

# Cost-Effective Deployment of Certified Cloud Composite Services

Marco Anisetti, Claudio A. Ardagna, Filippo Gaudenzi

*DI – Università degli Studi di Milano, Italy*

Ernesto Damiani

*EBTIC – Khalifa University, UAE*

Gwanggil Jeon

*Incheon National University, Department of Embedded Systems Engineering, Incheon, Republic of Korea*

---

## Abstract

The advent of cloud computing has radically changed the concept of distributed environments, where services can now be composed and reused at high rates. Today, service composition in the cloud is driven by the need of providing stable QoS, where non-functional properties of composite services are proven over time and composite services continuously adapt to both functional and non-functional changes of the component services. This scenario introduces substantial costs on the cloud providers that go beyond the cost of deploying component services, and require to consider the costs of continuously verifying non-functional properties of composite and component services. In this paper, we propose a cost-effective approach to certification-based cloud service composition. This approach is based, on one side, on a portable certification process for the cloud evaluating non-functional properties of composite services and, on the other side, on a cost-evaluation methodology aimed to produce the service composition that minimizes the total cost paid by the cloud providers, taking into

---

<sup>1</sup>Email: *firstname.lastname@unimi.it*

<sup>2</sup>Email: *ernesto.damiani@kustar.ac.ae*

<sup>3</sup>Email: *gjeon@inu.ac.kr*

account both deployment and certification/verification costs. Our service composition approach is driven by certificates awarded to single services and by a fuzzy-based cost evaluation methodology, and assumes certified properties as *must-have* requirements for service selection and composition.

*Keywords:* Cloud, Certification, Cost Optimization, Service composition, Security

---

## 1. Introduction

The maturity reached by cloud computing has fostered the implementation of a number of distributed infrastructure, platform, and application services available worldwide. Current trends in software distribution and provisioning envision services made available as commodities over distributed systems including the Internet or the cloud marketplace. At the same time, the trend towards coarse-granularity business services, which cannot be managed by a single entity, resulted in several approaches to service composition that maximize software re-use by dynamically composing single services on the basis of their functionalities [1].

A major challenge faced by distributed service-based systems deployed on the cloud goes beyond the ability to guarantee the functionality of composite services, and must consider the importance of guaranteeing stable Quality of Service (QoS) in the form of non-functional properties requirements such as security, performance, and trust [2]. Service compositions need to guarantee optimal and verifiable properties, managing different events that might change their structure such as component relocation, substitution, malfunctioning, versioning, adaptation [3]. Continuous monitoring and verification of service non-functional properties is needed and usually achieved by means of assurance techniques [4, 5]. Recently, certification-based assurance techniques have been introduced to guarantee stable QoS in the cloud [6, 7, 8, 9, 10, 11]. They are based on continuous collection of evidence on the behavior of the system, which is used to verify whether the considered system holds a specific

24 (set of) non-functional property and award a certificate proving it. To this  
25 aim, distributed agents are instrumented to connect to different endpoints in  
26 the cloud and retrieve evidence used to evaluate the non-functional status of  
27 the target cloud-based system. Current certification techniques mostly focus on  
28 the certification of single-service systems and often do not consider the cost of  
29 maintaining stable QoS. Even worse, a trend in service composition is to pro-  
30 vide an ad hoc composite service for each request, with high costs on the cloud  
31 providers (CPs).

32 In this scenario, two colliding requirements emerge. On one side, there is the  
33 need to guarantee non-functional properties of a service composition. This is a  
34 challenging task that requires continuous evaluation of compositions at cloud-  
35 provider side, to accomplish the dynamic and evolving nature of the cloud. On  
36 the other side, there is the need to take the costs observed by cloud providers for  
37 certified composition management under control. These costs, in fact, rapidly  
38 increase because the costs of continuous certification and verification become  
39 substantial. Current research on cloud computing has privileged solutions min-  
40 imizing costs on the final users [12, 13, 14], neglecting the costs on the cloud  
41 providers that often represent a major source of fee increase.

42 In this paper, we propose the first cost-effective approach to certification-  
43 based cloud service composition that addresses the above problems. It is inspired  
44 by our previous work in [4] and extends it according to the cloud challenges dis-  
45 cussed in [11] and [15]. Differently from existing work [12, 13], our service  
46 composition approach is driven by certificates awarded to single services and  
47 by a fuzzy-based cost evaluation methodology, and assumes certified properties  
48 as must-have requirements for service selection [16] and composition [17]. This  
49 methodology aims to decrease the costs of cloud providers, also analyzing those  
50 costs introduced by the need of keeping the composition continuously moni-  
51 tored and certified. More specifically, the cost of deploying a certified service  
52 composition includes *i) direct costs*, traditional costs of service deployment on  
53 the cloud, or costs of third-party services building the composition (i.e., multi-  
54 cloud composition), *ii) indirect costs*, the costs introduced by the certification

55 infrastructure to continuously monitor certificate validity, *iii) mismatch costs*,  
56 the costs modeling the discrepancy between what was agreed in terms of certi-  
57 fied properties and what was actually provided. The mismatch costs are often  
58 neglected by existing approaches. They evaluate the additional costs observed  
59 by a CP when sharing a service whose properties in the certificate are stronger  
60 than the properties requested by a composite service. For instance, providing a  
61 storage service ensuring end-to-end confidentiality, while just confidentiality of  
62 data at rest is requested, means that resources for confidentiality in transit are  
63 overspent without a real revenue.<sup>4</sup>

64 The contribution of this paper is twofold. First, we present a certification  
65 process for composite services that fits the dynamics of the cloud (Sections 3  
66 and 4). Our process guarantees continuous monitoring of certified properties,  
67 evaluating certificate validity over time and portability across different deploy-  
68 ments. Second, after introducing the cost factors and profiles affecting the  
69 costs of cloud providers (Section 5), we provide a fuzzy-based cost evaluation  
70 methodology at the basis of a cost-effective, certification-based cloud service  
71 composition approach (Section 6). Our approach selects component services  
72 on the basis of their certified properties and is implemented by means of two  
73 run-time heuristics for composition cost minimization, which are experimentally  
74 evaluated in terms of quality and performance (Section 7). It contributes to the  
75 resolution of the long-standing problem of managing non-functional properties  
76 of distributed applications and composite services in a cost-effective way. It  
77 provides an approach that effectively relocates and refines service compositions  
78 in the cloud at run time guaranteeing stable QoS.

## 79 **2. Problem Statement**

80 Our reference model is a cloud infrastructure where single services are com-  
81 posed to form complex services and certification-based assurance techniques

---

<sup>4</sup>We note that mismatch cost is a crucial metric for internal cost optimization, permitting a more effective monitoring of requests vs offers.

82 are deployed for continuous QoS evaluation. The participating entities are:  
83 *i) cloud provider*, providing functionalities for service delivery and composition;  
84 *ii) composite service owner*, managing a service composition; *iii) certification*  
85 *authority*, providing functionalities for continuous non-functional property certi-  
86 fication. Current approaches to service composition in the cloud are affected by  
87 a few limitations, which show a clear disalignment with the maturity reached  
88 by the cloud. These limitations, which are described in the following, must  
89 be addressed to provide a cost-effective service composition for the cloud with  
90 continuous QoS assessment.

- 91 • *Functional composition.* Service composition in the cloud puts great em-  
92 phasis on functionalities. Component services are selected on the basis of  
93 the implemented functionalities, while overall non-functional aspects are,  
94 in most of the cases, pushed aside. For example, a composite e-Health  
95 service composes services for planning for a visit, access clinical reports,  
96 and get medicine prescription, a payment service, and a database/storage  
97 service. This practice however increases the likelihood of composite ser-  
98 vices that, on one side, satisfy the expectations of the users, while on the  
99 other side increase risks of failures and misbehaviors (e.g., privacy risks in  
100 the e-Health service). A proper approach to service composition must not  
101 only focus on functional requirements, but also consider non-functional  
102 requirements from the outset. For example, non-functional requirements  
103 may refer to security, privacy, reliability requirements, and can be ad-  
104 dressed by proving specific properties such as confidentiality, integrity,  
105 availability, or showing compliance to specific standards/regulations, such  
106 as Payment Card Industry Data Security Standard (PCI-DSS), EU Gen-  
107 eral Data Protection Regulation (GDPR).
- 108 • *Ad hoc composite services.* Service composition in the cloud often consists  
109 of ad hoc workflows, where component services are designed and devel-  
110 oped for a specific composite service. For example, similarly to what hap-  
111 pens with reserved instances and dedicated hosts in the cloud, component

112 services are developed for and assigned to a specific, static service com-  
113 position and never shared with other composite services. This approach  
114 substantially decreases the utility of service composition, from both a flex-  
115 ibility and a cost point of view. Having no possibility of sharing a single  
116 service among multiple service compositions bound current approaches  
117 to hold fashion monolithic service deployments. This approach is often  
118 adopted at infrastructure layer, where the huge amount of available re-  
119 sources often point to single tenant scenarios, where a user is usually  
120 provided with isolated resources not shared with other tenants. If, on  
121 one side, ad hoc composite services lower complexity of QoS evaluation  
122 and management, on the other side, it substantially increases costs and  
123 reduces the benefits of service compositions.

124 • *QoS evaluation.* It mostly focuses on deployment-time evaluation and on  
125 composition adaptation in case of component service malfunctioning/fail-  
126 ure. QoS evaluation is however a more powerful concept that should rep-  
127 resents a first-class requirement driving composition operations. First, it  
128 should be based on assurance (e.g., certification) techniques guaranteeing  
129 stable and verifiable QoS; then, it should consider how the QoS of a single  
130 service contributes to the QoS of the whole composition; finally, it should  
131 implement a continuous process that evaluates non-functional properties  
132 over time and drives adaptation of service compositions to provide stable  
133 QoS.

134 • *Direct costs.* The evaluation of service composition costs, which mainly  
135 focuses on direct costs due to component service integration, does not fit a  
136 multi-tenant cloud environment where *i)* services can be shared, relocated  
137 and migrated among different compositions and *ii)* non-functional proper-  
138 ties are modeled as QoS requirements and integrated with the composite  
139 service life cycle. A proper cost evaluation at the basis of a cost-effective  
140 service composition must also consider the costs introduced by the infras-  
141 tructure responsible for continuous QoS evaluation, and the costs intro-

142           duced when QoS requested by the users are lower than the ones provided  
143           by the cloud infrastructure.

144       In the following of this paper, we provide a cost-effective, certification-based  
145 service composition approach for the cloud that fills in the above limitations.  
146 It is based on *i)* the concept of *portable certification*, supporting continuous  
147 QoS evaluation also in case of service migration and relocation and *ii)* a new  
148 cost evaluation methodology, considering direct, indirect, and mismatch costs  
149 on the cloud providers. We recall that our approach considers non-functional  
150 properties of composite services as must-have requirements; in other words, the  
151 QoS requirements of composite services are satisfied by design following our  
152 certification-based service composition. The design of an approach where QoS  
153 requirements in the form of non-functional properties in certificates are relaxed  
154 is out of the scope of this paper and will be the target of our future work.

### 155 **3. Basic Concepts**

A certification scheme for the cloud implements a continuous process whose goal is to verify whether a cloud service holds a given (set of) property [18]. The cloud service under evaluation is referred to as *Target of Certification (ToC)*. Properties  $p=(\hat{p},l)$ , as defined by the Cloud Security Alliance (CSA) [19], are composed of a controlled name  $\hat{p}$  (e.g., confidentiality of data in transit) and a level  $l$  modeling the strength of the supported property. Properties can be organized in a hierarchy based on their strength such that  $p_i \leq p_j$  (meaning  $p_i$  is weaker than  $p_j$ ) iff  $p_i.\hat{p}=p_j.\hat{p}$  and  $p_i.l < p_j.l$ . Based on levels  $l$ , a distance  $Dist(p_i, p_j)$  between two properties with the same  $\hat{p}$  is defined as:

$$Dist(p_i, p_j) = |p_i.l - p_j.l| \quad (1)$$

156 In this paper, without loss of generality, we consider security properties includ-  
157 ing, among the others, confidentiality, authentication, and data replication. For  
158 instance, property *confidentiality* can be further specified in properties *confi-*  
159 *dentiality at rest* and *confidentiality in transit*, each with three levels {AES128,  
160 AES192, AES256} and {TLS1.0, TLS1.1, TLS1.2}, respectively.

161 Our certification process is driven by a Certification Authority that manages  
 162 all certification activities leading to certification. It is composed of two sub-  
 163 processes: *i) evidence collection sub-process* and *ii) life cycle sub-process*. The  
 164 *evidence collection sub-process* collects the evidence at the basis of a trustworthy  
 165 certification and is carried out by the certification infrastructure. The *life cycle*  
 166 *sub-process* implements a continuous certification process that accomplishes the  
 167 evolution of the *ToC*, managing *ToC* migrations and versioning.

168 The certification process is based on two models, namely Certification Model  
 169 (CM) Template and Instance, driving certification activities [20]. The certifica-  
 170 tion authority defines CM Template  $\mathcal{T}$  specifying evidence collection activities  
 171 for a class of *ToC* and a (set of) property; the certification infrastructure im-  
 172 plements and executes the corresponding CM Instance  $\mathcal{I}$  specifying evidence  
 173 collection activities for a given *ToC* instance and a (set of) property. Collected  
 174 evidence is based on testing or monitoring, and permits to evaluate whether  
 175 the observed *ToC* behavior conforms to the expected one. Upon a positive  
 176 evaluation is retrieved following activities in  $\mathcal{I}$ , a certificate  $cert_{\mathcal{I}}$  is released.  
 177 Certificate  $cert_{\mathcal{I}}$  is signed by the certification authority and contains: *i)* a de-  
 178 scription of the property certified for a given service, *ii)* a link to the *ToC*, and  
 179 *iii)* a reference to the collected evidence and the relevant  $\mathcal{I}$ .

180 **Certification Model Template ( $\mathcal{T}$ ).** It is a declarative model that describes  
 181 the activities to be done to verify a set of properties according to the expected  
 182 behavior of a class of *ToC*. Formally, a CM Template  $\mathcal{T}_i$  is a triple  $(f_i, R_i, d-$   
 183  $eval_i)$ , where *i)*  $f_i$  is a functionality in the set  $F$  of functionalities offered by a  
 184 cloud provider, *ii)*  $r_k$  is a user requirement in the set  $R_i$  of requirements used  
 185 to annotate  $f_i$ , with  $r_k \in R_i$  a property  $(\hat{p}, l)$ , and *iii)*  $d-eval_i$  is a declarative  
 186 description of the evaluation activities to be carried out on the *ToC* to verify  
 187 requirements  $R_i$ .  $\mathcal{T}$  is built around  $d-eval$ , which is defined as a set of annotated  
 188 workflows.

189 **Definition 3.1 ( $d-eval$ ).**  $d-eval$  is a pair  $\langle \phi, \omega \rangle$ , where:

- 190 •  $\phi$  is a set of sequential workflows  $\{n_1, \dots, n_n\}$  for evidence collection, where



191 each node  $n_i$  defines an abstract action (e.g., test authentication interface)  
 192 and each edge  $(n_i, n_j)$  the flow between two actions.

193 •  $\omega$  is an annotation function on nodes  $n$ .  $\omega(\{n_i\})$  defines constraints (e.g.,  
 194 two factor authentication required) for a subset  $\{n_i\}$  of abstract actions.

195 We recall that *d-eval* refers to a generic class of *ToC* (e.g., an authentication system),  
 196 while it precisely pinpoints security and deployment requirements (e.g., a  
 197 given password strength policy). This means that, although there are a number  
 198 of different *ToC* for the selected class, their evaluation w.r.t. security/deploy-  
 199 ment requirements should follow the same declarative description.

200 **Certification Model Instance ( $\mathcal{I}$ ).** It is a procedural, executable model  
 201 generated by instantiating  $\mathcal{T}$  on a real *ToC*. It drives the certification process,  
 202 including the evidence collection process. Formally, a CM Instance  $\mathcal{I}_i$  is a triple  
 203  $(cs_i, \mathcal{P}_i, p\text{-eval}_i)$ , where *i*)  $cs_i$  is the *ToC*, *ii*)  $\mathcal{P}_i$  is the set of properties supported  
 204 by  $\mathcal{I}_i$ , and *iii*)  $p\text{-eval}_i$  defines certification activities as a concrete instantiation of  
 205 *d-eval* for a specific *ToC*.  $\mathcal{I}$  is built around  $p\text{-eval}$ , which covers the peculiarities  
 206 of the specific *ToC* w.r.t. the given properties.  $p\text{-eval}$  is an annotated workflow  
 207 defined as follows.

208 **Definition 3.2 (*p-eval*).** *p-eval* is a triple  $\langle \phi', \lambda \rangle$ , where:

- 209 •  $\phi'$  is a set of sequential workflows  $\{n_1, \dots, n_n\}$  for evidence collection,  
 210 where each node  $n_i$  defines an action implemented on the *ToC* instance  
 211 and each edge  $(n_i, n_j)$  the flow between two implemented actions.
- 212 •  $\lambda$  is an annotation function.  $\lambda(\{n_i\})$  defines the configuration settings of  
 213 each action, describes how to deploy *p-eval*, and describes possible depen-  
 214 dencies on its execution.

215 We note that CM Instance  $\mathcal{I}$  can be not unique for CM Template  $\mathcal{T}$ .

216 **Example 3.1.** Let us consider a Certification Model Template  $\mathcal{T}=(\text{Storage},$   
 217 *Confidentiality via encryption at rest, d-eval*), with  $d\text{-eval}=\langle \phi, \omega \rangle$ . For simplic-  
 218 ity, we assume  $\phi$  composed of a single sequential workflow  $\{n_1, n_2, n_3\}$ , where

219  $n_1 = \text{“ToC login”}$ ,  $n_2 = \text{“Test encryption”}$ ,  $n_3 = \text{“ToC logout”}$ , and annotations  
220  $\omega(\{n_1\}) = [\text{Administration credentials required}]$ ,  $\omega(\{n_2\}) = [\text{Resource URI}]$ .

221 The same Certification Model Template  $\mathcal{T}$  is instantiated in two different  
222 Certification Model Instances  $\mathcal{I}$  for a Linux file system and Amazon Simple  
223 Storage Service (S3). Both instances drive a certification process and evidence  
224 collection activity targeting the same property “Confidentiality at rest via en-  
225 cryption”.

226 Let us first consider a Linux file system using LUKS.  $\text{p-eval}_l = \langle \phi'_l, \lambda_l \rangle$  im-  
227 plementing the above d-eval is defined as follows:  $\phi'_l = \{\text{SSH login, Script testing}$   
228  $\text{encrypted volumes, SSH logout}\}$ ,  $\lambda_l(\{n_1\}) = [\text{root, cert}]$ ,  $\lambda(\{n_2\}) = \text{Volume path}$ .

229 Let us then consider Amazon S3.  $\text{p-eval}_{s3} = \langle \phi'_{s3}, \lambda_{s3} \rangle$  implementing the above  
230 d-eval is defined as follows:  $\phi'_{s3} = \{\text{Amazon login, API call for S3 configuration,}$   
231  $\text{Amazon logout}\}$ ,  $\lambda_l(\{n_1\}) = [\text{credentials, APIkey}]$ ,  $\lambda(\{n_2\}) = [\text{Config item}]$ .

#### 232 4. Portable Certification of Composite Services

233 We present a certification approach specifically tailored for cloud composite  
234 services, which is grounded on and extends the one in Section 3 to *i*) support  
235 service versioning, migration, and deployment changes (portability) and *ii*) ac-  
236 complish the dynamics of service orchestrations where component services can  
237 be replaced and migrated at run time according to contextual events. In the  
238 following, we first describe the portability of our certification process and then  
239 describe how we use it in the framework of composite service certification.

##### 240 4.1. Portability

241 A portable certification process is a certification process that is not bound  
242 to a specific *ToC* and can be easily applied to different service instances. It  
243 permits to apply the same certification process to different *ToC* with sufficient  
244 commonalities. Using our formalism, a certification process that derives from  
245 requirements in a template  $\mathcal{T}$  can be re-used (with or without minor modifi-  
246 cations) to certify all the services having an instance  $\mathcal{I}$  consistent with  $\mathcal{T}$ . To

247 verify this consistency we define a consistency check function, inspired by the  
 248 work in [20], as follows.

249 **Definition 4.1** ( $\xrightarrow{I}$ ). *CM Instance*  $\mathcal{I}_i=(cs_i, \mathcal{P}_i, \text{p-eval}_i)$  is consistent with CM  
 250 *Template*  $\mathcal{T}_i=(f_i, R_i, \text{d-eval}_i)$ , denoted as  $\mathcal{T}_i \xrightarrow{I} \mathcal{I}_i$ , iff i)  $cs_i$  implements  $f_i$ , ii)  
 251  $\mathcal{P}_i$  is such that  $R_i \leq \mathcal{P}_i$ , that is,  $\forall r_j \in R_i, p_j \in \mathcal{P}_i, r_j \leq p_j$ , meaning that the prop-  
 252 erties are stronger than the requirements according to property levels, and iii)  
 253  $\text{d-eval}_i \xrightarrow{i} \text{p-eval}_i$  (see Definition 4.2), meaning that  $\text{p-eval}_i$  is an instantiation of  
 254  $\text{d-eval}_i$ .

255 Consistency check  $\xrightarrow{I}$  is the cornerstone of process portability. A certification  
 256 process can be implemented and executed using different instances  $\mathcal{I}$ , thanks  
 257 to the decoupling between abstract definition ( $\mathcal{T}$ ) and concrete actuation ( $\mathcal{I}$ )  
 258 of the certification process. This decoupling also permits multiple consistent  
 259 instantiations ( $\mathcal{I}$ ) of the same process ( $\mathcal{T}$ ). We note that, having  $\mathcal{T}$  and  $\mathcal{I}$  the  
 260 same logical structure,  $\xrightarrow{I}$  can be used to verify the consistency between two  
 261 templates ( $\mathcal{T}_i \xrightarrow{I} \mathcal{T}_j$ ) or two instances ( $\mathcal{I}_i \xrightarrow{I} \mathcal{I}_j$ ).

262 As a complement to Definition 4.1, we detail how *p-eval* in  $\mathcal{I}$  is checked for  
 263 consistency against *d-eval* in  $\mathcal{T}$ .

264 **Definition 4.2** ( $\xrightarrow{i}$ ).  $\text{p-eval}_i = \langle \phi', \lambda \rangle$  is an instantiation of  $\text{d-eval}_i = \langle \phi, \omega \rangle$ ,  
 265 denoted as  $\text{d-eval}_i \xrightarrow{i} \text{p-eval}_i$ , iff i)  $\phi'$  implements  $\phi$ , ii) configurations  $\lambda(\{n_i\})$   
 266 in  $\text{p-eval}$  instantiate constraints  $\omega(\{n_i\})$  in  $\text{d-eval}$ , iii)  $\lambda$  permits the binding  
 267 between each action in  $\phi'$  and the corresponding end-point in the *ToC*.

268 Definition 4.1 ( $\xrightarrow{I}$ ) and Definition 4.2 ( $\xrightarrow{i}$ ) are at the basis of a portable  
 269 certification process that addresses two main scenarios: *service versioning* and  
 270 *service replacement*.

271 **Service versioning.** It considers a single service that either is migrated as it  
 272 to another location or evolves to a new version. It is defined as follows.

273 **Definition 4.3 (Process Portability (Versioning)).** *Let us consider a cer-*  
 274 *tification process driven by*  $\mathcal{I}_i=(cs_i, \mathcal{P}_i, \text{p-eval}_i)$  *for service*  $cs_i$ , *and a service*  $cs_k$

275 such that either i)  $cs_i=cs_k$  but they are deployed in different locations or ii)  $cs_k$   
 276 is the new version of  $cs_i$ . The certification process driven by  $\mathcal{I}_i$  can be ported to  
 277  $cs_k$  iff  $\lambda_i$  is modified to connect  $p\text{-eval}_i$  to  $cs_k$ .

278 Process portability (versioning) properly configures the certification model in-  
 279 stance in a way that permits the certification activities in  $p\text{-eval}_i$  to connect to  
 280 a different *ToC* (i.e., service  $cs_k$ ). To this aim,  $\lambda_i$  of  $p\text{-eval}_i$  must provide the  
 281 new configurations required to connect each action to  $cs_k$ .

282 **Service replacement.** It considers a migration of a service to another service  
 283 of the same class. For instance, a service implementing a MySQL database is  
 284 migrated to a service implementing an SQLServer database. Process portability  
 285 for service replacement is defined as described in the following definition.

286 **Definition 4.4 (Process Portability (Replacement)).** *Let us consider  $\mathcal{I}_i$*   
 287  *$= (cs_i, \mathcal{P}_i, p\text{-eval}_i)$  and  $\mathcal{I}_k = (cs_k, \mathcal{P}_k, p\text{-eval}_k)$  such that  $cs_i \neq cs_k$ . The certifica-*  
 288 *tion process driven by  $\mathcal{I}_i$  can be ported to  $\mathcal{I}_k$  according to the following condi-*  
 289 *tions:*

- 290 •  $\mathcal{I}_i \xrightarrow{I} \mathcal{I}_k$
- 291 •  $cs_i$  and  $cs_k$  provide the same functionality  $f_i$ .

292 Process portability (replacement) instantiates certification activities on dif-  
 293 ferent services  $cs_i$  and  $cs_k$ . To this aim, Condition 1 states that  $\mathcal{I}_i$  is consistent  
 294 with  $\mathcal{I}_k$ , and in turn their  $\mathcal{T}$  are consistent as well. We note that the con-  
 295 sistency at CM Instance level implies that  $p\text{-eval}_k$  is equivalent to  $p\text{-eval}_i$  (see  
 296 Definition 4.2). In other words the workflows for evidence collection in  $p\text{-eval}_i$   
 297 must be available also in  $p\text{-eval}_k$  possibly with different annotation functions  
 298 [17]. Details about this implications, and how to relax it are out of the scope of  
 299 this paper, and will be detailed in future work.

300 Condition 2 states that  $cs_i$  and  $cs_k$  provide the same functionality  $f_i$ , which  
 301 is specified in the corresponding templates  $\mathcal{T}_i$  and  $\mathcal{T}_k$ . In other words, a cer-  
 302 tification process can be ported to a service or in an environment where the

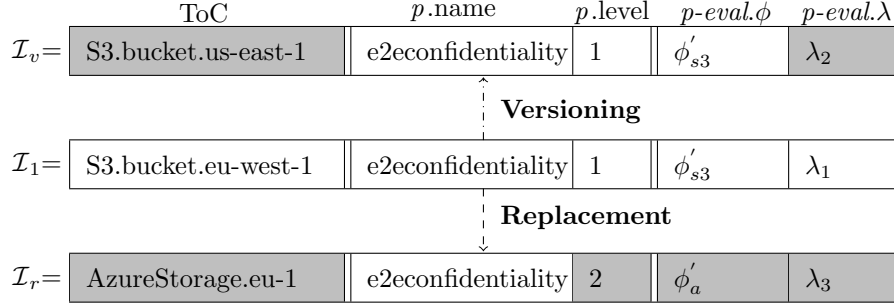


Figure 1: An example of Versioning and Replacement of a storage service

303 certification is driven by a different  $\mathcal{T}$  without the need to re-build the certifi-  
 304 cation process from scratch.

305 **Example 4.1.** Let us consider a CM template for a storage service defined as  
 306 follows  $\mathcal{T}_1 = \{Storage, \{(e2econfidentiality, 1)\}, d-eval_{storage}\}$ , where end-to-end  
 307 confidentiality (*e2econfidentiality*) is requested (i.e., both confidentiality in tran-  
 308 sit and confidentiality at rest). Let us consider a storage service based on Ama-  
 309 zon Simple Storage Service (S3) and specifically a bucket hosted on S3 eu-west-1  
 310 AWS region. Let us consider that this service has been certified according to  
 311 the CM Instance  $\mathcal{I}_1 = (S3.bucket.eu-west, \{(e2econfidentiality, 1)\}, p-eval_{s3})$  with  
 312  $\mathcal{T}_1 \xrightarrow{I} \mathcal{I}_1$ .

313 Figure 1 also shows the case where this service is moved to a different region  
 314 (versioning). In this scenario CM Instance  $\mathcal{I}_1$  is ported to  $\mathcal{I}_v = (S3.bucket.us-$   
 315  $east-1, \{(e2econfidentiality, 1)\}, p-eval_{s3})$ , where  $p-eval_{s3}$  is re-configured to ac-  
 316 cess the new bucket in the different region. We note that only parameters  $\lambda_2$   
 317 are modified since the service is exactly the same but in different location.

318 Figure 1 shows the case where this service is migrated to a different service  
 319 (replacement). More specifically the service is replaced with an Azure Storage  
 320 service which offers the same functionality and is certified for a given  $\mathcal{T}_2$  for  
 321 the same property but with higher level ( $\{(e2econfidentiality, 2)\}$ ) with  $(\mathcal{T}_1 \xrightarrow{I} \mathcal{T}_2)$ .  
 322 The corresponding CM Instance  $\mathcal{I}_r = (AzureStorage.eu-1, \{(e2econfidentiality,$   
 323  $2)\}, p-eval_a)$  is compatible with  $\mathcal{I}_1$  and can replace it. We note that, in a real

324 *environment, storage service replacement also implies functional compatibility*  
 325 *at orchestration level and application data migrations.*

#### 326 4.2. Composition

327 A certification process for composite services builds on our portable certifica-  
 328 tion process and is driven by a compositional CM Template  $\mathcal{T}^c$ , where functional  
 329 and certification requirements are specified for each component service.  $\mathcal{T}^c$  is  
 330 expressed as a set  $\{\mathcal{T}_1, \dots, \mathcal{T}_n\}$  of ordered templates, each to be linked to a  
 331 component service. A certified service  $cs_i$ , having  $cert_{\mathcal{I}_i}$ , can be selected as a  
 332 component service *iff* its  $\mathcal{I}_i$  is consistent with the corresponding  $\mathcal{T}_i$  in  $\mathcal{T}^c$ . We  
 333 note that templates, including compositional templates, are specified by a CA  
 334 and are the cornerstone of the certification chain of trust [21]. We extend Defi-  
 335 nition 4.1 to compositional instances ( $\mathcal{I}^c$ ) and compositional templates ( $\mathcal{T}^c$ ) as  
 336 follows.

337 **Definition 4.5** ( $\xrightarrow{\mathcal{I}^c}$ ). *A Compositional Instance  $\mathcal{I}_i^c$  is consistent with a Com-*  
 338 *positional Template  $\mathcal{T}_i^c$ , denoted as  $\mathcal{T}_i^c \xrightarrow{\mathcal{I}^c} \mathcal{I}_i^c$ , iff  $\forall \mathcal{T}_k \in \mathcal{T}_i^c, \exists \mathcal{I}_j \in \mathcal{I}_i^c$  such that*  
 339  *$\mathcal{T}_k \xrightarrow{\mathcal{I}} \mathcal{I}_j$ .*

340 The consistency check in Definition 4.5 supporting multiple consistent in-  
 341 stantiations ( $\mathcal{I}^c$ ) of the same certification process ( $\mathcal{T}^c$ ) is at the basis of com-  
 342 position portability. It provides higher flexibility and lower costs, supporting  
 343 automatic component substitution and reuse. Shared/reused components do  
 344 not need to be evaluated multiple times, saving certification effort, and their  
 345 management does not involve the certification authority.

346 **Example 4.2.** *Let us consider an e-Health service that allows patients to plan*  
 347 *and pay for a visit, access clinical reports, and get medicine prescription. The*  
 348 *e-Health service is a composite service that integrates i) a Web App providing ac-*  
 349 *cess to e-Health functionalities, ii) a Database that gives access to patients' doc-*  
 350 *uments, iii) a Payment service allowing patients to pay for visits, iv) a Storage*  
 351 *that stores all patients' documents. An e-Health composite service comes with*

352 *strong security requirements: it must guarantee confidentiality of data and com-*  
353 *munications, and robustness against known vulnerabilities. A compositional CM*  
354 *Template  $\mathcal{T}^c = \{\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_3, \mathcal{T}_4\}$  can then be defined as follow: i)  $\mathcal{T}_1 = \{WebApp,$   
355  $\{(confidentiality-in-transit, 3), (vulnerability-free, 10)\}, d - eval_{webapp}\}$  meaning  
356 *that the Web App must provide the service over an encrypted channel and be*  
357 *vulnerability free, ii)  $\mathcal{T}_2 = \{DB, \{(e2econfidentiality, 3)\}, d - eval_{db}\}$  and*  
358  $\mathcal{T}_4 = \{Storage, \{(e2econfidentiality, 3)\}, d - eval_{storage}\}$  meaning that the database  
359 *and storage must be encrypted and exchange data over an encrypted channel*  
360 *(e2econfidentiality) with highest level (i.e.,  $l=3$ ), iii)  $\mathcal{T}_3 = \{Payment, \{(con-$   
361  $identiality-in-transit, 2)\}, d - eval_{payment}\}$  meaning that the payment service  
362 *must provide the service over an encrypted channel (confidentiality-in-transit)*  
363 *with medium level (i.e.,  $l=2$ ).***

364 *A consistent Compositional Instance  $\mathcal{I}_i^c$  of the above CM Template  $\mathcal{T}^c$  can*  
365 *be defined as made by the following set of  $\mathcal{I}$ ,  $\{\{psql.h-6, confidentiality.0, p -$   
366  $eval_{psql}\}, \{nginx.h-2, vulnerability-free.10, p - eval_{nginx}\}, \{S3.h-2, e2e-confi-$   
367  $dentiality.4, p - eval_{s3}\}, \{pay.remote, confidentiality, p - eval_{pay}\}\}$ .*

368 *We note that, thanks to our portability (see Example 4.1), Amazon S3 (i.e.,*  
369  $S3.h-2)$  can be replaced with Azure Storage having instance  $\mathcal{I}_m = (AzureStorage.ue-$   
370  $1, e2econfidentiality.1, p - eval_{azurestorage})$  leading to another consistent compo-  
371 *sitional instance  $\mathcal{I}_j^c$ .*

## 372 5. Deployment of Certified Composite Services

373 *The enrichment of traditional composition solutions with certification tech-*  
374 *niques evaluating non-functional properties of composite services introduces the*  
375 *need of rethinking the algorithms driving selection of component services. If,*  
376 *on one side, service selection has been already renewed to accomplish selection*  
377 *of services that prove a set of non-functional properties [4], on the other side,*  
378 *solutions to cost-based service selection need to depart from the assumption*  
379 *that costs are only due to service deployment and resource consumption [12].*  
380 *The latter must consider costs introduced by certification processes, and by the*

381 need of keeping the composition continuously monitored and certified.

### 382 5.1. Deployment Composition Matrix

383 The status of a given CP at time  $t$  can be represented as a matrix  $D$  of size  
 384  $C \times F$  of deployed compositions  $\mathcal{I}_i^c$ , where  $C$  is the cardinality of deployed com-  
 385 positions at time  $t$  and  $F$  the cardinality of all possible functionalities provided  
 386 by service providers. Matrix  $D$  has the following structure

$$D = \begin{matrix} & f_1 & f_2 & f_3 & f_4 & \cdots & f_F \\ \mathcal{I}_1^c & \left( \begin{array}{cccccc} \mathcal{I}_{1,1} & \mathcal{I}_{1,2} & \mathcal{I}_{1,3} & \mathcal{I}_{1,4} & \cdots & \mathcal{I}_{1,F} \\ \mathcal{I}_2^c & \mathcal{I}_{2,1} & \mathcal{I}_{2,2} & \mathcal{I}_{2,3} & \mathcal{I}_{2,4} & \cdots & \mathcal{I}_{2,F} \\ \mathcal{I}_3^c & \mathcal{I}_{3,1} & \mathcal{I}_{3,2} & \mathcal{I}_{3,3} & \mathcal{I}_{3,4} & \cdots & \mathcal{I}_{3,F} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathcal{I}_C^c & \mathcal{I}_{C,1} & \mathcal{I}_{C,2} & \mathcal{I}_{C,3} & \mathcal{I}_{C,4} & \cdots & \mathcal{I}_{C,F} \end{array} \right) & (2) \end{matrix}$$

387 where each row represents a composite service  $\mathcal{I}_i^c$ , each column a function-  
 388 ality  $f_j$ , and each cell a component service of  $\mathcal{I}_i^c$  referred in the matrix with  
 389 the corresponding CM Instance  $\mathcal{I}_{i,j} = (cs_{i,j}, \mathcal{P}_{i,j}, p\text{-eval}_{i,j})$ . Each service  $\mathcal{I}_{i,j}$  is  
 390 annotated with a sharing level  $k \geq 1$ , specifying the number of compositions  $\mathcal{I}_i^c$   
 391 insisting on it. In the following, we denote the component service  $\mathcal{I}_{i,j}$  selected  
 392 as part of the composition  $\mathcal{I}_i^c$  as  $\mathcal{I}_i^c.\mathcal{I}_j$ .

393 **Example 5.1.** *Figure 2(a) shows an example of deployment composition matrix*  
 394  *$D$  with 4 functionalities ( $f$ ) and 8 cloud services ( $\mathcal{I}$ ), as follows:*

- 395 • *Functionality database (DB): mysql ( $\mathcal{I}_1$ ) and postgresql ( $\mathcal{I}_6$ ) are both cer-*  
 396 *tified for property confidentiality at different levels  $l$ .*
- 397 • *Functionality web application (WebApp): nginx ( $\mathcal{I}_2$ ,  $\mathcal{I}_4$ , and  $\mathcal{I}_7$ ) are cer-*  
 398 *tified for property vulnerability-free at different levels  $l$ . A level can refer*  
 399 *to the severity of the Common Vulnerability Scoring System (CVSS) score*  
 400 *related to the Common Vulnerabilities and Exposures (CVE) discovered*  
 401 *on the target; the highest the level the lower the severity discovered.*



- 402 • *Functionality storage (Storage): Amazon S3 ( $\mathcal{I}_8$ ) is certified for property*  
403 *end-to-end confidentiality.*
- 404 • *Functionality payment (Payment): a remote payment service ( $\mathcal{I}_3$ ) is cer-*  
405 *tified for property PCI-DSS compliance level 1 and the ENGPAY remote*  
406 *payment service ( $\mathcal{I}_5$ ) for PCI-DSS compliance level 3. Details about PCI-*  
407 *DSS compliance certification is available in [17] and in Section 7.1.*

408 *These services are composed in 3 composite services  $\mathcal{I}_i^c$  (Figure 2(a)): i) an*  
409 *e-commerce service  $\mathcal{I}_1^c$  based on Database  $\mathcal{I}_1$ , WebApp  $\mathcal{I}_2$ , payment service  $\mathcal{I}_3$ ;*  
410 *ii) a shared-economy app  $\mathcal{I}_2^c$  based on Database  $\mathcal{I}_1$ , WebApp  $\mathcal{I}_4$ , and payment*  
411 *service  $\mathcal{I}_5$ ; iii) the e-health service  $\mathcal{I}_3^c$  of Example 4.2 based on Database  $\mathcal{I}_6$ ,*  
412 *WebApp  $\mathcal{I}_7$ , storage  $\mathcal{I}_8$ , and payment service  $\mathcal{I}_3$ . Figure 2(b) shows a graphical*  
413 *overview of the composite services highlighting shared component services, that*  
414 *is,  $\mathcal{I}_1^c$  and  $\mathcal{I}_3^c$ .*

## 415 5.2. Cost Factors

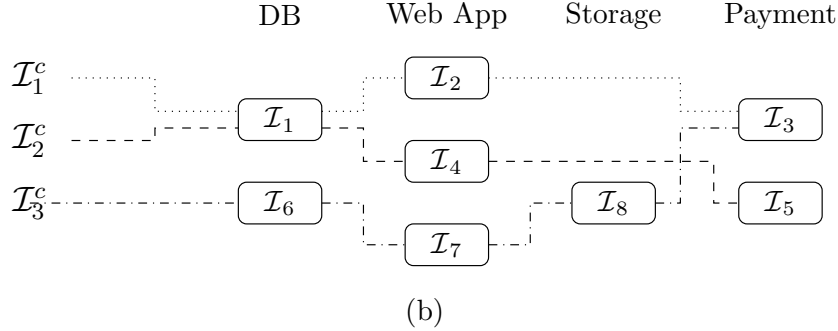
416 The management of a certified service introduces a cost on the cloud provider  
417 that can be evaluated using the Deployment Compositional Matrix and depends  
418 on three main factors: *direct (deployment) costs*, *indirect (certification) costs*,  
419 and *mismatch costs*. Each of these cost factors can be characterized using one  
420 of the four different cost behaviors identified by Horngren [22], and later used  
421 for cloud services by de Medeiros et al. [12]: *i) fixed costs*, a resource cost  
422 function that is completely independent from volume and time, indeed constant;  
423 *ii) variable costs*, a resource cost function that varies depending on volume or  
424 time, and is equal to zero for volume equal to zero; *iii) mixed costs*, a resource  
425 cost function that is the sum of a variable and a fixed cost function; and *iv) step*  
426 *costs*, a resource cost function that varies following different patterns.

427 **Direct (deployment) costs** ( $\alpha(\mathcal{I}.cs, \mathcal{I}.P, \mathcal{I}.k)$ ). They are defined by the  
428 cloud provider as the amount of resources to be allocated to a cloud service  $cs$   
429 w.r.t. the certified properties  $\mathcal{I}.P$ . They are usually estimated as a mixture of  
430 fixed deployment and variable usage costs [12]. Direct costs comprise servers,

$$\begin{matrix} \mathcal{I}_1^c \\ \mathcal{I}_2^c \\ \mathcal{I}_3^c \end{matrix} \begin{pmatrix} DB & WebApp & Storage & Payment \\ \mathcal{I}_1 & \mathcal{I}_2 & - & \mathcal{I}_3 \\ \mathcal{I}_1 & \mathcal{I}_4 & - & \mathcal{I}_5 \\ \mathcal{I}_6 & \mathcal{I}_7 & \mathcal{I}_8 & \mathcal{I}_3 \end{pmatrix}$$

$$\begin{aligned} \mathcal{I}_1 &= \{\text{mysql.h-1, confidentiality, 1, ...}\} \\ \mathcal{I}_2 &= \{\text{nginx.h-2, vulnerability-free, 4, ...}\} \\ \mathcal{I}_3 &= \{\text{pay.remote, PCI-DSS compliance, 1, ...}\} \\ \mathcal{I}_4 &= \{\text{nginx.h-4, vulnerability-free, 7, ...}\} \\ \mathcal{I}_5 &= \{\text{ENGPAY.remote, PCI-DSS compliance, 3, ...}\} \\ \mathcal{I}_6 &= \{\text{psql.h-6, confidentiality, 0, ...}\} \\ \mathcal{I}_7 &= \{\text{nginx.h-2, vulnerability-free, 10, ...}\} \\ \mathcal{I}_8 &= \{\text{S3.h-2, e2e-confidentiality, 4, ...}\} \end{aligned}$$

(a)



(b)

Figure 2: An example of CP Status Matrix  $D$  with 8 services and 3 compositions

431 infrastructure, power, networks, and personnel costs [23]. They also consider  
 432 the cost of orchestrating the composition, and managing service versioning and  
 433 migration. An appropriate cost behavior can be a step function that depends  
 434 on properties  $\mathcal{I.P}$  and the sharing level  $k$ , that is, the number of compositions  
 435 insisting on a given service. Figure 3(a) shows an example of a function of  
 436 direct costs; we observe a small cost increase between  $k=2$  and  $k=4$  due to  
 437 power consumption, and a more substantial increase from  $k=5$  when a vertical  
 438 scaling of resources is required to satisfy all requests.

439 **Indirect (certification) costs** ( $\beta(\mathcal{I.p-}eval, \mathcal{I.k})$ ). They are defined by the  
 440 cloud provider with the support of the certification authority as the amount

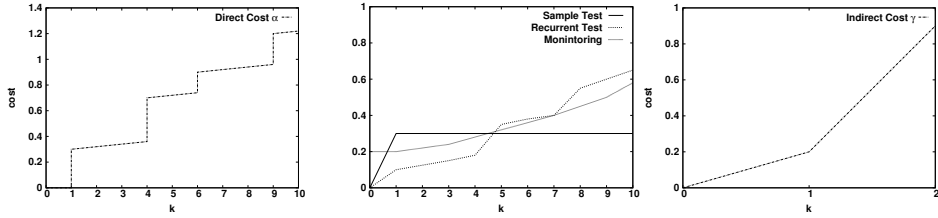


Figure 3: Cost functions

441 of resources to be allocated to the certification infrastructure (Section 7) for  
 442 continuous evaluation of service compositions. We note that, to execute evi-  
 443 dence collection in *p-eval*, our certification infrastructure considers three types  
 444 of collection activities (Figure 3(b)) having different costs.

- 445 • **Sample test:** one time evaluation of a specific part of the *ToC*. No need  
 446 to keep evaluation resources allocated after the evaluation.
- 447 • **Recurrent test:** a scheduled, repeatable evaluation of a specific target;  
 448 it is often part of a complex evaluation. The evaluation resources are  
 449 permanently allocated to re-execute the evaluation multiple times.
- 450 • **Monitoring:** continuous and permanent evaluation of the target.

451 **Mismatch costs** ( $\gamma(\mathcal{T}.R, \mathcal{I}.P)$ ). They are defined by the certification author-  
 452 ity, as inefficiency due to the distance between provided properties  $\mathcal{I}.P$  and re-  
 453 quired properties  $\mathcal{T}.R$ , with  $\mathcal{T} \xrightarrow{I} \mathcal{I}$ . Providing higher security properties means  
 454 in general allocating more resources than needed (e.g., more computational ef-  
 455 fort for encryption with 128-bit key when 32-bit key was required). This loss  
 456 of resources depends on the distance  $Dist(\mathcal{I}.p, \mathcal{T}.r)$  for each  $p \in \mathcal{P}$  and corre-  
 457 sponding  $r \in R$ , according to property levels in Section 3. Figure 3(c) shows an  
 458 example of mismatch cost function for a property/requirement distributed over  
 459 three levels. We note that the function is not necessarily homogeneous over  
 460 the number of property levels. An important boost for higher levels might be  
 461 observed, because high security levels may require more hardware facilities.

### 462 5.3. Cost Profiles

463 The CP behavior balances the contribution of direct, indirect, and mismatch  
464 costs to the computation of the total cost, while the total cost is the combination  
465 of each cost factor. We express the CP behavior as three cost profiles mapping  
466 to different CP strategies inspired by the deployment patterns in [24, 25].

- 467 • **Sharing profile** is typical of cloud providers that prefer to share re-  
468 sources, increasing the distance between requested and provided security  
469 properties (higher mismatch costs).
- 470 • **Fitting profile** is typical of cloud providers that prefer to achieve higher  
471 precision between requested and provided security properties, at the price  
472 of increasing the need of horizontal scaling. As a result, more compo-  
473 nent services are deployed precisely addressing users' needs, decreasing  
474 the average sharing level (higher direct and indirect costs).
- 475 • **Average profile** where direct, indirect, and mismatch costs equally con-  
476 tribute to the total cost.

477 We note that a degeneration of the fitting profile where an ad hoc composi-  
478 tion is deployed for each request is a good approximation of the actual strategy  
479 adopted by cloud providers.

## 480 6. Fuzzy-Based Cost-Effective Deployment of Service Compositions

481 We propose a fuzzy-based cost evaluation approach, which evaluates the cost  
482 of composite services in Matrix  $D$  on the basis of cost factors and profiles in  
483 Section 5. Our solution extends the one in [15] to provide a more accurate in-  
484 frastructure and easy to tune approach for cost evaluation and profile setting.  
485 In fact, the solution in [15] assumed uniform cost factors to balance their con-  
486 tribution to the final cost, and required a difficult and inaccurate tuning of the  
487 cost profiles on the needs of the cloud providers.

Fuzzy logic has been applied successfully in many fields including software  
cost estimation [26]. Let  $Y$  be the universe of discourse containing cost values

$y$  (i.e., cost factors in this paper). As customary, the membership degree of element  $y$  to a fuzzy set  $S$  is characterized by a membership function  $\mu_S(y)$ . A fuzzy set  $S$  in  $Y$  is denoted as follow.

$$S = \{(y, \mu_S(y)) | y \in Y\} \quad (3)$$

488 where  $\mu_S(y)$  is the membership function of the fuzzy set  $S$ . Let us then consider  
 489 each cost function  $\alpha, \beta, \gamma$  as a separate universe of discourse  $Y_\alpha, Y_\beta, Y_\gamma$ . We  
 490 define a standardized partition of each of them into different fuzzy sets. In the  
 491 following, we discuss  $Y_\alpha$ ; the same discussion holds for  $Y_\beta$  and  $Y_\gamma$ .

492 We define a standard fuzzy partition  $\{S_{\alpha_l}, S_{\alpha_m}, S_{\alpha_h}\}$  of  $Y_\alpha$  such that  
 493  $\forall y_\alpha \in Y_\alpha, \sum_{j \in \{l, m, h\}} \mu_{S_{\alpha_j}}(y_\alpha) = 1$ . Each fuzzy set corresponds to a linguistic  
 494 concept, that is, *LOW* for  $S_{\alpha_l}$ , *MEDIUM* for  $S_{\alpha_m}$ , and *HIGH* for  $S_{\alpha_h}$ . There  
 495 are a number of different shapes for the membership functions related to each  
 496 linguistic concept. In this paper, we use R-function and L-function for *LOW* and  
 497 *HIGH* membership functions and trapezoidal function for *MEDIUM* member-  
 498 ship function. Given these linguistic variables mapping the concepts of *LOW*,  
 499 *MEDIUM*, and *HIGH* costs for each cost factor, we define different sets of fuzzy  
 500 rules. The rules, one for each profile, are used by the fuzzy inference system to  
 501 infer the cost ( $Fc$ ) introduced by each single component service  $\mathcal{I}_i^c \cdot \mathcal{I}_j$  at time  
 502  $t$ .

503 **Example 6.1.** *Let us consider a component service  $\mathcal{I}_i^c \cdot \mathcal{I}_j$  ( $\mathcal{I}$  for brevity) to be*  
 504 *deployed at time  $t$ , following a sharing profile. Examples of fuzzy rules can be*  
 505 *defined as follows:*

506 **If**  $\alpha(\mathcal{I}.cs, \mathcal{I}.p, \mathcal{I}.k)$  *is HIGH and*  $\beta(\mathcal{I}.p\text{-eval}, \mathcal{I}.k)$  *is not LOW and*  $\gamma(\mathcal{T}.R, \mathcal{I}.P)$   
 507 *is LOW* **then**  $Fc_t(\mathcal{I})$  *is HIGH.*

508 **If**  $\alpha(\mathcal{I}.cs, \mathcal{I}.p, \mathcal{I}.k)$  *not HIGH and*  $\beta(\mathcal{I}.p\text{-eval}, \mathcal{I}.k)$  *is not LOW and*  $\gamma(\mathcal{T}.R, \mathcal{I}.P)$   
 509 *is LOW* **then**  $Fc_t(\mathcal{I})$  *is MEDIUM.*

510 **If**  $\alpha(\mathcal{I}.cs, \mathcal{I}.p, \mathcal{I}.k)$  *is LOW and*  $\beta(\mathcal{I}.p\text{-eval}, \mathcal{I}.k)$  *is LOW and*  $\gamma(\mathcal{T}.R, \mathcal{I}.P)$  *is*  
 511 *LOW* **then**  $Fc_t(\mathcal{I})$  *is LOW.*

512

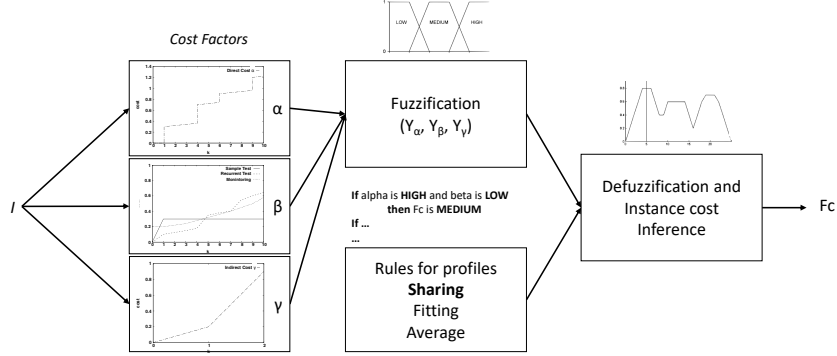


Figure 4: Fuzzy cost inference for a given  $\mathcal{I}$  and considering rules set for sharing profile (in bold).

513 Figure 4 shows the overview of our approach for fuzzy-based cost computa-  
 514 tion where, given a specific  $\mathcal{I}_i^c \cdot \mathcal{I}_j$  at time  $t$ , cost factors  $\alpha$ ,  $\beta$ , and  $\gamma$  are mapped  
 515 to fuzzy domains ( $\Theta$ ). This mapping is based on standard partitioning into dif-  
 516 ferent membership functions, which are specific for each cost function.<sup>5</sup> Then,  
 517 for each cost profile, a set of rules are defined and executed ( $\Xi$ ) to infer the  
 518 cost of each single  $\mathcal{I}_i^c \cdot \mathcal{I}_j$  ( $\mathcal{I}$  in Equation 4 for brevity), using a defuzzification  
 519 function ( $\Psi$ ), as follows.

$$F_{c_t}(\mathcal{I}) = \Psi(\Xi(\Theta(\alpha(\mathcal{I}.cs, \mathcal{I}.p, \mathcal{I}.k), \beta(\mathcal{I}.p\text{-eval}, \mathcal{I}.k), \gamma(\mathcal{T}.R, \mathcal{I}.P)))) \quad (4)$$

Considering the CP status Matrix in Equation 2, total fuzzy cost  $TF_{c_t}$  is the sum of the cost of each deployed composite service  $\mathcal{I}_i^c \in \{\mathcal{I}_1^c, \dots, \mathcal{I}_t^c\}$ , with  $t \leq C$ .  $TF_{c_t}$  can be formally expressed as

$$TF_{c_t} = \sum_{\substack{\mathcal{I}_i^c \\ 1 \leq i \leq C}} \sum_{\substack{\mathcal{I}_i^c \cdot \mathcal{I}_j \\ 1 \leq j \leq |\mathcal{I}_i^c|}} F_{c_t}(\mathcal{I}_i^c \cdot \mathcal{I}_j) \quad (5)$$

<sup>5</sup>Each CP tunes the membership functions of  $\alpha$  and contributes to the tuning of the ones of  $\beta$ , while the membership functions of  $\gamma$  are defined by the CA.

520 where  $\sum_{\substack{\mathcal{I}_i^c, \mathcal{I}_j \\ 1 \leq j \leq |\mathcal{I}_i^c|}} FC_t(\mathcal{I}_i^c, \mathcal{I}_j)$  is the cost of a composite service  $\mathcal{I}_i^c$  and calculated  
 521 as the sum of the costs of the corresponding component services  $\mathcal{I}_i^c, \mathcal{I}_j$ .

522 We note that total fuzzy cost is not the real cost incurred for a given de-  
 523 ployment. It represents the cost perceived by a cloud provider according to the  
 524 selected profile (i.e., sharing, fitting average) and the mixture of the different  
 525 cost sources (i.e., direct, indirect, mismatch).

### 526 6.1. Composition Cost Minimization

527 The aim of our solution is to find the best deployment  $\{\mathcal{I}_1^c, \dots, \mathcal{I}_n^c\}$  of com-  
 528 posite services that *i)* satisfies a set  $\{\mathcal{T}_1^c, \dots, \mathcal{T}_n^c\}$  of composition requests, that  
 529 is, guarantees  $\mathcal{T}_i^c \xrightarrow{\mathcal{I}_i^c} \mathcal{I}_i^c$ , with  $i=1, \dots, n$ , *ii)* minimize the cost  $TFc_n$  (Equation 5)  
 530 for the CP. We assume that a new composition request  $\mathcal{T}_i^c$  is received every time  
 531 instant  $t$ , introducing a uniform time of arrival for composition requests. Find-  
 532 ing the optimum deployment has however an exponential asymptotic behavior  
 533  $\mathcal{O}((|l| * F * t + t)^t)$ , in the worst case, with  $|l|$  the number of property levels,  $F$   
 534 the number of functionalities, and  $t$  the number of composition request. Being  $t$   
 535 the dominating factor that varies over time, the exponential asymptotic behav-  
 536 ior becomes  $\mathcal{O}(t^t)$ , which clearly does not fit the pseudo real-time requirement  
 537 in our paper. It is therefore necessary to design heuristic approaches for solving  
 538 the problem in polynomial time, even for relatively large composite services.

539 Many heuristic approaches balancing efficiency and quality in terms of pre-  
 540 cision and recall can be used for minimizing  $TFc_t$  at time  $t$ , though not all of  
 541 them are applicable in a cloud environment, where *i)* composition requests are  
 542 consecutive, *ii)* the requests may need to be served quickly. In our previous  
 543 paper [15], we considered a simpler approach that analyzes each request  $\mathcal{T}_t^c$   
 544 independently and aims to find  $\mathcal{I}_t^c$  providing the lowest cost increment with re-  
 545 spect to the total cost at time  $t-1$ . This approach however has some drawbacks,  
 546 which could affect the quality of the retrieved results. First, it provides a static  
 547 approach where each  $\mathcal{I}_i^c$  remains for its entire life-cycle deployed in the origi-  
 548 nal deployment slot; then, it supports real-time deployment of each incoming

549 request opening the door to wrong choices that could materialize only in the  
 550 subsequent requests.

551 We propose two heuristic algorithms: *i)* heuristic *sliding window* that selects  
 552 a cloud service deployment within a time-forwarding window  $w$  of composition  
 553 requests  $\mathcal{T}_i^c$  and *ii)* heuristic *sliding window with migration* that extends the  
 554 heuristic sliding window with the possibility of migrating component services  
 555  $\mathcal{I}_i^c, \mathcal{I}_j$ .

556 **Heuristic 1: Sliding Window.** It is based on the idea of finding the best  
 557 solution at time  $t$  using a time-forwarding window  $w$ . The heuristic selects the  
 558 best deployment  $\mathcal{I}_t^c$  at time  $t$  by evaluating a set of  $|w|$  consecutive requests  $\mathcal{T}_i^c$ ,  
 559 with  $t \leq i \leq t + |w|$ , received within a window  $w$  of size  $|w|$ . In other words, the  
 560 selected  $\mathcal{I}_t^c$  represents the composite service contributing to the global optimum  
 561 within window  $w$ .

562 At time  $t$ , the heuristic receives as input the CP status matrix  $D$ , which  
 563 contains all deployed  $\mathcal{I}_i^c$  with  $1 \leq i \leq t-1$ , all costs with related de-fuzzification  
 564 functions, window size  $|w|$ , and the total fuzzy cost  $TFc_{t-1}$ .

565 Upon collecting the last  $|w|$  requests  $\mathcal{T}_t^c, \dots, \mathcal{T}_{t+|w|}^c$ , the heuristic calculates  
 566 all possible candidate sets  $\mathcal{I}_t^c, \dots, \mathcal{I}_{t+|w|}^c$ . For each candidate  $\mathcal{I}_t^c, \dots, \mathcal{I}_{t+|w|}^c$ , our  
 567 heuristic calculates the total fuzzy cost  $TFc_{t+|w|}$  and chooses the one that entails  
 568 the minimum increase of cost. The minimum increase of cost is calculated as  
 569 the difference between the total fuzzy cost  $TFc_{t+|w|}$  within window  $w$  and the  
 570 current total fuzzy cost  $TFc_{t-1}$ . Both  $TFc_{t+|w|}$  and  $TFc_{t-1}$  are calculated using  
 571 Equation 4. Once the deployment  $\mathcal{I}_t^c, \dots, \mathcal{I}_{t+|w|}^c$  with minimum cost increase is  
 572 selected,  $\mathcal{I}_t^c$  is instantiated to satisfy request  $\mathcal{T}_t^c$ ; the window is then shifted of  
 573 one time interval and the process restarts to satisfy request  $\mathcal{T}_{t+1}^c$  when a new  
 574 request is received at time  $\mathcal{T}_{t+1+|w|}^c$ .

575 We note that  $|w|$  must be chosen carefully to balance the quality of the  
 576 retrieved solution and the performance/complexity of the overall heuristic. This  
 577 decision is left to the CP based on its requirements or preferences. We also note  
 578 that a degeneration of this approach with a sliding window of dimension  $|w|=1$



579 yields to the greedy approach presented in [15].

580 **Heuristic 2: Sliding Window with Migration.** It extends heuristic 1  
 581 with a better management of component deployment. Heuristic 2 supports  
 582 service versioning and replacement (see Definition 4.3 and Definition 4.4), and  
 583 in turn resource consolidation, for cost optimization. Migration in fact allows  
 584 CP to modify its status matrix, moving to a new deployment scenario with  
 585 lower costs. The global effect on the total cost, called *Migration Impact* ( $mi$ ), is  
 586 the difference between the total fuzzy cost  $TFc_t^{mi}$  after migration and the total  
 587 fuzzy cost  $TFc_t$  before migration:

$$mi = TFc_t^{mi} - TFc_t \quad (6)$$

588 Migration impact  $mi < 0$  introduces a cost saving; migration impact  $mi \geq 0$   
 589 introduces a cost increase.

590 A migration is triggered when a new composition request  $\mathcal{T}_t^c$  is processed  
 591 and results in the deployment of a new cloud service  $\mathcal{I}_i$  first instantiated in  $\mathcal{I}_t^c$   
 592 ( $\mathcal{I}_t^c.\mathcal{I}_i$ ). For clarity, we describe our heuristic using compositions with a single  
 593 component service, since every functionality  $f$  is independent and therefore can  
 594 be processed in parallel with the others. The migration process is composed of  
 595 two sequential phases and 4 steps as follows.

596 1. *Service migration:* this phase aims to optimize the cost of composite ser-  
 597 vices  $\mathcal{I}_i^c$  in  $D$  at time  $t-1$ . In particular, it migrates component service  
 598  $\mathcal{I}_i^c.\mathcal{I}_i$ , such that  $\mathcal{I}_i^c.\mathcal{I}_i \xrightarrow{\mathcal{I}} \mathcal{I}_t^c.\mathcal{I}_t$ , to the new cloud service  $\mathcal{I}_t^c.\mathcal{I}_t$  deployed at  
 599 time  $t$  (step 0) according to  $mi$ . Phase service migration starts with an or-  
 600 dering process, which introduces a migration priority among deployed ser-  
 601 vices. Services  $\mathcal{I}_i^c.\mathcal{I}_i$  are sorted in descending order according to function  
 602 property distance  $Dist(\mathcal{T}_i^c.\mathcal{T}_{i.r}, \mathcal{I}_i^c.\mathcal{I}_i.p)$  (see Equation 1), where  $\mathcal{T}_i^c.\mathcal{T}_{i.r}$   
 603 is the property originally requested for composition request  $\mathcal{T}_i^c$  and  $\mathcal{I}_i^c.\mathcal{I}_i.p$   
 604 is the property of the corresponding deployed composition  $\mathcal{I}_i^c$  (step 1).  
 605 Once all services are sorted, for each  $\mathcal{I}_i^c.\mathcal{I}_i$ , the migration impact  $mi$  is  
 606 calculated and, if  $mi < 0$ ,  $\mathcal{I}_i^c.\mathcal{I}_i$  is migrated to  $\mathcal{I}_t^c.\mathcal{I}_t$  (step 2).

607 2. *Resource consolidation*: this phase considers all component services  $\mathcal{I}_i^c \cdot \mathcal{I}_i$   
608 in the CP status matrix  $D$  migrated during the previous phase. Since  
609 each service instance  $\mathcal{I}_i$  is shared among different composite services, a  
610 migration changes the level of sharing of  $\mathcal{I}_i$  and introduces the need of  
611 a consolidation process to optimize the total fuzzy cost  $TFc$  (step 3).  
612 Resource consolidation is a *binary join* operation between two services  $\mathcal{I}_i$   
613 and  $\mathcal{I}_j$ , with  $\mathcal{I}_i.p < \mathcal{I}_j.p$ , which migrates all service composition  $\mathcal{I}_i^c$  that  
614 are deployed on  $\mathcal{I}_i$  to  $\mathcal{I}_j$ , that is,  $\mathcal{I}_i^c \cdot \mathcal{I}_i$  is migrated to  $\mathcal{I}_i^c \cdot \mathcal{I}_j$ . Among  
615 all possible pairs  $(\mathcal{I}_i, \mathcal{I}_j)$ , heuristic 2 chooses the one that offers the best  
616  $mi$ . Resource consolidation is recursively executed until no  $(\mathcal{I}_i, \mathcal{I}_j)$  offers  
617 a negative  $mi$ .

618 **Example 6.2.** *Let us consider a CSP offering compositions with a single func-*  
619 *tionality  $f$  and a property  $p$  with 3 levels. In the following, we describe the*  
620 *working of heuristic 2 as an extended version of heuristic 1.*

621 Step 0 – New request  $\mathcal{T}_t^c$  (Figure 5(a)). *A new request  $\mathcal{T}_t^c$  at time  $t$  triggers*  
622 *the execution of heuristic 2. The status of the CSP is depicted in the status*  
623 *Matrix  $D$  in Figures 5(a), where each row represents the deployed composition*  
624  *$\mathcal{I}_i^c$ , each column the deployed instance  $\mathcal{I}_j$  offering functionality  $f$  and property  $p$*   
625 *with level  $l$ , and each cell the request  $\mathcal{T}_i^c \cdot \mathcal{T}_j \cdot r$  to be satisfied by the correspond-*  
626 *ing property  $p$  of  $\mathcal{I}_j$ . For instance, in Figure 5(a), cloud service  $\mathcal{I}_1$  offers a*  
627 *property  $p$  at level 1 ( $p = (\hat{p}, 1)$ ) and is shared by composite services  $\mathcal{I}_1^c$  and*  
628  *$\mathcal{I}_3^c$ , whose templates  $\mathcal{T}_1^c$  and  $\mathcal{T}_3^c$  require property  $\mathcal{T}_1^c \cdot \mathcal{T}_1 \cdot r = (\hat{p}, 1)$  and property*  
629  *$\mathcal{T}_3^c \cdot \mathcal{T}_1 \cdot r = (\hat{p}, 1)$ , respectively. In the following, for simplicity, we consider the*  
630 *same  $\hat{p}$  for both  $r$  and  $p$ , and then refer to levels  $r.l$  and  $p.l$  only. At time  $t$ ,*  
631 *composition request  $\mathcal{T}_8^c$  with  $r.2$  is received. A composition instance  $\mathcal{I}_8^c$  of  $\mathcal{T}_8^c$*   
632 *is deployed on a new cloud service  $\mathcal{I}_4$  (denoted with a gray background in Fig-*  
633 *ure 5(a)) offering  $p.2$ . We note that the result of step 0 is both the final result*  
634 *of heuristic 1 and the initialization step of heuristic 2.*

635 Step 1 – Service ordering (Figure 5(b)). *All component services  $\mathcal{I}_i^c \cdot \mathcal{I}_j$ , with*  
636  *$1 \leq j \leq 3$  and with  $1 \leq i \leq 7$  in our example, are then sorted in descending order by*

$\mathcal{I}_1 (p.1)$	$\mathcal{I}_2 (p.3)$	$\mathcal{I}_3 (p.3)$	$\mathcal{I}_4 (p.2)$
$\mathcal{I}_1^c$	<i>r.1</i>		
$\mathcal{I}_2^c$		<i>r.3</i>	
$\mathcal{I}_3^c$	<i>r.1</i>		
$\mathcal{I}_4^c$		<i>r.1</i>	
$\mathcal{I}_5^c$			<i>r.3</i>
$\mathcal{I}_6^c$			<i>r.2</i>
$\mathcal{I}_7^c$		<i>r.2</i>	
$\mathcal{I}_8^c$			<i>r.2</i>

(a) Step 0

composition.cloud service	$\mathcal{I}_2^c \mathcal{I}_2$	$\mathcal{I}_3^c \mathcal{I}_3$	$\mathcal{I}_7^c \mathcal{I}_2$	$\mathcal{I}_6^c \mathcal{I}_3$	$\mathcal{I}_1^c \mathcal{I}_1$	$\mathcal{I}_3^c \mathcal{I}_1$	$\mathcal{I}_4^c \mathcal{I}_2$
property distance	(+1)	(+1)	(0)	(0)	(-1)	(-1)	(-1)

(b) Step 1

order	migration	migration impact	action	$\mathcal{I}_4(p.2)$
1	$\mathcal{I}_7^c \mathcal{I}_2 \xrightarrow{\mathcal{I}_7^c} \mathcal{I}_7^c \mathcal{I}_4$	$mi=-2$	migrate	$\mathcal{I}_1^c, \mathcal{I}_7^c$
2	$\mathcal{I}_6^c \mathcal{I}_3 \xrightarrow{\mathcal{I}_6^c} \mathcal{I}_6^c \mathcal{I}_4$	$mi=-1$	migrate	$\mathcal{I}_1^c, \mathcal{I}_7^c, \mathcal{I}_6^c$
3	$\mathcal{I}_1^c \mathcal{I}_1 \xrightarrow{\mathcal{I}_1^c} \mathcal{I}_1^c \mathcal{I}_4$	$mi=+2$	-	$\mathcal{I}_1^c, \mathcal{I}_7^c, \mathcal{I}_6^c$
4	$\mathcal{I}_3^c \mathcal{I}_1 \xrightarrow{\mathcal{I}_3^c} \mathcal{I}_3^c \mathcal{I}_4$	$mi=+3$	-	$\mathcal{I}_1^c, \mathcal{I}_7^c, \mathcal{I}_6^c$
5	$\mathcal{I}_4^c \mathcal{I}_2 \xrightarrow{\mathcal{I}_4^c} \mathcal{I}_4^c \mathcal{I}_4$	$mi=-2$	migrate	$\mathcal{I}_1^c, \mathcal{I}_7^c, \mathcal{I}_6^c, \mathcal{I}_4^c$

(c) Step 2

$\mathcal{I}_1 (p.1)$	$\mathcal{I}_2 (p.3)$	$\mathcal{I}_3 (p.3)$	$\mathcal{I}_4 (p.2)$	$\mathcal{I}_1 (p.1)$	$\mathcal{I}_2 \cup \mathcal{I}_3 (p.3)$	$\mathcal{I}_4 (p.2)$
$\mathcal{I}_1^c$	<i>r.1</i>			$\mathcal{I}_1^c$	<i>r.1</i>	
$\mathcal{I}_2^c$		<i>r.3</i>		$\mathcal{I}_2^c$		<i>r.3</i>
$\mathcal{I}_3^c$	<i>r.1</i>			$\mathcal{I}_3^c$	<i>r.1</i>	
$\mathcal{I}_4^c$			<i>r.1</i>	$\mathcal{I}_4^c$		<i>r.1</i>
$\mathcal{I}_5^c$		<i>r.3</i>		$\mathcal{I}_5^c$		<i>r.3</i>
$\mathcal{I}_6^c$			<i>r.2</i>	$\mathcal{I}_6^c$		<i>r.2</i>
$\mathcal{I}_7^c$			<i>r.2</i>	$\mathcal{I}_7^c$		<i>r.2</i>
$\mathcal{I}_8^c$			<i>r.2</i>	$\mathcal{I}_8^c$		<i>r.2</i>

(d) Step 3

(e) Step 3

Figure 5: An example of heuristic 2 execution

637 measuring distance  $Dist(\mathcal{T}_i^c, \mathcal{T}_j, r, \mathcal{I}_4.p)$ .

638 Step 2 – Service migration (Figure 5(c)). The migration impact  $mi$  in Equation 6  
639 is calculated for each of the  $\mathcal{I}_i^c$  (denoted with a gray background in Figure 5(b))  
640 showing a distance that is less or equal to zero (step 2). Figure 5(c) shows  
641 the results of our migration, that is,  $\mathcal{I}_6^c$  is migrated from  $\mathcal{I}_3$  to  $\mathcal{I}_4$ , and  $\mathcal{I}_4^c$  and  
642  $\mathcal{I}_7^c$  are migrated from  $\mathcal{I}_2$  to  $\mathcal{I}_4$ . We note that all these migrations are of type

643 replacement as presented in Definition 4.4.

644 Step 3 – Resource consolidation (Figures 5(d) and 5(e)). Upon phase service  
645 migration ends, phase resource consolidation is executed and considers all  $\mathcal{I}_j$   
646 such that at least one composite service  $\mathcal{I}_i^c$  insisting on it has been migrated to  $\mathcal{I}_4$ .  
647 In our example, service instances  $\mathcal{I}_2$  and  $\mathcal{I}_3$  are candidates for the binary join  
648 (denoted with a light grey background in Figure 5(d)). Since the join between  $\mathcal{I}_2$   
649 and  $\mathcal{I}_3$  has a negative  $mi=-2$ , resource consolidation is convenient and applied.  
650 Figure 5(e) finally shows the new CP status after the execution of heuristic 2,  
651 where the result of the join operation is denoted with a gray background. We  
652 note that all migrations due to consolidation are of type versioning as presented  
653 in Definition 4.3.

## 654 7. Experimental Evaluation

655 We experimentally evaluated the performance and quality of our approach  
656 for cost-effective deployment of service compositions, and the utility of our  
657 portable certification process.

### 658 7.1. Experimental Setup

659 We considered a scenario where a cloud provider hosts the three composi-  
660 tions depicted in Figure 2. For simplicity but with no lack of generality, we  
661 focused on the payment functionality only, which is used in all compositions.  
662 This choice was due to the fact that, as already discussed, considering the entire  
663 composition as a whole does not give any additional insights on the soundness  
664 of the proposed approach and its performance/cost. Our methodology in fact  
665 treats each functionality independently and the total cost is calculated as the  
666 sum of the cost of each functionality. CP offers two payment services, a standard  
667 payment service and ENGPay payment service offered by Engineering S.p.A.,  
668 one of the biggest system integrators in Italy, all certified for property PCI-  
669 DSS compliance. We considered three different certification levels for property

670 PCI-DSS compliance  $\mathcal{P}_c$  from basic confidentiality ( $\mathcal{P}_c.\text{level}=1$ ) to full PCI-  
 671 DSS ( $\mathcal{P}_c.\text{level}=3$ ), via generic CIA – Confidentiality, Integrity, Authentication  
 672 ( $\mathcal{P}_c.\text{level}=2$ ). Standard payment service is certified for property PCI-DSS com-  
 673 pliance at level 1; ENGPAY is offered with two levels of certification, level 2 and  
 674 level 3.

675 We developed a *request simulator* that randomly generates requests for a  
 676 payment service with a specific property level. We then built 10 data sets of  
 677 300 consecutive random requests  $\mathcal{T}^c$  submitted to the cloud provider. For all  
 678 data sets, we evaluated the deployment obtained using the sliding window and  
 679 migration heuristics in Section 6.1 with sharing and fitting profiles in Section 5.3.  
 680 We evaluated retrieved results according to *i*) a set of evaluation metrics, *ii*)  
 681 the fuzzy membership functions, and *iii*) the cost functions.

682 **Evaluation metrics.** We used three metrics to evaluate our approach.

683 *Metric 1* measures the execution time needed to deploy composite services  
 684 addressing composition requests.

685 *Metric 2*, called  $\Gamma_t(TFc, TFc')$ , is the relative cost increment. It is calculated  
 686 as the difference between the two areas identified by Total Fuzzy cost functions  
 687  $TFc$  and  $TFc'$  in the interval  $[1, t]$ . It is formally defined as follows:

$$\Gamma_t(TFc, TFc') = \frac{\sum_{i=1}^t (TFc_i - TFc'_i)}{\sum_{i=1}^t TFc_i} \quad (7)$$

688 where  $TFc_i$  and  $TFc'_i$  are the two Total Fuzzy cost functions evaluated at  
 689 time  $i$ . We used Total Fuzzy cost to calculate  $\Gamma$ , since our goal here is to evaluate  
 690 the overall cost increase and not the contribution of each cost factor ( $\alpha$ ,  $\beta$ , and  
 691  $\gamma$ ).

692 *Metric 3*, called  $\Delta_t$ , evaluates the cumulative number of portability events  
 693 (versioning or replacement) occurred until a given time  $t$ . It provides a measure  
 694 of how often a portability event and a consistency check are needed to support  
 695 our dynamic composition certification, and in turns a measure of its utility.

**Fuzzy membership functions.** Our fuzzy system is based on membership  
 functions and fuzzy rules that depend on the cost factors to be evaluated. We

adopted the generalized bell-based memberships  $f$  for all cost factors

$$f(x; a, b, c) = \frac{1}{1 + \left|\frac{x-c}{a}\right|^{2b}} \quad (8)$$

696 where  $c$  is the center of the curve,  $a$  controls the width of the curve, and  $b$   
 697 controls the slope of the curve. To optimize the membership function definition,  
 698 we evaluated the distribution of costs to adjust parameters  $a$ ,  $b$ ,  $c$  to the meaning  
 699 of the corresponding linguistic variable. More precisely, we used a fuzzy c-mean  
 700 approach to have an initial idea on the membership shapes using 100 requests  
 701 from each of the 10 data sets. Given this shape we tuned the membership  
 702 parameters to fit the fuzzy clusters with a *gbell* shape. We note that this  
 703 process, as well as the cost function definition, is a tuning process that may  
 704 depend on the cloud provider peculiarities. In general, the selected cost and  
 705 membership functions are suitable for a generic cloud provider working with  
 706 cloud service compositions, while the rule sets address the peculiarities of the  
 707 profiles.

708 **Cost functions.** We used cost functions  $\alpha$ ,  $\beta$ , and  $\gamma$  in Figure 3 for the three  
 709 property levels used in our experiments. We recall that cost function  $\gamma$  is defined  
 710 using property levels and ranges from 0 to 2.  $\alpha$ ,  $\beta$ , and  $\gamma$  have been used to  
 711 compare the cost retrieved by our heuristics (metric 2 and metric 3). Their  
 712 definition is CP specific and should reflect the costs of the CP infrastructure.  
 713 To fully evaluate our approach, we defined cost functions such that service  
 714 migrations are triggered also with low numbers of composition requests.

715 We run our experiments on a Blade server PowerEdge M630 (VRTX) 2 x  
 716 Intel(R) Xeon(R) CPU E5-2620 v4 2.10GHz 192GB of RAM 120GB SSD.

## 717 7.2. Performance evaluation

718 We evaluated the performance of our heuristics using *metric 1*, and cost  
 719 factors  $\alpha$ ,  $\beta$ , and  $\gamma$  in Figure 3. Similar to the exhaustive algorithm, our heuris-  
 720 tics have an exponential asymptotic behavior  $\mathcal{O}\left(\left(|l| * F * |w| + t\right)^{|w|}\right)$ , with  $|l|$   
 721 the number of property levels,  $F$  the number of functionalities,  $|w|$  the size of

722 window  $w$ , and  $t$  the number of received requests. We note that, since  $|w|$  is  
723 fixed a priori by our heuristics, the asymptotic behavior becomes polynomial  
724 as  $\mathcal{O}(t^{|w|})$ . Window  $w$  however makes our heuristics rapidly unusable given  
725 our assumption to serve requests in pseudo real time (in the order of minutes),  
726 though their complexity is far lower than the one of the exhaustive algorithm,  
727 which is  $\mathcal{O}(t^t)$ .

728 Figure 6 compares the average execution time of the heuristics and exhaus-  
729 tive algorithm on the 10 data sets, varying window size  $|w|$  from 1 to 7. We  
730 note that the execution time of all algorithms is reported only for configura-  
731 tions requiring less than 3-minutes. Sharing and fitting profiles show a similar  
732 performance trend just partially affected by optimizations based on branch cut.  
733 Heuristic 2 shows an additive execution time increment with respect to heuris-  
734 tic 1 due to the migration and consolidation algorithms, which require sorting of  
735 compatible requests ( $\mathcal{O}((t-1)^2)$  at time  $t$ , in the worst case). This additive fac-  
736 tor is not anyways substantial, since it depends on the presence and amount of  
737 possible migrations (e.g., between  $t = 70$  and  $t = 80$ ). Our results show that, as  
738 expected, the heuristics approximates polynomial execution time in the window  
739 size  $|w|$ , which can be taken under control by selecting proper  $|w|$ . For instance,  
740 when  $|w|=7$ , execution time exceeds the 3 minute limit with a number  $|\mathcal{T}^c|$  of  
741 composition requests equal to 22 for heuristic 1 and 20 for heuristic 2; when  
742  $|w|=6$ , execution time exceeds the 3-minute limit with  $|\mathcal{T}^c|=155$  for heuristic 1  
743 and  $|\mathcal{T}^c|=143$  for heuristic 2. The exhaustive algorithm shows the worst exe-  
744 cution time, exceeding the 3-minute limit with  $|\mathcal{T}^c|=12$ . We note that the two  
745 heuristics have comparable performance dominated by  $w$ . We also note that the  
746 exhaustive algorithm has better performance than our heuristics for a number  
747  $|\mathcal{T}^c| \leq |w|$  of requests, because the heuristics use a window size  $|w|$ , while ex-  
748 haustive algorithm uses the entire set of requests. Therefore, when the number  
749 of requests is less than  $|w|$ , it provides better or comparable performance.

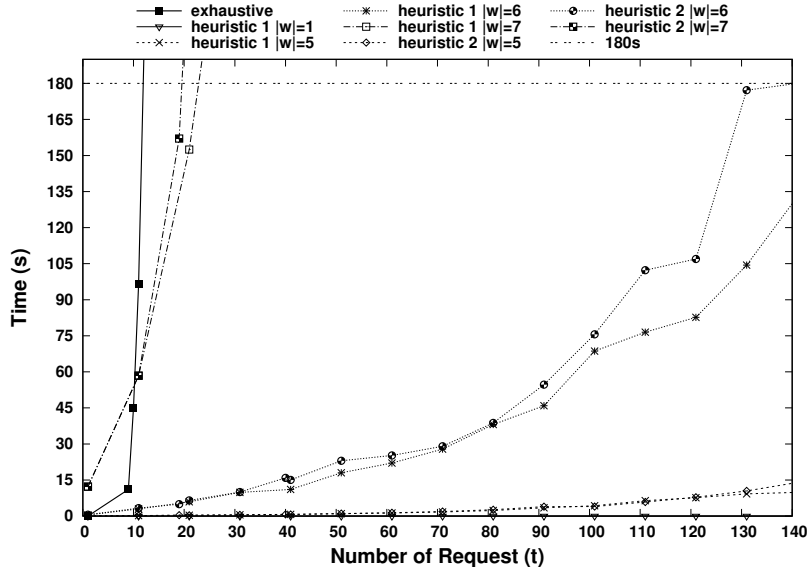


Figure 6: Performance evaluation: heuristic 1 and heuristic 2 varying window size  $|w|$ .

750 *7.3. Cost and Utility Evaluation*

751 Performance evaluation in Section 7.2 showed the unmanageable complexity  
752 of the exhaustive algorithm, which required 21 minutes for deploying 12 requests.  
753 We therefore compared the costs of our two heuristics on the 10 data sets and  
754 the utility of the portability underpinning them using *metric 2* and *metric 3*,  
755 and cost factors  $\alpha$ ,  $\beta$ , and  $\gamma$ . We first discuss the impact of windows size  
756 and then compare heuristic 1 and heuristic 2.

757 *7.3.1. Window size*

758 We measured the impact of window  $w$  on heuristic 1 and heuristic 2 using  
759 the relative cost increment (*metric 2* -  $\Gamma_t$ ) with sharing and fitting profiles  
760 and the entire data sets of 300 requests. Table 1 shows the average relative  
761 cost increment ( $\overline{\Gamma}_t$ ), expressed in percentage, over our 10 data sets, varying  
762  $|w|$ . Fitting profile shows a negligible variation of  $\overline{\Gamma}_t$  both for heuristic 1 and  
763 heuristic 2. Fitting profile in fact does not take significant advantages by looking  
764 forward in the incoming requests. In particular, Table 1 shows that the total



$\bar{\Gamma}_t$		$TFc_{w2}-TFc_{w1}$	$TFc_{w3}-TFc_{w2}$	$TFc_{w4}-TFc_{w3}$	$TFc_{w5}-TFc_{w4}$	$TFc_{w5}-TFc_{w1}$
heuristic 1	sharing	-0.6%	-1.5%	-1.31%	-2.08%	-5.48%
	fitting	0.6%	0.09%	-0.23%	-0.19%	0.29%
heuristic 2	sharing	-1.2%	-1.3%	-1.7%	-0.86%	-5.39%
	fitting	0.46%	0.06%	0.11%	0.007%	0.6%

Table 1: Average relative cost difference ( $\bar{\Gamma}_t$  in %) at  $t=300$  for heuristic 1 and heuristic 2, using sharing and fitting profiles, and varying window size  $|w|$ . We denote with  $TFc_{wi}$  the total fuzzy costs  $TFc$  with  $|w|=i$ .

765 fuzzy costs  $TFc_{w5}$  with  $|w|=5$  increases on average of 0.29% with respect to the  
766 total fuzzy costs  $TFc_{w1}$  with  $|w|=1$  (denoted as  $TFc_{w5}-TFc_{w1}$  in Table 1) for  
767 heuristic 1 and of 0.6% for heuristic 2. When the sharing profile is adopted,  
768 the average cost decreases as the window size increases. Table 1 shows that the  
769 total fuzzy costs  $TFc_{w5}$  with  $|w|=5$  decrease on average of -5.48% with respect  
770 to the total fuzzy costs  $TFc_{w1}$  with window  $|w|=1$  for heuristic 1 and of -5.39%  
771 for heuristic 2.

772 Figure 7 shows an excerpt of the total fuzzy cost  $TFc$  of heuristic 1 for 4  
773 representative data sets for sharing profile, varying the window size from  $|w|=1$   
774 to  $|w|=5$ . We note that an increase in the window size  $|w|$  does not always  
775 result in a cost decrease. Figure 7(d) shows a data set where heuristic 1 with  
776  $|w|=4$  has lower total fuzzy cost than the one with  $|w|=5$ . This mainly depends  
777 on the bias introduced by the random generation of the data sets and by the  
778 random selection of the best deployment when different candidate deployments  
779 have the same total fuzzy cost. This latter scenario may lead to a sub-optimal  
780 deployment drifting from the optimal total cost and is more probable at the  
781 beginning of the deployment process where there are more deployments with  
782 the same cost.

783 Figure 8 shows an excerpt of the total fuzzy cost  $TFc$  of heuristic 2 for 4  
784 representative data sets for sharing profile, varying the window size from  $|w|=1$   
785 to  $|w|=5$ . We note that the behavior of heuristic 2 (Figure 8) is similar to the  
786 one of heuristic 1 (Figure 7), which is reasonable considering the refinement  
787 nature of heuristic 2. We also note that *i*) the drifting effects causing sub-

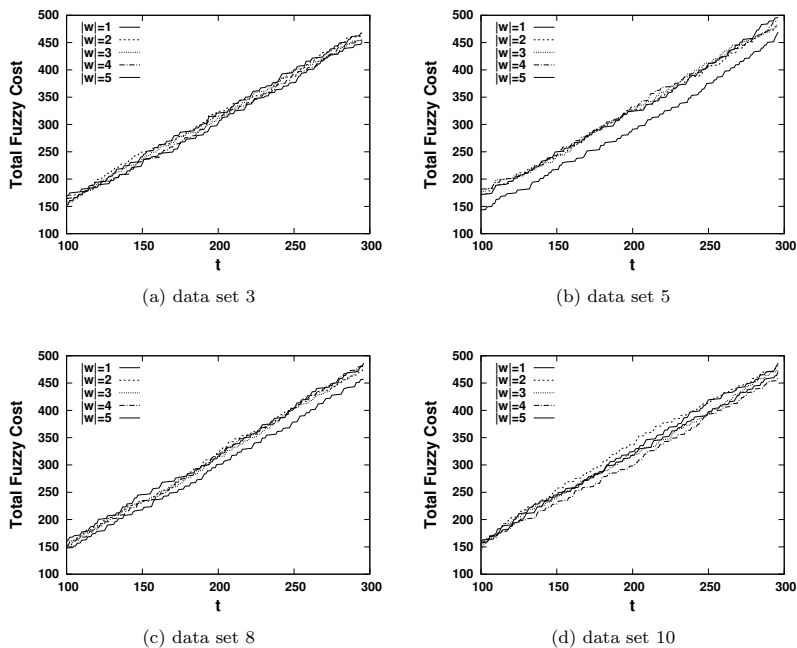


Figure 7: Heuristic 1 cost evaluation (TFc) for sharing profile varying window size  $|w|$ . Time frames 100–300 have been plotted to improve figure readability.

788 optimal deployments are reduced especially for bigger windows (greater than  
 789  $|w|=3$ ) and *ii*) the variance of the total fuzzy cost observed with  $|w|=3$ ,  $|w|=4$ ,  
 790 and  $|w|=5$  using heuristic 2 is lower than the one observed using heuristic 1,  
 791 meaning that heuristic 2 reduces the gap between different window sizes. This  
 792 effect is also visible in Table 1 where the average improvement between  $|w|=4$   
 793 and  $|w|=5$  is lower compared to the others.

### 794 7.3.2. Heuristic 1 vs. Heuristic 2

795 We compared heuristic 1 and heuristic 2 using the relative cost increment  
 796 ( $metric\ 2 - \Gamma_t$ ), the total fuzzy cost  $TFc$  with sharing and fitting profiles, and  
 797 the entire data sets of 300 requests. Table 2 shows the average relative cost  
 798 difference ( $\overline{\Gamma_t}$ ) between heuristic 2 and heuristic 1, expressed in percentage,  
 799 over our 10 data sets, varying  $|w|$ . Negative values indicate a cost decrease  
 800 in heuristic 2 with respect to heuristic 1. We note that, even with a minimal

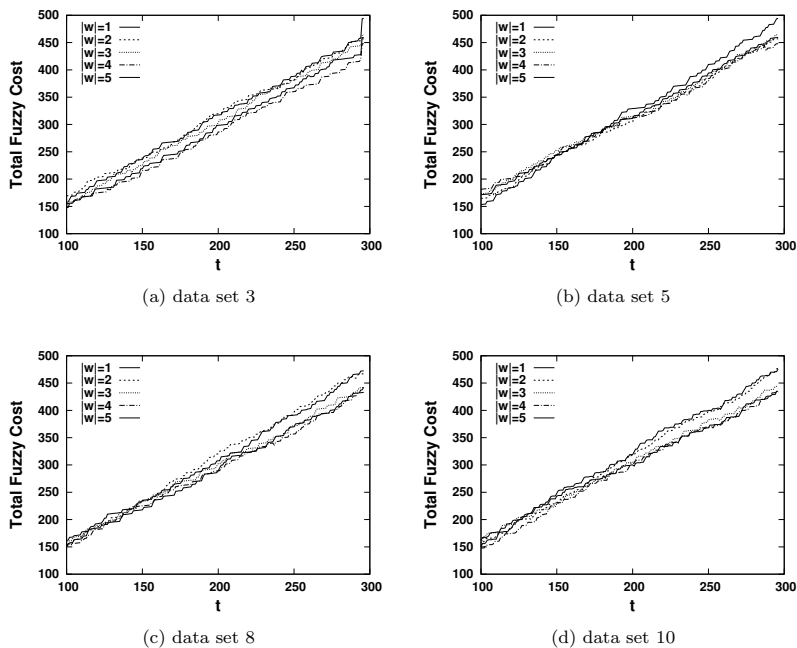


Figure 8: Heuristic 2 cost (TFc) evaluation for sharing profile varying window size  $|w|$ . Time frames 100–300 have been plotted to improve figure readability.

$\bar{\Gamma}_t$	$ w =1$	$ w =2$	$ w =3$	$ w =4$	$ w =5$
sharing	−0.54%	−1.20%	−1.16%	−1.89%	−0.46%
fitting	−3.96%	−3.81%	−3.88%	−3.08%	−2.78%

Table 2: Average relative cost difference ( $\bar{\Gamma}_t$  in %) at  $t=300$  between heuristic 2 and heuristic 1, using sharing and fitting profiles, and varying window size  $|w|$ . Negative values indicate a cost decrease in heuristic 2 with respect to heuristic 1.

801 bias effect introduced by different windows sizes, heuristic 2 outperforms, on  
802 average, heuristic 1 regardless the used profile. More specifically, considering  
803 the sharing profile, our results show that the average relative cost difference be-  
804 tween the two heuristics is around  $-1\%$ . Considering fitting profile, the average  
805 relative cost difference is more than three times higher for all window sizes and  
806 is around  $-3.5\%$ . For this reason, in the following, we further discuss the effect  
807 of heuristic 2 on fitting profile only.

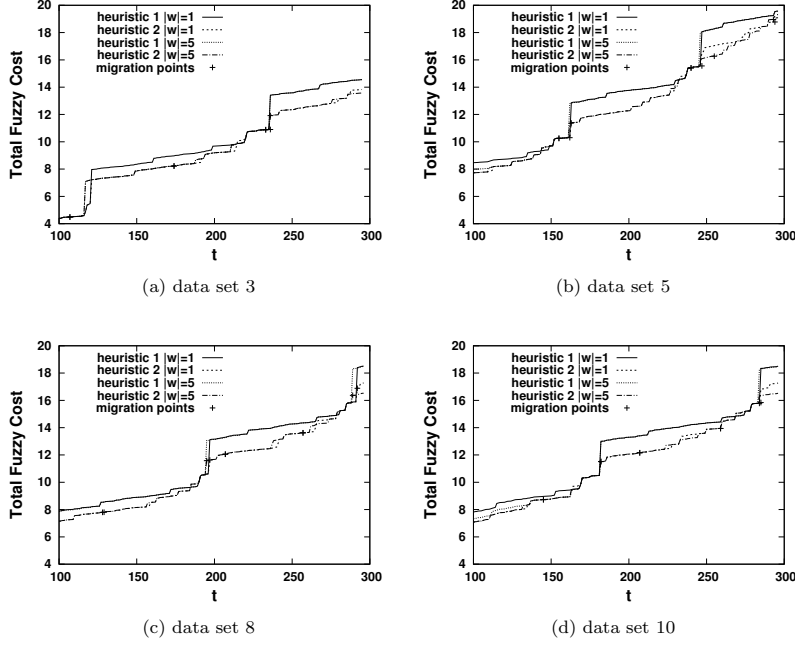


Figure 9: Comparison between heuristics 1 and 2 with fitting profile (TFc), using  $|w|=1$  and  $|w|=5$ . Time frames 100–300 have been plotted to improve figure readability. Migration events are marked with “+”.

808 Figure 9 shows the comparison of the total fuzzy costs  $TFc$  of heuristic 1  
 809 and heuristic 2, using 4 representative data sets and window size  $|w|=1$  and  
 810  $|w|=5$ . We note that, when a migration is triggered in heuristic 2 (denoted with  
 811 “+” in the figure), a substantial decrease in  $TFc$  is observed (e.g., at  $t=161$   
 812 and  $t=244$  for data set 5 in Figure 9(b)). We also note that, when  $|w|=1$ , the  
 813 migrations are synchronous with the cost increment, while, when  $|w|=5$  and in  
 814 general with  $|w|>1$ , the migrations are triggered before the cost increment due  
 815 to the look-forward effect of window  $w$ . We finally note that the total fuzzy  
 816 costs  $TFc$  in Figure 9 presents a “step behavior” that mimics the one of cost  
 817 factors  $\alpha$  and  $\beta$  and confirms the findings in Table 1 about the negligible impact  
 818 of window size on both heuristic 1 and heuristic 2 for fitting profile.

819 To further evaluate the contribution of migrations in heuristic 2, Figure 10  
 820 presents the distribution of the difference between  $\Gamma_t$  of heuristic 2 and  $\Gamma_t$  of

821 heuristic 1, where a negative value indicates a cost saving when heuristic 2 is  
 822 used, with  $|w|=5$  for 4 representative data sets and fitting profile. Here, a recur-  
 823 rent pattern “a negative peak followed by uphill steps” can be easily identified  
 824 in all data sets and reflects the negative peak (cost decrease) introduced by a  
 825 migration followed by a step-like cost degradation showing a convergence be-  
 826 tween the two heuristics after migration. For instance, in case of data set 5,  
 827 after migration at  $t=162$ , the cost shows a negative peak followed by a cost  
 828 degradation (cost increase), until the following migration at  $t=189$ . We note  
 829 that degradation can lead to a cost increase (e.g., positive difference for data  
 830 set 8 and data set 10), where heuristic 1 is less costly than heuristic 2. This  
 831 effect is however limited to a very short period of time and the cost difference  
 832 is very small. Considering all 10 data sets and  $|w|=5$ , heuristic 1 outperforms  
 833 heuristic 2 only 7 times for data set 8 and 6 times for data sets 1, 9, 10, showing  
 834 a maximum cost difference of 0.23%. Quantitatively ( $metric\ 3 - \overline{\Delta}_t$ ), heuristic 2  
 835 produced 8 migrations on average on the 10 data sets for sharing profile and  
 836 window  $|w|=5$ , 5 migrations for fitting profile and window  $|w|=5$ .

837 We note that, the comparison in this section, being based on total fuzzy  
 838 costs, is affected by the “normalization” introduced by defuzzification, and re-  
 839 flects a *perceived cost* more than a *concrete cost*. To provide a more tangible  
 840 quantitative analysis of the cost decrease introduced by heuristic 2, Figure 11  
 841 presents a comparison in terms of  $\alpha$  and  $\beta$  for sharing and fitting profiles over the  
 842 10 data sets. In particular, *i*) Figure 11(a) compares heuristic 1 (fitting profile,  
 843  $|w|=1$ ) with heuristic 2 (fitting profile,  $|w|=5$ ), *ii*) Figure 11(b) compares heuris-  
 844 tic 1 (sharing profile,  $|w|=1$ ) with heuristic 2 (sharing profile,  $w=5$ ). Heuristic 2  
 845 provides an average cost reduction over heuristic 1 on the real cost  $\alpha + \beta$  of  
 846  $-18.5\%$  for sharing profile and  $-8.2\%$  for fitting profile.

847 To conclude, while proving the effectiveness of both heuristics, our exper-  
 848 iments show the utility of certification portability supporting migrations (re-  
 849 placement), as well as new instantiations of the same services (versioning). We  
 850 remark that, while migrations generate intrinsic costs [27], they are in most of  
 851 the cases “operational costs” due to computation or bandwidth degradation [28].

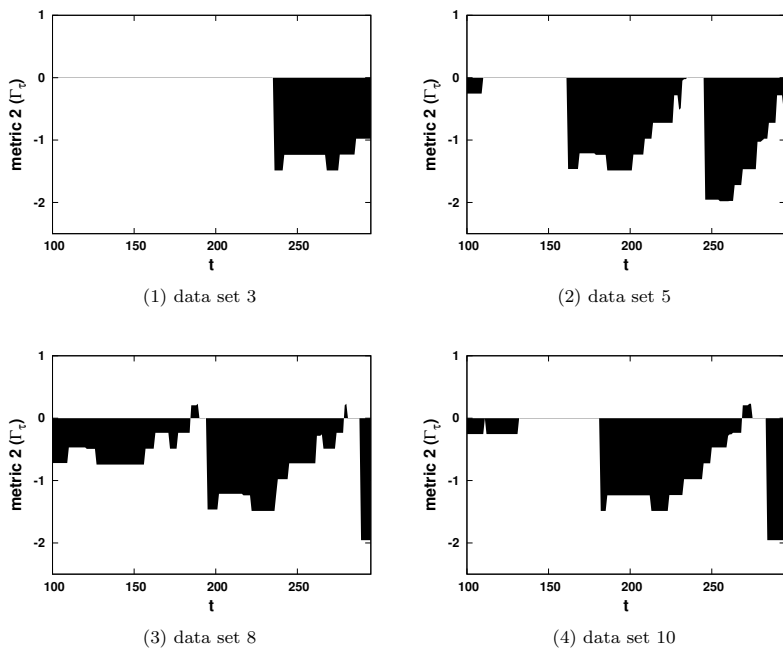
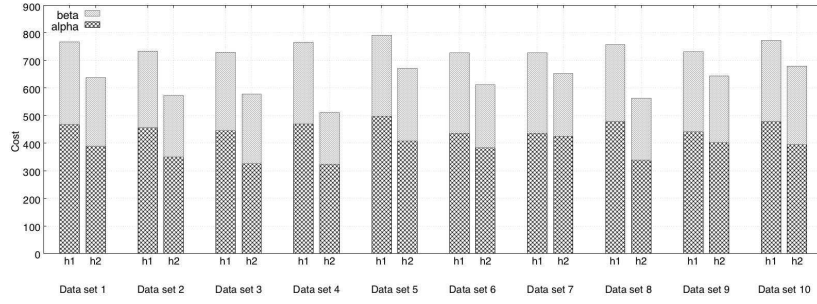


Figure 10: Filled curve plot of the distribution of the difference between  $\Gamma_t$  of heuristic 2 and  $\Gamma_t$  of heuristic 1. Time frames 100–300 have been plotted to improve figure readability.

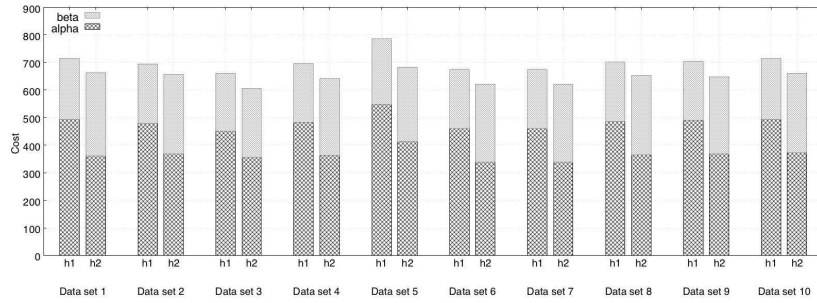
852 These costs can be normally mitigated by a number of well-established tech-  
 853 niques (e.g., cold start servers, microservice architectures), as well as elasticity-  
 854 based solutions [29]. In case of no mitigations, our approach can deal with  
 855 migration events as additional CP costs to be added within the general cost  
 856 functions. However, this scenario is not in the scope of this experimental eval-  
 857 uation and will be considered in our future work.

## 858 8. Related Work

859 Research on cloud service composition has recently focused on the prob-  
 860 lem of selecting component services on the basis of non-functional (including  
 861 security) requirements [4, 7, 8, 9, 30, 31, 32]. Wang et al. [9] propose a network-  
 862 aware service composition, which builds on candidate services geolocation to  
 863 keep stable network performance. Qi et al. [7] propose a QoS-aware composi-



(a)



(b)

Figure 11: Comparing  $\alpha$ ,  $\beta$  total cost for the 10 data sets for sharing (a) and fitting (b) profiles using heuristic 1 window  $|w|=1$  and heuristic 2  $|w|=5$ , respectively.

864 tion method supporting cross-platform service invocation in cloud environment,  
 865 using the execution time of single components. Manshan et al. [31] propose a  
 866 fuzzy way solution to represent and solve QoS-based web services composition.  
 867 Wu et al. [30] discuss a composition method providing a trustworthy selection  
 868 of component services and guaranteeing trust in the composition. Kurdi et  
 869 al. [8] focus on multiple cloud composition providing a combinatorial optimiza-  
 870 tion algorithm for cloud service composition, aimed to maximize the fulfilment  
 871 of clients' requests with minimal overhead. Arman et al. [32] propose a solu-  
 872 tion for moving application to the cloud, aiming to select the cloud service that  
 873 matches at best application requirements and plan characteristics. Another line  
 874 of research relevant for the work in this paper evaluated the costs of service  
 875 composition algorithms [12, 13, 14, 33, 34]. Greenberg et al. [23] analyzed the  
 876 most relevant direct and indirect costs in cloud service data centers, identi-

877 fying the following: *i*) server costs (45% of the total costs) depend on server  
878 utilization and optimization, *ii*) infrastructure costs (25% of the total costs)  
879 comprise all facilities for power delivery, air conditioning, ups, and the like, *iii*)  
880 power draw (15% of the total costs) consists of all power consumption, and  
881 *iv*) network costs (15% of the total costs) comprises costs for switches, routers,  
882 agreements and traffic with ISP. Patel et al. [35] extend the above costs with  
883 the cost of personnel per rack and license costs, and identify model and func-  
884 tions representing the whole data center costs. We added on top of these works  
885 the costs of managing service certification and continuous service verification,  
886 as well as the costs due to *service versioning* and *service replacement*. These  
887 costs are due to the need of guaranteeing business continuity or, in other words,  
888 the cost of keeping the replica of components up and running while migrating  
889 a service. Jiang et al. [36], identified CPs physical resources utilization as an  
890 emerging problem and proposed a cloud capacity planning based on an ensem-  
891 ble time-series prediction method. He et al. [13] propose three novel QoS-aware  
892 service selection approaches for composing multi-tenant service-based systems.  
893 Li et al. [34] compare costs and service behaviours from different CPs, while  
894 Medeiros et al. [12] provide different cost patterns which may fit different types  
895 of services and service composition. Singh et al. [37] propose an agent-based  
896 and autonomous framework able to optimize the resource provisioning cost; the  
897 approach only focuses on virtual machine composition.

898 The solution in this paper extends the the cost-based certification approach  
899 in [15]. It provides a certification-based service composition for the cloud, which  
900 continuously evaluates non-functional properties using declarative and procedu-  
901 ral modeling. It also investigates the composition costs from a cloud provider  
902 point of view [11, 14, 33], providing a cost-effective approach, using a fuzzy-based  
903 approach, for composite service deployment. The cost evaluation considers the  
904 costs introduced by the certification infrastructure and mismatch costs of main-  
905 taining services providing more than what is strictly needed for ensuring clients  
906 non-functional requirements.



## 907 **9. Conclusions**

908 We proposed an approach to cost-effective deployment of certified cloud com-  
909 posite services. The proposed solution provides a composition approach that  
910 comparatively evaluates service certificates to build a composite service address-  
911 ing clients' requirements. It also provides a fuzzy-based cost evaluation method-  
912 ology aimed to minimize the CP total costs of composition, identifying the best  
913 composition among possible alternative deployments. Our cost minimization  
914 approach has been implemented using two heuristics algorithms and assuming  
915 certified properties as *must-have* requirements. It has then been extensively ex-  
916 perimentally evaluated simulating composition requests and comparing average  
917 composition costs of the two heuristics varying specific settings like window size  
918 and profile. This paper leaves some space for future work. First, our approach to  
919 cost optimization will consider scenarios distinguishing between non-functional  
920 requirements annotated by the users as soft (should-have) or hard (must-have),  
921 where cost improvements can be achieved by relaxing soft properties. Second,  
922 possible applications of the migration concept will be investigated to the aim of  
923 improving sub-optimal deployments in the cloud.

## 924 **Acknowledgments**

925 This work was partly supported by the program “piano sostegno alla ricerca  
926 2015-17” funded by Università degli Studi di Milano, and the program FFABR  
927 “Fondo per il Finanziamento delle Attività Base di Ricerca”, funded by the  
928 Italian MIUR.

## 929 **References**

- 930 [1] S.-Y. Hwang, E.-P. Lim, C.-H. Lee, C.-H. Chen, Dynamic web service selec-  
931 tion for reliable web service composition, *IEEE TSC 1 (2)* (2008) 104–116.
- 932 [2] H. Wang, P. Ma, Q. Yu, D. Yang, J. Li, H. Fei, Combining quantitative con-  
933 straints with qualitative preferences for effective non-functional properties-  
934 aware service composition, *JPDC 100* (2017) 71 – 84.

- 935 [3] D. Ye, Q. He, Y. Wang, Y. Yang, An agent-based service adaptation ap-  
936 proach in distributed multi-tenant service-based systems, JPDC 122 (2018)  
937 11 – 25.
- 938 [4] M. Anisetti, C. Ardagna, E. Damiani, Security certification of composite  
939 services: A test-based approach, in: Proc. of ICWS 2013, San Francisco,  
940 CA, USA, 2013.
- 941 [5] F. Tao, D. Zhao, Y. Hu, Z. Zhou, Resource service composition and its  
942 optimal-selection based on particle swarm optimization in manufacturing  
943 grid system, IEEE TII 4 (4) (2008) 315–327.
- 944 [6] T. Xiang, X. Li, F. Chen, Y. Yang, S. Zhang, Achieving verifiable, dynamic  
945 and efficient auditing for outsourced database in cloud, JPDC 112 (2018)  
946 97 – 107.
- 947 [7] L. Qi, W. Dou, X. Zhang, J. Chen, A qos-aware composition method sup-  
948 porting cross-platform service invocation in cloud environment, J. Comput.  
949 Syst. Sci. 78 (5) (2012) 1316–1329.
- 950 [8] H. Kurdi, A. Al-Anazi, C. Campbell, A. Al Faries, A combinatorial opti-  
951 mization algorithm for multiple cloud service composition, Comput. Electr.  
952 Eng. 42 (C) (2015) 107–113.
- 953 [9] X. Wang, J. Zhu, Y. Shen, Network-aware qos prediction for service com-  
954 position using geolocation, IEEE TSC 8 (4) (2015) 630–643.
- 955 [10] K. Kofler, I. ul Haq, E. Schikuta, User-centric, heuristic optimization of  
956 service composition in clouds, in: Proc. of Euro-Par 2010, Ischia, Italy,  
957 2010.
- 958 [11] A. Jula, E. Sundararajan, Z. Othman, Cloud computing service compo-  
959 sition: A systematic literature review, Expert Systems with Applications  
960 41 (8) (2014) 3809–3824.

- 961 [12] R. Medeiros, N. S. Rosa, L. F. Pires, Predicting service composition costs  
962 with complex cost behavior, in: Proc. of IEEE SCC, New York, NY, USA,  
963 2015.
- 964 [13] Q. He, J. Han, F. Chen, Y. Wang, R. Vasa, Y. Yang, H. Jin, Qos-aware  
965 service selection for customisable multi-tenant service-based systems: Ma-  
966 turity and approaches, in: Proc. of IEEE CLOUD, New York, NY, USA,  
967 2015.
- 968 [14] P. Leitner, W. Hummer, S. Dustdar, Cost-based optimization of service  
969 compositions, IEEE TSC 6 (2) (2013) 239–251.
- 970 [15] M. Anisetti, C. Ardagna, E. Damiani, F. Gaudenzi, A cost-effective  
971 certification-based service composition for the cloud, in: Proc of IEEE  
972 SCC 2016, San Francisco, CA, USA, 2016.
- 973 [16] M. Anisetti, C. A. Ardagna, E. Damiani, J. Maggesi, Security certification-  
974 aware service discovery and selection, in: Proc. of IEEE SOCA, 2012, pp.  
975 1–8.
- 976 [17] M. Anisetti, C. Ardagna, E. Damiani, G. Polegri, Test-based security cer-  
977 tification of composite services, ACM TWEB 13 (1) (2018) 3.
- 978 [18] M. Anisetti, C. Ardagna, E. Damiani, F. Gaudenzi, A certification frame-  
979 work for cloud-based services, in: Proc. of ACM SAC 2016, Pisa, Italy,  
980 2016.
- 981 [19] C. Consortium, D2.1: Security-aware SLA specification lan-  
982 guage and cloud security dependency model, [https://cordis.](https://cordis.europa.eu/docs/projects/cnect/0/318580/080/deliverables/001-D21SecurityawareSLAspecificationlanguageandcloudsecuritydependencymodelv101.pdf)  
983 [europa.eu/docs/projects/cnect/0/318580/080/deliverables/](https://cordis.europa.eu/docs/projects/cnect/0/318580/080/deliverables/001-D21SecurityawareSLAspecificationlanguageandcloudsecuritydependencymodelv101.pdf)  
984 [001-D21SecurityawareSLAspecificationlanguageandcloudsecuritydependencymodelv101.](https://cordis.europa.eu/docs/projects/cnect/0/318580/080/deliverables/001-D21SecurityawareSLAspecificationlanguageandcloudsecuritydependencymodelv101.pdf)  
985 [pdf](https://cordis.europa.eu/docs/projects/cnect/0/318580/080/deliverables/001-D21SecurityawareSLAspecificationlanguageandcloudsecuritydependencymodelv101.pdf), Accessed in Date June 2019.
- 986 [20] M. Anisetti, C. Ardagna, E. Damiani, F. Gaudenzi, A semi-automatic and  
987 trustworthy scheme for continuous cloud service certification, IEEE TSC  
988 (2017).

- 989 [21] M. Anisetti, C. Ardagna, E. Damiani, A certification-based trust model for  
990 autonomic cloud computing systems, in: Proc. of IEEE ICCAC, 2014, pp.  
991 212–219.
- 992 [22] C. T. Horngren, G. Foster, S. M. Datar, M. Rajan, C. Ittner, A. A. Baldwin,  
993 Cost accounting: A managerial emphasis, *Issues in Accounting Education*  
994 25 (4) (2010) 789–790.
- 995 [23] A. Greenberg, J. Hamilton, D. A. Maltz, P. Patel, The cost of a cloud:  
996 Research problems in data center networks, *SIGCOMM Comput. Commun.*  
997 *Rev.* 39 (1) (2008) 68–73.
- 998 [24] S. Newman, *Building Microservices*, O’Reilly Media, Inc., 2015.
- 999 [25] C. Sadtler, Z. X. Chen, S. Imazeki, M. Kelm, S. Kofkin-Hansen, Z. Q. Kou,  
1000 B. McChesney, et al., *IBM Workload Deployer: Pattern-based Application*  
1001 *and Middleware Deployments in a Private Cloud*, IBM Redbooks, 2012.
- 1002 [26] I. Maleki, L. Ebrahimi, S. Jodati, I. Ramesh, Analysis of software cost  
1003 estimation using fuzzy logic, *IJFCST* 4 (3) (2014) 27–41.
- 1004 [27] S. Kikuchi, Y. Matsumoto, Performance modeling of concurrent live mi-  
1005 gration operations in cloud computing systems using prism probabilistic  
1006 model checker, in: Proc of IEEE CLOUD, 2011, pp. 49–56.
- 1007 [28] D. Breitgand, G. Kutiel, D. Raz, Cost-aware live migration of services in  
1008 the cloud., in: SYSTOR, 2010.
- 1009 [29] U. Sharma, P. Shenoy, S. Sahu, A. Shaikh, A cost-aware elasticity provi-  
1010 sioning system for the cloud, in: Proc. of ICDCS, IEEE, 2011, pp. 559–570.
- 1011 [30] X. Wu, B. Li, R. Song, C. Liu, S. Qi, Trust-based service composition and  
1012 optimization, in: Proc of APSEC 2012, Hong Kong, 2012.
- 1013 [31] M. Lin, J. Xie, H. Guo, H. Wang, Solving qos-driven web service dynamic  
1014 composition as fuzzy constraint satisfaction, in: Proc. of IEEE EEE, 2005,  
1015 pp. 9–14.

- 1016 [32] A. Arman, S. Foresti, G. Livraga, P. Samarati, A consensus-based approach  
1017 for selecting cloud plans, in: Proc. of IEEE RTSI, 2016, pp. 1–6.
- 1018 [33] D. Worm, M. Zivkovic, H. van den Berg, R. van der Mei, Revenue max-  
1019 imization with quality assurance for composite web services, in: Proc. of  
1020 IEEE SOCA, Taipei, Taiwan, 2012.
- 1021 [34] A. Li, X. Yang, S. Kandula, M. Zhang, Cloudcmp: Comparing public cloud  
1022 providers, in: Proc. of IMC 2010, Melbourne, Australia, 2010.
- 1023 [35] C. D. Patel, J. S. Amip, Cost model for planning, development and  
1024 operation of a data center,[http://hpl.hp.com/techreports/2005/  
1025 HPL-2005-107R1.pdf](http://hpl.hp.com/techreports/2005/HPL-2005-107R1.pdf).
- 1026 [36] Y. Jiang, C.-s. Perng, T. Li, R. Chang, Self-adaptive cloud capacity plan-  
1027 ning, in: Proc. of IEEE SOCA, 2012, pp. 73–80.
- 1028 [37] A. Singh, D. Juneja, M. Malhotra, A novel agent based autonomous and  
1029 service composition framework for cost optimization of resource provision-  
1030 ing in cloud computing, Journal of King Saud University - Computer and  
1031 Information Sciences 29 (1) (2017) 19 – 28.