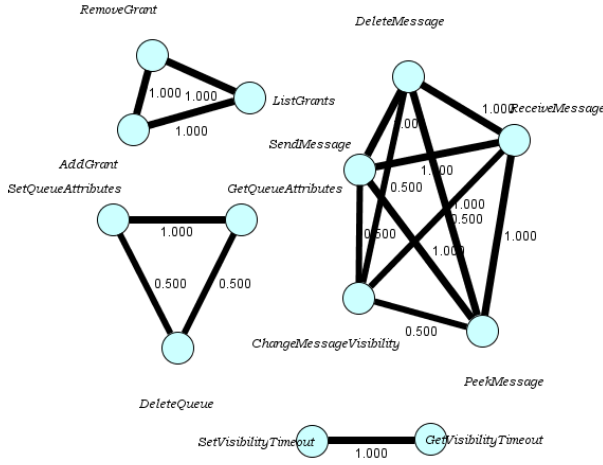hesion of r_{r_1} is 0.72. In the second step, r_{r_1} is decomposed. In particular, the timeout management operations (`GetVisibilityTimeout` and `SetVisibilityTimeout`) are removed from r_{r_1}, and the splinter interface, r_{s_2}, is formed (in Figure 2, this interface is called `TimeoutMgt`). The rest of the operations of r_{r_1}, form r_{r_2}. The lack of cohesion of r_{s_2} is 0, while the lack of cohesion of r_{r_2} is 0.61. In the last step, r_{r_2} is decomposed, by removing the access rights management operations (`AddGrant`, `RemoveGrant` and `ListGrants`), which constitute the splinter interface, r_{s_3} (in Figure 2, this interface is named `GrantsMgt`). The rest of the operations are related to messaging and form r_{r_3} (in Figure 2, this interface appears as `MsgMgt`). The lack of cohesion of r_{s_3} is 0, while the lack of cohesion of r_{r_3} is 0.2. To sum up, the results of the decomposition of `MessageQueue` are $R_I = \{r_{s_1}, r_{s_2}, r_{s_3}, r_{r_3}\}$.

5　VALIDATION

Amazon and Yahoo are two major service providers that offer a wide variety of Web services. To validate the proposed approach, we selected services that provide interfaces with at least 10 operations. Overall, we used 11 Amazon services and 11 Yahoo services[9]. Hereafter, we use identifiers A1-A11 and Y1-Y11 to refer, respectively, to the interfaces of the Amazon and the Yahoo services that we used. Table 3 provides the mapping between the identifiers and the service interfaces, along with the sizes of the interfaces (i.e., the number of provided operations) and the values of LoC_{msg}, LoC_{conv} and LoC_{dom} for the interfaces.

In the rest of this section we detail our findings. In Section 5.1 we concentrate on the effectiveness of the proposed approach from a quantitative perspective. In Section 5.2, we discuss the usefulness of the approach from the developers' perspective. Finally, in Section 5.3 we discuss threats to validity.

TABLE 3
Amazon & Yahoo case studies.

(a) *Amazon: aws.amazon.com/*

Service Interface					
Name	Size	ID	LoC_{msg}	LoC_{conv}	LoC_{dom}
AmazonEC2PortType	87	A1	0.98	0.99	0.94
MechanicalTurkRequesterPortType	27	A2	0.92	0.84	0.83
AmazonFPSPortType	27	A3	0.92	0.97	0.96
AmazonRDSv2PortType	23	A4	0.91	0.96	0.56
AmazonVPCPortType	21	A5	0.95	0.98	0.82
AmazonFWSInboundPortType	18	A6	0.93	0.96	0.73
AmazonS3	16	A7	0.89	0.97	0.75
AmazonSNSPortType	13	A8	0.96	0.97	0.84
ElasticLoadBalancingPortType	13	A9	0.93	0.97	0.72
MessageQueue	13	A10	0.98	0.98	0.81
AutoScalingPortType	13	A11	0.96	0.98	0.79

(b) *Yahoo:developer.searchmarketing.yahoo.com/docs/V6/reference/*

Service Interface					
Name	Size	ID	LoC_{msg}	LoC_{conv}	LoC_{dom}
KeywordService	34	Y1	0.84	0.93	0.91
AdGroupService	28	Y2	0.84	0.94	0.65
UserManagementService	28	Y3	0.96	0.97	0.91
TargetingService	23	Y4	0.74	0.96	0.74
AccountService	20	Y5	0.92	0.98	0.88
AdService	20	Y6	0.79	0.89	0.88
CampaignService	19	Y7	0.83	0.91	0.91
BasicReportService	12	Y8	0.91	0.99	0.92
TargetingConverterService	12	Y9	0.84	0.80	0.53
ExcludedWordsService	10	Y10	0.72	0.81	0.54
GeographicalDictionaryService	10	Y11	0.79	0.99	0.65

5.1　Effectiveness

To assess the effectiveness of the approach from a quantitative perspective we focus on the following research questions:

RQ1: To what extent is cohesion improved by adopting the proposed method ?

RQ2: Is the number of produced interfaces reasonable

9. The WSDL specifications of the Amazon and the Yahoo services can be found at: www.cs.uoi.gr/~zarras/WS-Decomp-Material/

with respect to the size of the decomposed interface ?

To respond to these questions we decomposed the examined service interfaces, based on the metrics that we defined in Section 3, and the method that we detailed in Section 4. To address RQ1, we measured the *cohesion improvement*, $CI(si)$, that is achieved for a service interface si. Formally, for a set of interfaces, R_I, produced by the proposed method for si, the cohesion improvement is: $CI(si) = \frac{LoC_*(si,OpS_*) - \frac{\sum_{\forall r \in R_I}(LoC_*(r,OpS_*))}{|R_I|}}{LoC_*(si,OpS_*)} * 100\%$. To address RQ2, we measured the number of interfaces, $DS(si) = |R_I|$, produced by the proposed method for si. We further examined the relation between the number of operations offered by si (the independent variable) and $DS(si)$ (the dependent variable), using ordinary least squares regression (OLS). Hereafter, we use the term, *decomposition of si*, to refer to the set of interfaces, R_I, that is produced by the proposed method for si. Moreover, we use to term, *size of decomposition*, to refer to $DS(si)$.

RQ1: Figure 12(left col.), gives the values of CI that we obtained for the examined service interfaces. Concerning our first question, the combination of the proposed method with the domain-level cohesion metric (i.e., LoC_{dom}) was effective in all cases. The cohesion improvement for the domain-level decompositions is medium-high (CI ranges from 38% to 100%). The combination of the proposed method with the message-level cohesion metric (i.e., LoC_{msg}) was also effective in all cases. The cohesion improvement for the message-level decompositions is medium (CI is up to 41.9%). Finally, the combination of the proposed method with the conversation-level cohesion metric (i.e., LoC_{conv}) was effective in 77% of the cases. The cohesion improvement for the conversation-level decompositions is low. In 5 cases (A9, Y2, Y4, Y7, Y10), the similarities between the operations of the examined interfaces were such that the conversation-level cohesion of the initial interfaces could not be further improved.

RQ2: Figure 12(middle col.), gives the values of DS that resulted for the examined interfaces. Moreover, Figure 12(right col.) gives the results of the OLS analysis; in the x-axis of the scatter plots we have the number of the operations that are offered by the examined interfaces, in the y-axis we have the values of DS, and at the lower left corner of the scatter plots we have the regression equations and the values of the R^2 statistic. In general, the values of the R^2 statistic range from 0 to 1; high R^2 values indicate that a regression equation explains well the relationship between the variables involved in the equation. In our analysis, the values of the R^2 statistic are quite high (ranging from 0.71 to 0.89). Thus, the size of the decompositions, produced by the proposed method, linearly increases

with the number of operations that are offered by the decomposed interfaces. The regression equations that we obtained for the different cohesion metrics are similar. The maximum value of the regression coefficients that could result from the OLS analysis is 1. A regression coefficient that equals to 1, would mean that the number of interfaces that are produced by the decomposition method equals to the number of operations of the decomposed interface. In our analysis, the regression coefficients are quite small, ranging from 0.33 to 0.35. Hence, the size of the produced decompositions is reasonable, with respect to the number of operations of the decomposed interface. Nevertheless, there are certain cases where the size of the produced decompositions is relatively high – see Fig. 12(middle col.). For instance, for the combination of the decomposition method with the domain-level cohesion metric we have the cases of A3 and Y3. Similarly, for the combination of the decomposition method with the message-level cohesion metric we have the cases of A3 and Y7. Finally, for the combination of the decomposition method with the conversation-level cohesion metric we have the cases of A2 and Y5.

5.2 The Developers' Opinions

To evaluate the usefulness of the approach from the developers' perspective we investigate the following research questions:

RQ1: Does the proposed approach produce useful results for the developers ?

RQ2: What are the developers' preferences (if any) concerning the metrics that are employed ?

RQ3: To what extent should the results be refined to fully satisfy the developers' needs ?

To address the aforementioned questions we looked for volunteers with the following skills: software development experience; knowledge of the service-oriented computing paradigm, related technologies and standards. Overall, 10 volunteers participated in our study. The participants had 3 to 15 years experience in software development. They were all familiar with the service-oriented computing paradigm.

We organized the participants in two groups. The first group assessed the decompositions of the Amazon service interfaces, while the second group assessed the decompositions of the Yahoo service interfaces.

In a first meeting with the participants, we explained the overall purpose of the study, without giving any details, concerning the metrics and the method used for the decomposition of the examined service interfaces. Following, we gave to each participant a document [10] that contained the following information for each one of the examined interfaces: (a)

10. The documents can be found at www.cs.uoi.gr/~zarras/WS-Decomp-Material/
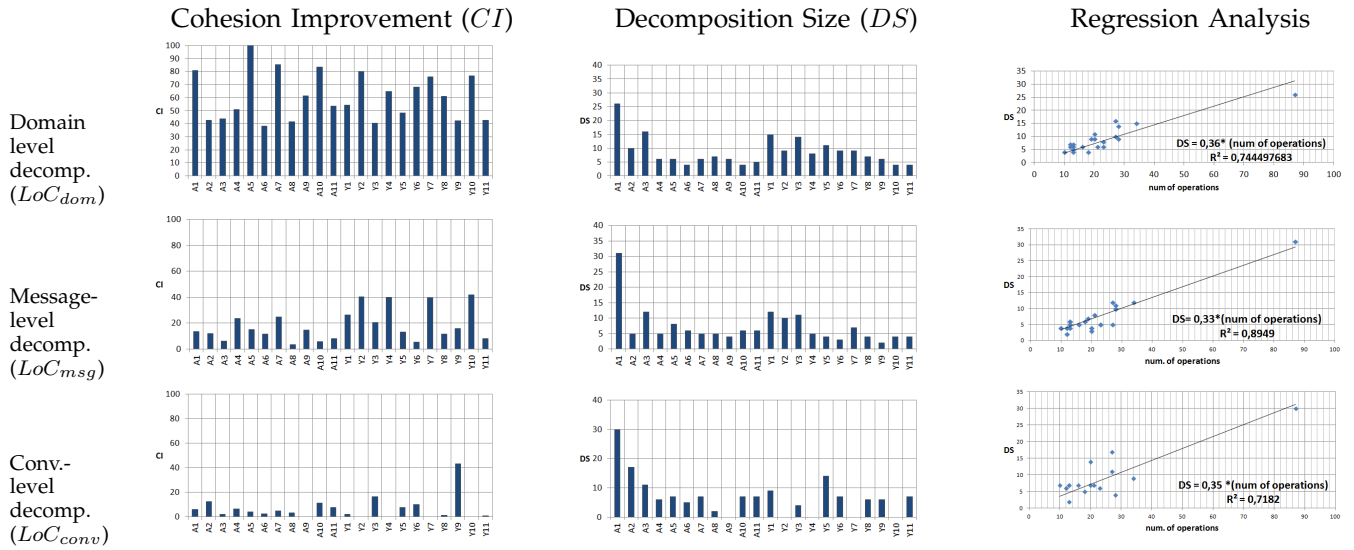
Fig. 12. Effectiveness Assessment

a high-level description (represented as a UML class) of the interface; (b) the domain-level, the message-level and the conversation-level decompositions of the interface. The decompositions were given in random order. The document further contained detailed instructions concerning the assessment tasks that should be performed for each one of the examined service interfaces:

- In the *first task*, the participants had to choose whether a service interface should be decomposed, or remain as is.
- In the *second task*, the participants had to report which of the provided decompositions is closest to their preferences; in this task the participants could also report that none of the provided decompositions is satisfactory.
- The *third task* was to suggest, if necessary, further changes on a selected decomposition.

In a second meeting with each of the participants, we collected the documents and we analyzed the participants' feedback. The participants' feedback is summarized in Figure 13. In this figure we use the following notations that correspond to the possible choices that could be made by a participant for a particular service interface: NO-SPLIT - the participant suggested that the interface should not be decomposed; NONE - none of the provided decompositions was selected by the participant; Msg - the participant selected the message-level decomposition; Conv - the participant selected the conversation-level decomposition; Dom - the participant selected the domain-level decomposition. Figure 13(a), gives for each service the percentage of the participants that made a particular choice. Finally, Figure 13(b) gives for each participant the percentage of the services for which he/she made a particular choice.

(a) The choices made for each service.

(b) The choices made by each participant.
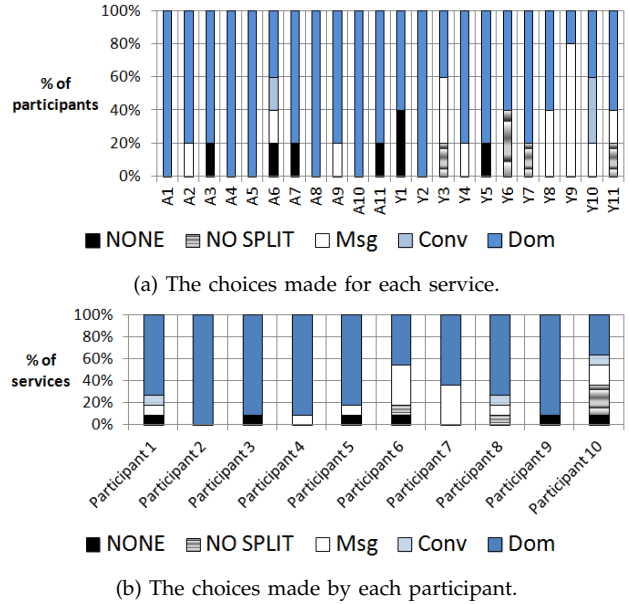
Fig. 13. Usefulness from the developers' perspective.

RQ1: Concerning the first question, the participants suggested to decompose most of the examined interfaces. The only exceptions are Y3, Y6, Y7 and Y11 (Figure 13(a)). For Y3, Y7 and Y11, one of the participants suggested to leave the interface as is, while the others were in favor of decomposing the interface. For Y6, two of the participants suggested to leave the interface as is. For most of the service interfaces, the participants selected decompositions that were among the ones that we provided. In the Amazon services, we have 4 cases (A3, A6, A7, A11 in Figure 13(a)), for which one of the participants was not satisfied by any of the provided decompositions. For A6 and A7, the participants proposed their own

decompositions. Specifically, the proposed decomposition for A6 was: "3 interfaces to manage fulfilment, items and shipments". For A7 the suggestion was: "3 interfaces for objects, buckets, access policies". In the Yahoo services, we have the case of Y1 (Figure 13(a)), for which two of the participants were not satisfied by any of the provided decompositions. The participants pointed out that the proposed decompositions do not separate clearly the underlying concepts (keywords, bids, adGroups, optimization guidelines). Moreover, in the Yahoo services we have the case of Y5 (Figure 13(a)), for which one of the participants proposed his own decomposition: "three sets of operations one for the standard accounts, one for the mobile accounts and one for the credit cards".

RQ2: Regarding the second question, the domain-level cohesion metric worked very well for the participants. Specifically, in 63% of the services, more than 80% of the participants selected the domain-level decomposition (Figure 13(a)). Concerning each one of the participants, the percentage of the services for which the domain-level decomposition was selected ranges from 36% to 100% (Figure 13(b)). On the other hand, the percentage of the services for which the message-level decomposition was selected ranges from 0% to 36% (Figure 13(b)). Finally, the percentage of the services for which the conversation-level decomposition was selected ranges from 0% to 9% (Figure 13(b)).

RQ3: Concerning the third question, in several cases the participants did not suggest any changes in the selected decompositions. In the 55 decompositions that have been chosen for the Amazon services (5 participants × 11 services) there were 21 such occurrences; in the 55 decompositions of the Yahoo services this amounted to 18 occurrences. However, we also have several cases for which the participants moved certain operations between interfaces. In the 55 decompositions that have been chosen for the Amazon services there were 20 such occurrences; in the 55 Yahoo decompositions, this amounted to 14 occurrences. Moreover in several cases the participants decreased the size of the decompositions by merging certain interfaces. Specifically, we had 21 occurrences for the Amazon services and 24 occurrences for the Yahoo services. The details of the individual participants suggestions are found in Appendix D.

To conclude this study, we performed a X^2 test, so as to check the statistical significance of the results. The goal of the test was to examine the following null hypothesis:

H_0: The choices that have been made by the participants are not significantly different from the ones that we would have by chance alone.

Table 4 provides the details for the X^2 test that we performed. In particular, the first row of Table 4 gives the percentages of NO-SPLIT, NONE, Msg, Conv and Dom that we observed in the study in the overall 110

TABLE 4
X^2 test for the overall results.

	NO-SPLIT	NONE	Msg	Conv	Dom
Observed	4.50%	13.51%	2.70%	73.87%	5.40%
Expected	20.00%	20.00%	20.00%	20.00%	20.00%
Squared diffs.	12.00	2.10	14.95	145.11	10.65
X^2	184.83				
ρ **value**	6.81E-39				

choices that have been made by the participants (22 services × 5 participants per service). The second row of Table 4 gives the percentages of NO-SPLIT, NONE, Msg, Conv and Dom, that we would have by chance alone. Based on the squared differences between the expected and the observed percentages, the overall X^2 value that we got is 184.83. Then, according to the X^2 distribution, the probability of having a X^2 as large as 184.83, by chance alone, is too small ($\rho \ll 0.001$). Therefore, we rejected H_0.

5.3 Threats to Validity

A possible threat to the internal validity of the results that we obtained from the developers' involved in the validation is the developers' fatigue or boredom [33]. To reduce this threat we arranged our study according to the developers' availability, instead of imposing a strict schedule. To avoid effects caused by interactions between the developers [33], we made clear that the required tasks should not be performed in a collaborative manner. Finally, to avoid learning effects [33], the different decompositions of each interface were provided to the developers in a random order.

External validity concerns whether the results of a study can be generalized to a wide population [33]. Regarding external validity, the positive aspects of our validation are the following:

- It is among the very few ones [34], [6] that involve real services. Specifically, we used a representative set of services, provided by two major service providers; the services offer diverse functionalities and their interfaces vary in size and complexity.
- It was based on a representative set of developers that have knowledge of the service-oriented computing paradigm, related technologies and standards.

On the other hand, a possible limitation is that the validation was not based on a large number of developers. Nevertheless, the number of developers that we considered is comparable with other similar studies [5], [9], [7].

6 CONCLUSION

In this paper, we have proposed an approach that enables the cohesion-driven decomposition of service interfaces, without information of how the services

are implemented. The proposed approach iteratively decomposes a given service interface in a set of more cohesive interfaces. We validated the approach in 22 real-world services provided by Amazon and Yahoo. Our findings showed that the proposed approach is able to improve cohesion. The number of interfaces produced by the approach linearly increases with the size of the decomposed interface. In general, the developers found the proposed approach useful. As anticipated, the decompositions produced by the method are not perfectly adjusted to the developers' needs. In certain cases the developers would prefer smaller and even more cohesive decompositions. Future work can be pursued towards avoiding unnecessary splits, interactiveness with the user and working with semantically annotated services.

REFERENCES

[1] W. Stevens, G. Myers, and L. Constantine, "Structured Design," *IBM Systems Journal*, vol. 13, no. 2, pp. 115–139, 1974.
[2] M. Papazoglou and W.-J. van den Heuvel, "Service-Oriented Design and Development Methodology," *International Journal on Web Engineering and Technology*, vol. 2, no. 4, pp. 412–442, 2006.
[3] C. Legner and T. Vogel, "Design Principles for B2B Services - An Evaluation of two Alternative Service Designs," in *Proceedings of the IEEE International Conference on Service Computing (SCC)*, 2007, pp. 372–379.
[4] T. Kohlborn, A. KortHaus, T. Chan, and M. Rosemann, "Identification and Analysis of Business and Software Services - A Consolidated Apporach," *IEEE Transactions on Services Computing*, vol. 2, no. 1, pp. 1–15, 2009.
[5] M. Perepletchikov, C. Ryan, and Z. Tari, "The Impact of Service Cohesion on the Analyzability of Service-Oriented Software," *IEEE Transactions on Services Computing*, vol. 3, no. 2, pp. 89–103, 2010.
[6] D. Athanasopoulos and A. Zarras, "Fine-Grained Metrics of Cohesion Lack for Service Interfaces," in *Proceedings of the 9th IEEE International Conference on Web Services (ICWS)*, 2011, pp. 588–595.
[7] N. Tsantalis and A. Chatzigeorgiou, "Identification of Move Method Refactoring Opportunities," *IEEE Transactions on Software Engineering*, vol. 99, no. 3, pp. 347–367, 2009.
[8] F. Simon, F. Steinbrückner, and C. Lewerentz, "Metrics Based Refactoring," in *Proceedings of the 5th IEEE European Conference on Software Maintenance and Reengineering (CSMR)*, 2001, pp. 30–39.
[9] M. Fokaefs, N. Tsantalis, A. Chatzigeorgiou, and J. Sander, "Decomposing Object-Oriented Class Modules Using an Agglomerative Clustering Technique," in *Proceedings of the IEEE International Conference on Software Maintenance (ICSM)*, 2009, pp. 93–101.
[10] D. Doval, S. Mancoridis, and B. S. Mitchell, "Automatic Clustering of Software Systems Using a Genetic Algorithm," in *Proceedings of the 9th IEEE Software Technology and Engineering Practice (STEP)*, 1999, pp. 73–81.
[11] B. D. Bois, S. Demeyer, and J. Verelst, "Refactoring: Improving Coupling and Cohesion of Existing Code," in *Proceedings of the 11th IEEE Working Conference on Reverse Engineering (WCRE)*, 2004, pp. 144–151.
[12] O. Seng, J. Stammel, and D. Burkhart, "Search-Based Determination of Refactorings for Improving the Class Structure of Object-Oriented Systems," in *Proceedings of the 8th Annual Conference on Genetic and Evolutionary Computation (GECCO)*, 2006, pp. 1909–1916.
[13] L. Tahvildari and K. Kontogiannis, "Improving Design Quality Using Meta-Pattern Transformations: A Metric-Based Approach," *Journal of Software Maintenance*, vol. 16, no. 4-5, pp. 331–361, 2004.

[14] M. O'Keeffe and M. í Cinnéide, "Search-Based Refactoring for Software Maintenance," *Journal of Systems and Software*, vol. 81, no. 4, pp. 502–516, 2008.
[15] M. Bowman, L. C. Briand, and Y. Labiche, "Solving the Class Responsibility Assignment Problem in Object-Oriented Analysis with Multi-Objective Genetic Algorithms," *IEEE Transactions on Software Engineering*, vol. 36, no. 6, pp. 817–837, 2010.
[16] W. F. Opdyke, "Refactoring Object-Oriented Frameworks," Ph.D. dissertation, Univ. of Illinois - Urbana Champaign, 1992.
[17] Y. Kataoka, M. D. Ernst, W. G. Griswold, and D. Notkin, "Automated Support for Program Refactoring using Invariants," in *Proceedings of the IEEE International Conference on Software Maintenance (ICSM)*, 2001, pp. 736–746.
[18] E. Tilevich and Y. Smaragdakis, "Binary Refactoring: Improving Code Behind the Scenes," in *Proceedings of the 27th International Conference on Software Engineering (ICSE)*, 2005, pp. 264–273.
[19] L. Vidács, M. Gogolla, and R. Ferenc, "From C++ Refactorings to Graph Transformations," in *Proceedings of the 3rd Workshop on Software Evolution through Transformations*, 2006.
[20] D. Dig, C. Comertoglu, D. Marinov, and R. E. Johnson, "Automated Detection of Refactorings in Evolving Components," in *Proceedings of 20th European Conference on Object-Oriented Programming (ECOOP)*, 2006, pp. 404–428.
[21] P. Anbalagan and T. Xie, "Automated Inference of Pointcuts in Aspect-Oriented Refactoring," in *Proceedings of the 29th international conference on Software Engineering (ICSE)*, 2007, pp. 127–136.
[22] D. Binkley, M. Ceccato, M. Harman, F. Ricca, and P. Tonella, "Tool-Supported Refactoring of Existing Object-Oriented Code into Aspects," *IEEE Transactions on Software Engineering*, vol. 32, no. 9, pp. 698–717, 2006.
[23] T. Mens and T. Tourwé, "A Survey of Software Refactoring," *IEEE Transactions on Software Engineering*, vol. 30, no. 2, pp. 126–139, 2004.
[24] M. Harman and L. Tratt, "Pareto Optimal Search Based Refactoring at the Design Level," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2007, pp. 1106–1113.
[25] S. R. Chidamber and C. F. Kemerer, "A Metrics Suite for Object Oriented Design," *IEEE Transactions on Software Engineering*, vol. 20, no. 6, pp. 476–493, 1994.
[26] M. Harman, "Search-Based Software Engineering," *Information and Software Technology*, vol. 43, no. 14, pp. 833–839, 2001.
[27] J. Clarke, J. J. Dolado, M. Harman, R. Hierons, B. Jones, M. Lumkin, B. Mitchell, S. Mancoridis, K. Rees, M. Roper, and M. Shepperd, "Reformulating Software Engineering as a Search Problem," *IEE Proceedings Software*, vol. 150, no. 3, pp. 161–175, 2003.
[28] L. C. Briand, J. W. Daly, and J. Wüst, "A Unified Framework for Cohesion Measurement in Object-Oriented Systems," *Empirical Software Engineering*, vol. 3, no. 1, pp. 65–117, 1998.
[29] J. A. Dallal and L. Briand, "A Precise Method-Method Interaction-Based Cohesion Metric for Object-Oriented Classes," *ACM Transactions on Software Engineering and Methodology*, vol. 21, no. 2, 2012.
[30] A. Kazemi, A. Rostampour, A. Zamiri, P. Jamshidi, H. Haghighi, and F. Shams, "An Information Retrieval Based Approach for Measuring Service Conceptual Cohesion," in *Proceedings of the 11th IEEE International Conference on Quality Software (QSIC)*, pp. 102–111.
[31] L. Briand, S. Morasca, and V. R. Basili, "Property-Based Software Engineering Measurement," *IEEE Transactions on Software Engineering*, vol. 22, no. 1, pp. 68–86, 1996.
[32] W3C, "Web Services Architecture," W3C, Tech. Rep., http://www.w3c.org/TR/ws-arch.
[33] W. R. Shadish, T. D. Cook, and D. T. Campbell, *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Houghten Mifflin, 2002.
[34] M. Fokaefs, R. Mikhaiel, N. Tsantalis, E. Stroulia, and A. Lau, "An Empirical Study on Web Service Evolution," in *Proceedings of the 9th IEEE International Conference on Web Services (ICWS)*, 2011, pp. 49–56.

Appendices to:
**Cohesion-Driven Decomposition of Service
Interfaces Without Access to Source Code**

Dionysis Athanasopoulos, Apostolos V. Zarras,
George Miskos, Valerie Issarny, and Panos Vassiliadis

APPENDIX A
SUMMARY OF RELATED WORK

Table 5, provides a summary of metrics-driven refactoring approaches that have been proposed in the past and highlights the contribution of our approach for the cohesion-driven decomposition of service interfaces. Moreover, Table 6, briefly summarizes the relation between the cohesion metrics that we employ in our approach, the object-oriented cohesion metrics surveyed in [29], and the service-oriented cohesion metrics that have been proposed in [5], [6], [30].

TABLE 5
A summary of metrics-driven refactoring approaches.

Refactoring Approach	Purpose	Type of relations
[24]	Class coupling	Implementation-level
[7]	Class cohesion	Implementation-level
[8]	Class cohesion	Implementation-level
[9]	Class cohesion	Implementation-level
[10]	Class coupling, cohesion	Implementation-level
[11]	Class coupling, cohesion	Implementation-level
[12]	Class coupling, cohesion, code complexity	Implementation-level
[13]	Class coupling, cohesion, code complexity	Implementation-level
[14]	Class coupling, cohesion, code size	Implementation-level
[15]	Class coupling, cohesion, code size, code complexity	Implementation-level
Proposed approach	Service interfaces cohesion	Interface-level

TABLE 6
A summary of cohesion metrics [29], [5], [6], [30].

	Implementation-level	Interface-level
Class cohesion	$LCOM1$, $LCOM2$, $LCOM3$, $LCOM4$, $LCOM5$, TCC, LCC, DCD, DCI, CC, $SCOM$, DMC, $CBMC$, $LSCC$	$CAMC$, NHD, $SNHD$, $MMAC$ $CAMC$,
Service cohesion	$SIIC$, SCV	$SIDC$, $SISC$, $SIUC$
		Proposed metrics: LoC_{msg}, LoC_{conv}, LoC_{dom}

APPENDIX B
ANALYTIC VALIDATION OF COHESION METRICS

In the 90's, Briand et al. [31] proposed a mathematical framework for the theoretical validation of cohesion metrics. Here, we rely on this framework for the

validation of the metrics that we employ for the cohesion-driven decomposition of service interfaces. Briefly, in [31], a software system is represented by a graph $S = (E, R)$, where E is the set of elements that constitute the system and $R \subseteq E \times E$ is a set of relations between elements. A module of the system is represented by a graph $m = (E_m, R_m)$, where $E_m \subseteq E$, and $R_m \subseteq R$. According to Briand et al., a cohesion metric has to satisfy the following properties:

- *Nonnegativity and normalization*: The cohesion of a module $m = (E_m, R_m)$ belongs to a specified interval, i.e., $cohesion(m) \in [0, M]$.
- *Null value*: The cohesion of a module $m = (E_m, R_m)$ is null if R_m is empty, i.e., $R_m = \emptyset \Rightarrow cohesion(m) = 0$.
- *Monotonicity*: Let $m = (E_m, R_m)$ and $m' = (E_m, R_{m'})$ be two modules (with the same set of elements), such that $R_m \subseteq R_{m'}$. Then, $cohesion(m) \leq cohesion(m')$.
- *Cohesive modules*: Let $m_1 = (E_{m_1}, R_{m_1})$ and $m_2 = (E_{m_2}, R_{m_2})$ be two unrelated modules and $m_{1 \cup 2}$ is the union of m_1, m_2. Then, $cohesion(m_{1 \cup 2}) \leq \max(cohesion(m_2), cohesion(m_2))$.

For brevity, we focus our validation on LoC_*. The three different refinements of LoC_* can be validated identically. Differently from [31], the metrics that we consider measure the *lack of cohesion* of service interfaces. Moreover, the interface-level graphs that we employ are *weighted*. Hence, we appropriately adapt the properties that should hold for the proposed metrics. We begin our validation with the proof of a supportive lemma, which concerns the similarity functions that we employ. Then, we prove that LoC_* satisfies the properties of the Briand et al. framework.

Lemma 1: The similarity functions that we use for the different notions of cohesion belong to the interval $[0, 1]$.

Proof: In the case of *message-level cohesion*, the similarity, $OpS_{msg}(op_i, op_j)$, between two operations is the average of the similarities between the input/output messages of op_i, op_j. The similarity, $MsgS(m_i, m_j)$, between two messages is measured based on the message-level graphs G_{m_i}, G_{m_j} of the messages. On the one extreme, the maximum common subgraph $G_{m_i \cap m_j}$ of G_{m_i}, G_{m_j} may be trivial if G_{m_i}, G_{m_j} have nothing in common. In this case we have:

$$|V_{m_i \cap m_j}| = 0 \tag{1}$$

On the other extreme, G_{m_i}, G_{m_j} may be identical, in which case we have:

$$|V_{m_i}| = |V_{m_j}| = |V_{m_i \cap m_j}| = |V_{m_i \cup m_j}| \tag{2}$$

From (1) and (2) we have:

$$0 \leq MsgS(m_i, m_j) \leq 1 \Rightarrow$$
$$0 \leq OpS_{msg}(op_i, op_j) \leq 1 \tag{3}$$

In *conversation-level cohesion*, the similarity $OpS_{conv}(op_i, op_j)$ is also an average of message similarities. Hence:

$$0 \leq OpS_{conv}(op_i, op_j) \leq 1 \qquad (4)$$

In *domain-level cohesion*, the similarity $OpS_{dom}(op_i, op_j)$ is measured based on the sets of domain-level terms T_{op_i} and T_{op_j}. On the one extreme, T_{op_i} and T_{op_j} may have nothing in common. In this case we have:

$$|T_{op_i} \cap T_{op_j}| = 0 \qquad (5)$$

On the other extreme, T_{op_i} and T_{op_j} may be identical, in which case we have:

$$|T_{op_i}| = |T_{op_j}| = |T_{op_i} \cup T_{op_j}| = |T_{op_i} \cap T_{op_j}| \quad (6)$$

(5) and (6) imply that:

$$0 \leq OpS_{dom}(op_i, op_j) \leq 1 \qquad (7)$$

□

Theorem 1: For a service interface, si, and the interface-level graph, $G^*_{si} = (V_{si}, E_{si})$, that represents the interface, $0 \leq LoC_*(si, OpS_*) \leq 1$.

Proof: Based on Lemma 1, for any two operations op_i, op_j of si we have:

$$0 \leq OpS_*(op_i, op_j) \leq 1 \qquad (8)$$

From graph theory we further have:

$$|E_{si}| \leq \frac{|V_{si}| * (|V_{si}| - 1)}{2} \qquad (9)$$

Based on (8) and (9) the following holds:

$$0 \leq \frac{\sum_{\forall (op_i, op_j) \in E_{si}} OpS_*(op_i, op_j)}{\frac{|V_{si}| * (|V_{si}| - 1)}{2}} \leq 1 \Rightarrow$$
$$0 \leq LoC_*(si, OpS_*) \leq 1 \qquad (10)$$

□

Theorem 2: Let si be a service interface, represented by the interface-level graph, $G^*_{si} = (V_{si}, E_{si})$. If $E_{si} = \emptyset$, then $LoC_*(si, OpS_*) = 1$.

Proof: Given that E_{si} is empty we have:

$$E_{si} = \emptyset \Rightarrow \sum_{\forall (op_i, op_j) \in E_{si}} OpS_*(op_i, op_j) = 0 \Rightarrow$$
$$LoC_*(si, OpS_*) = 1 \qquad (11)$$

□

Theorem 3: Let si, si' be two service interfaces, represented by the interface-level graphs, $G^*_{si} = (V_{si}, E_{si})$, $G^*_{si} = (V_{si}, E_{si'})$. G^*_{si} and $G^*_{si'}$ have the same nodes. Moreover, $E_{si} \subseteq E_{si'}$ and $\sum_{\forall (op_i, op_j) \in E_{si}} OpS_*(op_i, op_j) \leq \sum_{\forall (op_i, op_j) \in E_{si'}} OpS_*(op_i, op_j)$. Then, $LoC_*(si, OpS_*) \geq LoC_*(si', OpS_*)$.

Proof: Given the initial assumptions of the theorem for si and si' (i.e., G^*_{si} and $G^*_{si'}$ have the same nodes, $E_{si} \subseteq E_{si'}$, and $\sum_{\forall (op_i, op_j) \in E_{si}} OpS_*(op_i, op_j) \leq \sum_{\forall (op_i, op_j) \in E_{si'}} OpS_*(op_i, op_j)$), the following implications hold:

$$\frac{\sum_{\forall (op_i, op_j) \in E_{si}} OpS_*(op_i, op_j)}{\frac{|V_{si}| * (|V_{si}| - 1)}{2}} \leq$$
$$\frac{\sum_{\forall (op_i, op_j) \in E_{si'}} OpS_*(op_i, op_j)}{\frac{|V_{si}| * (|V_{si}| - 1)}{2}} \Rightarrow$$
$$1 - \frac{\sum_{\forall (op_i, op_j) \in E_{si}} OpS_*(op_i, op_j)}{\frac{|V_{si}| * (|V_{si}| - 1)}{2}} \geq$$
$$1 - \frac{\sum_{\forall (op_i, op_j) \in E_{si'}} OpS_*(op_i, op_j)}{\frac{|V_{si}| * (|V_{si}| - 1)}{2}} \Rightarrow$$
$$LoC_*(si, OpS_*) \geq LoC_*(si', OpS_*) \qquad (12)$$

□

Theorem 4: Let si_1, si_2 be two unrelated service interfaces, represented by the interface-level graphs, $G^*_{si_1} = (V_{si_1}, E_{si_1})$, $G^*_{si_2} = (V_{si_2}, E_{si_2})$. Let $si_{1 \cup 2}$ be the union of si_1, si_2, represented by, $G^*_{si_{1 \cup 2}} = (V_{si_{1 \cup 2}}, E_{si_{1 \cup 2}})$. Then, $LoC_*(si_{1 \cup 2}, OpS_*) \geq \max(LoC_*(si_1, OpS_*), LoC_*(si_2, OpS_*))$.

Proof: Without loss of generality, we assume that si_2 is more cohesive than si_1. Based on this assumption, we have:

$$LoC_*(si_1, OpS_*) \geq LoC_*(si_2, OpS_*) \Rightarrow$$
$$\frac{\sum_{(op_i, op_j) \in E_{si_1}} OpS_*(op_i, op_j)}{\frac{|V_{si_1}| * (|V_{si_1}| - 1)}{2}} \geq$$
$$\frac{\sum_{(op_i, op_j) \in E_{si_2}} OpS_*(op_i, op_j)}{\frac{|V_{si_2}| * (|V_{si_2}| - 1)}{2}} \Rightarrow$$
$$|V_{si_2}| * (|V_{si_2}| - 1) * \sum_{E_{si_1}} OpS_*(op_i, op_j) \geq$$
$$|V_{si_1}| * (|V_{si_1}| - 1) * \sum_{E_{si_2}} OpS_*(op_i, op_j) \quad (13)$$

Given that si_1, si_2 are unrelated we have that:

$$\forall (op_{si_1}, op_{si_2}) \in V_{si_1} \times V_{si_2},$$
$$OpS_*(op_{si_1}, op_{si_2}) = 0 \qquad (14)$$

From (14) we derive the following for the interface-level graph $G^*_{si_{1 \cup 2}}$ that represents the union of si_1, si_2:

$$V_{si_{1 \cup 2}} = V_{si_1} \cup V_{si_2} \qquad (15)$$
$$E_{si_{1 \cup 2}} = E_{si_1} \cup E_{si_2} \qquad (16)$$

From (15), (16) we have that:

$$LoC_*(si_{1\cup2}, OpS_*) = \tag{17}$$
$$1 - \frac{\sum_{E_{si_1}} OpS_*(op_i, op_j) + \sum_{E_{si_2}} OpS_*(op_i, op_j)}{\frac{(|V_{si_1}|+|V_{si_2}|)*(|V_{si_1}|+|V_{si_2}|-1)}{2}}$$

Given (17), to prove the theorem we have to show that the following inequality holds:

$$1 - \frac{\sum_{E_{si_1}} OpS_*(op_i, op_j) + \sum_{E_{si_2}} OpS_*(op_i, op_j)}{\frac{(|V_{si_1}|+|V_{si_2}|)*(|V_{si_1}|+|V_{si_2}|-1)}{2}} \geq$$
$$1 - \frac{\sum_{(op_i,op_j)\in E_{si_1}} OpS_*(op_i, op_j)}{\frac{|V_{si_1}|*(|V_{si_1}|-1)}{2}} \tag{18}$$

From (18), with trivial algebraic operations, we derive the following inequality that must hold to prove the theorem:

$$2 * |V_{si_1}| * |V_{si_2}| + \tag{19}$$
$$|V_{si_2}| * (|V_{si_2}| - 1) * \sum_{(op_i,op_j)\in E_{si_1}} OpS_*(op_i, op_j) \geq$$
$$|V_{si_1}| * (|V_{si_1}| - 1) * \sum_{(op_i,op_j)\in E_{si_2}} OpS_*(op_i, op_j)$$

Given that (13) holds, (19) is also true.

□

APPENDIX C
DECOMPOSITION METHOD TERMINATION & COMPLEXITY

The analysis of the proposed decomposition method focuses on two issues. First, we prove that the cohesion-driven decomposition of service interfaces terminates. Second, we show that the complexity of decomposing a given interface with the proposed method is, in the worst case, cubic to the number of operations, offered by the given interface.

Theorem 5: Given a service interface si, $decomposeInterface$ (Algorithm 1) terminates.

Proof: decomposeInterface performs a number of iterations, until the size of Q is 0. During each iteration, $decomposeInterface$ picks a service interface r_i from Q. If the cohesion of r_i can not be improved, r_i is put in the results set R_I. Otherwise, the cohesion of r_i is improved by splitting it in two new interfaces r_r, r_s, which are stored in Q. $decomposeInterface$ can not perform infinite splits because:

- The lower bound for the lack of cohesion of a service interface is 0 (Theorem 1).
- The lower bound for the number of operations of a service interface is 1.

Therefore, the size of Q eventually becomes 0 and $decomposeInterface$ terminates.

□

Theorem 6: In the worst case, the complexity of decomposing a service interface, si, is cubic to the number of operations of si.

Proof: In the worst case scenario, $decomposeInterface$ (Algorithm 1) starts with si that provides $|si.O|$ operations and results in $|si.O|$ new interfaces, one per operation. To achieve this, the algorithm starts with si and splits it in two new interfaces r_r and r_s. The new interfaces are enqueued in Q. In the i-th iteration, one of the intermediate interfaces, r_i is chosen and split again. Hence, at the end of the i-th iteration, Q contains $i+1$ new interfaces. Once, the size of Q becomes $|si.O|$, there are another $|si.O|$ iterations to dequeue the interfaces that are held in Q (again in the worst case). Therefore, in the worst case $decomposeInterface$ performs $2 * |si.O|$ iterations.

The two factors that affect the complexity of splitting an intermediate interface r_i in two interfaces is the creation (Algorithm 2) and population of the splinter interface r_s (Algorithm 3).

- $createSplinter$ performs $|r_i.O|$ iterations to find the operation, op_s, whose removal maximizes the cohesion improvement of r_i. Then, it creates r_s that contains op_s, and r_r that contains the rest of the r_i operations.
- $populateSplinter$, takes as input the interfaces r_r, r_s that result from $createSplinter$. Hence, $|r_r.O| = |r_i.O| - 1$ and $|r_s.O| = 1$. To improve the cohesion of the two interfaces $populateSplinter$ moves operations from r_r to r_s. In the worst case, we can have $|r_r.O| - 1 = |r_i.O| - 2$ operations moved, i.e., the outer loop of $populateSplinter$ performs $|r_i.O| - 1$ iterations. To find the first operation, the inner loop of $populateSplinter$ performs $|r_i.O| - 1$ iterations. To find the i-th operation, the inner loop of $populateSplinter$ performs $|r_i.O| - 1 - i + 1$ iterations, and so on. Therefore, the overall number of iterations performed is $\sum_{i=1}^{|r_i.O|-1} |r_i.O| - i = \sum_{i=1}^{|r_i.O|-1} i = \frac{|r_i.O|*(|r_i.O|-1)}{2}$.

Based on the previous analysis, in the worst case the complexity of decomposing si is $O(|si.O|^3)$.

□

APPENDIX D
INDIVIDUAL PARTICIPANTS' SUGGESTIONS FOR IMPROVEMENT

Tables 7, 8 give a detailed summary of the changes that have been performed by the participants on the decompositions that they selected. In particular, for each one of the examined interfaces and each participant we provide the percentage of the moved operations (over the size of the examined interface) and the percentage of the decomposition size decrease.

Overall, the percentage of moved operations ranged from 1.15% to 15.38% whenever this happened in

TABLE 7
Amazon services: Changes per participant: % of moved operations and % of decomposition size decrease.

ID	Participant 1		Participant 2		Participant 3		Participant 4		Participant 5	
	% move oper.	% DS decr.	% move oper.	% DS decr.	% move oper.	% DS decr.	% move oper.	% DS decr.	% move oper.	% DS decr.
A1	01.15	14.81	02.30	03.70	03.45	11.11	02.30	14.81	00.00	07.41
A2	03.70	00.00	00.00	10.00	03.70	10.00	03.70	10.00	00.00	30.00
A3	00.00	00.00	07.41	06.25	07.41	12.50	00.00	00.00	00.00	16.67
A4	04.35	16.67	00.00	16.67	04.35	16.67	04.35	16.67	00.00	00.00
A5	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00
A6	00.00	00.00	00.00	00.00	-	-	00.00	00.00	00.00	25.00
A7	-	-	06.25	00.00	06.25	00.00	12.50	16.67	-	-
A8	00.00	16.67	00.00	16.67	00.00	00.00	00.00	16.67	00.00	16.67
A9	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00
A10	07.69	00.00	07.69	00.00	15.38	00.00	07.69	00.00	00.00	00.00
A11	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	15.38	00.00

TABLE 8
Yahoo services: Changes per participant: % of moved operations and % of decomposition size decrease.

ID	Participant 6		Participant 7		Participant 8		Participant 9		Participant 10	
	% move oper.	% DS decr.	% move oper.	% DS decr.	% move oper.	% DS decr.	% move oper.	% DS decr.	% move oper.	% DS decr.
Y1	00.00	20.00	00.00	00.00	00.00	13.33	-	-	-	-
Y2	00.00	11.11	00.00	00.00	00.00	11.11	03.57	22.22	00.00	33.33
Y3	14.29	00.00	07.14	18.18	10.71	14.29	00.00	14.29	00.00	28.57
Y4	00.00	00.00	13.04	00.00	17.39	12.50	13.04	12.50	00.00	20.00
Y5	-	-	00.00	00.00	05.00	18.18	00.00	09.09	-	-
Y6	-	-	00.00	11.11	00.00	22.22	00.00	11.11	-	-
Y7	00.00	00.00	00.00	00.00	00.00	11.11	00.00	11.11	00.00	33.33
Y8	08.33	00.00	16.67	00.00	25.00	14.29	00.00	00.00	08.33	25.00
Y9	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00	00.00
Y10	00.00	25.00	20.00	00.00	00.00	00.00	20.00	00.00	00.00	00.00
Y11	00.00	00.00	00.00	00.00	-	-	00.00	00.00	00.00	00.00

Amazon services and 03.57% to 25% for the Yahoo services. The percentage of the decomposition size decrease ranged from 3.70% to 25% for Amazon and 11.11% to 33.33% for Yahoo services.