

Capturing the functionality of Web services with functional descriptions

Ruben Verborgh · Thomas Steiner · Davy Van Deursen ·
Jos De Roo · Rik Van de Walle · Joaquim Gabarró Vallés

Published online: 5 February 2012
© Springer Science+Business Media, LLC 2012

Abstract Many have left their footprints on the field of semantic RESTful Web service description. Albeit some of the propositions are even W3C Recommendations, none of the proposed standards could gain significant adoption with Web service providers. Some approaches were supposedly too complex and verbose, others were considered not RESTful, and some failed to reach a significant majority of API providers for a combination of the reasons above. While we neither have the silver bullet for universal Web service description, with this paper, we want to suggest a lightweight approach called RESTdesc. It expresses the semantics of Web services by pre- and postconditions in simple N3 rules, and integrates existing standards and conventions such as Link headers, HTTP OPTIONS, and URI templates for

R. Verborgh (✉) · D. Van Deursen · R. Van de Walle
ELIS – Multimedia Lab, Ghent University – IBBT, Gaston Crommenlaan 8 bus 201, 9050
Ledeborg-Ghent, Belgium
e-mail: ruben.verborgh@ugent.be

D. Van Deursen
e-mail: davy.vandeursen@ugent.be

R. Van de Walle
e-mail: rik.vandewalle@ugent.be

T. Steiner · J. Gabarró Vallés
Universitat Politècnica de Catalunya – Department LSI, 08034 Barcelona, Spain

T. Steiner
e-mail: tsteiner@lsi.upc.edu

J. Gabarró Vallés
e-mail: gabarro@lsi.upc.edu

J. De Roo
Agfa Healthcare, Moutstraat 100, 9000 Gent, Belgium
e-mail: jos.deroo@agfa.com

discovery and interaction. This approach keeps the complexity to a minimum, yet still enables service descriptions with full semantic expressiveness. A sample implementation on the topic of multimedia Web services verifies the effectiveness of our approach.

Keywords Semantic Web · Service description · Service discovery

1 Introduction

The immense diversity of various multimedia analysis and processing algorithms makes it difficult to integrate them in an automated platform to perform compound tasks. Yet, recent research has indicated the importance of the fusion of different techniques [2]. It is impossible to make different algorithms interoperate if there are no agreements or guidelines on how communication should happen. A coordinating platform can only select algorithms based on their capabilities in presence of a formal description detailing their preconditions and postconditions.

In this paper, we show how to lift multimedia algorithms to the level of Semantic Web services¹ with a formal description mechanism that follows a pragmatic approach. Rather than reinventing the existing methodologies, which focus on *technical* process details, we want to express an algorithm's *functionality* in a way that captures its functionality without requiring lengthy specifications. Our intention is to use existing standards such as the HTTP protocol, Link headers, and URI templates and apply common best practices for implementing multimedia algorithms as true Semantic Web services. The aim is a versatile description and communication model, enabling fully automated service discovery and execution, even under changing conditions. The sole starting point is a Web address of a server, required additional information is gathered at runtime.

Can a client just follow its nose—like humans do—and access the right service by reasoning? We will explain our approach by three real-world multimedia use cases, each of which represents challenges that are currently not fully addressed by alternative techniques. The aim of this work is to provide a simple, flexible, and dynamic solution to semantically describe multimedia services and the associated communication models, thereby enabling their implementation as Semantic Web services.

The remainder of this paper is structured as follows: Section 2 gives an overview on related work. Section 3 describes our RESTdesc approach for Semantic Web service description. Section 4 details how our approach can be used for automatic service discovery. Section 5 shows how our approach is able to adapt to change and react dynamically on errors. The paper terminates with Section 6, which provides a conclusion and gives an outlook on future work.

We have implemented a sample multimedia Web service with mock data that follows our description approach. It is available at our RESTdesc testing web-

¹We use the terms “API” and “(Web) service” synonymously throughout this paper.

site.² Additional RESTdesc information can also be found on the official website <http://restdesc.org/>.

2 Related work

2.1 Web Service Description Language

The description of Web services has a long history. The XML-based Web Service Description Language (WSDL, [9, 10]) provided one of the first models. WSDL focuses on the communicational aspect of Web services, looking from a message-oriented point of view. The details of the message format are written down in a very verbose way and concretized to actual bindings such as SOAP [20] or plain HTTP [14]. Finally, the description can contain endpoints, which implement the specified bindings.

For our use case, we spot two major problems with the use of WSDL. First, WSDL only provides the mechanisms to characterize the technical implementation of Web services. It does not provide the means to capture the functionality of a service. For example, a service that counts the number of words in a text will be described by WSDL as an interface, which accepts a string and outputs an integer. Clearly, an infinite number of algorithms share those input and output properties, so this information is insufficient to infer any meaning or functionality.

Secondly, in practice, a WSDL description is used to generate module source code automatically, which is then compiled into a larger program. If the description changes, the program no longer works, even if such a change leaves the functionality intact. A concrete example of such brittleness is the switch from 32 to 64 bit integer identifiers that occurred at some point in Google's AdWords API, a small change in the service's WSDL file that required the complete recompilation of the relevant pieces of source code [35]. This indicates that WSDL is not well adapted to real-world circumstantial changes. Therefore, WSDL cannot offer automatic service discovery at runtime and why we should investigate other possibilities.

2.2 Semantic annotations for WSDL

The W3C Recommendation named Semantic Annotations for WSDL and XML Schema (SAWSDL, [26]) describes a way how to add semantic annotations to various parts of a WSDL document such as interfaces and operations, and input and output message structures. In addition to that, Web services can be assigned a category with the objective of making them discoverable in a central registry of Web services. SAWSDL also defines an annotation mechanism for specifying the data mapping of XML Schema [15, 36] types to and from ontologies, often referred to as *up-* and *down-lifting*.

While the standard fulfills parts of our requirements, it inherits all the disadvantages from WSDL, specifically its brittleness and verbosity. Although SAWSDL provides semantic descriptions that can be used at runtime, similarly in intent to our

²RESTdesc testing website: <http://restdesc.no.de/>.

aim, we deliberately chose to start from a different perspective. This allows us to provide an alternative for the legacy structures in SAWSDL.

2.3 REST services

A REST service or RESTful service is built on the following principles [13]:

- Servers and clients are separated from each other by a uniform interface. Both servers and clients have well-defined responsibilities, also referred to as *separation of concerns*. This is to guarantee maximum independence from the one and the other.
- All client requests are *stateless*, this means that each request from a client has all the information that the server needs to process it.
- Responses must define themselves as *cacheable* or not using standard HTTP caching techniques.
- When layered systems (like load-balancing) are used, this fact must not be exposed to the API user.

In RESTful APIs, resources are identified by URIs [17]. A resource is to be differentiated from its representation. For example, a set of RDF triples (the resource) might be represented in different serializations (syntaxes), such as RDF/XML or Turtle. The manipulation of any of the representations should carry sufficient information to manipulate the original resource. All messages need to be self-descriptive, for example, the media type of a message needs to make clear what can be done with this message. Each representation needs to communicate relevant related resources, or next steps the client can take at each state.

2.4 Web Application Description Language

The Web Application Description Language (WADL, [21]) is another Web service description format, also XML-based, which does not degrade HTTP to a tunneling mechanism for SOAP, but advocates proper use of all the aspects of the HTTP protocol. While WSDL 2.0 is also capable of specifying bindings to RESTful endpoints, it still requires the abstractions that enable bindings to SOAP and others. WADL, on the other hand, was tailored to the needs of RESTful services, but only exists as a W3C Member Submission and will most likely never reach the W3C Recommendation status of WSDL 2.0 [29].

In addition to that, WADL still suffers from the same problem: it does emphasize the technical properties of the underlying service and does not leave any room for the semantics of the task it performs. This also means that there is no way to automatically discover services based on the desired functionality. Therefore, there is no reason why WADL would be used any differently than WSDL, as also argued by Joe Gregorio in [19].

The main criticism by the REST community, however, is that WADL does document beforehand what, according to the REST principles [13], should be discovered dynamically at run-time. One of the fundamental properties of REST is the so-called *hypermedia constraint*, which basically can be summarized as the constraint that each server response should contain the possible next steps the client can take, since the

application state is not stored on the client. It should be noted that WADL could be used in this way at run-time, yet most current usage continues to be beforehand.

2.5 Semantic markup for Web services

OWL-S [31] is an OWL [34] ontology for describing Semantic Web services in RDF [23]. A service description consists of three parts: a profile, a model and a grounding. Some aspects of profile and model are very similar, in the sense that they both describe input, output, preconditions and effects. The difference is that the profile is used for high-level discovery, while the model is used for more detailed condition matching. Finally, the grounding part specifies the implementation of the service, for instance to WSDL, but other groundings are possible (e.g., in SPARQL [41]).

This marks the first time that there is a focus on the functionality of a service (profile/model), separate from how the interaction (grounding) happens. However, there is no way to enforce the consistency of profile, model, and grounding of a single service. OWL-S input and output types provide more or less the equivalent of what a WSDL message format contains, albeit with RDF types, so there is only a minimal added semantic value on that level. The real possibilities lie in the use of preconditions and postconditions (the latter under the form of result effects), which allow to express complex relationships between input and output values, finally capturing the semantics and functionality of the service.

Unfortunately, these conditions are not expressed within the RDF document that carries the OWL-S description. Instead, the existing inside string literals inside that document, effectively forming another document context. The semantics that connect the OWL-S RDF document and the expression literals are not inherent to neither RDF nor the expression language. Furthermore, a variety of languages to create these expressions are possible: KIF [16], DRS [32], and SWRL [22] are mentioned in the submission, but other languages can be supported through extensions e.g., N3Logic, [7]). While this is a clear benefit for description authors, description interpreters are now faced with a broad spectrum of languages they should a) support and b) be able to integrate with the initial OWL-S RDF document. We believe this is one of the main reasons why the conditions mechanism of OWL-S is seldom used, leaving the interpreter with a parameter-only description.

Furthermore, while OWL-S offers functional descriptions capable of automatic discovery of the capabilities of a single service, it does not provide mechanisms to express its relation to other services. Also, descriptions contain redundancies and require a fair amount of vocabulary, even to express conceptually simple services, and rely on external groundings for technical implementations.

2.6 Linked Open Services

The obligation to make explicit the relation between input and output is present within the Linked Open Service (LOS, [28]) principles. LOS does this by expressing preconditions and postconditions with SPARQL [37] query graph patterns, because RDF currently cannot express quantification, as we also argue in Section 3.4. The drawback of this approach is that these patterns also have to be contained inside string literals, like the OWL-S expression languages. This similarly results

in the expression of the conditions residing in a different document level from the remainder of the service description.

2.7 Resource Linking Language

The Resource Linking Language (ReLL, [1]) aims to provide a natural mapping from RESTful services to RDF. The authors recognize the issues regarding RESTful service descriptions in general and provide an excellent discussion thereof. ReLL differs from our approach in that it only offers “*static description of RESTful services that does not cover [...] new resources or identification and access schemes*” [1], whereas we specifically aim to address these cases in the context of automated service discovery and consumption.

2.8 Universal Description, Discovery, and Integration

The XML-based OASIS standard Universal Description, Discovery, and Integration (UDDI, [4]) was developed to enable the definition of a set of services supporting the discovery and description of (i) businesses, organizations, and other Web service providers, (ii) the Web services that those institutions offer, and finally (iii) the technical interfaces, which may be used to access those services. UDDI was based on a common set of industry standards at that time, including HTTP, XML, XML Schema, and SOAP. The standard was designed to allow for the description and discovery of both public services and non-public in-house services. It was meant to be used as a service broker where parties interested in a special service could go to and retrieve a list of service providers offering the desired service (for example, shipping address verification). Such services would be described in the so-called Green Pages, including not only technical details, but also contact details of the Web service provider.

While for various reasons out of scope of this paper UDDI could not gain the adoption its creators had hoped for, the overall idea of automatically being able to select a service from a (not necessarily central) registry of services still seems useful to us. We will show in Section 4 how we imagine this idea to work decentralized and dynamically using our approach.

3 RESTdesc semantic description

3.1 Motivation

The answer to the question whether Semantic Web service description and Web service discovery are necessary needs to be split up in two parts.

On the one hand there is the question whether Web service description is needed. In RESTful systems, the common opinion is that each message should be self-descriptive enough so that user agents can make sense of each message, given a documented media type that the message is serialized in. On a pure technical layer this works well. For example, let us imagine a very simple image search engine that

simply returns the most adequate image of media type `image/gif` as the result to a query, similar to Google's "*I'm feeling lucky*" functionality. This gives the user agent enough information to process the response with its Graphics Interchange Format (GIF) library, however, a priori it is not clear that the image stands in a relation to a search query that the user agent has used as an input. Therefore OpenSearch [11] defines a description format, which can be used to describe a search engine so that it can be used by search client applications [24, 38]. While we could perfectly use OpenSearch to describe this search API, even slight variations of the API semantics render its use impossible. For instance, let us imagine a Web font preview API where you give the name of a Web font as an input, and get a GIF image with a preview of the text "*The quick brown fox jumps over the lazy dog*" in that very Web font as an output. There is currently no universal way to describe the exact functionality of such API, and yet it might be desirable for a Web font vendor to announce its availability.

The second question is whether automatic discovery of Web services is needed. The first approach for automatic service discovery was UDDI, outlined in Section 2.8. It was driven by the vision that central service registries would serve as so-called Green Pages for parties interested in a specific service. The problem with this approach, however, is that companies do not work this way. There is always a human being involved in the process [27]. We see the potential of service discovery in the generation and run-time supervision of automatic execution plans as outlined in [42], a task that can highly profit from discoverable service descriptions.

3.2 Multimedia example

To make the explanation more concrete, we introduce two related multimedia services, one for face detection, and the other for face recognition. A user agent can upload a photo to the face detection service and use it to check for the existence of faces in the uploaded image. If faces are found, the user agent can use the face recognition service to try to find out more details on the persons whose faces are contained in the image. Each image is considered a resource, for example represented by a binary image file (like `/photos/1`). Each face is a resource, for example represented by an RDF document serialized in Turtle, or a cropped version of the entire image showing only the particular face (like `/photos/1/faces/1`). Each person is a resource (like `/photos/1/persons/1`), for example represented by a string of the person's name. Some of the potential next steps after detecting faces could be, to follow a link to a Web service that allows for recognizing these faces, or starting from the first person on an image, to follow a link to the next person on the image. We will use these two sample Web services, namely a face detection and a face recognition Web service, throughout the paper.

3.3 Introducing RESTdesc

By now, it is clear that we aim to provide a semantic method to express the functionality of a service—as well as its communication—in a concise way that appeals to humans and can be processed automatically. The word "semantic" obviously hints at

```

@prefix http: <http://www.w3.org/2006/http#>.

{
  Preconditions about a certain resource...
}
=> ... imply ...
{
  ... that a certain request exists:
  _:request http:methodName [...];
            http:requestURI [...];
            http:resp [...].
  This request then effectuates postconditions on the resource.
}.

```

Listing 1 The RESTdesc description skeleton

the Semantic Web [6] and its core language RDF [23], upon which our solution will be based.

Listing 1 shows the general skeleton for RESTdesc descriptions. The expression language is Notation3 (N3, [5]), which is based on RDF. The justification for this choice is explained in Section 3.4. Here, we want to focus on the essential elements of the description format:

1. the *preconditions*, which indicate the state a certain resource should have before being able to take part in the interaction;
2. the *postconditions*, which describe the new state for that resource (or related resource);
3. the *request details*, which explain exactly what HTTP request should be made to achieve the postconditions.

These elements are brought together in the form of a rule, which takes care of correct quantification and variable instantiation. Since it is difficult to reason about such an abstract skeleton, we will provide the derivation of the RESTdesc description method with a concrete example.

3.4 Deriving a functional description

We will now formulate the logic basis of RESTdesc, by applying it to the aforementioned example. Let us first revise what we actually want to express. In the example, an informal expression for photo retrieval could be:

I can retrieve a photo by going to /photos/ and appending its identifier. (1)

An intuitive formalization of the above would be:

$$\begin{aligned}
 &hasUri(request, \{"/photos/", id\}) \wedge photoId(photo, id) \wedge \\
 &hasResponse(request, resp) \wedge represents(resp, photo)
 \end{aligned}
 \tag{2}$$

This is straightforward to represent in RDF:

$$\begin{aligned}
 &:request :uri ("/photos/" :id); \\
 &:response [:represents [:photoId :id]].
 \end{aligned}
 \tag{3}$$

Upon closer inspection, it is clear that the formalization (2)—and thus its RDF counterpart (3)—does not contain all the semantics of the informal expression (1). While (1) implies (2), the opposite implication (2) \Rightarrow (1) is broken, and thus the equivalence does not hold. Indeed, fragment (3) states that there exists *one specific* request which returns the photograph with the identifier specified in its URI. It does however not convey the intention of (1) that *all* requests with this URI structure behave the same way. This is a problem of existential \exists versus universal \forall quantification, which has important consequences that should be dealt with formally.

Revising (4) with quantifiers gives:

$$\begin{aligned} \forall photo : \exists id, request, uri, resp : \\ hasUri(request, \{"/photos/", id\}) \wedge photoId(photo, id) \wedge \\ hasResponse(request, resp) \wedge represents(resp, photo) \end{aligned} \quad (4)$$

However, this still remains insufficient, because the universal quantification introduces the claim that *every* photograph in the world possesses an identifier—a false statement for the majority of photographs, with the exception of those uploaded to the server. Similarly, requests exist for such photographs only. Looking back at the informal expression (1), we now spot the (again, implicit) assumption that the photograph we want to retrieve has a known identifier.

Therefore, our last revision of the formal expression takes into account this notion as follows:

$$\begin{aligned} \forall photo, id : photoId(photo, id) \Rightarrow \exists request, uri, resp : \\ hasUri(request, \{"/photos/", id\}) \\ \wedge hasResponse(request, resp) \wedge represents(resp, photo) \end{aligned} \quad (5)$$

The above expression now corresponds to the intended meaning of (1): that a representation of every photograph with an identifier can be retrieved by following the constructed URI. Now the issue of expressing (5) in RDF remains. The original RDF specification [23] does not include a form of quantifiers. Although some attempts have been made in the past (e.g., [33]), the most successful initiative is the W3C submission Notation3 (N3, [5]), which also includes syntactical support for implications as an added benefit.

Expressing (5) in Notation3 gives:

```
@forall :photo, :id.
@forSome :request.
{ :photo :photoId :id. }
:implies
{ :request :uri ("/photos/" :id);
  :response [ :represents :photo ]. }. \quad (6)
```

Note the automatic existential quantification of blank nodes. By turning the request also in a blank node and using the full expressive power of Notation3, we can conveniently write (6) as:

$$\begin{aligned} & \{?photo :photoId ?id.\} \\ => \\ & \{_:request :uri ("/photos/" ?id); \\ & \quad :response [:represents ?photo].\}. \end{aligned} \quad (7)$$

This minimal syntax fully reflects the functionality of the service as intended by the original equation (1). The uniqueness of this approach lies in the fact that the logical underpinnings or Notation3 were so far only used in pure reasoning contexts, where the accent is on the execution of the rule. Here, we use the descriptive part of the rule paradigm to introduce service descriptions, while their executional semantics provide automated composition possibilities.

3.5 RESTdesc description format

With the syntax and required concepts in mind, we now look at existing recommendations, proposals, and vocabularies that we can integrate to obtain an interchangeable description format.

Since RESTful services are centered around the correct use of the HTTP protocol, one of the obvious elements we need is a way to describe HTTP requests. The *HTTP Vocabulary in RDF* [25] has already registered widespread use and, in addition to that, it has the status of a W3C Working Draft. It defines all the necessary concepts to rigorously describe HTTP messages, their structure, and their relationships.

The resource-oriented nature of RESTful services implies the use of descriptive URIs, based on a structure specific to each server. Therefore, we should be able to express the relationship between a resource and its URI. In an Internet-Draft, the IETF describes the concept of URI templates [18] to refer to a category of resources.

Below is an example of a URI template for a person in a photograph:

```
http://example.org/photos/{photoId}/persons/{personId}
```

The identifiers between the curly braces are variables, which can be assigned a value. For example, to get the person with identifier 3 on photograph 241, the URI gets expanded to:

```
http://example.org/photos/241/persons/3
```

While URI templates are still in draft status, many implementations and applications exist. In consequence we decided that we should include them in our design.

Finally, we need a way to tie the URI templates to HTTP request parameters such as the request URI. Also, some additional template semantics are required, for instance to describe what the response body contains. Since such a vocabulary was not available yet, we created the *HTTP template ontology*.³

³Located at <http://purl.org/restdesc/uri-template>.

Listing 2 RESTdesc

description of photo retrieval

```

@prefix : <http://restdesc.no.de/ontology#>.
@prefix http: <http://www.w3.org/2006/http#>.
@prefix tmpl: <http://purl.org/restdesc/http-template#>.

{
  ?photo :photoId ?id.
}
=>
{
  _:request http:methodName "GET";
            tmpl:requestURI ("/photos/" ?photoId);
            http:resp [ tmpl:represents ?photo ].
}.

```

Listing 3 Listing 2 with explicit quantifiers

```

@prefix log: <http://www.w3.org/2000/10/swap/log#>.

@forall :photo, :id.          #  $\forall photo, id :$ 
@forSome :request, :response.
{
  :photo :photoId :id.      #  $photoId(photo, id)$ 
}
log:implies                  #  $\Rightarrow \exists request, r : resp(request, r)$ 
{
  :request [...] http:resp :response.
  :response tmpl:represents :photo.
}.

```

Listing 2 shows the final description of the photo retrieval service. On a high level, we see the precondition, followed by the request and the postcondition. Concepts detailing precise semantics of the service are expressed in a *server-specific vocabulary*⁴ (in this case, photo identifiers) or by reusing publicly available vocabularies (here, for people and depictions). The precondition thus states that an object with a photo identifier is required. In the postcondition, we use the HTTP vocabulary to describe a GET request and its associated response. Finally, we use the HTTP template ontology to specify the URI template, and the contents of the response.

Contrary to its appearance, this short description conveys a vast amount of semantic information. Of course, most importantly, there is the explicit relation expressing precisely how the input relates to the output. An alternative way to look at the implication is to state that the specified request only exists in presence of a photograph. The semantics of the quantification have been highlighted in Listing 3, which contains the same description with the explicit quantifier syntax (prefixes from this and further listings omitted for clarity). The incorporation of the URI template is also particularly strong: the variables in the URI have been bound to the actual values that will be present during execution. Interesting here is that these variables, due to the server-specific ontology, do not only have an *associated data type*, but *fully linked semantics*. For instance, if the server describes the `photoId` predicate by specifying its range as integers and its domain as photographs, this information is propagated into the URI template. Also note that we do not need an ontology for services: the description is complete by the expression of its functionality.

⁴It is not obligatory to detail the server-specific vocabulary in an ontology. Its consistent use across different descriptions may suffice for interpretation and composition.

Listings 4–6 show example descriptions of other services on the same server. For photo upload (Listing 4), we see the prerequisite is to have an image. Note that the service description author is free to use any vocabulary, in this case the *FOAF* ontology [8]. Since the request URI is fixed, no URI template was used. The response, in contrast, will have a `Location` header with a URI containing the photo identifier. For the request, we specify the format of the `POST` body. Note how the precondition of photo retrieval (Listing 2) naturally follows from the postcondition of photo upload, hinting at a possible causality.

This effect is also visible in Listings 5 and 6, which both demonstrate the ease of expressing complex conditions. The required expressions involve a complicated

```
@prefix foaf: <http://xmlns.com/foaf/>.

{
  ?photo a foaf:Image.
}
=>
{
  _:request http:methodName "POST";
            http:requestURI "/photos";
            http:body [ tmpl:formData ("photo=" ?photo) ];
            http:resp [ tmpl:location ("/photos/" ?photoId) ].

  ?photo :photoId ?photoId.
}.
```

Listing 4 RESTdesc description of photo upload

Listing 5 RESTdesc description of face detection

```
{
  ?photo :photoId ?photoId.
}
=>
{
  _:request http:methodName "GET";
            tmpl:requestURI ("/photos/" ?photoId "/faces");
            http:resp [ tmpl:representsMultiple ?region ].

  ?region foaf:depicts [ a foaf:Person ];
            :regionId _:regionId;
            :belongsTo ?photo.
}.
```

```
{
  _:region foaf:depicts ?person;
            :regionId ?regionId;
            :belongsTo [ :photoId ?photoId].
}
=>
{
  _:request http:methodName "GET";
            tmpl:requestURI ("/photos/" ?photoId "/people/" ?regionId);
            http:resp [ tmpl:represents ?person ].

  ?person foaf:name _:personName.
}.
```

Listing 6 RESTdesc description of face recognition

indirection (e.g., “the photograph contains a region that depicts a person”), yet they can be understood quite easily, while the formal semantics are sound.

When we overlook all of the above, it becomes apparent that RESTdesc descriptions are a simple and elegant way of describing Web services in an integrated semantic manner. They capture the functional aspects formally without resorting to complex artifices. The use of the HTTP vocabulary and semantic identifiers was taken from previous work [40], as well as the use of Notation3 conditions [41], both which were extended and combined into a single method. The resulting RESTdesc descriptions can be used for automatic discovery, service composition [39, 42], and execution. These and other aspects will be demonstrated in Section 4.2.

3.6 Automated interpretation and composition

An interesting fact about Notation3 implications is that, besides the *descriptive/declarative* semantics we have used so far, they also entail *operational* semantics. This means that, given a reasoner that is able to make *modus ponens* inferences, the following action takes place:

$$\frac{P \Rightarrow Q, P}{Q} \quad (8)$$

This is a very relevant property for RESTdesc descriptions, which *enables context-based service discovery*.

For example, we might want to know what we can do on a server given the situation where we have an image. RESTdesc makes this a trivial task. The triple (9) below expresses our current condition:

$$\langle \text{http://example.org/photo.jpg} \rangle \text{ a foaf:Image.} \quad (9)$$

It is also the precondition of photo upload (Listing 4). Consequently, using *modus ponens* (8), we can derive the postcondition of photo upload. Yet it does not stop there. The statements of the postcondition can also trigger other inferences. In the end, the result chain is:

- we can upload the photo, upon which it will receive an identifier;
- we can use this identifier to receive the photo;
- we can use this identifier to detect faces within it;
- we can then ask the server to recognize these faces.

In addition to *what* steps we can take, the inference process also tells us *how* to take this steps by listing the concrete HTTP requests.

These and other examples can be verified online using the EYE Semantic Web reasoner [12] on the interactive RESTdesc test website².

An even more interesting approach is to add a goal, in addition to a starting point (9). If we indeed want to know who is depicted in the photograph, our query might be:

$$\langle \text{http://example.org/photo.jpg} \rangle \text{ foaf:depicts ?person.}$$

The proof of the reasoner for this query forms a list of ordered steps to obtain the desired results, again with detailed instructions on how to execute these steps. This differs from the previous output, which was just an unordered list of possible actions.

Here, the result is an actual execution plan, instructing to first upload the photo, then ask for detected faces, and finally find out the associated persons [42].

It is apparent that RESTdesc descriptions provide a powerful and instant way to deal with automated interpretation and composition of different Web services.

3.7 Compatibility and automatic translation

One outstanding property of RESTdesc is that it, capturing the complete functionality of the service, can be converted into a multitude of other formats. After all, the functionality description, combined with invocation information, contains the maximal superset of what needs to be known to machines about a service.

Is it important to realize that RESTdesc starts from RESTful principles from which the invocation aspects follow. For instance, the fact that GET and HEAD operations are safe and idempotent, whereas POST is not, conveys necessary information to generate certain descriptions.

One relevant example is the translation to the Composition of Identifier Names vocabulary (CoIN, [30]). The CoIN vocabulary “*defines a set of classes and properties used to describe what properties of a resource constitute components of a URI*”. This information is present in the RESTdesc service descriptions, due to its use of URI templates which are bound to concrete variables. Automatic translation can be seen in action on our website².

It is also possible to generate WSDL, WADL, or OWL-S descriptions, as all the information—besides human-targeted elements such as labels and textual information—are readily available. In fact, human-centric textual descriptions could also be generated based on RESTdesc and ontological information. Listing 2 could translate to “given an identifier, retrieves a representation of the photograph with that identifier.”

By means of content negotiation, a description format understood by the client can be returned upon request.

4 RESTdesc service discovery

In this section we will show with the help of the concrete example of a face detection and recognition API introduced before how starting from a single URI one can follow one’s nose to explore the capabilities of an API.

4.1 On Web Service discovery

Web Service discovery can be seen as the process of locating a suitable Web service for a given task. Typically in classic WS-* architectures there are several options for this process, which can involve Web services registering themselves with a central registry (as in Section 2.8 with UDDI), or Web services exposing their capabilities using the Web Service Description Language (WSDL, [9, 10]). On a related note, the Web Services Inspection Language (WSIL, [3]) has been proposed to list groups of Web services and their endpoints. When we say Web service discovery, we currently limit ourselves to enable discovery of Web services by following one’s nose from a given start URI by resolving links and making sense of Notation3 service invocation

Listing 7 An OPTIONS call on an API's base URI

```
$ curl -i -H "Accept: text/n3; charset=utf-8" \
-X OPTIONS http://restdesc.no.de
HTTP/1.1 200 OK
Link: <./>; rel=self,
      <./photos>; rel=index; type=text/n3; charset=utf-8
Allow: GET, OPTIONS, HEAD
```

descriptions. However, RESTdesc is a very powerful concept, as given just one starting URI, the full reasoning chain of available Web services is enabled, not constrained to Web services on the same domain.

4.2 Learning about one's OPTIONS

An execution plan can be dynamically created by a user agent that is given a concrete task like “identify all persons in a certain photo” (see Section 3.6). The user agent first starts to check out its options on the Web service's base URI /, as shown in Listing 7. From there, the user agent discovers that there is a link of type “index” to /photos, which in turn it checks its options on. As can be seen in Listing 8 (edited slightly for clearness), the user agent has indicated that it accepts responses of type “text/n3; charset=utf-8”, and therefore is given instructions that new items can be added to the index by means of a POST request. Next, the user agent uploads the photo to the server, and is notified that the photo has been stored at the location /photos/1, as stated in the Location header of the 201 Created-type response. The user agent then executes an OPTIONS call to that location in order to find out what it can do with the uploaded photo. The response contains Notation3 instructions on how to detect potentially contained faces in the photo by performing a GET request to /photos/1/faces. Upon execution of that GET call, each of the detected faces has its own URI, for example /photos/1/faces/1, where a cropped image region of just that face is available, as can be derived from the Link header. In addition to that, a different Link header reveals that the persons behind each of the faces can be recognized by navigating to the particular person's URI, for example /photos/1/persons/1.

5 Adapting to change and errors

In this section, we describe how our approach reacts to change and errors in a forgiving and tolerant way. We investigate how RESTdesc descriptions ensure clients can adapt to long-term changes and possible errors.

5.1 Focus on runtime decisions

RESTdesc is designed from the start to be consumed at runtime and to make decisions only at the moment this becomes necessary. We want to mimic the flexibility of human beings browsing the Web, who follow hyperlinks to achieve a predefined goal—which is perhaps adjusted along the way. Mostly, humans have a high-level plan, that is refined as each step becomes more and more concrete, and if necessary, steps can be taken back.

Listing 8 An OPTIONS call on a discovered index path from an API's base URI while accepting text/n3 responses

```
$ curl -i -H "Accept: text/n3; charset=utf-8" \
-X OPTIONS http://restdesc.no.de/photos
HTTP/1.1 200 OK
Link: <./>;
      rel=index;
      rel=self
Allow: GET, OPTIONS, HEAD, POST
Content-Type: text/n3; charset=utf-8

@prefix : <http://restdesc.no.de/ontology#>.
@prefix http: <http://www.w3.org/2006/http#>.
@prefix tmpl: <http://purl.org/restdesc/http-template#>.
@prefix foaf: <http://xmlns.com/foaf/>.

{
  ?photo a foaf:Image.
} => {
  _:request http:methodName "POST";
  http:requestURI "/photos";
  http:body [ tmpl:formData ("photo=" ?photo) ];
  http:resp [ tmpl:location ("/photos/" ?photoId) ].
  ?photo :photoId _:photoId.
}

{
  ?photo :photoId ?id.
} => {
  _:request http:methodName "GET";
  tmpl:requestURI ("/photos/" ?photoId);
  http:resp [ tmpl:represents ?photo ].
}.
```

Listing 9 A POST call in order to upload a photo discovered via the Notation3 data from the previous request

```
$ curl -i -F "photo=@./obama-gillard.jpg" \
http://restdesc.no.de/photos
HTTP/1.1 100 Continue

HTTP/1.1 201 Created
Location: http://restdesc.no.de/photos/1
Content-Type: text/html

Your photo was uploaded:
<a href="http://restdesc.no.de/photos/1">
  http://restdesc.no.de/photos/1
</a>
```

5.2 Fluent change coping

This focus on the runtime aspect makes RESTdesc well adapted to changes. The key to that functionality is offered by the operational semantics of the integrated pre- and postconditions: in order for a RESTdesc description to apply, its preconditions must be satisfied. This is inherently different from static descriptions, where the description can be interpreted separately. This adaptive behavior does not only work for small interface changes, even more complex situations can be handled gracefully.

We will briefly consider some examples. For instance, suppose the server changes its URI structure (which is similar to the change of data format presented in Section 2.1). This does not pose a problem, since the URI templating mechanism fills out the parameters dynamically. A more subtle change, for instance, if the server only wants to accept images with maximum dimensions $1,000 \times 1,000$, can be handled on

Listing 10 An OPTIONS call to find out about one's options with a concrete photo

```
$ curl -i -H "Accept: */*" -X OPTIONS \\  
      http://restdesc.no.de/photos/1  
HTTP/1.1 200 OK  
Content-Type: image/jpeg  
Allow: GET, OPTIONS, HEAD  
Link: <http://restdesc.no.de/restdesc/photos/1/faces>;  
      rel="http://dbpedia.org/resource/Face_detection";  
      title="contained faces";  
      type="text/n3"  
  
{  
  ?photo :photoId ?photoId.  
} => {  
  _:request http:methodName "GET";  
  tmpl:requestURI ("/photos/" ?photoId "/faces");  
  http:resp [ tmpl:representsMultiple ?region ].  
  
  ?region foaf:depicts [ a foaf:Person ];  
  :regionId _:regionId;  
  :belongsTo ?photo.  
}.  
  
{  
  _:region foaf:depicts ?person;  
  :regionId ?regionId;  
  :belongsTo [ :photoId ?photoId ].  
}  
=>  
{  
  _:request http:methodName "GET";  
  tmpl:requestURI  
    ("/photos/" ?photoId "/people/" ?regionId);  
  http:resp [ tmpl:represents ?person ].  
  ?person foaf:name _:personName.  
}.  
}
```

two levels. The preconditions will state this requirement on the image, and should the client attempt a larger image, the server will return an error code. More interestingly, the server can also return hyperlinks to image resizing services, which can help the client to work out a solution on its own. Even changes that affect the process structure can be handled transparently: for example, if the face recognition algorithm needs grayscale input images, the preconditions can list this requirement and the server could return service links in a similar way.

The central idea is that the client uses descriptions in a dynamic way: “*Given a certain input, how can the service descriptions reach my predefined goal?*”. The server furthermore aims to support the client by providing information on how to reach subsequent steps. This vision differs completely from the traditional static approach, which cannot deal with changing contexts.

5.3 Adaptive error handling

WSDL and OWL-S provided very detailed ways to specify error conditions and faults. This does not correspond to the human strategy when browsing the Web: we just try, and if something does not work out as expected, we continue, possibly aided by hyperlinks on last visited pages. The underlying rationale is simple: if we had to anticipate *every* possible error (page not found, irrelevant information, network failure, ...), we might as well give up before we start. Consequently, our approach

is to handle errors dynamically as they arise,⁵ guided by the service itself where applicable.

An important benefit of this pragmatic error handling is that all causes can be dealt with in an uniform manner. Clients assume services will handle their request as described. If an exception or error should occur, it is detected and remedied, irrespective of whether it could have been expected. For example, a face detection request can fail for numerous reasons: the image does not exist or has been deleted, no faces detected in the image, the server is unavailable or crashes... The central idea is that there is no point in anticipating foreseeable errors, since errors can always occur. A RESTdesc description details necessary preconditions for executing a request, but it does not strive to handle exceptional situations because it can never cover all of them.

The REST practice of correctly using HTTP status codes forms the corner stone of error detection. They can precisely identify the source of the problem (client, request, or server), its temporal scope (temporary or permanent), and offer additional information (even in case of success). What we suggest is that the service should supply hyperlinks that can help the client to remedy the problem, similar to service discovery in Section 4. For example, depending on the error, the server could list the photo upload API (image does not exist), or an alternative API with a different face detection algorithm (no faces detected), or even another server (server unavailable) in its responses.

6 Conclusion and future work

In this paper, we have shown a proposal for a Web service description and interaction approach for automatic Web service discovery and execution called RESTdesc. Our approach builds on top of RESTful principles and consists of a semantic markup model, offering a formal description of a service's functionality, with extensive flexibility, and an HTTP-based discovery method of services, both within a domain of related services, and also beyond. It is to be noted that in order for our approach to work, obeying to REST principles is essential for the APIs that RESTdesc should be applied to. We have demonstrated the feasibility and the pragmatism of our proposal with a concrete implementation. In addition to that, and unlike OWL-S, our approach is integrated in the normal Web service data flow.

Future work will be to prove the applicability of the approach to a broad family of existing RESTful Web services. We are also planning to investigate ways to link to external services that not necessarily follow our approach, including multi-domain-spanning Web services. In addition to that, we want to perform an in-depth study of compatibility and exchangeability with other standards and practices (namely with WSDL, WADL, and OWL-S). Currently we are at the very beginnings of our

⁵This approach works well with actions that do not involve commitment, i.e. information exchange, which we primarily focus on. Data fetching and even state-changing actions, such as image upload, are thus perfectly possible. In case of irreversible actions with binding consequences (e.g., reservations), special care might be necessary. But then again, not all errors can be foreseen, which renders this topic inherently difficult and not generally solvable.

work towards allowing for complex automated execution plan creation including the creation of automated clients against RESTdesc-described services. With this paper we have laid a humble foundation stone for semantic Web service description. Future versions of RESTdesc will encourage the decoupled use of the method, meaning that instead of relying on URI templates (which allow for a certain degree of freedom, but still introduce a form of tight coupling) we shift the URI descriptions into the Link headers, and only specify the relation of those Link headers to the result in the server response.

Acknowledgements The authors wish to thank Niklas Lindström for his help with the CoIN vocabulary and Mark Nottingham for his suggestion on HTTP header line folding.

The research activities as described in this paper were funded by Ghent University, the Interdisciplinary Institute for Broadband Technology (IBBT), the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT), the Fund for Scientific Research Flanders (FWO Flanders), and the European Union.

This work was partially supported by the European Commission under Grant No. 248296 FP7 I-SEARCH project. Joaquim Gabarró is partially supported by TIN-2007-66523 (FORMALISM), and SGR 2009-2015 (ALBCOM).

References

- Alarcón R, Wilde E (2010) Linking data from RESTful services. In: Proceedings of the 3rd international workshop on linked data on the Web
- Atrey P, Hossain M, El Saddik A, Kankanhalli MS (2010) Multimodal fusion for multimedia analysis: a survey. In: Multimedia systems. <http://www.springerlink.com/index/E31M71152774R630.pdf>
- Ballinger K, Brittenham P, Malhotra A, Nagy WA, Pharies S (2001) Web services inspection language (WS-Inspection) 1.0. <http://www-106.ibm.com/developerworks/webservices/library/ws-wsinspect.html>
- Bellwood T, Capell S, Clement L, Colgrave J, Dovey MJ, Feygin D, Hately A, Kochman R, Macias P, Novotny M, Paolucci M, von Riegen C, Rogers T, Sycara K, Wenzel P, Wu Z (2004) UDDI version 3.0.2. <http://www.oasis-open.org/committees/uddi-spec/doc/spec/v3/uddi-v3.0.2-20041019.htm>
- Berners-Lee T, Connolly D (2011) Notation3 (N3): a readable RDF syntax. W3C Team Submission. <http://www.w3.org/TeamSubmission/n3/>
- Berners-Lee T, Hendler J, Lassila O (2001) The Semantic Web. *Sci Am* 284(5):34
- Berners-Lee T, Connolly D, Kagal L, Scharf Y, Hendler J (2008) N3Logic: a logical framework for the World Wide Web. *Theory and Practice of Logic Programming* 8(3):249–269
- Brickley D, Miller L (2010) FOAF vocabulary specification 0.97. Namespace document. <http://xmlns.com/foaf/spec/20100101.html>
- Chinnici R, Moreau J-J, Ryman A, Weerawarana S (2007) Web Services Description Language (WSDL) version 2.0 part 1: core language. W3C Recommendation. <http://xml.coverpages.org/wSDL20000929.html>
- Christensen E, Curbera F, Meredith G, Weerawarana S (2000) Web Services Description Language (WSDL) 1.0. <http://xml.coverpages.org/wSDL20000929.html>
- Clinton D (2007) OpenSearch specification 1.1 draft 3. <http://www.opensearch.org/Specifications/OpenSearch/1.1>
- De Roo J (2012) Euler proof mechanism. <http://eulerssharp.sourceforge.net/>
- Fielding RT, Taylor RN (2002) Principled design of the modern Web architecture. *ACM Trans on Int Tech* 2(2):115–150. doi:10.1145/514183.514185
- Fielding R, Gettys J, Mogul J, Frystyk H, Masinter L, Leach P, Berners-Lee T (1999) Hypertext transfer protocol—HTTP/1.1. Request for comments: 2616. <http://www.ietf.org/rfc/rfc2616.txt>
- Gao S, Sperberg-McQueen CM, Thompson HS, Mendelsohn N, Beech D, Maloney M (2008) W3C XML Schema Definition Language (XSD) 1.1 part 1: structures. World Wide Web consortium, working draft WD-xmlschema11-1-20080620.


16. Generereth MR (1998) Knowledge interchange format. Draft Proposed American National Standard. <http://logic.stanford.edu/kif/dpans.html>
17. Gonzalez JL, Marcelnez R (2011) Phoenix: fault-tolerant distributed Web storage based on URLs. *JoC* 2(1):79–86
18. Gregorio J, Fielding R, Hadley M, Nottingham M (2010) URI template. <http://tools.ietf.org/html/draft-gregorio-uritemplate-04>
19. Gregorio J (2007) Do we need WADL?. <http://bitworking.org/news/193/Do-we-need-WADL>
20. Gudgin M, Hadley M, Mendelsohn N, Moreau J-J (2007) SOAP version 1.2 part 1: messaging framework, 2nd edn. W3C recommendation. <http://www.w3.org/TR/2007/REC-soap12-part1-20070427/>
21. Hadley M (2009) Web Application Description Language. W3C Member Submission. <http://www.w3.org/Submission/wadl/>
22. Horrocks I, Patel-Schneider PF, Boley H, Tabet S (2004) Swrl: a Semantic Web rule language combining OWL and RuleML. W3C Member Submission. <http://www.w3.org/Submission/SWRL/>
23. Klyne G, Carrol JJ (2004) Resource Description Framework (RDF): concepts and abstract syntax. W3C Recommendation. <http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/>
24. Klyuev V, Oleshchuk V (2011) Semantic retrieval: an approach to representing, searching and summarising text documents. *IJITCC* 1(2):221–234
25. Koch J, Velasco CA (2009) HTTP vocabulary in RDF 1.0. W3C working draft. <http://www.w3.org/TR/HTTP-in-RDF10/>
26. Kopecký J, Vitvar T, Bournez C, Farrell J (2007) SAWSDL: semantic annotations for WSDL and XML Schema. *IEEE Internet Comput* 11:60–67
27. Krill P (2005) Microsoft, IBM, SAP discontinue UDDI registry effort. *InfoWorld*. <http://www.infoworld.com/d/architecture/microsoft-ibm-sap-discontinue-uddi-registry-effort-777>
28. Kruppenacher R, Norton B, Marte A (2010) Towards linked open services and processes. In: *Proceedings of FIS'2010*, pp 68–77
29. Lafon Y (2009) Team comment on the “Web Application Description Language” Submission. <http://www.w3.org/Submission/2009/03/Comment>
30. Lindström N (2011). The CoIN vocabulary. <http://court.googlecode.com/hg/resources/docs/coin/spec.html>
31. Martin D, Burstein M, Hobbs J, Lassila O (2004) OWL-S: semantic markup for Web services. W3C Member Submission. <http://www.w3.org/Submission/OWL-S/>
32. McDermott D (2004) DRS: a set of conventions for representing logical languages in RDF. <http://cs-www.cs.yale.edu/homes/dvm/daml/DRSGuide.pdf>
33. McDermott D, Dou D (2002) Representing disjunction and quantifiers in RDF. In: *Proceedings of international Semantic Web conference 2002*, pp 250–263
34. McGuinness DL, van Harmelen F (2004) OWL Web ontology language overview. W3C Recommendation. <http://www.w3.org/TR/2004/REC-owl-features-20040210/>
35. Minar N (2006) Why SOAP sucks. <http://www.somebits.com/weblog/tech/bad/whySoapSucks.html>
36. Peterson D, Gao S, Malhotra A, Sperberg-McQueen CM, Thompson HS (2008) W3C XML Schema Definition Language (XSD) 1.1 part 2: datatypes. In: *World Wide Web Consortium, working draft WD-xmlschema11-2-20080620*.
37. Prud'hommeaux E, Seaborne A (2008) SPARQL query language for RDF. W3C Recommendation. <http://www.w3.org/TR/rdf-sparql-query/>
38. Pyshkin E, Kuznetsov A (2011) Approaches for Web search user interfaces: how to improve the search quality for various types of information. *JoC* 1(1):1–8
39. Sathappan OL, Chitra P, Venkatesh P, Prabhu M (2011) Modified genetic algorithm for multi-objective task scheduling on heterogeneous computing system. *IJITCC* 1(2):146–158
40. Steiner T, Algermissen J (2011) Fulfilling the hypermedia constraint via HTTP OPTIONS, the HTTP vocabulary In RDF, and link headers. In: *Proceedings of the second international workshop on RESTful design*
41. Verborgh R, Van Deursen D, De Roo J, Mannens E, Van de Walle R (2010) SPARQL endpoints as front-end for multimedia processing algorithms. In: *Proceedings of the fourth international workshop on service matchmaking and resource retrieval in the Semantic Web*
42. Verborgh R, Van Deursen D, Mannens E, Poppe C, Van de Walle R (2011) Enabling context-aware multimedia annotation by a novel generic semantic problem-solving platform. *Multimedia Tools and Applications special issue on Multimedia and Semantic Technologies for Future Computing Environments*. doi:10.1007/s11042-010-0709-6



Ruben Verborgh received his Master degree in Computer Science Engineering from Ghent University, Belgium in 2010. He is a researcher with the Multimedia Lab - Ghent University, Department of Electronics and Information Systems (ELIS) and the Interdisciplinary institute for BroadBand Technology (IBBT). His interests include Semantic Web technologies, multimedia annotation, artificial intelligence, and their relation to multimedia processing. Currently, he is working on moving multimedia algorithms to the Web and associated problems, such as semantic service descriptions.



Thomas Steiner is a proud dad-of-two, a Research Scientist at Google, and a PhD student at UPC.edu. His main research interests these days are the Semantic Web, Linked Data, and the architectural style REST. He holds two Master of Computer Science degrees, one from the Technical University of Karlsruhe, Germany, and the other from the École Nationale Supérieure d'Informatique et de Mathématiques Appliquées de Grenoble, France. In addition to that he works on making the Internet a better place, tweets as @tomayac and blogs at blog.tomayac.com.



Davy Van Deursen received his M.Sc. and Ph.D. degrees in computer science from Ghent University, Belgium in 2005 and 2009, respectively. He is currently a post-doctoral researcher at Ghent University - IBBT (Ghent, Belgium), where he is working for the Multimedia Lab research group (<http://multimedialab.elis.ugent.be>). Since 2008, he is also actively participating within the W3C Media Fragments Working Group, whose mission is to specify a URI scheme for the addressing of media fragments. His research interests and areas of publication include video coding, media content adaptation and delivery, media content presentation, Semantic Web technologies, metadata modeling, standardization, and multichannel publishing.



Jos De Roo graduated in 1979 as Master of Science and did a doctoral study on image computing from 1979 till 1982. He was co-founder of ICOS (Image Computer Systems) and worked there until 1985. Since then he works at Agfa. He is researcher at Agfa Healthcare and doing advanced development of components to realize semantic web oriented architectures for healthcare imaging applications, healthcare enterprise integration and medical decision support. He is also involved in Euler proof mechanism open source development and in W3C Semantic Web Activity Working Groups.



Rik Van de Walle received his M.Sc. and PhD degrees in Engineering from Ghent University, Belgium in 1994 and 1998 respectively. After a visiting scholarship at the University of Arizona (Tucson, USA), he returned to Ghent University, where he became professor of multimedia systems and applications, and head of the Multimedia Lab. His current research interests include multimedia content delivery, presentation and archiving, coding and description of multimedia data, content adaptation, and interactive (mobile) multimedia applications.



Joaquim Gabarró Vallés was born in Barcelona (Spain) in 1953 and studied Physics and Computer Science, obtained his Ph.D. from Universitat Politècnica de Calatunya in 1983. From 1985 is Associate Professor (Profesor Titular) in Universitat Politècnica de Calatunya. Currently represents Spain in the TC1 (Foundations of Computer Sciences) of IFIP. Main research topics are Concurrency, Complexity and Web algorithms.