

Process-aware web programming with Jolie

Fabrizio Montesi

Department of Mathematics and Computer Science, University of Southern Denmark, Campusvej 55, 5230 Odense M, Denmark. Phone number: +4565507171

Abstract

We extend the Jolie programming language to capture the native modelling of process-aware web information systems, i.e., web information systems based upon the execution of business processes. Our main contribution is to offer a unifying approach for the programming of distributed architectures on the web, which can capture web servers, stateful process execution, and the composition of services via mediation. We discuss applications of this approach through a series of examples that cover, e.g., static content serving, multiparty sessions, and the evolution of web systems. Finally, we present a performance evaluation that includes a comparison of Jolie-based web systems to other frameworks and a measurement of its scalability.

Keywords: Business Processes, Programming Languages, Sessions, Web Services

1 Introduction

A Process-Aware Information System (PAIS) is an information system based upon the execution of business processes. These systems are required in many application scenarios, from inter-process communication to automated business integration [24]. Processes are typically expressed as structures that determine the order in which communications should be performed in a system. These structures can be complex and of different kinds; a systematic account can be found at [11]. For this reason, many formal methods [63, 29, 40, 22], tools [65, 35, 27, 33, 49], and standards [51, 66, 1] have been developed to provide languages for the definition, verification, and execution of processes. In these works, compositionality plays a key role to make the development of processes manageable. For example, in approaches based on process calculi, complex process structures are obtained by composing simpler ones through the usage of standard composition operators such as sequence, choice, and parallel (see, e.g., [41]). Other approaches follow similar ideas using graphical formal models, e.g., Petri Nets [64, 56].

In the last two decades, web applications have become increasingly process-aware. Web processes – i.e., processes inside of a web information system – are usually implemented server-side on top of *sessions*, which track incoming messages related to

Email address: `fmontesi@imada.sdu.dk` (Fabrizio Montesi)

the same conversation. Sessions are supported by a shared memory space that lives through different client invocations. Differently from the aforementioned approaches for designing processes, the major languages and platforms for developing web applications (e.g., PHP, Ruby on Rails, and Java EE) do not support the explicit programming of process structures. As a workaround, programmers have to simulate processes using bookkeeping variables in the shared memory space of sessions. For example, consider a process in a Research Information Service (RIS) where a user has to authenticate through a `login` operation before accessing another operation, say `addPub`, for registering a publication. This would be implemented by defining the `login` and the `addPub` operations separately. The code for `login` would update a bookkeeping variable in the session state and the implementation for `addPub` would check that variable when it is invoked by the user. Although this approach is widely adopted, it is also error-prone: since processes can assume quite complex structures, simulating them through bookkeeping variables soon becomes cumbersome. Consequently, the produced code may be poorly readable and hard to maintain.

The limitations described above can be avoided by adopting a multi-layered architecture. For example, it is possible to stratify an application by employing: a web server technology (e.g., Apache Tomcat) for serving content to web browsers; a web scripting language (e.g., PHP) for programmable request processing; a process-oriented language (e.g., WS-BPEL [51]) for modelling the application processes; and, finally, mediation technologies such as proxies and ESB [23] for integrating the web application within larger systems. Such an architecture would offer a good separation of concerns. However, the resulting system would be highly heterogeneous, requiring a specific know-how for handling each part. Thus, it would be hard to maintain and potentially prone to breakage in case of modifications.

The aim of this paper is to simplify the programming of process-aware web information systems. We build our results on top of Jolie, a general-purpose service-oriented programming language that can handle both the structured modelling of processes and their integration within larger distributed systems [49, 37]. Jolie is briefly introduced in § 2.

1.1 Contributions

Our main contribution is the extension of Jolie to obtain a unifying technology for the programming of processes, web technologies (web servers and scripting) and mediation services (e.g., proxies), which facilitates the development of heterogeneous information systems that involve the Web. We then investigate the applicability of this framework, showing that our extended version of Jolie captures the different components of web systems and their integration using a homogeneous set of concepts. We proceed as described below.

Web processes. We extend the Jolie interpreter to support the HTTP protocol, enabling processes written in Jolie to send and receive HTTP messages (§ 3). The integration is *seamless*, meaning that the processes defined in Jolie remain abstract from the underlying data formats used in the Web: data structures in Jolie are transparently transformed to HTTP message payloads and vice versa (§ 3.1). These transformations can be configured using *port parameters*, an extension of the Jolie language that we develop to

allow processes to map information from HTTP headers to application data and vice versa (§ 3.3). Parameters support mobility: they can be transparently transmitted from service registries to clients, allowing clients to be transparently configured at runtime with the right binding information (§ 3.4, Example 3.3).

Web servers as processes. We develop a web server, called Leonardo (§ 4), using our approach. The web server is given as a simple process that (i) receives the name of the resource a client wants to access, then (ii) reads the content of such resource, and (iii) sends the content back to the client. Leonardo is an example of the fact that, in our framework, a web server is not a separate technology but it is instead a simple case of process. We also show how to extend Leonardo to handle simple CRUD operations over HTTP.

Sessions. We combine our HTTP extension for Jolie with message correlation, a mechanism used in service-oriented technologies to route incoming messages to their respective processes running inside of a service [51, 49]. We first show that this combination is adequate wrt existing practice: it enables Jolie processes to use the standard methodology of tracking client-server web sessions using unique session identifiers (§ 5.1). Then, we generalise such methodology to program multiparty sessions, i.e., structured conversations among a process and multiple external participants [32] (§ 5.2).

Architectural programming. We present how to obtain separation of concerns in a web architecture implemented with our approach, by combining HTTP with aggregation, a Jolie primitive for programming the structure of service networks [47, 57, 49] (§ 6). We demonstrate the usefulness of this combination by implementing a multi-layered system that integrates different components. We also discuss how to deal with the evolution of software architectures obtained with our approach (§ 6.2).

RESTful services. We discuss how to develop web systems based on the REST style [26] with our framework, using URI templates to bridge resource-oriented interactions to processes (§ 7). The standard separation of concerns between routing and business logic can be achieved by developing a routing service that routes REST requests to other services, which required developing a reflection library for the Jolie language (§ 7.3). We then show how structured processes can be combined with our REST router to obtain RESTful multiparty sessions (§ 7.4). Since, in these scenarios, requests typically have to be validated both at client- and server-side, we provide an integration between Jolie and JavaScript to be able to run the same validation code, achieving the same de-duplication benefits as in frameworks based on JavaScript (§ 7.5).

Performance. A performance evaluation of our approach is given (§ 8). This evaluation covers two main aspects. First, our data shows that our solution has comparable performance to that of other existing frameworks in the execution of basic tasks, such as offering static files or templated web pages. Second, we analyse the scalability of our approach wrt the number of services involved in the computation of responses to clients.

2 Overview of Jolie

Jolie [45] is a general-purpose service-oriented programming language, released as an open-source project [37] and formally specified as a process calculus [29, 48]. In this section, we briefly describe some aspects of Jolie that are relevant for our discussion. We refer the interested reader to [49] and [37] for a more comprehensive presentation of the language. Readers who are already familiar with the Jolie language may skip this section and resume reading from § 3.

2.1 Jolie programs

Every Jolie program defines a service and consists of two parts: *behaviour* and *deployment*. A behaviour defines the implementation of the operations offered by the service; it consists of communication and computation instructions, composed into a structured process (a workflow) using constructs such as sequences, parallels, and internal/external choices. Behaviours rely on *communication ports* to perform communications, which are to be defined in the deployment part. The latter can also make use of architectural primitives for handling the structure of an information system. Formally, a Jolie program is structured as:

$$D \text{ main } \{ B \}$$

Above, D represents the deployment and B the behavior of the program.

2.2 Behaviour

We report (a selection of) the syntax of behaviours in Figure 1. A behaviour B can use primitives for performing communications, computation, and their composition in processes. We briefly comment the syntax. Terms (*input*), (*output*), and (*input choice*) implement communications. An input η can either be a one-way or a request-response, following the WSDL standard [9]. Statement (*one-way*) receives a message for operation \circ and stores its content in variable x . Term (*request-response*) receives a message for operation \circ in variable x , executes behaviour B (called the *body* of the request-response input), and then sends the value of the evaluation of expression e to the invoker. Dual to input statements, an output $\bar{\eta}$ can be a (*notification*) or a (*solicit-response*). A (*notification*) sends a message to OP containing the value of the evaluation of expression e . Term (*solicit-response*) sends a message to OP containing the evaluation of e and then waits for a response from the invoked service, storing it afterwards in variable y . In both (*notification*) and (*solicit-response*), OP is the name of an *output port*, which acts as a reference to an external service. Output ports are concretely defined in the deployment part of a program; we will present them in § 2.3.

Term (*input choice*) implements an input-guarded choice; it is similar to the `pick` primitive in WS-BPEL [51]. Specifically, the construct waits for a message for any of the inputs in η_1, \dots, η_n . When a message for one of these inputs is received, say for η_i where $1 \leq i \leq n$, then the statement is executed as follows, in order: (i) all the other branches in the choice (i.e., all the $[\eta_j] \{ B_j \}$ such that $j \neq i$) are discarded; (ii) η_i is executed; and, finally, B_i is executed.

Terms (*cond*) and (*while*) implement, respectively, the standard conditional and iteration constructs. Term (*seq*) models sequential execution and reads as: execute B ,

$B ::= \eta$		(input)
$\bar{\eta}$		(output)
$[\eta_1] \{B_1\} \dots [\eta_n] \{B_n\}$		(input choice)
if (e) B_1 else B_2		(cond)
while (e) B		(while)
$B ; B'$		(seq)
$B \mid B'$		(par)
throw (f)		(throw)
$x = e$		(assign)
$x \rightarrow y$		(alias)
nullProcess		(inact)

$\eta ::= o(x)$		(one-way)
$o(x) (e) \{B\}$		(request-response)

$\bar{\eta} ::= o@OP (e)$		(notification)
$o@OP (e) (y)$		(solicit-response)

Figure 1: Jolie, syntax of behaviours (selection).

wait for its termination, and then run B' . In term (*par*), instead, B and B' are run in parallel. Term (*throw*) throws a fault signal f , interrupting execution. If a fault signal is thrown from inside a request-response body, the invoker of the request-response statement is automatically notified of the fault [44]. We omit the syntax for handling faults, which is not necessary for reading this paper.

Term (*assign*) stores the result of the evaluation of expression e in variable x . Term (*alias*) makes variable x an alias for variable y , i.e., after its execution accessing x will be equivalent to accessing y . Term (*inact*) denotes the empty behaviour (no-op).

Example 2.1 (Structured data). *Jolie natively supports the manipulation of structured data. In Jolie's memory model the program state is a tree (possibly with arrays as nodes, see [47]), and every variable, say x , can be a path to a node in the memory tree. Paths are constructed through the dot "." operator; for instance, the following sequence of assignments*

```
1 person.name = "John"; person.age = 42
```

would lead to a state containing a tree with root label `person`. For the reader familiar with XML, a corresponding XML representation would be:

```
1 <person> <name>John</name> <age>42</age> </person>
```

2.3 Deployment

We introduce now (a selection of) the syntax of deployments. A deployment includes definitions of *input ports*, denoted by IP , and *output ports*, denoted by OP ,

```

IP ::= inputPort Port    OP ::= outputPort Port
Port ::= id {
    Location: Loc
    Protocol: Proto
    Interfaces: iface1, ..., ifacen
}

```

Figure 2: Jolie, syntax of ports (selection).

which respectively support input and output communications with other services. Input and output ports are one the dual concept of the other, and their respective syntaxes are quite similar. Both kinds of ports are based on the three basic elements of *location*, *protocol* and *interface*. Their syntax is reported in Figure 2. In the syntax of ports, i.e., term *Port*, *Loc* is a URI (Uniform Resource Identifier) that defines the location of the port; *Proto* is an identifier referring to the data protocol to use in the port, which specifies how input or output messages through the port should be respectively decoded or encoded; the identifiers *iface*₁, ..., *iface*_n are references to the interfaces accessible through the port.

Jolie supports several kinds of locations and protocols. For instance, a valid *Loc* for accepting TCP/IP connections on TCP port 8000 would be "**socket://localhost:8000**". Other supported locations are based, respectively, on Unix sockets, Bluetooth communication channels, and local in-memory channels (channels implemented using shared memory). Some supported instances of *Proto* are *sodep* [37] (a binary protocol, optimised for performance), *soap* [8], and *xmlrpc* [12].

The interfaces referred to by a communication port define the operations that can be accessed through that port. Each interface defines a set of operations, along with their respective (i) operation types, defining if an operation is to be used as a one-way or a request-response, and (ii) types of carried messages. For example, the following code

```

1 interface SumIface { RequestResponse: sum(SumT) (int) }

```

defines an interface named *SumIface* with a request-response operation, called *sum*, that expects input messages of type *SumT* and replies with messages of type **int** (integers). Data types for messages follow a tree-like structure; for example, we could define *SumT* as follows:

```

1 type SumT:void { .x:int .y:int }

```

We read the code above as: a message of type *SumT* is a tree with an empty root node (**void**) and two subnodes, *x* and *y*, that have both type **int**.

Example 2.2 (A complete Jolie program). *We give an example of how to combine behaviour and deployment definitions, by showing a simple service defined in Jolie. The code follows:*

```

1 type SumT:void { .x:int .y:int }
2
3 interface SumIface { RequestResponse: sum(SumT) (int) }
4
5 inputPort SumInput {
6 Location: "socket://localhost:8000"
7 Protocol: soap
8 Interfaces: SumIface
9 }
10
11 main
12 {
13     sum( req )( resp ) {
14         resp = req.x + req.y
15     }
16 }

```

Above, input port `MyInput` deploys the interface `SumIface` (and thus the `sum` operation) on TCP port 8000, waiting for TCP/IP socket connections by invokers using the `soap` protocol. The behaviour of the service is contained in the `main` procedure, the entry point of execution in Jolie. The behaviour in `main` defines a request-response input on operation `sum`. In this paper, we implicitly assume that all services are deployed with the **concurrent** execution modality for supporting multiple session executions, from [47]. This means that whenever the first input of the behavioural definition of a service receives a message from the network, Jolie will spawn a dedicated process with a local memory state to execute the rest of the behaviour. This process will be equipped with a local variable state and will proceed in parallel to all the others. Therefore, in our example, whenever our service receives a request for operation `sum` it will spawn a new parallel process instance. The latter will enter into the body of `sum`, assign to variable `resp` the result of adding the subnodes `x` and `y` of the request message `req`, and finally send back this result to the original invoker. \square

3 Extending Jolie to HTTP

We extend Jolie to support web applications by introducing a new protocol for communication ports, named `http`, and by extending the language of deployments to support configuration parameters for protocols. The protocol follows the specifications of HTTP, and integrates the message semantics of Jolie to that of HTTP and its different content encodings. In this section, we discuss the main aspects of our implementation.

3.1 Message transformation

The central issue to address for integrating Jolie with the HTTP protocol is establishing how to transform HTTP messages in messages for the input and output primitives of Jolie and vice versa. Our objective is twofold: on the one hand, we aim at having transparent transformations between data payloads inside of HTTP messages (e.g., XML documents or JSON structures) and Jolie values, so that the programmer

does not have to deal with them; on the other hand, we also need to give Jolie programs enough low-level control on HTTP messages such that implementing standard components found in web systems is possible (e.g., web servers or REST routers, as we will discuss respectively in § 4 and § 7). Hereby we discuss primarily how our implementation manages request messages; response messages are similarly handled. The (abstract) structure of a *request message* in HTTP is:

Method Resource **HTTP**/*Version Headers Body*

Above, *Method* specifies the action that the client intends to perform and can be picked by a static set of keywords, such as GET, PUT, POST, etc. Term *Resource* is a URI path telling which resource the client is requesting. Term *Version* is the HTTP protocol version of the message. The term *Headers* may include descriptive information on the message *Body*, e.g., the type of its content (Content-Type), or parameters that influence the behaviour of the receiver, e.g., the wish to keep the underlying connection open for future requests (Connection:keep-alive). Finally, *Body* contains the content (payload) of the HTTP message.

A Jolie message consists of an operation name (the operation the message is meant for) and a structured value (the content of the message) [47]. Hence, we need to establish where to read or write these elements in an HTTP message. For operation names, we interpret the path part of the *Resource* URI as the operation name. The *Method* of an HTTP message, instead, is read and written by Jolie programs through a configuration parameter of our extension, described later in § 3.3. The value of a Jolie message is obtained from *Body* and the rest of the *Resource* URI (query and fragment parts). We use the latter to decode querystring parameters as Jolie values.

The content of an HTTP message may be encoded in one of different formats. Our `http` extension handles querystrings, form encodings (simple and multipart), XML, JSON [4], and GWT-RPC¹ [3]. Programmers can use the `format` parameter (§ 3.3) to control the data format for encoding and decoding messages. Most of the times, however, this decision is performed automatically via standard HTTP content negotiation and the programmer does not need to know which format is used (`format` is a fallback parameter in the case that the client does not ask for a content type). As an example of message translation, the HTTP message:

1 **GET** /sum?x=2&y=3 **HTTP**/1.1

would be interpreted as a Jolie message for operation `sum`. The querystring `x=2&y=3` would be translated to a structured value with subnodes `x` and `y`, containing respectively the strings `"2"` and `"3"`.

3.2 Automatic type casting

Querystrings and other common message formats used in web applications, such as HTML form encodings, do not carry type information. Instead, they carry only string

¹We also developed a companion GWT-RPC client library, called `jolie-gwt`, for a more convenient access to web services written in Jolie by integrating with the standard GWT development tools.

representations of values; the information on the types that these values may have had in the code of the sender (e.g., in Javascript) is therefore lost. However, type information is necessary for supporting services such as the sum service in Example 2.2, which specifically requires its input values to be integers. To handle such cases, we introduce the mechanism of *automatic type casting*. Automatic type casting reads incoming messages that do not carry type information (such as querystrings or HTML forms) and tries to cast their content values to the types expected by the service interface for the message operation. As an example, consider the querystring `x=2&y=3` that we discussed before. Since its HTTP message is a request for operation `sum`, the automatic type casting mechanism would retrieve the typing for the operation and see that nodes `x` and `y` should have type `int`. Therefore, it would try to re-interpret the strings `"2"` and `"3"` as integers before giving the message to the behaviour of the Jolie program. There are cases that type casting may fail to handle; for example, in `x=hello` the string `hello` cannot be cast to an integer for `x`. In such cases, our `http` protocol will send a `TypeMismatch` fault back to the invoker with HTTP status code 400 (Bad Request).

3.3 Configuration Parameters

We augment the deployment syntax of Jolie to support *configuration parameters* for our `http` protocol. Specifically, these can be accessed through (*assign*) and (*alias*) statements that can be written inside a code block immediately after declaring the `http` protocol in a port. For instance, consider the following input port definition:

```
1 inputPort MyInput {
2   /* ... */
3   Protocol: http {
4     .default = "d"; .debug = true;
5     .method -> m
6   }
7 }
```

The code above would set the `default` parameter to `"d"`, set the `debug` parameter to `true`, and bind the `method` parameter to the value of variable `m` in the current Jolie process instance.

We briefly describe some configuration parameters. All of them can be modified at runtime using the standard Jolie constructs for dynamic port binding, from [47], which we omit here. Parameter `default` allows to mark an operation as a special fallback operation for receiving messages that cannot be handled by any other operation defined in the interface of the enclosing input port. Parameter `cookies` allows to store and retrieve data from browser cookies, by mapping cookie values in HTTP messages to subnodes in Jolie messages. Parameter `method` allows to read/write the *Method* field of HTTP messages. Parameter `statusCode` gives read/write access to HTTP status codes (default status codes apply depending on the request method, e.g., the successful response to a GET request containing data has status code 200). Parameter `format` can be used as a default setting for the data format to use as previously discussed (e.g., `json`, `xml`). The parameter `alias` allows to map values inside a Jolie

message to resource paths in the HTTP message, to support interactions with REST services. Parameter `redirect` gives access to the **Location** header in HTTP, allowing to redirect clients to other locations. The parameter `cacheControl` allows to send directives to the client on how the responses sent to it should be cached. Finally, parameter `debug` allows to print the HTTP messages sent and received through the network on screen.

3.4 Examples

We report some examples about how our `http` protocol implementation integrates with some standard mechanisms of web technologies.

Example 3.1 (Access from web browsers). *Let us consider a modification of the `sum` service from Example 2.2, where we change the input port to use the `http` protocol that we developed:*

```
1 type SumT: void { .x:int .y:int }
2
3 interface SumIface { RequestResponse: sum(SumT) (int) }
4
5 inputPort SumInput {
6   Location: "socket://localhost:8000"
7   Protocol: http
8   Interfaces: SumIface
9 }
10
11 main
12 {
13   sum( req ) ( resp ) {
14     resp = req.x + req.y
15   }
16 }
```

Now, our implementation of `http` allows us to access the service above in multiple ways. The most obvious is to write a Jolie client using an output port with the `http` protocol. A more interesting way is to use a web browser. For example, we can use the service by passing parameters through a querystring; navigating to the following URL is valid:

`http://localhost:8000/sum?x=2&y=3`

Accessing the URL above would show the following content on the browser:

```
1 <sumResponse>5</sumResponse>
```

The standard format used for responses is XML, as above. Responses from Jolie services using `http` can of course also be themed using, e.g., HTML and Javascript (we refer to the online documentation for more information about this aspect [37]).

Another possibility is to use HTML forms, such as the one that follows:

```
1 <form action="sum" method="GET">
2   <input type="text" name="x"/>
3   <input type="text" name="y"/>
4   <input type="submit"/>
5 </form>
```

The content displayed as a response in the web browser would be the same XML document as before.

We also offer support for AJAX programming. The following Javascript snippet calls the `sum` operation using jQuery [61]: first, it reads the values for `x` and `y` from two text fields (respectively identified in the DOM by the names `x` and `y`); then, it sends their values to the Jolie service by encoding them as a JSON structure; and, finally, it displays the response from the server in the DOM element with id `result`:

```
1 $.ajax(
2   'sum', { x: $('#x').val(), y: $('#y').val() },
3   function( response ) { $('#result').html( response ); }
4 );
```

Our implementation of the `http` protocol for Jolie auto-detects the format of messages sent by clients, so the `sum` service does not need to distinguish among all the different access methods shown above: they are all handled using the same Jolie code.

Example 3.2 (Accessing REST services). We exemplify how to access REST services, where resources are identified by URLs, using our configuration parameters. In this example we invoke the DBLP server, which provides bibliographic information on computer science articles [62]. We use DBLP to retrieve the BibTeX entry of an article, given the `dblp` key of the latter (i.e., the identifier of such article in `dblp`). The code follows:

```
1 include "console.iol"
2
3 type FetchBib:void { .dblpKey:string }
4
5 interface DBLPiface {
6   RequestResponse: fetchBib( FetchBib )( string )
7 }
8
9 outputPort DBLP {
10  Location: "socket://dblp.uni-trier.de:80/"
11  Protocol: http {
12    .osc.fetchBib.alias = "rec/bib2/!{dblpKey}.bib";
13    .format = "html" }
14  Interfaces: DBLPiface
15 }
16
17 main
```

```

18 {
19   r.dblpKey = args[0];
20   fetchBib@DBLP( r )( bibtex );
21   println@Console( bibtex )()
22 }

```

In the example above, we start by importing the `Console` service from the Jolie standard library. We then declare an output port towards the DBLP server. The interesting part here is the usage of parameter `osc.fetchBib.alias`, which passes to our implementation a configuration for parameter `alias` for operation `fetchBib` (`osc` stands for operation-specific configuration and is used for configuration parameters that make sense when associated to an operation). The value of the alias for operation `fetchBib` specifies how to map calls for that operation to resource paths that the DBLP server understands. The interface offered by DBLP for retrieving bibtex entries is REST-based, with paths rooted at “`rec/bib2/`”. As an example, assume that we wanted to retrieve the bibtex entry for the book “The C Programming Language” by Kernighan and Ritchie [38]. Its `dblp` key is “`books/ph/KernighanR78`”; therefore, the bibtex entry can be accessed at the URL:

`http://dblp.uni-trier.de/rec/bib2/books/ph/KernighanR78.bib`

In our implementation, we capture this kind of patterns for REST paths by providing a syntax for replacing parts of paths with the value of a subnode in a request message, based on URI templates. For instance, the term `%!{dblpKey}` in the alias for operation `fetchBib` means that that part of the path will be replaced with value of the sub node `dblpKey` in messages sent for that operation on port `DBLP`. The behaviour of the service is simple: we invoke operation `fetchBib` reading the `dblp` key we want from the first command line argument that Jolie is invoked with; then, we print the received bibtex entry on screen.

An extended version of this example is deployed as a tool at [46].

Example 3.3 (Parameter mobility). Configuration parameters for ports can be dynamically transmitted and used for binding services at runtime. For example, our DBLP client may invoke a service registry to get the correct location and parameters to invoke the DBLP service as follows:

```

1  /* Interface information is as before */
2
3  outputPort DBLP {
4  Interfaces: DBLPInterface
5  }
6
7  outputPort Registry { /* ... */ }
8
9  main
10 {
11   getBinding@Registry( "DBLP" )( DBLP );
12   r.dblpKey = args[0];

```

```
13  fetchBib@DBLP( r )( bibtex );
14  println@Console( bibtex )()
15 }
```

In this example, the DBLP client starts by invoke an external Registry service on operation `getBinding` to dynamically discover how the DBLP service can be contacted. The registry is programmed to return the same data structure that was previously statically defined inside of the client. We refer the reader to [49] for a more detailed description of dynamic binding; here, the point is to show that our extension to configuration parameters does not lead to any additional complexity.

4 Web Servers

In Example 3.1 we have seen how to make operations in a Jolie service accessible by invokers using HTTP (in the example, operation `sum`): any operation in a Jolie service can be exposed to HTTP clients just by changing the protocol of its related input port(s) to `http`. This technique covers scenarios in which the interface that we want to expose over HTTP is statically defined as a finite set of operations, which is the typical situation when using service-oriented technologies such as Jolie or WS-BPEL [51]. However, web servers do not fall into this category. A web server allows clients to access files, e.g., web pages, images, and JavaScript libraries. Since files can be created and deleted during execution, we cannot statically map each single file to an operation name as would be required by our methodology in § 3. In this section, we discuss how to deal with this kind of situations by introducing default operations; these will form the basis for our development of more structured REST-oriented routers in § 7.

Default operations bridge the mutating nature of dynamic resource sets that web servers have to offer (such as parts of a filesystem) to the static operation names used in processes. Specifically, a default operation is a special operation marked as a fallback in case a client sends a request message for an operation that is not statically defined by the service. In this case, the message is wrapped in the following data structure (we omit some subnodes not relevant for this discussion):

```
1  type DefaultOperationHttpRequest: void {
2      .requestUri: string
3      .data: undefined
4  }
```

where `requestUri` is the URI of the resource that has been requested by the client and `data` is the data content of the message.

A default operation is set through the parameter `default` of the `http` protocol, and can either be associated to the `Method` field of incoming HTTP messages or be defined as a “catch-all” operation in case no other more specific operation can be found (in the latter case, the method used in the message can then be retrieved as a variable). For example, the following configuration states that requests for undefined operations with HTTP method `PUT` should be handled by operation `put`, requests for

undefined operations with method GET should be handled by operation `get`, and all other requests for undefined operations should be handled by operation `d`.

```
1 Protocol: http {
2   .default = "d";
3   .default.get = "get";
4   .default.put = "put"
5 }
```

Example 4.1 (Leonardo Web Server). Parameter `default` allows us to easily model a simple web server: whenever we receive a request for the default operation, we try to find a file in the local filesystem that has the same name as the operation originally requested by the client. We have used this mechanism to implement Leonardo [6], a web server implementation written in pure Jolie. For clarity, here we report a simplified version. The entire implementation of Leonardo consists of only about 80 LOCs, showing that our language pushes many of the details of dealing with HTTP to the underlying implementation; many of these details can be accessed through configuration parameters when needed. Leonardo can be downloaded at [6].

```
1 /* ... */
2
3 interface MyInterface {
4   RequestResponse:
5     d( DefaultOperationHttpRequest )( undefined )
6 }
7
8 inputPort HTTPInput {
9   Location: "socket://localhost:80/"
10  Protocol: http { .default = "d" /* ... */ }
11  Interfaces: MyInterface
12 }
13
14 main {
15   d( req )( resp ) {
16     /* ... */
17     readFile@File( req.requestUri )( resp )
18   }
19 }
```

Above, we have set the `default` parameter for the `http` protocol in input port `HTTPInput` to operation `d`. Therefore, when a message for an unhandled operation is received through input port `HTTPInput`, it will be managed by the implementation of operation `d`. The body of the latter invokes operation `readFile` of the `File` service from the Jolie standard library, which reads the file with the same name as the originally requested resource (`req.requestUri`). Finally, the data read from the file (`resp`) is returned back to the client.

Example 4.2 (CRUD Web Servers). We extend *Leonardo* to a simple web server supporting *CRUD* operations (Create, Read, Update, Delete). As usual, we map create and update to *PUT* requests, read to *GET*, and delete to *DELETE*. The code follows:

```
1  /* ... */
2
3  interface MyInterface {
4  RequestResponse:
5      get( DefaultOperationHttpRequest )( undefined )
6      put( DefaultOperationHttpRequest )( void )
7      delete( DefaultOperationHttpRequest )( bool )
8  }
9
10 inputPort HTTPInput {
11 Location: "socket://localhost:80/"
12 Protocol: http {
13     .default.get = "get";
14     .default.put = "put";
15     .default.delete = "delete"
16 }
17 Interfaces: MyInterface
18 }
19
20 main {
21     [ get( req )( resp ) {
22         readFile@File( req.requestUri )( resp )
23     } ] { nullProcess }
24
25     [ put( req )() {
26         f.filename = req.requestUri;
27         f.content -> req.data;
28         writeFile@File( f )( resp )
29     } ] { nullProcess }
30
31     [ delete( req )( resp ) {
32         delete@File( req.requestUri )( resp )
33     } ] { nullProcess }
34 }
```

In the code above, *GET* requests are served by operation `get`, which reads the requested file and replies with its content. Similarly, operation `put` uses the Jolie standard library to write a file with the data sent by the invoker, and operation `delete` deletes a file from the filesystem.

5 Sessions

A main aspect of web-based information systems is the modelling of *sessions*, which allow to relate different incoming messages to the same logical “conversation”. In this section, we present how to program sessions over HTTP with our extension of the Jolie language. A major benefit is that sessions are process-aware: the order in which messages are sent and received over different operations is syntactically explicit, and it is enforced without requiring bookkeeping variables.

5.1 Binary sessions

We start by addressing binary sessions, i.e., sessions with exactly two participants [34]. Consider the scenario mentioned in the Introduction about a Research Information Service (RIS), where the RIS allows users to add a publication to a repository after having successfully logged in. This structure is expressed by the following behaviour:

```
1 login( cred )( r ) { checkCredentials };
2 addPub( pub )
```

Above, `login` is a request-response operation that, when invoked, checks the received credentials by calling the subprocedure `checkCredentials`. If the latter does not throw a fault, the process proceeds by making operation `addPub` available.

Suppose now that, e.g., two users are logged in at the same time in a service with the behaviour above. The service would have then two separate process instances, respectively dedicated to handle the two clients. When a message for operation `addPub` arrives in this situation, how can we know if it is from the first user or the second? We address this kind of issues by using correlation sets, as defined in [48]. A correlation set declares special variables that identify an internal service process from the others. In our example we use the following correlation set declaration:

```
1 cset { userKey: addPub.userKey }
```

Above, we used the **cset** keyword to declare a correlation set consisting of variable `userKey`. We will use `userKey` to distinguish users that have logged in. Variable `userKey` is associated to the subnode `userKey` in incoming messages for operation `addPub`. This means that whenever a message for operation `addPub` is received from the network, Jolie will assign the message to the internal running process with the same value for the correlation variable `userKey`. We can now write a working implementation of the service:

```
1 inputPort RISInput {
2   /* ... */
3   Protocol: http
4 }
5
6 cset { userKey: addPub.userKey }
7
8 define checkCredentials { /* ... */ }
```

```

9  define updateDB { /* ... */ }
10
11  main
12  {
13    login( cred )( r ) {
14      checkCredentials;
15      r.userKey = csets.userKey = new
16    };
17    addPub( pub );
18    updateDB
19  }

```

Our RIS allows the creation of new processes by invoking operation `login`. If the procedure `checkCredentials` does not throw a fault, then the process creates a fresh value for the correlation variable `csets.userKey` using the `new` keyword. The process sends the value of `csets.userKey` back to the client through variable `r`. Then, the process waits for an invocation of operation `addPub` and stores the incoming message in variable `pub`. The correlation set declaration in the program guarantees that only invocations with the same user key as that returned by operation `login` will be given to this process. We finally update the internal database of the RIS using the (unspecified) procedure `updateDB`.

5.1.1 Integrating cookies with correlation sets

Our implementation of the RIS requires clients to write the `userKey` as a subnode in the messages they send to operation `addPub`. Since this may be cumbersome in the case of many operations that require correlation, web applications typically use HTTP cookies to store this kind of information. Our `http` protocol integrates cookies with message correlation through the `cookies` parameter, which allows to map cookies to subnodes in Jolie variables. We change the definition of input port `RISInput` to the following:

```

1  inputPort RISInput {
2  /* ... */
3  Protocol: http { .cookies.userKeyCookie = "userKey" }
4  }

```

The parameter assignment `.cookies.userKeyCookie = "userKey"` instructs our `http` protocol implementation to store the value of the cookie `userKeyCookie` in subnode `userKey` for incoming messages, and vice versa for outgoing messages.

In general, our `http` extension allows developers to abstract from where correlation data is encoded when programming a service behaviour. For instance, an important disadvantage of using cookies to store correlation data is that this breaks the statelessness constraint of REST interactions [26]. Instead of using a cookie, the web user interface may also send the value for a correlation variable in other ways, e.g., embedded a hyperlink, a JSON or XML subnode, or an element in an HTML form encoding. Our extension transparently support these different methods without requiring changes

in the behaviour of a service. We discuss the usage of hyperlinks to keep track of process execution in § 7.

5.2 Multiparty Sessions

As far as binary sessions are concerned, there is not much difference between standard session identifiers as used, e.g., in PHP, and correlation sets, aside from the fact that the generation and sending of correlation variables is explicit programmed in Jolie behaviours. However, correlation sets are more expressive when it comes to providing (i) compound session identifier based on multiple values, as in BPEL [51], and (ii) multiple identifiers for the same process. We are particularly interested in the second aspect, since it allows us to model *multiparty* sessions, i.e., sessions with more than two participants.

Multiparty sessions are useful when considering scenarios with multiple actors that need to be coordinated to reach a common goal. As an example, we extend our RIS implementation to deal with a use case from the Pure software by Elsevier [25]. In Pure, when a user (e.g., a research scientist) adds a publication, a moderator (e.g., the head of the scientist's department) has to be notified of the change. Then, the moderator has to choose whether to approve or reject the newly added publication for confirmation in the database, after reviewing the data inserted by the user. We show the code for this multiparty version of our RIS implementation in the following:

```
1 inputPort RISInput {
2 /* ... */
3 Protocol: http { .cookies.userKeyCookie = "userKey" }
4 }
5
6 outputPort Logger { /* ... */ }
7 outputPort Moderator { /* ... */ }
8
9 cset { userKey: addPub.userKey }
10 cset { modKey: approve.modKey reject.modKey }
11
12 define checkCredentials { /* ... */ }
13 define updateDB { /* ... */ }
14
15 main
16 {
17   login( cred )( r ) {
18     checkCredentials;
19     r.userKey = csets.userKey = new
20   };
21   addPub( pub );
22   noti.bibtex = pub.bibtex;
23   noti.modKey = csets.modKey = new;
24   { log@Logger( pub.bibtex ) | notify@Moderator( noti ) };
25   [ approve() ] {
```

```

26     log@Logger( "Accepted " + pub.bibtex );
27     updateDB
28   }
29   [ reject() ] {
30     log@Logger( "Rejected " + pub.bibtex )
31   }
32 }

```

Above, we have added the output ports `Logger`, an external service that maintains a log of actions that we assume the user can read, and `Moderator`, an external service playing the role of the moderator in our scenario. We have also added a new correlation set for variable `modKey` (moderator key), which we use to track incoming messages from the moderator of a session. The correlation set declares also that the moderator may use `modKey` to invoke operations `approve` and `reject`. In the behaviour, the code is unchanged until after we receive an invocation for operation `addPub`. Now, after we receive a request for operation `addPub`, we prepare a notification `noti` for the moderator containing (i) the descriptor of the publication (we assume that it is given by the user in BibTeX format), and (ii) the moderation key `modKey` (which is instantiated as a fresh value with the keyword `new`). Then we use the parallel construct of Jolie to concurrently send a message to, respectively, the `Logger` on operation `log` (to log the user's request) and the `Moderator` on operation `notify` (to notify the moderator of the user's request). The process now enters into an input choice on operations `approve` and `reject`, which can be invoked only by the moderator; this is because the correlation set declaration of variable `modKey` requires it to be present for invocations of these operations, and we sent the value of `modKey` only to the moderator. If `approve` is invoked, then we log the approval and we update the database of publications. Otherwise, if `reject` is invoked, we log the rejection only.

6 Layering

In the previous sections, we focused separately on how to use our extension of Jolie to program web servers (§ 4) and structured process-aware sessions (§ 5). Typically, a real-world web architecture has to deal with both aspects. In this section, we show how they can be combined in our context by building multi-layered architectures.

6.1 Aggregation

A simple way of designing a service that serves content *and* provides process-aware sessions is to combine the respective operations in the same behaviour as an input choice. Consider the following code:

```

1  /* ... */
2  main
3  {
4    [ get( req )( resp ) { /* ... */ } ] { nullProcess }
5    [ login( cred )( r ) { /* ... */ } ] { /* ... */ }
6  }

```

Above, we assume that the ports and correlation sets are configured by merging the configurations found in Example 4.2 and the RIS implementation in § 5.2. Then, operation `get` would serve HTML and JavaScript files to clients, which could also invoke operation `login` to access the behaviour of the RIS.

While combining the code for a web server with that of sessions with complex structures as done above is simple, in the long term it also leads to code that is hard to maintain due to poor separation of concerns: all concerns are mixed in the same service. Ideally, separate concerns should be addressed by separate services. This methodology, however, raises the question of how services addressing separate concerns can be composed together as a single system that clients can access without knowing the inner complexity of the system. We tackle this issue by integrating our `http` protocol implementation with the notion of service aggregation found in Jolie [47, 49].

Aggregation is a Jolie primitive that allows a service to expose the interfaces of other services on one of its input ports, in addition to its own interfaces. In the remainder, we refer to the service using aggregation as aggregator and to the services it aggregates as aggregated services. The semantics of aggregation is a simple generalisation of the mechanism used in proxy services: when a message from the network reaches an aggregator, the aggregator checks whether the message is for an operation in (i) one of its own interfaces or (ii) the interfaces of an aggregated service. In the first case, the message is given to the behaviour of the aggregator; in the second case, the aggregator forwards the message to the aggregated service providing the operation requested in the message.

Using aggregation in combination with our `http` protocol we can easily build a multi-layered web architecture for our RIS scenario, where services communicate using different protocols as needed. We depict the architecture in Figure 3, where circles represent services, rectangles represent the interfaces exposed by services, full arrows represent dependencies from actors (users or services) to services, and dashed arrows represent aggregations; each arrow is annotated with the protocol used for communications. We comment the architecture. Users can access the web server using a web browser, through the `http` protocol. Requests for files, intended to be for the user interface (e.g., HTML pages), are handled directly by the web server through operation `get`. Instead, invocations of operations `login` and `addPub` are forwarded to the RIS by aggregation. The web server and the RIS communicate using the `sodep` protocol, for performance (`sodep` is a binary protocol). As in § 5.2, the RIS uses an additional service, Moderator, to decide whether publications should be accepted into the system. The RIS and the Moderator services communicate using the `soap` protocol. Below, we exemplify how our architecture can be implemented. We assume that the Moderator service is externally provided, and focus instead on the web server and the RIS.

Web server. The code of the web server follows:

```
1 /* ... */
2
3 outputPort RIS {
4 Location: "socket://www.ris-example.com:8090/"
5 Protocol: sodep
```




Figure 3: Architecture of the RIS scenario.

```

6 Interfaces: RISIface
7 }
8
9 inputPort WebServerInput {
10 Location: "socket://www.webserver-example.com:80/"
11 Protocol: http {
12   .default.get = "get";
13   .cookies.userKeyCookie = "userKey"
14   /* ... */
15 }
16 Interfaces: GetIface
17 Aggregates: RIS
18 }
19
20 main {
21   get( req )( resp ) {
22     /* ... */
23     readFile@File( req.requestUri )( resp )
24   }
25 }

```

Our web server implements only the operation `get`, which serves static files to clients. It also aggregates the RIS by using the aggregation instruction **Aggregates**:RIS in its input port, where RIS is an output port pointing to the RIS. Therefore, all invocations from users for the operations offered by the RIS will be automatically forwarded to the latter. Observe that output port RIS uses the `sodep` protocol: our implementation automatically takes care of translating incoming HTTP messages from users destined to the RIS into binary `sodep` messages. In general, the programmer does not need to worry about data format transformations in our extension of the Jolie language: messages are implicitly converted to/from the HTTP format as needed.

Remark 6.1. Here, we showed the code for the web server modified to aggregate the RIS for clarity purposes. In real-world production environments, the practice of rewriting the web server component every time comes with the unnecessary risk of introducing bugs. Therefore, in such environments, the web server is deployed as an

autonomous service, such that important updates to it from the Leonardo project (or any other web server project based on Jolie) can be immediately applied. In this kind of set-up, the configuration information on which services should be aggregated by Leonardo is kept in a separate configuration file.

RIS (Research Information Service). The code for the RIS is the same as that shown in § 5.2, with the exception that we now use sodep as communication protocol and that we removed the usage of the external service Logger for simplicity:

```
1 inputPort RISInput {
2 Location: "socket://www.ris-example.com:8090/"
3 Protocol: sodep
4 Interfaces: RISIface
5 }
6
7 outputPort Moderator {
8 Location: "socket://www.moderator-example.com:8080/"
9 Protocol: soap
10 Interfaces: ModeratorIface
11 }
12
13 cset { userKey: addPub.userKey }
14 cset { modKey: approve.modKey reject.modKey }
15
16 define checkCredentials { /* ... */ }
17 define updateDB { /* ... */ }
18
19 main
20 {
21   login( cred )( r ) {
22     checkCredentials;
23     r.userKey = csets.userKey = new
24   };
25   addPub( pub );
26   noti.bibtex = pub.bibtex;
27   noti.modKey = csets.modKey = new;
28   notify@Moderator( noti );
29   [ approve() ] {
30     updateDB
31   }
32   [ reject() ] {
33     /* ... */
34   }
35 }
```




Figure 4: Architecture of the RIS scenario with DBLP importer.

6.2 Evolvability

Implementing multi-layered web architectures using our approach, i.e., combining `http` with aggregation, results in systems that are robust wrt future modifications, or their *evolution*. We distinguish between *vertical* and *horizontal modifications*, which respectively represent modifications that influence an existing chain of aggregations or new ones.

A vertical modification is a modification of an interface aggregated by another service. In our example, changing the code of the RIS to add, remove, or change the type of an operation in interface `RISInterface` would be a vertical modification, because `RISInterface` is aggregated by the web server. Vertical modifications do not require any intervention on the rest of the architecture, as aggregation is a parametric mechanism: the web server simply needs to be restarted to read the new definition of interface `RISInterface`.

Horizontal modifications deal with the addition or removal of operations without requiring an intervention on the behaviour of existing services. Assume that, as an example, we wanted to add the possibility to import publications from the DBLP bibliography service to our RIS by offering a new operation called `import`. We could implement this new feature by changing the code of the RIS, both its interface and behaviour. However, in some scenarios this may not be possible, e.g., the RIS may be a black box to which we do not have access (third-party proprietary code), or the RIS cannot be modified due to quality or security regulations. We deal with this kind of situations by developing the new operations we need in a new service, and then by aggregating this service together with the RIS in the web server. The resulting situation in our example scenario is depicted in Figure 4. The only difference between Figure 4 and our previous architecture from Figure 3 is the presence of a new service, called `Importer`, offering operation `import`; the web server now aggregates `Importer` together with the RIS, to make operation `import` accessible by users through web browsers. We report the updated code for the web server and the importer.

Web server. For the web server, we simply need to add an output port towards the importer service and aggregate it in the input port of the server. We report only the code interested by our changes, the rest remains the same as in § 6.1.

```
1 /* ... */
2
3 outputPort Importer {
4 Location: "socket://localhost:8009/"
5 Protocol: sodep
6 Interfaces: ImporterIface
7 }
8
9 outputPort RIS { /* ... */ }
10
11 inputPort WebServerInput {
12 /* ... */
13 Aggregates: RIS, Importer
14 }
15
16 /* ... */
```

By changing the web server as done above, invocations for operation `import` will now be redirected to the importer service.

Observe that, by using aggregation, all the invocations from the web client to the aggregated services pass through the web server. This implies that our programming methodology respects the standard Same Origin Policy by design, allowing the web application run by users to access the aggregated services regardless of where the latter are located.

Importer service. The code for the importer service follows:

```
1 /* ... */
2
3 outputPort DBLP {
4 Location: "socket://dblp.uni-trier.de:80/"
5 Protocol: http {
6   .osc.fetchBib.alias = "rec/bib2/%!{dblpKey}.bib"
7   /* ... */
8 }
9 Interfaces: DBLPInterface
10 }
11
12 outputPort RIS { /* ... */ }
13
14 inputPort ImporterInput {
15 Location: "socket://localhost:8009/"
16 Protocol: sodep
```



```

17 Interfaces: ImporterIface
18 }
19
20 main
21 {
22     import( request );
23
24     dblpReq.dblpKey = request.dblpKey;
25     fetchBib@DBLP( dblpReq )( result );
26
27     addReq.bibtex = result;
28     addReq.userKey = request.userKey;
29
30     addPub@RIS( addReq )
31 }

```

The importer service offers a single operation, `import`, which takes as input a message containing two subnodes: `dblpKey`, the dblp key of the publication to import from DBLP, and `userKey`, which must be a valid user key for a session inside of the RIS. The idea is that a user has to invoke operation `login` before using operation `import`, thus opening a session in the RIS, and that she then invokes `import` with the `userKey` it got as a response from `login`. After receiving a message for operation `import`, the importer service proceeds by invoking the DBLP service to retrieve the BibTeX record stored therein for the dblp key passed by the user. Finally, after retrieving the BibTeX record, the importer asks the RIS to add it through operation `addPub`.

7 RESTful Services

So far, we have considered architectures based on the style of Web Services [13, 10], where our `http` extension is used to bind HTTP messages to verb-based services. In this section, we explore how to use the Jolie-HTTP binding to follow the REpresentational State Transfer architectural style (REST) [26].

7.1 REpresentational State Transfer (REST)

The REST architectural style is a collection of principles for the design and development of web systems, based on the key abstract concept of *resource* [26]. These principles are aimed at promoting the scalability and reusability of services by constraining how they should interact. Such constraints reduce coupling and allow for the use of intermediary services to improve, among other aspects, performance, security, and integration. We briefly present the principles of the REST style in the following (see [54] for a more comprehensive overview of REST and its adoption).

1. *Resource Identification.* Interactions with a RESTful service happen by referring to the resources that it exposes. Resources are globally identified by URIs [42].

Path	Actions			
	GET	POST	PUT	DELETE
/poll	Get list of polls	Create poll	-	-
/poll/{pid}	Get poll <i>pid</i>	-	-	Delete poll <i>pid</i>
/poll/{pid}/vote	Get votes of poll <i>pid</i>	Create vote in poll <i>pid</i>	-	-
/poll/{pid}/vote/{vid}	Get vote <i>vid</i> of poll <i>pid</i>	-	Set state of vote <i>vid</i>	Delete vote <i>vid</i>

Table 1: Resources offered by the poll service.

2. *Uniform Interface.* The operations that can be used on resources are fixed. For example, the uniform interface offered by HTTP includes the operations (called methods) GET, POST, PUT, and DELETE. GET retrieves a representation of the state of a resource. POST submits some data to modify the state of a resource. PUT creates a new resource. Dually, DELETE is used for deleting a resource. The idea is that the uniform interface should be a small set of operations, yet generic enough to implement all the desired functionalities in a web system.
3. *Self-descriptive Messages.* Each message must contain all the necessary data and metadata for the message to be processed. For example, metadata can be used to indicate the format used to encode the message payload, informing the receiver of how it should be unmarshalled.
4. *Hyperlinks as the engine of application state.* State transitions in a RESTful service are explicitly communicated. After the service processes a request that causes a state transition, it sends hyperlinks in the response to inform the client of what resources it can now use.

7.2 Routing via URI templates

We show how to program a RESTful service in our framework by implementing a service for online polls, inspired by the reference example used in [54] to introduce the REST style. This poll service allows clients to manage polls, where each poll has a number of options that can be chosen from. Technically, this is achieved by offering the resources reported in Table 1, where we also describe the semantics of each HTTP method for a resource.

A naive but immediate solution for obtaining a RESTful service is to write a Jolie service that implements the Uniform Interface (here we limit our discussion to the operations GET, POST, PUT, and DELETE). We call this interface `WebIFace`:

```

1 interface WebIface {
2 RequestResponse:
3 get (WebRequest) (undefined), post (WebRequest) (undefined),
4 put (WebRequest) (undefined), delete (WebRequest) (undefined)
5 }

```

Each operation receives a message of type `WebReq`, from § 4, and replies with a message of type `undefined`, since the type of the responses depends on the resource that the operation is applied to. We can implement the interface using our default operations:

```

1 inputPort WebIn {
2 Location: "socket://www.mysrv.com:80"
3 Protocol: http {
4   .default.get = "get"; .default.post = "post";
5   .default.put = "put"; .default.delete = "delete"
6 }
7 Interfaces: WebIface
8 }
9
10 main {
11   [ get( request )( response ) { /* ... */ } ]
12   [ post( request )( response ) { /* ... */ } ]
13   [ put( request )( response ) { /* ... */ } ]
14   [ delete( request )( response ) { /* ... */ } ]
15 }

```

In the service above, each operation implements its respective HTTP method. We now need a way of differentiating the behaviour of each operation, e.g., `get`, depending on the resource requested by the client. For this purpose, we use standard URI templates [28]. Specifically, we introduce a Jolie standard library service, called `UriTemplates`, that we can use to match the resource URI sent by a client to a URI template and extract the parameters embedded in the resource URI (e.g., `pid` in Table 1). As URI templates for the poll service, we use those reported in the first column of Table 1. A naive use of service `UriTemplates` is the following, where we exemplify how to support GET invocations for individual polls and votes:

```

1 main {
2   get( request )( response ) {
3     match@UriTemplates( {
4       .uri = request.requestUri,
5       .template = "/poll/{pid}"
6     } )( m );
7     if ( m ) { // GET /poll/{pid}
8       response << global.polls.(m.pid)
9     } else { // GET /poll/{pid}/vote/{vid}
10      match@UriTemplates( {

```




Figure 5: Router architecture for RESTful Jolie services.

```

11     .uri = request.requestUri,
12     .template = "/poll/{pid}/vote/{vid}"
13   } ) ( m );
14   response << global.polls.(m.pid).votes.(m.vid)
15 }
16 }
17 }

```

As an example, in the code above, if we receive an HTTP request targeting URI `/poll/5/vote/2`, we would get a positive match in Lines 10-13 and the subnodes `pid` and `vid` in variable `m` would be set to the respective values 5 and 2 from the URI. In Line 14, we use these values to find the right vote to return to the invoker.

7.3 The Router Service

While functional, the approach we followed in programming the poll service suffers from poor separation of concerns and readability. This is because the code for determining which action we should perform, by matching URIs with URI templates, is mixed with the code for implementing each action. We solve this problem architecturally, by introducing a standard router service that can be autonomously deployed and therefore modularly reused in the implementation of any RESTful system. The idea is to separate a RESTful service into two services: a *router*, which deals with matching URIs to actions, and a *controller*, which deals with the implementations of actions. (The same router may actually manage multiple controllers, but here we will focus on just one.)

We depict the resulting architecture for a RESTful Jolie service in Figure 5. In the Figure, the router service is the service that exposes the uniform interface to clients. The implementation of actions is then delegated to the controller, which offers a custom interface defined by the programmer. We remark two operations that the router service offers to the controller: `config` and `makeLink`. Operation `config` is used to tell the router how client invocations for its uniform interface should be matched with the user-defined operations offered by the controller. Operation `makeLink`, instead, is used by the controller to build hyperlinks that respect the configuration of the router, ensuring that clients receive hyperlinks that can be correctly routed in the future. We impose no restriction on the transport between the router and the controller: it can be anything supported by Jolie, including, e.g., shared memory communications (the

default, obtained using local locations [49]) or Web Services protocols (e.g., SOAP). The architecture can be nested, allowing for the composition of RESTful and/or non-RESTful services: the controller may be a router itself, or orchestrate operations offered by other routers.

In our poll service example, the poll service is the controller. We can use the router to rewrite the poll service as follows.

```
1 init {
2   config.host = "localhost:8080";
3   // Resource poll
4   config.resources[0] << {
5     .name = "poll",
6     .id = "pid",
7     .template = "/poll"
8   };
9   // vote sub-resources of poll
10  config.resources[0].resources[0] << {
11    .name = "vote",
12    .id = "vid",
13    .template = "/vote"
14  };
15  config@Router( config )()
16 }
17
18 main {
19   [ poll_index()( response ) { // GET /poll
20     foreach( pid : global.polls ) {
21       makeLink@Router( {
22         .operation = "poll_show",
23         .params.pid = pid
24       } )( response.href[i++] )
25     }
26   } ]
27
28   [ poll_show( request )( response ) { // GET /poll/{pid}
29     findPoll;
30     response.options << poll.options;
31     makeLink@Router( {
32       .operation = "vote_index",
33       .params.pid = request.id
34     } )( response.votes.href );
35     for( i = 0, i < #poll.vote, i++ ) {
36       response.votes.vote[i] << poll.vote[i]
37     }
38   } ]
39 }
```

```

40 [ vote_index( request )( response ) { // GET /poll/{pid}/vote
41     findPoll;
42     response.vote -> poll.vote
43 } ]
44
45 // etc.
46 }

```

In Lines 2–15, the router is configured to target the address of our controller and then to offer poll resources and their sub-resources vote. The router will then route requests to corresponding operations suffixed with special identifier to distinguish what HTTP method they should receive from. These operations are provided in the **main** block of the controller; we put comments to indicate what kind of calls correspond to their invocations. For example, operation `poll_index` is invoked by GET requests for resource `/poll`, which returns a list of hyperlinks towards existing polls. Observe that we build hyperlinks by invoking operation `makeLink`, offered by the router, which allows us to abstract from the URI templates that we are using and avoid the error-prone manual writing of hyperlinks. The rest of the service implements the behaviour described in Table 1.

We now move to the implementation of the router service, focusing on its most important behaviour:

```

1  /* ... */
2
3  define route {
4      findRoute;
5      if ( !found ) {
6          setErrorStatusCode
7      } else {
8          /* prepare invokeReq */
9          invokeReq.operation = found.operation;
10         invoke@Reflection( invokeReq )( response );
11         statusCode = 200;
12         /* set headers and status code */
13     }
14 }
15
16 main
17 {
18     [ get( request )( response ) {
19         method = "get";
20         route
21     } ]
22
23     [ post( request )( response ) {
24         method = "post";

```

```

25     route
26   } ]
27
28   [ makeLink( request )( response ) { /* ... */ } ]
29
30   /* ... */
31 }

```

The router offers the Uniform Interface and the `config` and `makeLink` operations (we omit code that works as in other standard frameworks, not relevant for our discussion). The most important part is how messages are routed from the Uniform Interface to the controller. For each operation in the Uniform Interface, e.g., `get` and `post`, the `route` procedure is called and looks for the appropriate route for that method in the configuration. When this is found, variable `invokeReq` is programmed to contain the location of the controller, the data to send, and the operation to invoke (Lines 229–230). Now we need a way to invoke an operation that we do not statically know; for this, we introduce a new `Reflection` service to the Jolie standard library that supports typical reflection operations, as in Java. In Line 231, `Reflection` is used to invoke the operation decided at runtime on the controller. Depending on the configuration and whether the operation replies with a response or a fault, the headers and status code in the reply are set accordingly.

7.4 RESTful Processes

In our poll service example, the life cycle of resources is handled with a shared memory space. Hereby, we describe how a more process-oriented approach can be used by combining correlation sets with REST routing. We also report on how custom operations (without suffixes) can be used in route configurations. The program below is a service that manages quizzes under path `/quiz`. The code follows:

```

1  /* ... */
2
3  cset {
4  id: Start.id Show.id Answer.id
5     Confirm.id Giveup.id Timeout.id
6  }
7
8  init {
9    config.routes[0] << {
10     .method = "post",
11     .template = "/quiz",
12     .operation = "start"
13   };
14   config.routes[1] << {
15     .method = "get",
16     .template = "/quiz/{id}",
17     .operation = "show"

```

```

18     };
19     config.routes[2] << {
20         .method = "delete",
21         .template = "/quiz/{id}?reason=confirm",
22         .operation = "confirm"
23     };
24     /* other routes ... */
25     config@Router( config )()
26 }
27
28 main
29 {
30     start( request )( response ) { // POST /quiz
31         csets.id = new;
32         makeLink@Router( {
33             .operation = "show",
34             .params.id = "id"
35         } )( response.href );
36         send@Mail( /* send link to request.player */ )();
37         quiz -> request.quiz
38     };
39     setNextTimeout@Time( TO { .message.id = csets.id } );
40     provide
41     [ show()( state ) { // GET /quiz/{id}
42         /* reply with quiz text and hyperlinks */
43     } ]
44     [ answer( quiz.answers ) ] // PUT /quiz/{id}/answers
45     until
46     [ confirm() () ] // DELETE /quiz/{id}?reason=confirm
47     [ giveup() () ] // DELETE /quiz/{id}?reason=giveup
48     [ timeout() ]; // Local memory call from Time
49
50     /*
51     Send e-mail with results to
52     request.quizmaster and request.player
53     */
54 }

```

The service starts with the configuration of its own correlation mechanism and that of the router (Lines 3–26). In the **main** procedure (the behaviour of the service), we start by offering to clients the possibility of creating a new quiz. A quiz creation request must include some text (`request.quiz`) and the e-mail address of the player the quiz is intended for (`request.player`). After the correlation token for the quiz is created (Line 31), we send the hyperlink to access the quiz (asked to the router in Lines 32–35) to the player (Line 36). The process now created to handle the quiz registers a timeout (of duration `TO`) using the `Time` service from the Jolie standard

library. We then enter a **provide-until** block: the operations `show` and `answer` are repeatedly provided to clients until either `confirm`, `giveup`, or `timeout` is invoked. Each operation is commented with how it can be reached from the Web (aside from `timeout`, which is local). Operation `show` is meant to be invoked by the player, who received the hyperlink to access it via e-mail when the quiz was created. Operation `answer` is used to create (or replace) the sub-resource answers of the quiz. When she is satisfied, the player can either `confirm` her answers or `giveup`, using different hyperlinks. Finally, the quizmaster and the player are notified of the outcome.

7.5 Integrating Javascript

The client input for manipulating or accessing resources typically has to be validated to check for errors (e.g., invalid syntax). This must be done both on the client and on the server: the former is to offer rapid feedback to the user and avoid sending wrong requests in the first place, whereas the latter is aimed at avoiding malicious client requests. Since validation code in the client is typically written in JavaScript (unless it is trivially supported by HTML or similar declarative mechanisms), we developed an integration mechanism between Jolie and JavaScript to be able to reuse JavaScript programs in Jolie services and avoid error-prone duplication of logic. The basic idea is to enable the invocation of JavaScript programs as if they were services. We obtain this property by extending the *embedding* mechanism of Jolie [49] to support the execution of JavaScript programs as sub-services of a Jolie program.

Concretely, our extension allows to bind output ports to functions offered by a JavaScript program via an **embedded** block in the deployment part of a Jolie service. In our quiz example, assume that the JavaScript code used by the client for validating the quiz creation request is in a function called `validate` that resides in a file called `script.js`. We can bind this function to an output port as follows:

```
1 interface JSIface {
2   RequestResponse: validate(QuizRequest) (bool) }
3 outputPort JS { Interfaces: JSIface }
4 embedded { Javascript: "script.js" in JS }
```

Our extension takes care of converting Jolie values to JavaScript objects and vice versa. Now, we can invoke the `validate` function before Line 31 in our example:

```
1 validate@JS( request )( ok );
2 if ( !ok ) { throw( MalformedRequest ) }
```

8 Performance

In this Section, we present the results of some representative performance experiments executed with our framework. The aim of these experiments is to obtain indicative information on the applicability of our work. All experiments were run on a server machine equipped with an i5-2500k CPU and 8GB of RAM memory.

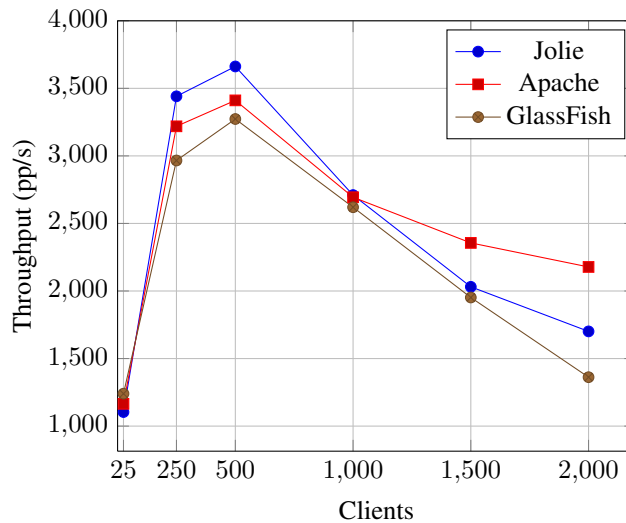


Figure 6: Static page download experiment.

Serving static content. A central feature of any web system is serving static content, e.g., an HTML page, an image, or a CSS file. In this experiment, we measure the throughput (pages downloaded per second) of each framework when serving a growing number of clients a simple web page of middle size (about 1,500 bytes). The results are shown in Figure 6. In the graph, our solution is labelled Jolie. The other frameworks used in the comparison are the Apache HTTP server² and the GlassFish server³, the reference implementation of Java EE (Enterprise Edition) by Oracle. It is clear from the data that Jolie offers comparable performance wrt the other two frameworks.

Templated content. In this experiment, we measure the speed in serving dynamic pages, i.e., pages whose content is computed at runtime for each request rather than being statically stored in a file. Specifically, here the server receives the request for a page, retrieves the main page content and then sets it inside of a template containing a navigation menu and graphical layout⁴. To compute the final result for the clients in Apache and GlassFish, we used PHP (version 5) and Java Server Pages (JSP) respectively. The results, given in Figure 7, show again that the three frameworks perform with relatively close speeds.

Scalability. In our setting, replying to a client request may involve the use of many services (and hence many processes). It is therefore interesting to observe the impact that the number of services involved in processing a client request has on scalability.

²<https://httpd.apache.org/>

³<https://glassfish.java.net/>

⁴The template used is that of the Jolie website, with a simple header and footer: <http://www.jolie-lang.org/>




Figure 7: Templated page download experiment.

To this aim, we run an experiment where a client can ask a Jolie web server for a list of news; the web server then asks a collector service to retrieve (in parallel) news from a number of autonomous news services; finally, the aggregated list of news is returned to the client. We measure the performance of the system (as throughput) using 20, 40, 60, and 80 news services that the collector service must contact. The results are shown in Figure 8, where each curve represents the behaviour of the system when continuously receiving requests from 250, 500, 750, or 1000 clients respectively. As first observation, we note that the system scales well wrt the number of clients (as for the templated page experiment). When we increase the number of news services, throughput decreases more significantly. This is to be expected since each client request requires all news services to be contacted in order to be served. However, since news services are contacted in parallel by the collector service, performance degrades better than directly proportional to the number of news services. (Doubling the number of services does not halve the throughput.) More interestingly, it seems that the number of clients does not have a strong effect on performance in the interval we considered. To understand this better, in Figure 9 we show the result of normalising the data points of each curve in Figure 8 wrt their maximum speed (in the interval we considered) at 20 services. This gives us an indication of how much performance degrades in percentage wrt the best data point at 20 services, for each client load. We observe that all curves follow similar behaviour, confirming the impression that performance degrades gracefully independently of the number of clients in our interval.

9 Related Work

To the best of our knowledge, our work is the first to propose a unified language for dealing with the programming of web servers, scripting, and the architecture of service




Figure 8: Scalability experiment.

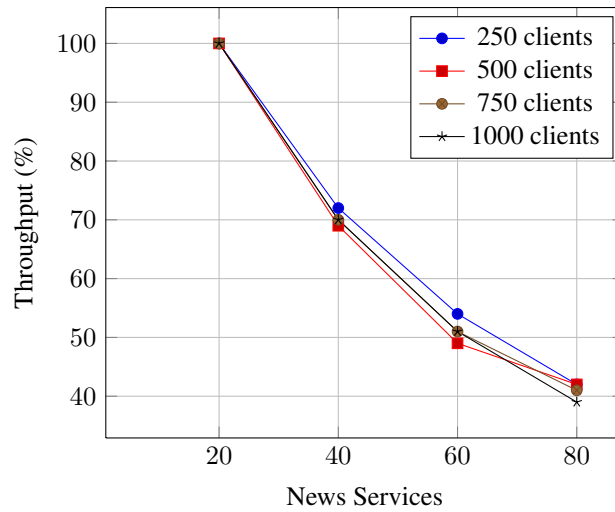


Figure 9: Scalability experiment (normalised curves).

systems in the Web by means of mediator services. Our development was inspired by related work in these areas, described below.

The frameworks most similar to ours are those for modelling business processes, such as WS-BPEL [51], WS-CDL [66], and YAWL [65]. Differently from our approach, these tools are integrated with web applications through third-party tools. Some of the ideas presented in this paper (e.g., the **default** parameters for implementing web servers) may be easily applied to WS-BPEL, making our work a potential reference.

The idea of using a router component to obtain RESTful applications is used also in other frameworks, e.g., Ruby on Rails [7]; we can find similar methodologies also in .NET and Java EE. The difference in our approach is that a router is a service. (In general, in Jolie every component in Jolie is a service.) Therefore, it can be independently configured and deployed (or even replicated, if there is a need to scale wrt requests).

Other works offer tools for supporting the development of process-aware web applications. The papers [20, 21] propose a formally-specified language, implemented in Java, for defining processes that can transparently access resources on the web using a fixed set of primitive operations; the language supports similar process structured as those found in Jolie behaviours, although in our case operations are user-defined. In [58], the authors present a process-based approach to deal with user actions through web interfaces using EPML; like Jolie, EPML is formally specified and comes with an execution engine. JOpera comes with an integration layer for offering REST-based interfaces to business processes [55]. These solutions are formed by integrating separate modules for process modelling, computation, and system integration. In contrast, our framework addresses all these aspects using the same language. JOpera also supports the composition of RESTful services using a graphical notation [53]. EPML can integrate with other languages to integrate user interfaces with process execution; we are currently investigating in a similar direction (see § 10, *Scaffolding of User Interfaces*). The modelling language IFML [19] captures both processes and the modelling of user front-ends; IFML offers an expressive behavioural model, but is focused on modelling rather than implementation as in here, so it may be an interesting complement to the implementation framework presented in this work. The S scripting language is a domain-specific language for writing high-performance RESTful web services [16]. S is natively based on the Uniform Interface defined in REST and the resource abstraction. This allows for some language-based analyses and optimisations based on the semantics of the Uniform Interface. In our case, we could implement similar features in the router service (§ 7.3), since its configuration tells us which methods are going to be used for accessing which operations/resources. We leave such aspects as an interesting future work, since it may also require an analysis of state modifications as indicated in [16]. In [52], an extension to the BPEL language is proposed to capture the publishing and composition of resources. This is a more high-level approach than that proposed in this paper, where we leave the modelling of resource semantics to the programmer using our more low-level HTTP extension. We see the two approaches as complementary: Jolie may be used as the lower layer of similar high-level abstractions, implemented via a compiler or similar techniques. Indeed, the methodology given in [52] should be straightforward to port to a DSL based on our implementation.

Hop [60, 17] and GWT [3] are programming frameworks that deal with the pro-

programming of both the user interface and the server-side application logic using a single codebase, which gets then compiled in the code for the client interface and the services. Differently, in this paper we do not deal with the generation of client code. Instead, we developed an integration between existing technologies (HTML, AJAX calls, JSON, etc.) and our services, by using our `http` protocol to convert the data structures handled by these technologies to/from those handled by Jolie. This leaves the choice of which framework to use for implementing the web user interface to the developer. The client code compiled from GWT projects can be reused with our `http` extension, which is able to parse GWT requests. HipHop [15] is an extension of Hop based on the synchronous language Esterel [14], which introduces orchestration primitives to Hop. The major difference between HipHop and our solution is that behavioural code in Jolie is kept separated from deployment information, making it reusable in different environments, whereas HipHop code mixes the two aspects (for example, cookies in Hop are handled in behavioural code).

Another work that shares some of our aims is the Bigwig project, which offers a language for the programming of session-aware web applications [18]. Our language for behaviours is more expressive than that of Bigwig, which does not support, e.g., the programming of processes using multiparty sessions; however, in our setting we obtain this expressiveness by requiring the programmer to manually handle session identifiers in processes, whereas in Bigwig these are handled automatically. Bigwig is based on the Apache web server, whereas our approach is self-contained: web servers, services, and service mediators (which Bigwig does not handle) are all written in Jolie.

Our **default** configuration parameters for `http` allows a service implementation to catch and reply to invocations for operations that were not known at design time. The same aspect has been previously theoretically modelled through mobility mechanisms for names in process calculi, e.g., in [43, 59, 30]. Our approach is less powerful because these theoretical models elevate the received operation names at the language level: a service may receive an operation name, store it in a variable, and then use the latter in (*input*) and (*output*) primitives as an operation. This is not possible in our behavioural language, since operations in input and output statements are statically defined. We chose not to support this kind of mobility, since it would make the definitions of Jolie interfaces change at runtime. This would break the basic assumption of statically defined operations used in the formal model and implementation of the Jolie language, which goes out of the scope of this paper. It would also make Jolie fundamentally different from other standards for web services, such as WSDL [9], with unclear consequences on their integration. Static operation names are also used in many formal models for the verification of concurrent programming languages (e.g., session types [34]), which we are interested in adopting for Jolie in the future.

10 Discussion and Future Extensions

We discuss some aspects of web programming with Jolie and future extensions related to the work that we presented in this article.

Holystic Approach. The main motivation of this work is to lower the complexity of programming web-based systems by offering a unified language to capture their differ-

ent aspects. However, the current widespread approach of having a specialised technology for each of such aspects may have an advantage when it comes to the required knowledge to use them, as each technology can be studied in isolation. For example, the administrator of a web server in a larger system has to learn only how to use the web server software she uses, abstracting from the other technologies in the rest of the system (where, e.g., WS-BPEL or ESB technologies may be present).

When dealing with only one aspect of web programming, learning how to use a specific software to deal with such aspect may be less time consuming than learning the Jolie language, which is more general. A possible solution for this problem could be to develop Domain Specific Languages (DSLs), supported by Integrated Development Environments (IDEs), that are compiled to Jolie code. The idea is that a specific DSL would deal with one aspect of web programming, while retaining the benefits of having a single underlying language for the different components of a web system. It is still uncertain whether this step would really be necessary, for two reasons. The first reason is that for simple tasks, such as serving static content, we can offer a reference distributable implementation such as the Leonardo web server in § 4, as an alternative to other standard implementations such as the Apache Web Server. The second reason is that many web systems require dealing with multiple aspects of web programming. In those cases, it can take less time to learn Jolie than learning about the available specific technologies to cover all the use cases that the work we presented can address; this would amount to learning, at least, a web server, an orchestration (e.g., WS-BPEL), and a service mediation technologies.

Adoption. How and when should a solution such as that proposed in this paper be adopted in real-world software projects? We discuss an answer by distinguishing between two main cases: the development of new systems and the extension of existing systems.

When dealing with the development of an entirely new web system, Jolie offers a simple and unified language for dealing with the architectural aspects (layering, deployment), the behavioural aspects (application logic), and the serving of static content (web servers). Therefore, Jolie is now a candidate for the rapid prototyping of a web system. Since Jolie integrates with other technologies, starting with our framework does not imply that the final system must be written entirely in Jolie: different parts may be refined later either by using Jolie or other technologies (e.g., Java, WS-BPEL).

When dealing with the extension of an existing system, or even the development of a new system that has to integrate with other existing systems, Jolie can be considered as a glue framework for bridging services based on different technologies. In particular, it is convenient to use the simple syntax of Jolie for writing processes that direct the behaviour of other services in a system. In general, the integration capabilities of Jolie allow for its introduction in a development team by starting from a single service in a larger system, which can be used by the team to assess whether Jolie should be adopted in other parts of the system after seeing how it performs. We conjecture that this step-by-step introduction of web services written in Jolie will be key for its adoption by expert web developers. We are currently following this development methodology in some software projects at the University of Southern Denmark, for the improvement of the web-based tools provided to students and staff.

Since Jolie is a relatively new language, most programmers are still unfamiliar with it and therefore their training must be taken into account in a project. An advantage of a unified framework such as ours, though, is that it allows to understand multiple aspects of web-based systems by learning a single language. With the rise of more complex and structured web-based systems, we believe that there can be a motivation to learn Jolie even for developers who are expert in more established technologies.

Reversibility. Reversibility techniques (see, e.g., [39]) deal with the automatic reversal to previous states in a distributed system. In the context of the Web, reversibility could play a role in allowing users to revert the effects of unsafe operations, e.g., to “un-delete” resources. In complex sessions involving multiple parties, this requires inferring which parties should be notified of such a reversal event. The formal semantics of Jolie [48] should play an important role in enabling the formal study of such notions.

Scaffolding of User Interfaces. The explicit structure of processes written in Jolie allows us to statically see the workflow that a user interface should follow when interacting with a Jolie service. We could use this aspect to develop a scaffolding tool for user interfaces, starting from the process structure of a service. Specifically, given a behaviour in Jolie, it would be possible to automatically generate a user interface that follows the communication structure of the behaviour. This would be in line with the notions of *duality* formalised in [34, 31].

Behavioural analyses. Since our framework makes the process logic of a web application explicit, it would be possible to develop a tool for checking that the invocations performed by a web user interface written in, e.g., Javascript, match the structure of their corresponding Jolie service. The techniques presented in [35, 27, 50] may offer useful first steps towards this aim.

Declarative data validation. Our framework exploits the message data types declared in the interfaces of a Jolie service to *validate* the content of incoming messages from web user interfaces (§ 3.2). We plan to extend this declarative support for data validation by introducing an assertion language for message types that can check more complex properties (e.g., integer ranges and regular expressions).

Extensions to other web protocols. Our work lays the foundations for using Jolie as a fully-fledged language to handle HTTP-based systems. By following the same approach, it would be possible to develop support for new emerging protocols for the web, such as WebSockets [36] and SPDY [2].

11 Conclusions

We have presented a framework for the programming of process-aware web systems, where processes are used as a holistic approach to capture the development of the different components of such systems, such as web servers, orchestrators, and service mediators. Through examples, we have shown how our solution subsumes useful web design patterns and how it captures complex scenarios involving, e.g., multiparty

sessions and evolvability. Our `http` extension is open source and is included in the standard distribution of Jolie, along with the language additions that we introduced to support protocol configurations [5, 37]. Our integration is seamless wrt data formats, meaning that existing Jolie code can be ported to HTTP without having to deal explicitly with the typical data formats used in the Web (e.g., JSON). Since our alterations to the Jolie language targeted only the configuration of communication ports, all the techniques developed for the verification and execution of Jolie programs (as the typing system in [48] for correlation sets) can be transparently applied to the process-aware web application logic written in our framework.

Acknowledgements

The author thanks Claudio Guidi, Saverio Giallorenzo, and the anonymous referees for their useful comments. This work was supported by CRC (Choreographies for Reliable and efficient Communication software), grant no. DFF-4005-00304 from the Danish Council for Independent Research.

References

- [1] Business Process Model and Notation. <http://www.omg.org/spec/BPMN/2.0/>.
- [2] Google SPDY. <https://developers.google.com/speed/spdy/>.
- [3] Google Web Toolkit. <http://code.google.com/webtoolkit/>.
- [4] JavaScript Object Notation. <http://www.json.org/>.
- [5] Jolie HTTP extension. <https://jolie.svn.sourceforge.net/svnroot/jolie/trunk/extensions/http>.
- [6] Leonardo Web Server. <http://www.sourceforge.net/projects/leonardo/>.
- [7] Ruby on Rails. <http://rubyonrails.org/>.
- [8] SOAP Specifications. <http://www.w3.org/TR/soap/>.
- [9] Web Services Description Language. <http://www.w3.org/TR/wsdl>.
- [10] Web Services Interoperability (WS-I). <http://www.ws-i.org>.
- [11] Workflow Patterns. <http://www.workflowpatterns.com/>.
- [12] XML-RPC. <http://www.xmlrpc.com/>.
- [13] Gustavo Alonso, Fabio Casati, Harumi A. Kuno, and Vijay Machiraju. *Web Services - Concepts, Architectures and Applications*. Data-Centric Systems and Applications. Springer, 2004.

- [14] Gérard Berry. The foundations of esterel. In *Proof, Language, and Interaction, Essays in Honour of Robin Milner*, pages 425–454, 2000.
- [15] Gérard Berry and Manuel Serrano. Hop and hiphop: Multitier web orchestration. In *Distributed Computing and Internet Technology - 10th International Conference, ICDCIT 2014, Bhubaneswar, India, February 6-9, 2014. Proceedings*, pages 1–13, 2014.
- [16] Daniele Bonetta, Achille Peternier, Cesare Pautasso, and Walter Binder. S: a scripting language for high-performance restful web services. In *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2012, New Orleans, LA, USA, February 25-29, 2012*, pages 97–106, 2012.
- [17] Gérard Boudol, Zhengqin Luo, Tamara Rezk, and Manuel Serrano. Reasoning about web applications: An operational semantics for hop. *ACM Trans. Program. Lang. Syst.*, 34(2):10, 2012.
- [18] Claus Brabrand, Anders Møller, and Michael I. Schwartzbach. The ¡Bigwig¿ Project. *ACM Trans. Internet Technol.*, 2(2):79–114, May 2002.
- [19] Marco Brambilla and Piero Fraternali. *Interaction Flow Modeling Language: Model-Driven UI Engineering of Web and Mobile Apps with IFML*. Morgan Kaufmann, 2014.
- [20] Mario Bravetti. File managing and program execution in web operating systems. *CoRR*, abs/1005.5045, 2010.
- [21] Mario Bravetti. Formalizing restful services and web-os middleware. In *Web Services and Formal Methods - 10th International Workshop, WS-FM 2013, Beijing, China, August 2013, Revised Selected Papers*, pages 48–68, 2013.
- [22] Marco Carbone and Fabrizio Montesi. Deadlock-freedom-by-design: multiparty asynchronous global programming. In Roberto Giacobazzi and Radhia Cousot, editors, *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, pages 263–274. ACM, 2013.
- [23] David A. Chappell. *Enterprise Service Bus - Theory in practice*. O’Reilly, 2004.
- [24] Marlon Dumas, Wil M. P. van der Aalst, and Arthur H. M. ter Hofstede. *Process-Aware Information Systems: Bridging People and Software Through Process Technology*. Wiley, 2005.
- [25] Elsevier. Pure. <http://www.elsevier.com/online-tools/research-intelligence/products-and-services/pure>.
- [26] Roy Thomas Fielding. *Architectural styles and the design of network-based software architectures*. PhD thesis, University of California, Irvine, 2000.

- [27] Simon J. Gay, Vasco Thudichum Vasconcelos, António Ravara, Nils Gesbert, and Alexandre Z. Caldeira. Modular session types for distributed object-oriented programming. In *POPL*, pages 299–312, 2010.
- [28] Joe Gregorio, R Fielding, Marc Hadley, Mark Nottingham, and David Orchard. URI Template. *IETF RFC 6570*, 2012.
- [29] Claudio Guidi. *Formalizing languages for Service Oriented Computing*. PhD. thesis, University of Bologna, 2007. <http://www.cs.unibo.it/pub/TR/UBLCS/2007/2007-07.pdf>.
- [30] Claudio Guidi and Roberto Lucchi. Formalizing mobility in service oriented computing. *JSW*, 2(1):1–13, 2007.
- [31] Daniel Hirschhoff, Jean-Marie Madiot, and Davide Sangiorgi. Duality and i/o-types in the π -calculus. In *CONCUR*, pages 302–316, 2012.
- [32] K. Honda, N. Yoshida, and M. Carbone. Multiparty asynchronous session types. In *POPL*, volume 43(1), pages 273–284. ACM, 2008.
- [33] Kohei Honda, Aybek Mukhamedov, Gary Brown, Tzu-Chun Chen, and Nobuko Yoshida. Scribbling interactions with a formal foundation. In *ICDCIT*, volume 6536 of *LNCS*, pages 55–75. Springer, 2011.
- [34] Kohei Honda, Vasco Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In *ESOP’98*, volume 1381 of *LNCS*, pages 22–138, Heidelberg, Germany, 1998. Springer-Verlag.
- [35] Raymond Hu, Dimitrios Kouzapas, Olivier Pernet, Nobuko Yoshida, and Kohei Honda. Type-safe eventful sessions in java. In *ECOOP*, pages 329–353, 2010.
- [36] IETF. WebSocket protocol. <http://tools.ietf.org/html/rfc6455>.
- [37] Jolie. Programming Language. <http://www.jolie-lang.org/>.
- [38] Brian W. Kernighan and Dennis Ritchie. *The C Programming Language*. Prentice-Hall, 1978.
- [39] Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversing higher-order π . In *CONCUR*, pages 478–493, 2010.
- [40] Alessandro Lapadula, Rosario Pugliese, and Francesco Tiezzi. A calculus for orchestration of web services. In *ESOP*, pages 33–47, 2007.
- [41] Roberto Lucchi and Manuel Mazzara. A π -calculus based semantics for WS-BPEL. *J. Log. Algebr. Program.*, 70(1):96–118, 2007.
- [42] Larry Masinter, Tim Berners-Lee, and Roy T Fielding. Uniform resource identifier (URI): Generic syntax. *IETF RFC 3986*, 2005.

- [43] Robin Milner, Joachim Parrow, and David Walker. A calculus of mobile processes, I and II. *Information and Computation*, 100(1):1–40,41–77, September 1992.
- [44] F. Montesi, C. Guidi, I. Lanese, and G. Zavattaro. Dynamic Fault Handling Mechanisms for Service-Oriented Applications. In *ECOWS*, pages 225–234, 2008.
- [45] F. Montesi, C. Guidi, and G. Zavattaro. Composing Services with JOLIE. In *ECOWS*, pages 13–22, 2007.
- [46] Fabrizio Montesi. DBLP Tools. <http://www.fabriziomontesi.com/dblp/>.
- [47] Fabrizio Montesi. Jolie: a Service-oriented Programming Language. Master’s thesis, University of Bologna, Department of Computer Science, 2010.
- [48] Fabrizio Montesi and Marco Carbone. Programming services with correlation sets. In *ICSOC*, pages 125–141, 2011.
- [49] Fabrizio Montesi, Claudio Guidi, and Gianluigi Zavattaro. Service-oriented programming with jolie. In *Web Services Foundations*, pages 81–107. 2014.
- [50] Fabrizio Montesi and Nobuko Yoshida. Compositional choreographies. In *CONCUR 2013 - Concurrency Theory - 24th International Conference, CONCUR 2013, Buenos Aires, Argentina, August 27-30, 2013. Proceedings*, pages 425–439, 2013.
- [51] OASIS. Web Services Business Process Execution Language. <http://docs.oasis-open.org/wsbpel/2.0/wsbpel-v2.0.html>.
- [52] Cesare Pautasso. BPEL for REST. In *Business Process Management, 6th International Conference, BPM 2008, Milan, Italy, September 2-4, 2008. Proceedings*, pages 278–293, 2008.
- [53] Cesare Pautasso. Composing restful services with jopera. In *Software Composition, 8th International Conference, SC 2009, Zurich, Switzerland, July 2-3, 2009. Proceedings*, pages 142–159, 2009.
- [54] Cesare Pautasso. Restful web services: Principles, patterns, emerging technologies. In *Web Services Foundations*, pages 31–51. 2014.
- [55] Cesare Pautasso and Erik Wilde. Push-enabling restful business processes. In *ICSOC*, pages 32–46, 2011.
- [56] James L. Peterson. Petri nets. *ACM Comput. Surv.*, 9(3):223–252, September 1977.
- [57] Mila Dalla Preda, Maurizio Gabbrielli, Claudio Guidi, Jacopo Mauro, and Fabrizio Montesi. Interface-based service composition with aggregation. In *ESOCC*, pages 48–63, 2012.

- [58] Davide Rossi and Elisa Turrini. Designing and architecting process-aware web applications with epml. In *SAC*, pages 2409–2414, 2008.
- [59] D. Sangiorgi and D. Walker. *The π -calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001.
- [60] Manuel Serrano, Erick Gallesio, and Florian Loitsch. Hop: a language for programming the web 2.0. In *OOPSLA Companion*, pages 975–985, 2006.
- [61] The jQuery Foundation. jQuery. <http://www.jquery.com/>.
- [62] University of Trier and Schloss Dagstuhl. dblp: computer science bibliography. <http://www.jquery.com/>.
- [63] Wil M. P. van der Aalst. Verification of workflow nets. In *ICATPN*, pages 407–426, 1997.
- [64] Wil M. P. van der Aalst. The application of petri nets to workflow management. *Journal of Circuits, Systems, and Computers*, 8(1):21–66, 1998.
- [65] Wil M. P. van der Aalst and Arthur H. M. ter Hofstede. Yawl: yet another workflow language. *Inf. Syst.*, 30(4):245–275, 2005.
- [66] W3C WS-CDL Working Group. Web services choreography description language version 1.0. <http://www.w3.org/TR/ws-cdl-10/>, 2004.