

---

Basic Research in Computer Science

BRICS RS-02-11 Christensen et al.: Extending Java for High-Level Web Service Construction

## Extending Java for High-Level Web Service Construction

Aske Simon Christensen  
Anders Møller  
Michael I. Schwartzbach

BRICS Report Series

ISSN 0909-0878

RS-02-11

March 2002

**Copyright © 2002, Aske Simon Christensen & Anders Møller & Michael I. Schwartzbach.  
BRICS, Department of Computer Science  
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work  
is permitted for educational or research use  
on condition that this copyright notice is  
included in any copy.**

**See back inner page for a list of recent BRICS Report Series publications.  
Copies may be obtained by contacting:**

**BRICS  
Department of Computer Science  
University of Aarhus  
Ny Munkegade, building 540  
DK-8000 Aarhus C  
Denmark  
Telephone: +45 8942 3360  
Telefax: +45 8942 3255  
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through the World Wide  
Web and anonymous FTP through these URLs:**

`http://www.brics.dk`  
`ftp://ftp.brics.dk`  
**This document in subdirectory RS/02/11/**

# Extending Java for High-Level Web Service Construction

Aske Simon Christensen, Anders Møller, and Michael I. Schwartzbach  
BRICS, Department of Computer Science  
University of Aarhus, Denmark

---

We incorporate innovations from the <bigwig> project into the Java language to provide high-level features for Web service programming. The resulting language, JWIG, contains an advanced session model and a flexible mechanism for dynamic construction of XML documents, in particular XHTML. To support program development we provide a suite of program analyses that at compile-time verify for a given program that no run-time errors can occur while building documents or receiving form input, and that all documents being shown are valid according to the document type definition for XHTML 1.0.

We compare JWIG with Servlets and JSP which are widely used Web service development platforms. Our implementation and evaluation of JWIG indicate that the language extensions can simplify the program structure and that the analyses are sufficiently fast and precise to be practically useful.

Categories and Subject Descriptors: D.3.3 [**Programming Languages**]: Language Constructs and Features; D.2.4 [**Software Engineering**]: Software/Program Verification

General Terms: Languages, Design, Verification

Additional Key Words and Phrases: Interactive Web Services, Data-Flow Analysis, XML

---

## 1. INTRODUCTION

The Java language is a natural choice for developing modern Web services. Its built-in network support, strong security guarantees, concurrency control, and widespread deployment in both browsers and servers, together with popular development tools make it relatively easy to create Web services. In particular, JavaServer Pages (JSP) [34] and Servlets [33], which are both Java based technologies, have become immensely popular. However, both JSP, Servlets, and many other similar and widely used technologies, such as ASP, PHP, and CGI/Perl, suffer from some problematic shortcomings, as we will argue in the following and try to address.

### 1.1 Sessions and Web Pages

In general, JSP, Servlets, and related approaches provide only low-level solutions to two central aspects of Web service design: *sessions* and *dynamic construction of Web pages*.

A session is conceptually a sequential thread on the server that has local data, may access data shared with other threads, and can perform several interactions with a client. With standard technologies, sessions must be encoded by hand which is tedious and error prone. More significantly, it is difficult to understand the

---

Authors' address: Ny Munkegade, Bldg. 540, DK-8000 Aarhus C, Denmark.  
Email: {aske,amoeller,mis}@brics.dk

control-flow of an entire service from the program source since the interactions between the client and the server are distributed among several seemingly unrelated code fragments. This makes maintenance harder for the programmer if the service has a complicated control flow. Also, it prevents compilers from getting a global view of the code to perform whole-service program analyses.

The dynamic construction of Web pages is typically achieved by print statements that piece by piece construct HTML fragments from text strings. Java is a general purpose language with no inherent knowledge of HTML, so there are no compile-time guarantees that the resulting documents are valid HTML. For static pages, it is easy to verify validity [29], but for pages that are dynamically generated with a full programming language, the problem is in general undecidable. Not even the much simpler property of being well-formed can be guaranteed in this way. Instead of using string concatenation, document fragments may be built in a more controlled manner with libraries of tree constructor functions. This automatically ensures well-formedness, but it is more tedious to use and the problem of guaranteeing validity is still not solved.

The inability to automatically extract the control-flow of the sessions in a service raises another problem. Typically, the dynamically generated HTML pages contain input forms allowing the client to submit information back to the server. However, the HTML page with the form is constructed by one piece of the service code, while a different piece takes care of receiving the form input. These two pieces must agree on the input fields that are transmitted, but when the control-flow is unknown to the compiler, this property cannot be statically checked. Thorough and expensive run-time testing is then required, and that still cannot give any guarantees.

The `<bigwig>` language [9] is a research language designed to overcome these problems. Its core is a strongly typed C-like language. On top is a high-level notion of sessions where client interactions resemble remote procedure calls such that the control-flow is explicit in the source code [7]. Also, XHTML [30], the XML version of HTML, is a built-in data type with operations for dynamically constructing documents. The values of this data type are well-formed XHTML fragments which may contain “named gaps”. Such fragments can be combined with a special “plug” operation which inserts one fragment into a gap in another fragment. This proves to be a highly flexible but controlled way of building documents. The client interactions and the dynamic document construction are checked at compile-time using a specialized type system [32] and a program analysis [8] performing a conservative approximation of the program behavior to attack the problems mentioned above. More specifically, `<bigwig>` services are verified at compile-time to ensure that 1) a plug operation always finds a gap with the specified name in the given fragment, 2) the code that receives form input is presented with the expected fields, and 3) only valid XHTML 1.0 is ever sent to the clients.

## 1.2 Contributions

In this paper we obtain similar benefits for Java applications. Our specific contributions are to show:

- how the session model and the dynamic document model of `<bigwig>` can be integrated smoothly into Java;

- how the type system from [32] and the program analysis from [8] can be combined, generalized, and applied to Java to provide even stronger static guarantees than known from `<bigwig>`; and
- how our service model subsumes and extends both the Servlet and the JSP style of defining Web services.

The integration into Java is achieved using a class library together with some extensions of the language syntax. The resulting language is called JWIG. When running a JWIG service without applying the static analyses, a number of special run-time errors may occur: If one of the three correctness properties mentioned in the previous section is violated, an exception is thrown. The goal of the static analyses is to ensure at compile-time that these exceptions never occur. In addition to having this certainty, we can eliminate the overhead of performing run-time checks.

Such guarantees cannot be given for general Servlet or JSP programs. However, we show that the structures of such programs are special cases in JWIG: Both the script-centered and the page-centered styles can be emulated by the session-centered, so none of their benefits are lost.

Our current implementation uses XHTML 1.0, but the approach generalizes in a straightforward manner to an arbitrary interaction language described by an XML schema, such as WML or VoiceXML.

A cornerstone in our program analyses is a novel notion of *summary graphs* which provides a suitable abstraction of the sets of XML fragments that appear at run-time. We show how these graphs can be obtained from a data-flow analysis and that they comprise a precise description of the information needed to verify the correctness properties mentioned above.

Throughout each phase of our program analysis, we will formally define in what sense the phase is correct and we will give a theoretical bound on the worst-case complexity. We expect the reader to be familiar with Java and monotone data-flow analysis, and to have a basic understanding of HTML and XML.

### 1.3 Problems with Existing Approaches

In the following we give a more thorough explanation of the support for sessions and dynamic documents in JSP and Servlets and point out some related problems.

The overall structure of a Web service written as a Servlet resembles that of a CGI script. When a request is received from a client, a thread is started on the server. This thread generates a response, usually an HTML page, and perhaps has some side-effects such as updating a database. Before terminating it sends to the client the response, which is dynamically created by printing strings to an output stream. We call this a *script-centered* approach. The main advantages of Servlets compared to CGI scripts are higher performance and a convenient API for accessing the HTTP layer. A Servlet engine typically uses a thread pool to avoid the overhead of constantly starting threads. Also, Servlets have the general Java benefits of a sandboxed execution model, support for concurrency control, and access to the large set of available Java packages.

A small Servlet program is shown in Figure 1. A client running this service is guided through a sequence of interactions which we call a *session*: First, the service prompts for the client's name, then the name and the total number of invocations is

```

import javax.servlet.*;
public class SessionServlet extends HttpServlet {
    public void doGet(HttpServletRequest request,
        HttpServletResponse response)
        throws ServletException, IOException {
        ServletContext context = getServletContext();
        HttpSession session = request.getSession(true);
        response.setContentType("text/html");
        PrintWriter out = response.getWriter();
        out.println("<html><head><title>Servlet</title></head><body>");
        if (session.isNew()) {
            out.println("<form action=\"SessionServlet\">" +
                "Enter your name: <input name=\"handle\">" +
                "<input type=\"submit\"></form>");
            session.putValue("state", "1");
        } else {
            String state = (String) session.getValue("state");
            if (state.equals("1")) {
                String name = (String) request.getParameter("handle");
                int users =
                    ((Integer) context.getAttribute("users")).intValue() + 1;
                context.setAttribute("users", new Integer(users));
                session.putValue("name", name);
                out.println("Hello " + name +
                    ", you are user number " + users);
                session.putValue("state", "2");
            } else /* state.equals("2") */ {
                String name = (String) session.getValue("name");
                out.println("Goodbye " + name);
                session.invalidate();
            }
        }
        out.println("</body></html>");
    }
}

```

Fig. 1. Example Servlet code. The session flow is encoded into the `HttpSession` object, and HTML documents are constructed by printing string fragments to a stream.

shown, and finally a “goodbye” page is shown. The `ServletContext` object contains information shared to all sessions, while the `HttpSession` object is local to each session. Both kinds of state are accessed via a dictionary interface. The overall structure of the code resembles a `switch` statement that branches according to the current interaction. The Servlet API hides the details of cookies and URL rewriting which is used to track the client throughout the session. HTML documents are generated by printing strings to the output stream.

A JSP service turns the picture inside out by being defined by an HTML document with embedded code snippets. We call this a *page-centered* approach. Figure 2 shows a simple JSP program which dynamically inserts the current time together with a title and a user name based on the input parameters. This approach is quite similar to ASP and PHP, except that the underlying language and its run-time model is safer and better designed. An implementation of JSP typically performs a simple translation into Servlets. This model fits into situations where the service presents pages that are essentially static but with a few dynamic fragments

```

<html><head><title>JSP Demo</title></head><body>
Hello <%
    String name = request.getParameter("who ");
    if (name==null) name = "stranger";
    out.print(name);
%>!
<p>
This page was last updated: <%= new Date() %>
</body></html>

```

Fig. 2. Example JSP page. Code snippets are embedded within the HTML code using special `<%...%>` tags. When a client requests the page, the code snippets are replaced by the strings that result from the evaluation.

inserted. For more complicated services the code tends to dominate the pages, such that they converge towards straight Servlet implementations.

Both JSP and Servlets allow only a single interaction with the client before termination. To simulate a sequential thread, the client is given a session identifier that is stored either as a cookie, as an SSL session key, or in a hidden form field. The local state associated with the thread is stored as attributes in an `HttpSession` object from which it must be recovered when execution is later resumed. Thus, a sequence of interactions must be encoded by the programmer in a state machine fashion where transitions correspond to individual interactions. This is somewhat cumbersome and precludes cases where the resident local state includes objects in the heap or a stack of pending method invocations. However, it should be noted that the state machine model implies other advantages when it is applicable. Specifically, it allows robust and efficient implementations as evidenced by J2EE engines that ensure scalability and transaction safety by dividing session interactions into atomic transitions. If a Web service runs on a J2SE engine, then only the disadvantages exist since every interaction is then typically handled by a new thread anyway.

Data that is shared between several session threads must in JSP and Servlets be stored in a dictionary structure or an external databases. This is in many cases adequate, but still conceptually unsatisfying in a language that otherwise supports advanced scoping mechanisms, such as nested classes.

Finally, JSP and Servlets offer little support for the dynamic construction of XHTML or general XML documents. JSP allows the use of static templates in which gaps are filled with strings generated by code fragments. Servlets generate all documents as strings. These techniques cannot capture well-formedness of the generated documents, let alone validation according to some schema definition. Well-formedness can be ensured by relying on libraries such as JDOM [19], where XML documents are constructed as tree data structures. However, this loses the advantages from JSP of using human readable templates and validity can still only be guaranteed by expensive run-time checks.

Another concern in developing Web services is to provide a separation between the tasks of programmers and designers. With JSP and Servlets, the designer must work through the programmer who maintains exclusive control of the markup tags being generated. This creates a bottleneck in the development process.

The Jtwig language is designed to attack all of these problems within the Java

framework. For a more extensive treatment of Servlets, JSP, and the many related languages, we refer to the overview article on `<bigwig>` [9]. The central aspects of Jwig are an explicit session model and a notion of higher-order XML templates, which we will explain in detail in the following sections.

#### 1.4 Outline

We begin by describing the special Jwig language constructs and packages in Section 2. These extensions are designed to allow flexible, dynamic construction of documents and coherent sessions of client interactions. In Section 3 we explain how to obtain Jwig program *flow graphs* that abstractly describe the flow of strings and XML templates through programs. These flow graphs form the basis of the data-flow analyses that are described in Section 4. They culminate in the inference of *summary graphs* which model the sets of XML documents that variables or expressions may evaluate to at given program points. To compute this information we need a preliminary *string analysis*. In Section 5 we describe how the results from the summary graph analysis are used to verify that the run-time errors mentioned earlier cannot occur for a given program. This involves the use of a novel XML schema language, *Document Structure Description 2.0*. We provide an overview of our implementation in Section 6, and evaluate it on some benchmark programs to show that the analysis techniques are sufficiently fast and precise to be practically useful. Finally, we describe ideas for future work in Section 7.

## 2. THE Jwig LANGUAGE

The Jwig language is designed as an extension of Java. This extension consists of a service framework for handling session execution, client interaction and server integration, and some syntactic constructs and support classes for dynamic construction of XML documents.

### 2.1 Program Structure

A Jwig application is a subclass of the class `Service` containing a number of fields and inner classes. An instance of this class plays the role of a running *service*.

A service contains a number of different *sessions*, which are defined by inner classes. Each session class contains a `main` method, which is invoked when a client initiates a session. The fields of a service object are then accessible from all sessions. This provides a simple shared state that is useful for many applications. Concurrency control is not automatic but can be obtained in the usual manner through `synchronized` methods. Of course, external databases can also be applied, for instance using JDBC, if larger data sets are in use. The fields of a session object as well as local variables in methods are private to each session thread. This approach of applying the standard scope mechanisms of Java for expressing both shared state and per-session state is obviously simpler than using the `ServletContext` and `HttpSession` dictionaries in Servlets and JSP.

Figure 3 shows a Jwig service which is equivalent to the Servlet service from Figure 1. In the following, we describe the new language constructs for defining XML templates, showing documents to the clients, and receiving form input.

Session classes come in a number of different flavors, each with different purposes and capabilities, as indicated by its superclass:



```

import dk.brics.jwig.runtime.*;

public class MyService extends Service {
    int counter = 0;
    synchronized int next() { return ++counter; }

    public class ExampleSession extends Session {
        XML wrapper =
            [[ <html><head><title>JWIG</title></head>
              <body><[body]></body></html> ]];
        XML hello =
            [[ <form>Enter your name: <input name="handle">
              <input type="submit"></form> ]];
        XML greeting =
            [[ Hello <[who]>, you are user number <[count]> ]];
        XML goodbye =
            [[ Goodbye <[who]> ]];

        public void main() throws IOException {
            show wrapper<[body=hello];
            String name = receive handle;
            show wrapper<[body=greeting<[who=name,count=next()]];
            exit wrapper<[body=goodbye<[who=name]];
        }}

```

Fig. 3. Example JWIG program. The `MyService` service contains one session type named `ExampleSession` which has the same functionality as the Servlet in Figure 1.

- `Service.Session` is the most general kind of interaction pattern, allowing any number of interactions with the client while retaining an arbitrary session state. When a `Service.Session` subclass is initiated, a new thread is started on the server, which lives throughout all interactions with the client. At all intermediate interactions, after supplying the XML to be sent to the client, the thread simply sleeps, waiting for the client to respond.
- `Service.Page` is used for simple requests from the client. A `Service.Page` is similar to a `Service.Session`, except that it allows no intermediate client interactions. This is conceptually similar to the mechanism used in Servlet and JSP applications, where a short-lived thread is also initiated for each interaction.
- `Service.Seslet` is a special kind of session, called a *seslet*, that is used to interact with the service from applets residing on the client or with other Web services. A `Service.Seslet` does not interact with the client directly in the form of input forms and e.g. XHTML output; instead, it is parameterized with an `InputStream` and an `OutputStream` which are used for communication with the applet or Web service on the client-side. The notion of seslets was introduced in [9].

For the remainder of this article we focus on `Service.Session`. In contrast to the session API in Servlets and JSP, it provides a clear view of the service flow because the sequences of interactions constituting sessions are explicit in the program control-flow. This *session-centered* approach originates from the MAWL project [24].

We specify `Service.Page` by a separate class in order to illustrate that the script- and page-centered approaches are special cases of the session-centered approach,




Fig. 4. Client-server sessions in Web services. On the left is the client's browser, on the right is a session thread running on the server. The thread is initiated by a client request and controls the sequence of session interactions.

and to identify applications of these simpler interaction models to allow implementations to perform special optimizations.

## 2.2 Client Interaction

Communication with the client is performed through the `show` statement which takes as argument an XML template to be transmitted to the client. A session terminates by executing the `exit` statement whose argument is an XML template that becomes the final page shown. Intermediate XML templates need not conform to any XML schema, but when they are used as arguments to `show` or `exit` statements they must be valid XHTML 1.0 [30]. Otherwise, a `ValidateException` is thrown.

During client interactions, the session thread is suspended on the server. Thus, the execution of the `show` statement behaves as a remote procedure call to the client. Return values are specified by means of form fields in the document. For all `form` elements, a default `action` attribute is automatically inserted with a URL pointing to the session thread on the server. The responses from the client are subsequently obtained using the `receive` expression, which takes as argument the name of an input field in the XHTML document that was last shown to the client and returns the corresponding value provided by the client as a `String`. If no such input field was shown to the client, then no corresponding value is transmitted and a `ReceiveException` is thrown. There may be several occurrences of a given input field. In this case, all the corresponding values may be received in order of occurrence in the document into a `String` array using the expression `receive[]`. The non-array version is only permitted if the input field occurs exactly once; otherwise, a `ReceiveException` is thrown. The array version cannot fail: In case there are no fields, the empty array is produced.

Figure 4 illustrates the interactions of a session. In the Jwig example service in Figure 3, the `main` method contains three client interactions: two `show` statements and one `exit` statement. Clearly, the session flow is more explicit than in the corresponding Servlet code in Figure 1.

### 2.3 Dynamic Document Construction

In Servlets and JSP, document fragments are generated dynamically by printing strings to a stream. In Jwig, we instead use a notion of XML *templates*. A template is a well-formed XML fragment which may contain named *gaps*. A special *plug* operation is used to construct new templates by inserting existing templates or strings into gaps in other templates. These templates are higher-order, because we allow the inserted templates to contain gaps which can be filled in later, in a way that resembles higher-order functions in functional programming languages. Templates are identified by a special data type, `XML`, and may be stored in variables and passed around as any other type. Once a complete XHTML document has been built, it can be used in a `show` statement.

Syntactically, the Jwig language introduces the following new expressions for dynamic XML document construction:

<code>[[ xml ]]</code>	(template constant)
<code>exp<sub>1</sub> &lt;[g = exp<sub>2</sub>]</code>	(the plug operator)
<code>( [[ xml ] ] ) exp</code>	(XML cast)
<code>get url</code>	(run-time template inclusion)

These expressions are used to define template constants, plug templates together, cast values to the `XML` type, and to include template constants at run-time, respectively. The *url* denotes a URL of a template constant located in a separate file, and *xml* is a well-formed XML template according to the following grammar:

<i>xml</i> :	<i>str</i>	(character data)
	<code>&lt;name atts&gt; xml &lt;/name&gt;</code>	(element)
	<code>&lt;[g]&gt;</code>	(template gap)
	<code>&lt;{stm}&gt;</code>	(code gap)
	<i>xml xml</i>	(sequence)
<i>atts</i> :	$\varepsilon$	(empty)
	<code>name="str"</code>	(attribute constant)
	<code>name=[g]</code>	(attribute gap)
	<i>atts atts</i>	(sequence)

Here, *str* denotes an arbitrary Unicode string, *name* an arbitrary identifier, *g* a gap name, and *stm* a statement block that returns a value of type `String` or `XML`. Actual XML values must of course be further constrained to be well-formed according to the XML 1.0 specification [12]. Moreover, in this description we abstract away all DTD information, comments, processing instructions, etc. In Figure 3, there are four template constants: `wrapper`, `hello`, `greeting`, and `goodbye`. The `hello` template, for instance, contains two gaps named *who* and *count*, respectively.

XML templates can be composed using the plug operation `exp1 <[g = exp2]`. The result of this expression is a copy of `exp1` with all occurrences of the gap named *g* replaced by copies of `exp2`. This is illustrated in Figure 5. If `exp2` is a string, all special XML characters (`<`, `>`, `&`, `'`, and `"`) are automatically escaped by the corresponding XML character references. If `exp1` contains no gaps named *g*, a `PlugException` is thrown. A gap that has not been plugged is said to be *open*. The `exp2` expression is required to be of type `String` or `XML`. If it is not one




Fig. 5. The plug operation. The two XHTML templates in the top are combined to produce the one below by plugging into the *who* gap.

of these, it is coerced to `String`. There are three kinds of gaps: template gaps, attribute gaps, and code gaps. Both strings and templates may be plugged into template gaps, but only strings may be plugged into attribute gaps. Attempts to plug templates into attribute gaps will cause a `PlugException` to be thrown. When a template is shown, all remaining open gaps are removed in the following way: Each template gap is replaced by the empty string, and for each attribute gap, the entire attribute containing the gap is removed. As an example, this removal of attributes is particularly useful for checkboxes and radio buttons where `checked` attributes either must have the value `checked` or be absent. One can use a template containing the attribute gap `checked=[c]` and then plug `checked` into `c` whenever the field should be checked and simply omit the plug otherwise.

The code gaps are not filled in using the plug operation: Instead, when a template containing code gaps is shown, the code blocks in the code gaps are executed in document order. The resulting strings or templates are then inserted in place of the code. Because this does not happen until the template is shown, the code in code gaps can only access variables that are declared in the service class or locally within the code gap.

The plug operation is the only way of manipulating XML templates. Thus, our XML type is quite different from both the string output streams in Servlets and JSP and the explicit tree structures provided by JDOM. We exploit this to obtain a compact and efficient run-time representation, and as a foundation for performing our program analyses. The notion of code gaps allows us to directly emulate the JSP style of writing services, which is often convenient, but while still having the explicit notion of sessions and the guarantees provided by the program analyses.

In order to be able to perform the static analyses later, we need a simple restriction on the use of forms and input fields in the XML templates: In `input` and `button` elements we require syntactically that attributes named `type` and `multiple` cannot occur as attribute gaps in elements defining input fields. The same restriction holds for `name` attributes, unless the `type` attribute has value `submit` or `image`. In addition, we make a number of simple modifications of the documents being shown and of the field values being received:

- (1) In HTML and XHTML, lists, tables, and select menus are not allowed to have zero entries. However, it is often inconvenient to be required to perform special actions in those cases. Just before a document is shown, we therefore remove all occurrences of `<ul></ul>` and similar constructs. For select menus, we add a dummy option in case none are present.

- (2) If attempting to receive the selected value of a `select` menu that is not declared as `multiple` or the value of a `radio` button set, then, if no option is either pre-selected using `checked` or selected by the client, the null value is received instead of throwing a `ReceiveException` — even though no name–value pair is sent according to the XHTML 1.0/HTML 4.01 specification [31].
- (3) For `submit` and `image` fields, we change the corresponding returned name–value pair from `X=Y` to `submit=X`. This makes it easier to recognize the chosen submit button in case of graphical buttons.
- (4) For `submit` buttons, a single name–value pair is produced from the `name` and the `value` attributes. However, for graphical submit buttons, that is, fields with `type="image"`, HTML/XHTML produces two name–value pairs, `X.x` and `X.y` for the click coordinates, where `X` is the value of the `name` attribute, but the `value` attribute is ignored. To obtain a clean semantics, we ensure by patching the returned list of pairs that in every case, all three name–value pairs are produced. For instance, with graphical submit buttons we add a `submit=X` pair, and `submit.x` and `submit.y` contain the click coordinates. For normal submit buttons, `submit.x` and `submit.y` contain the value `-1`.

Clearly, these restrictions and modifications do not impose any practical limitations on the flexibility of the template mechanism. In fact, they serve as a convenience to the programmer since many special cases in XHTML need not be considered.

As for other Java types, casting is required when generic containers or methods are used. An XML template may be cast using the special syntax (`[[xml]] exp`). This is a promise from the JWIG programmer that any value that will ever be contained in `exp` at this point will be legal to use in all situations where `xml` is legal. If this is the case, the cast is said to be *valid*. At run-time, a `CastException` is thrown if the sets of gaps and input fields in `exp` do not match those in `xml`. If casting to a template with exactly one occurrence of a given field, then it is required that the actual values also has exactly one occurrence of that field—except for radio buttons, where multiple occurrences count as one. For the gaps, two properties must be satisfied: If the template being cast to has any gaps of a given name, then at least one such gap must also exist in the actual value; and, if there is a template gap in the template being cast to, then the actual value cannot contain any attribute gaps of that name. Note that this run-time check is not complete since it only considers gaps and input fields and not XHTML validity. However, if invalid XHTML is produced, it will eventually result in a `ValidateException` at a `show` statement.

Alternatively, the ordinary cast `(XML) exp` may be used. It generates a `CastException` if the actual type of `exp` is not `XML`, but no promises are made about the gaps or input fields.

Large JWIG applications may easily involve hundreds of template constants. For this reason, there is support for specifying these externally, instead of inlined in the program source. The construct `get url` loads the XML template located at `url` at run-time. This template can then later be modified and reloaded by the running service.

When the service is analyzed, the template constant referred to by the `get url` construct is loaded and treated as a constant in the analysis. The analysis is then of course only valid as long as the template is unchanged. However, validity will be




Fig. 6. The Jwig program translation process in our current implementation. The Jwig program is desugared into Java code which is compiled to class files. These class files are used both in the Web server and to perform the program analyses.

preserved if the template remains structurally the same. To obtain fresh security guarantees, it is simply required to reinvoke the program analyzer.

These features can also be used to support cooperation between the programmers and the Web page designers. For a first draft, the programmers can create some templates that have the correct structure but only a primitive design. While the program is being developed using these templates, the designers can work on improving them to give them a more sophisticated design. The program analyzer will ensure that the structure of gaps and fields is preserved.

In addition to the main features mentioned above, a session object contains a number of fields and methods that control the current interaction, such as, access control using HTTP authentication, cookies, environment variables, SSL encryption, and various timeout settings. The service object additionally contains a `checkpoint` method which serializes the shared data and stores it on disk. This is complemented by a `rollback` method which can be used in case of server crashes to improve robustness.

To summarize, we have now added special classes and language constructs to support session management, client interactions, and dynamic construction of XHTML documents. By themselves, we believe that these high-level language extensions aid development of Web services. The extensions may cause various exceptions: a `ValidateException` if one attempts to show an XML document which is not valid XHTML 1.0; a `ReceiveException` if trying to receive an input field that occurs an incompatible number of times; a `PlugException` if a plug operation fails because no gaps of the given name and type exist in the template; and a `CastException` if an illegal cast is performed because the gaps or fields do not match. In the following sections, we show that it is possible to statically check whether the first three kinds of exceptions can occur. This is possible only because of the program structure that the new language constructs enforce. For instance, it is far from obvious that similar guarantees could be given for Servlet or JSP services.

#### 2.4 The Jwig Program Translation Process

The steps involved in compiling, analyzing, and running a Jwig program are depicted in Figure 6. First, the special syntactic constructs of Jwig are translated into appropriate Java constructs by a simple source-to-source desugaring transformation. The resulting Java source files are compiled into Java class files as usual.




Fig. 7. Structure of the program analyzer.

These class files together with the accompanying externally specified XML template constants constitute the Web service. Of course, an implementation is not forced to have this structure: For instance, one could imagine a Jwig compiler that directly produces class files instead of going via Java code.

The analysis works on the class file level. When the analyzer is invoked, it is given a collection of class files to analyze. We call this collection the *application classes*; all others constitute the *non-application classes*. Exactly one of the application classes must be a subclass of `Service`. For efficiency reasons, the application classes can be just the few classes that actually constitute the Jwig service, not including all the standard Java classes that the program uses. Our analyses are designed to cope with this limited view of the program as an open system.

The soundness of the analyses that we describe in the following sections is based on a set of well-formedness assumptions:

- all invocation sites in the application classes must either always invoke methods in the application classes or always invoke methods in the non-application classes;
- no fields or methods of application classes are accessed by a non-application class;
- no XML operations are performed in non-application classes; and
- XML casts are always valid, according to the definition in the previous section.

These assumptions usually do not limit expressibility in practice. In some cases, the second assumption can be relaxed slightly, for instance if some method called from a non-application class does not modify any `String` or XML value that will ever reach other application class methods. This makes it possible to safely use callback mechanisms such as the `Comparator` interface. The assumption about casts is deliberately quite strong: As ordinary casts, XML casts provide a back-door to the programmer to bypass the static type system.

The structure of the program analyzer is shown in Figure 7. From the class files, we first generate *flow graphs*. From these, we generate *summary graphs*, which we analyze in three different ways corresponding to the properties mentioned in the previous section.

## 2.5 An Example Jwig Program

We will use the following Jwig program throughout the remaining sections to illustrate the various phases of the analysis. This admittedly rather artificial service applies most Jwig-specific language constructs:

```
import dk.brics.jwig.runtime.*;
```

```

public class Greetings extends Service {
    String greeting = null;

    public class Welcome extends Session {
        XML cover = [[ <html>
            <head><title>Welcome</title></head>
            <body bgcolor=[color]>
                <{
                    if (greeting==null)
                        return [[ <em>Hello World!</em> ]];
                    else
                        return [[ <b><[g]></b> ]] <[g=greeting];
                }>
                <[contents]>
            </body>
        </html> ]];

        XML getinput = [[ <form>Enter today's greeting:
            <input type="text" name="salutation">
            <input type="submit"></form> ]];

        XML message = [[ Welcome to <[what]>. ]];

        public void main() {
            XML h = cover[<[color="white", contents=message];
            if (greeting==null) {
                show cover[<[color="red", contents=getinput];
                greeting = receive salutation;
            }
            exit h[<[what=[[<b>BRICS</b>]]];
        }
    }
}

```

The first time the `Welcome` session is initiated, the client is prompted for a greeting text in one interaction, and then in the next interaction the greeting is shown together with a “Welcome to BRICS” message. For subsequent sessions, only the second interaction is performed.

### 3. FLOW GRAPH CONSTRUCTION

Given a Jwig program, we first construct an abstract flow graph as a basis for the subsequent data-flow analyses. The flow graph captures the flow of string and XML template values through the program and their uses in `show`, `plug`, and `receive` operations.

#### 3.1 Structure of Flow Graphs

The *nodes* in this graph correspond to abstract statements:

<code>x = exp;</code>	(assignment)
<code>show x;</code>	(client interaction)
<code>receive f;</code>	(receive field)



<b>receive[]</b> <i>f</i> ;	(receive field array)
<b>nop</b> ;	(no operation)

where *exp* denotes an expression of one of the following kinds:

<i>x</i>	(variable read)
" <i>str</i> "	(string constant)
[ <i>xml</i> ]	(XML template constant)
<i>x</i> <[ <i>g</i> = <i>y</i> ]	(plug operation)
<b>null</b>	(null value)
<b>anystring</b>	(arbitrary string)

and *x* and *y* are program variables, *g* is a gap name, *f* is a field name, *str* is a string constant, and *xml* is an XML template constant that does not contain any code gaps. All code gaps in the original Jwig program are expressed using normal gaps and plug operations in the flow graph, as will be explained in Section 3.3.

We assume that every variable occurring in the flow graph has a declared type: **STRING** representing strings, or **XML** representing XML templates. These types are extended to expressions as one would expect, and **null** has the special type **NULL**. Let  $EXP_{STRING}$  denote the expressions of type **STRING** or **NULL**, and  $EXP_{XML}$  denote those of type **XML** or **NULL**.

The assignment statement evaluates its expression and assigns the value to the given variable. The variable and the expression must have the same type. All flow graph variables are assumed to be declared with a global scope. Evaluating expressions cannot have side-effects. The argument to **show** statements is always of type **XML**. As described later, we model **receive** expressions from the Jwig program as pairs of statements, each consisting of a **receive** statement and an assignment. The **receive** and **receive[]** statements record program locations where input field values are received. The last kind of flow-graph statement, **nop**, is the no-operation statement which we use to model all operations that are irrelevant to our analyses, except for recording split and join points.

The expressions basically correspond to those in concrete Jwig programs, except **anystring** which is used to model standard library string operations where we do not know the exact result. In the plug operation, the first variable always has type **XML**, and the second has type **XML** or **STRING**.

Each node in the graph is assigned an *entry label* and an *exit label* as in [28], and additionally each XML template constant has a *template label*. All labels are assumed to be unique. The union of entry labels and exit labels constitute the *program points* of the program.

The graph has two kinds of edges: *flow edges* and *receive edges*. A flow edge models the flow of data values between the program points. Each edge has associated one source node and one destination node, and is labeled with a set of program variables indicating which values that are allowed to flow along that edge. A receive edge goes from a **receive** node to a **show** node. Its presence indicates that the control-flow of the program may lead from the corresponding **show** statement to the **receive** statement without reaching another **show** first. We use these edges to describe from which **show** statements the received field can originate.

### 3.2 Semantics of Flow Graphs

Formally, the semantics of a flow graph is defined by a constraint system. Let  $V$  be the set of variables that occur in the flow graph,  $XML$  be the set of all XML templates, and  $STRING$  be the set of all strings over the Unicode alphabet. Each program point  $\ell$  is associated an environment  $\mathcal{E}_\ell$ :

$$\mathcal{E}_\ell : V \rightarrow 2^{XML \cup STRING}$$

The entire set of environments forms a lattice ordered by pointwise set inclusion. For each node in the graph we generate a constraint. Let *entry* and *exit* denote the entry and exit labels of a given node. If the statement of the node is an assignment,  $x = exp$ , then the constraint is:

$$\mathcal{E}_{exit}(y) = \begin{cases} \widehat{\mathcal{E}}_{entry}(exp) & \text{if } x = y \\ \mathcal{E}_{entry}(y) & \text{if } x \neq y \end{cases}$$

For all other nodes, the constraint is:

$$\mathcal{E}_{exit} = \mathcal{E}_{entry}$$

The map  $\widehat{\mathcal{E}}_\ell : EXP \rightarrow 2^{XML \cup STRING}$  defines the semantics of flow graph expressions given an environment  $\mathcal{E}_\ell$ :

$$\widehat{\mathcal{E}}_\ell(exp) = \begin{cases} \mathcal{E}_\ell(x) & \text{if } exp = x \\ \{str\} & \text{if } exp = "str" \\ \{xml\} & \text{if } exp = \llbracket xml \rrbracket \\ \pi(\mathcal{E}_\ell(x), g, \mathcal{E}_\ell(y)) & \text{if } exp = x \llbracket g = y \rrbracket \\ \emptyset & \text{if } exp = \mathbf{null} \\ STRING & \text{if } exp = \mathbf{anystring} \end{cases}$$

The function  $\pi$  captures the meaning of the plug operation. Due to the previously mentioned type requirements, the first argument to  $\pi$  is always a set of XML template values. The function is defined by:

$$\pi(A, g, B) = \bigcup_{xml \in A, b \in B} \{\bar{\pi}(xml, g, b)\}$$

where  $\bar{\pi}$  is defined by induction in the XML template according to the definition in Section 2.3:

$$\bar{\pi}(xml, g, b) = \begin{cases} str & \text{if } xml = str \\ \langle name \bar{\pi}(attr, g, b) \rangle & \\ \bar{\pi}(xml', g, b) \langle /name \rangle & \text{if } xml = \langle name attr \rangle xml' \langle /name \rangle \\ b & \text{if } xml = \llbracket g \rrbracket \\ \langle \llbracket h \rrbracket \rangle & \text{if } xml = \langle \llbracket h \rrbracket \rangle \text{ and } h \neq g \\ \bar{\pi}(xml_1, g, b) \bar{\pi}(xml_2, g, b) & \text{if } xml = xml_1 xml_2 \end{cases}$$

$$\bar{\pi}(attr, g, b) = \begin{cases} \epsilon & \text{if } attr = \epsilon \\ name="str" & \text{if } attr = name="str" \\ name="b" & \text{if } attr = name=[g] \text{ and } b \in STRING \\ name=[h] & \text{if } attr = name=[h] \\ & \text{and } (h \neq g \vee b \in XML) \\ \bar{\pi}(attr_1, g, b) \bar{\pi}(attr_2, g, b) & \text{if } attr = attr_1 attr_2 \end{cases}$$

This defines plug as a substitution operation where template gaps may contain both strings and templates and string gaps may contain only strings.

For each flow edge from  $\ell$  to  $\ell'$  labeled with a variable  $x$  we add the following constraint:

$$\mathcal{E}_\ell(x) \subseteq \mathcal{E}_{\ell'}(x)$$

to model that the value of  $x$  may flow from  $\ell$  to  $\ell'$ .

We now define the semantics of the flow graph as the least solution to the constraint system. This is well-defined because all the constraints are continuous. Note that the environment lattice is not finite, but we do not need to actually compute the solution for any concrete flow graph.

In the following section we specify a translation from Jwig programs into flow graphs. In this translation, each **show** statement, plug expression, and **receive** expression occurring in the Jwig program has a corresponding node in the flow graph. Also, each operand of a **show** or plug operation has a corresponding variable in the flow graph. Correctness of such a translation is expressed as two requirements: 1) Let  $env$  be the least solution to the flow graph constraint system. If we observe the store of a Jwig program at either a **show** or a plug operation during some execution, then the value of each operand is contained in  $env_\ell(x)$  where  $\ell$  is the node corresponding to the Jwig operation and  $x$  is the variable corresponding to the operand. 2) If some session thread of an execution of the Jwig program passes a **show** statement and later a **receive** expression without passing any other **show** statement in between, then the flow graph contains a receive edge from the node corresponding to the **receive** expression to the node corresponding to the **show** statement.

### 3.3 From Jwig Programs to Flow Graphs

The flow graph must capture the flow of string and XML values in the original Jwig program. Compared to `<bigwig>` and the flow graph construction in [8] this is substantially more involved due to the many language features of Java. We divide this data flow into three categories: 1) per-method flow of data in local variables, 2) data flow to and from field variables, and 3) flow of argument and return values for method invocations. Since local variables are guaranteed to be private to each method invocation, we model the first kind of data flow in a control-flow sensitive manner. With field variables, this cannot be done because they may be accessed by other concurrently running session threads, and because we are not able to distinguish between different instantiations of the same class. The second kind of data flow is therefore modeled in a control-flow insensitive manner.

The translation ignores variables whose type is not `String`, `XML`, or an array of any dimension of these two. For each of the two analyzed types, a unique flow graph

*pool variable* is created for representing all the values of that type that cannot be tracked by the analysis. Pooled values include those assigned to and from `Object` variables and arrays, and arguments and results of methods outside the application classes. We add an assignment of `anystring` to the pool variable of type `STRING` to be maximally pessimistic about the string operations in the non-application classes. Something similar is not done for the `XML` type since we have assumed that `XML` values are produced only inside the analyzed classes.

In addition to capturing data flow, the flow graph must contain receive edges that reflect the correspondence between show and receive operations in the Jwig program. This requires knowledge of the control flow in addition to the data flow.

Before the actual translation into flow graphs begin, each code gap is converted to a template gap with a unique name, and the code inside the gap is moved to a new method in the service class.

The whole translation of Jwig programs into flow graphs proceeds through a sequence of phases, as described in the following. Since Jwig includes the entire Java language we attempt to give a brief overview of the translation rather than explain all its details. We claim that this translation is correct according to the definition in the previous section, however it is beyond the scope of this article to state the proof.

*1. Individual methods.* In the first phase, each method in the application classes is translated individually into a flow graph. Each statement produces one or more nodes, and edges are added to reflect the control flow inside the method. Each edge is labeled with all local variables of the method. Nested expressions are flattened using fresh local variables and assignments. The special Jwig operations are translated into the corresponding flow graph statement or expression, and all irrelevant operations are modeled with `nop` nodes. Each `receive` expression is translated into two nodes: a `receive` node and an assignment of `anystring`, since we need to model the locations of these operations but have no knowledge of the values being received. The control structures, `if`, `switch`, `for`, etc., are modeled with `nop` nodes and flow edges, while ignoring the results of the branch conditions. XML casts are translated into XML template constants. This is sound since we have assumed that all casts are valid. Figure 8 shows the flow graph for the `main` method of the example Jwig program in Section 2.5.

*2. Code gaps.* As mentioned, each code gap has been converted into a template gap whose name uniquely identifies the method containing its code. Before every `show` statement, a sequence of method calls and plug operations is inserted to ensure that all code gaps that occur in the program are executed and that their results are inserted. To handle code gaps that generate templates that themselves contain code gaps, an extra flow edge is added from the end of the plug sequence to the start. The analysis in [8] does not support code gaps.

*3. Method invocations.* The methods are combined *monovariantly*: each method is represented only once in the flow graph for the whole program. This means that the subsequent analyses that build on the flow graphs also are monovariant. To estimate which methods can be called at each invocation site, a call graph of the Jwig program is constructed using a *class hierarchy analysis* (CHA) [15; 35]. This

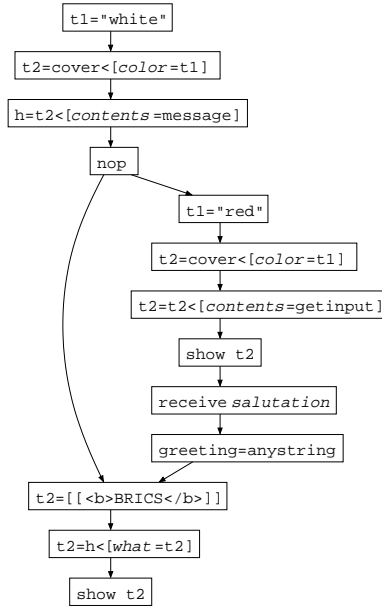


Fig. 8. Flow graph for the `main` method after Phase 1. All edges are here implicitly labeled with the set of variables  $\{h, t1, t2\}$ .

gives us for each invocation site a set of possible target methods. Of course, other call graph analyses could be applied instead, but CHA has proven to be fast and sufficiently precise for these purposes.

For each method invocation, we need to transfer the arguments and the return value, and to add flow edges to and from the possible target methods. The caller method uses a set of special local variables for collecting the arguments and the return value. We first insert a chain of assignments to these caller argument variables. Then we branch for each possible target method and add a chain of assignments from the caller argument variables to the target method parameters, followed by a flow edge to the target method entry point. Similarly, we add flow edges from the method exit points and transfer the return value via a caller result variable. For target methods in non-application classes, we use the pool variables in place of the argument and return value variables of the target method.

Figure 9 shows an example of a flow graph for a method invocation where the CHA has determined that there are two possible targets.

4. *Exceptions.* For every `try-catch-finally` construct, we add edges from all nodes corresponding to statements in the `try` block to the entry node of the `catch` block. These edges are labeled with all local variables of the method. This ensures that the data flow for local variables is modeled correctly. Adding edges from all nodes of the `try` blocks may seem to cause imprecision. A more complex analysis would only add edges from the nodes that correspond to statements that actually may throw exceptions that are caught by the `catch` block. However, our simple approach appears to be sufficiently precise in practice.

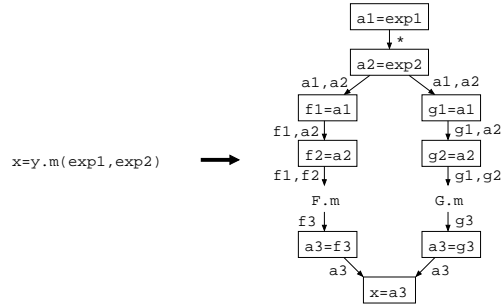


Fig. 9. Modeling method invocations. Assuming that the invocation of `m` in the expression on the left may lead to the classes `F` or `G`, the flow graph on the right is generated where `F.m` and `G.m` are the flow graphs for the target methods. First, the actual parameters are evaluated, then they are passed to the formal parameters for each method and the method bodies are processed, and finally, the return value is collected and the flow is merged. The `*` label denotes all local variables in the caller method.

In order to be able to set up the receive edges in a later phase, we also need to capture the intraprocedural control flow of exceptions. For this purpose, we add a special *drain node* for each method. For each statement, we add a flow edge with an empty label to the drain node of its method. This represents the control flow for uncaught exceptions within each method. This flow may subsequently lead either to drain nodes for caller methods or to `catch` blocks. To model this, we use the CHA information: For each target method of an invocation site, an edge is added from the drain node of the target method to the drain node of the method containing the invocation site. If the invocation site is covered by an exception handler within the method, an extra edge is added from the drain node of the target method to the entry of the handler.

5. *Show and receive operations.* The preceding phases have set up flow edges representing all possible control flow in the program. This means that we at this point are able to infer the correspondence between `show` and `receive` nodes to create the receive edges by looking at the flow graph alone, without considering the original Jwig program.

We treat the entry points of `main` methods of session classes as if they were `show` statements that show a document with a single form containing no input fields. This models the fact that no input fields can be read with `receive` until a document has been shown.

For each program point  $\ell$  in the flow graph we compute the set of `show` nodes that according to the flow edges may lead to  $\ell$  without passing another `show` node in between. This is done with a fixed point iteration, starting with empty sets at each program point. To each exit label of a `show` node  $n$ , we associate the singleton set containing just  $n$ . For all other nodes, the sets at the entry points are propagated to the exit points, and at join points, we take the union of the incoming sets. This is essentially a simple forward data-flow analysis on the flow graph. Since there are finitely many `show` nodes and we only apply monotone operations, the iteration eventually terminates. For each `receive` and `receive[]` node, we then add a receive edge to each node in the associated set.




Fig. 10. Shortcutting array updates. Assignments into arrays are modeled using weak update where all entries are merged.

6. *Arrays.* Array variables are translated into variables of their base type. An array is treated like a single entity whose possible values are the union of its entries. Construction of arrays using `new` is modeled with `null` values to reflect that they are initially empty.

An assignment to an array entry is modeled using weak updating [13] where the existing values of the array are merged with the new value. This is done by inserting two `nop` nodes around the assignment and adding an edge bypassing it labeled by the updated variable. This process is shown in Figure 10.

When one array variable is assigned to another, these variables become aliases. Such aliased variables are joined into one variable. This variable will be treated as a field variable and handled as described below, if at least one of its original variables was a field variable. This joining is similar to the technique used in [35].

7. *Field variables.* As mentioned, we model the use of field variables in a flow-insensitive manner where all instances of a given class are merged. This is done for each field simply by adding flow edges labeled by the name of the field from all its definitions to all its uses. To avoid constructing a quadratic number of edges, we add a dummy “ $x=x$ ” node to collect the definitions and the uses for each variable  $x$ .

In `<bigwig>`, a simpler and more restrictive approach was chosen: All global string variables were modeled with `anystring`, and for the global HTML variables, which correspond to the XML field variables in JWIG, the initializer expressions would dictate the types [8].

8. *Graph simplification.* Finally, we perform some reductions of the flow graph. This is not necessary for correctness or precision of the subsequent analyses, but as we show in Section 6.2, it substantially decreases the time and space requirements.

First, we remove all code that is unreachable from session entry points according to the flow edges. We ignore edges that originate from method returns since these edges do not indicate reachability.

Using a standard reaching definitions analysis on the flow graph [1; 28], we then find for each assignment all possible uses of that definition. This gives us a set of pairs of nodes where the first node is an assignment to some variable and the second node contains an expression which uses that variable. Once this information is obtained, we remove every flow edge and `nop` node in the graph, and then add new flow edges corresponding to the definition–use pairs. Each new edge is labeled with the single variable of the pair. Finally, a copy propagation optimization is

performed to compress chains of copying statements [1].

These transformations all preserve the data flow. In the resulting flow graphs, there are no `nop` nodes and all edges are labeled with a single variable, which is crucial for the performance of the subsequent analyses.

This construction of flow graphs for Jwig programs is correct in the sense defined in Section 3.2, both with and without the simplification phase.

### 3.4 Complexity

During the construction of the flow graph, we have performed two forward data-flow analyses on the intermediate graphs: one for setting up the receive edges in Phase 5 and the reaching definitions analysis in Phase 8. In the following sections, we will describe two more forward data-flow analyses on flow graphs. To bound the worst-case time requirements for all these analyses, we make some general observations. By implementing the analyses using a standard work-list algorithm rather than the chaotic iteration algorithm [28], the time can be bound by:

$$O\left(t \cdot \sum_{m \in \text{nodes}} |\text{var}(m)| \cdot h \cdot \sum_{m' \in \text{succ}(m)} |\text{var}(m')|\right)$$

where

- $t$  is the maximum cost of computing one binary least-upper-bound operation or one transfer function for a single variable;
- $\text{nodes}$  is the set of flow-graph nodes;
- $\text{var}(m)$  denotes the union of the labels of edges adjacent to the node  $m$ ;
- $h$  is the height of the lattice for a single variable; and
- $\text{succ}(m)$  is the set of successor nodes of  $m$ .

For each node  $m$ , an environment associates two lattice elements to each variable  $v$  in  $\text{var}(m)$ , one for the entry label and one for the exit label. Each can change at most  $h$  times. Because of the work-list, each change for an exit label can result in at most  $\sum_{m' \in \text{succ}(m)} |\text{var}(m')|$  binary least-upper-bound operations and transfer function computations, that is, one for each variable in each successor node, without any other environment changes for exit labels.

Phases 1-7 create at most  $O(n^2)$  flow-graph nodes and  $O(n^2)$  edges where  $n$  is the textual size of the program. The reason for the quadratic increase in the number of nodes is the encoding of argument transferring for method invocations in Phase 3 and the encoding of code gap execution in Phase 2.

The receive edge analysis in Phase 5 does not consider the variables at all. This is equivalent to setting  $\text{var}(m)$  to contain just one variable. Since there are at most  $O(n^2)$  edges,  $\sum_{m \in \text{nodes}} |\text{succ}(m)|$  is  $O(n^2)$ . The lattice height,  $h$ , is the number of `show` statements, which is  $O(n)$ , and the time  $t$  is  $O(n)$ . Therefore, this analysis runs in time  $O(n^4)$ .

The complexity of the reaching definitions analysis in Phase 8 can also be bound by the formula above. There are  $O(n)$  variables, the lattice height is  $O(n)$ , and the time  $t$  is  $O(n)$ . Again, since there are  $O(n^2)$  edges,  $\sum_{m \in \text{nodes}} |\text{succ}(m)|$  is  $O(n^2)$ . Together, we get that this analysis runs in time  $O(n^6)$ . Since the other phases of the flow-graph construction are linear in the size of the flow graph, the total



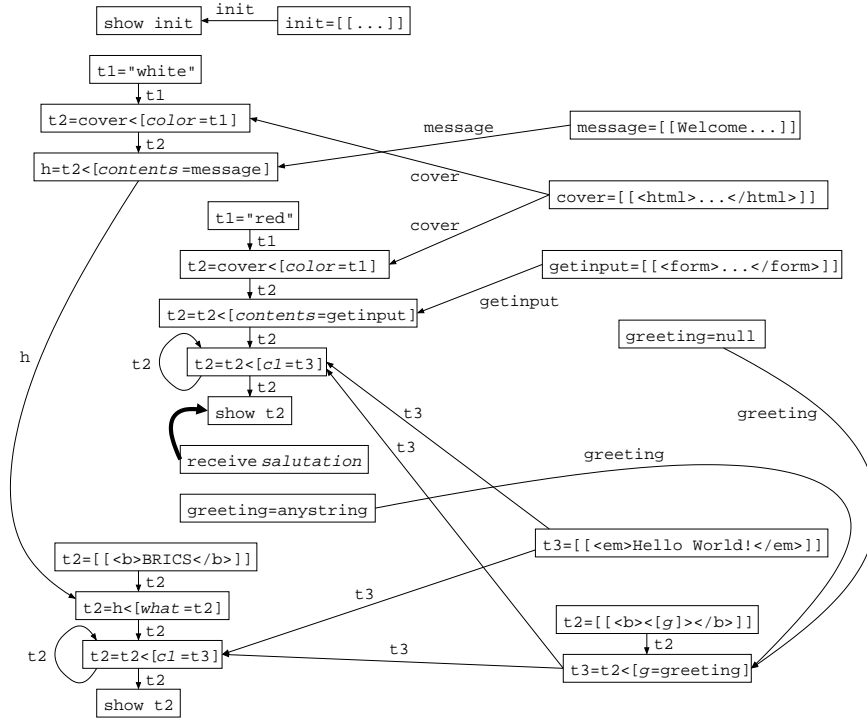


Fig. 11. Flow graph for the Jwig example program.

construction of the flow graph of a given Jwig program requires worst-case  $O(n^6)$  time. As shown in Section 6.2, this bound is rarely encountered in practice.

After the simplification phase, an extra property is satisfied: Since all flow edges are definition–use edges,  $|var(m)|$  is  $O(1)$  for all nodes  $m$ . Since the flow graph still contains only  $O(n^2)$  nodes and edges, the formula above then reduces to:

$$O(n^2 \cdot h \cdot t)$$

This will be used later to estimate the complexities of the remaining analyses.

### 3.5 Flow Graph for the Example

Figure 11 shows the flow graph that is generated for the Jwig program from Section 2.5. The left part of the graph corresponds to the `main` method, the top right part is the initialization of the field variables, and the bottom right part corresponds to the code in the code gap. The thin edges are flow edges, and the single thick edge is a receive edge. The `show` node in the top left corresponds to the entry point of the session. The `init` template is a simple valid XHTML document with a form that contains no fields. This models the fact that no input values are receivable when session threads are initiated. Note that edges in the part corresponding to the `main` method have changed compared to Figure 8 because of the graph simplification phase.

#### 4. SUMMARY GRAPH ANALYSIS

To statically verify that a given Jwig program will never throw any of the special Jwig exceptions we perform a *summary graph analysis* based on the flow graph which contains all the information we need from the original program. Summary graphs model how templates are constructed and used at run-time. This analysis depends on a preliminary *string analysis* that for each string expression finds a regular language that approximates the set of strings it may evaluate to. For each analysis we define a lattice expressing an abstraction of the data values along with a corresponding abstract semantics of the expressions and statements, and then apply standard data-flow analysis techniques to find the least solutions.

##### 4.1 String Analysis

Given a flow graph of a Jwig program, we must statically model which strings are constructed and used at run-time. In [8] the corresponding analysis is mixed into the summary graph analysis. Separating these analyses leads to a simpler specification and implementation without damaging the analysis precision. We describe here a rather simple analysis which is adequate for all our benchmarks. However, it should be clear that a more precise string analysis capturing relevant string operations easily could be applied instead, as explained in Section 7.

We first define a *string environment* lattice:

$$SE = Y \rightarrow REG$$

where  $Y$  is the set of string variables that occur in the program and  $REG$  is the family of regular languages over the Unicode alphabet. We choose regular languages for modeling string sets because they fit elegantly into the validity analysis algorithm in Section 5.5. The ordering on  $REG$  is language inclusion and  $SE$  inherits this ordering pointwise. We compute an element of this lattice for every program point with a forward data-flow analysis using standard techniques, such as the monotone frameworks of [28; 20]: For every statement that can appear in the flow graph we define a monotone transfer function  $SE \rightarrow SE$  and then compute the least fixed point by iteration. First, for each flow-graph string expression we define its abstract denotation by extending every environment map  $\Sigma \in SE$  from variables to string expressions,  $\widehat{\Sigma} : EXP\_STRING \rightarrow REG$ :

$$\widehat{\Sigma}(exp) = \begin{cases} \Sigma(x) & \text{if } exp = x, \\ \{str\} & \text{if } exp = "str" \\ U^* & \text{if } exp = \mathbf{anystring} \\ \emptyset & \text{if } exp = \mathbf{null} \end{cases}$$

where  $U$  denotes the Unicode alphabet.

For every string assignment statement  $x = exp$ ; the transfer function is defined by:

$$\Sigma \mapsto \Sigma[x \mapsto \widehat{\Sigma}(exp)]$$

that is, the string environment is updated for  $x$  to the environment value of  $exp$ . Clearly, this is a monotone operation. For all other statements the transfer function is the identity function, since they do not modify any string variables. The lattice

is not finite, but by observing that the only languages that occur are either total or subsets of the finitely many string constants that appear in the program, termination of the fixed point iteration is ensured. A more advanced string analyses, for instance one modeling concatenations more precisely, would require widening for ensuring termination.

The worst-case complexity of this analysis can be estimated by the formula from the previous section. By the observation above, we only use a part of the lattice. This part has height  $O(n)$  since there are at most  $O(n)$  string constants in the program. The time  $t$  for performing a least-upper-bound or transfer function computation is  $O(n)$ . Thus, this particularly simple string analysis runs in worst-case time  $O(n^4)$  where  $n$  is the size of the original Jwig program.

The result of this analysis is for each program point  $\ell$  a map:  $string_\ell : Y \rightarrow REG$ . This analysis is correct in the following sense: For any execution of the program, any program point  $\ell$ , and any string variable  $x$ , the regular language  $string_\ell(x)$  will always contain the value of  $x$  at  $\ell$ . That is, the analysis result is a conservative upper approximation of the string flow.

## 4.2 Summary Graphs

Like the string analysis, the summary graph analysis fits into standard data-flow frameworks, but it uses a significantly more complex lattice which we define in the following. Let  $X$ ,  $G$ , and  $N$  be respectively the sets of template variables, gap names, and template labels that occur in the program. A *summary graph*  $SG$  is a finite representation of a set of XML documents defined as:

$$SG = (R, T, S, P)$$

where

- $R \subseteq N$  is a set of *root nodes*,
- $T \subseteq N \times G \times N$  is a set of *template edges*,
- $S : N \times G \rightarrow REG$  is a *string edge* map, and
- $P : N \times G \rightarrow 2^{\{OPEN, CLOSED\}}$  is a *gap presence map*.

Intuitively, the language  $\mathcal{L}(SG)$  of a summary graph  $SG$  is the set of XML documents that can be obtained by unfolding its templates, starting from a root node and plugging templates and strings into gaps according to the edges. Assume that  $t : N \rightarrow xml$  maps every template label to the associated template constant. The presence of a template edge  $(n_1, g, n_2) \in T$  informally means that the  $t(n_2)$  template may be plugged into the  $g$  gaps in  $t(n_1)$ , and a string edge  $S(n, g) = L$  means that every string in the regular language  $L$  may be plugged into the  $g$  gaps in  $t(n)$ . The gap presence map,  $P$ , specifies for each gap named  $g$  originating from the template constant with label  $n$  whether or not  $g$  gaps may remain open in  $t(n)$ . The value OPEN means that the gaps may be open, and CLOSED means that they may be closed or never have occurred. Recall from Section 2.3 that attributes with gaps that remain open are removed when shown. An attribute gap labeled OPEN thus means that the attribute may be removed.

This unfolding is explained more precisely with the following formalization:

$$unfold(SG) = \{d \in XML \mid \exists r \in R : SG, r \vdash t(r) \Rightarrow d \text{ where } SG = (R, T, S, P)\}$$

The *unfolding relation*,  $\Rightarrow$ , is defined by induction in the structure of the XML template. For the parts that do not involve gaps the definition is a simple recursive traversal:

$$\frac{}{SG, n \vdash str \Rightarrow str}$$

$$\frac{SG, n \vdash xml_1 \Rightarrow xml'_1 \quad SG, n \vdash xml_2 \Rightarrow xml'_2}{SG, n \vdash xml_1 xml_2 \Rightarrow xml'_1 xml'_2}$$

$$\frac{SG, n \vdash atts \Rightarrow atts' \quad SG, n \vdash xml \Rightarrow xml'}{SG, n \vdash \langle name atts \rangle xml \langle /name \rangle \Rightarrow \langle name atts' \rangle xml' \langle /name \rangle}$$

$$\frac{}{SG, n \vdash \epsilon \Rightarrow \epsilon}$$

$$\frac{}{SG, n \vdash name="str" \Rightarrow name="str"}$$

$$\frac{SG, n \vdash atts_1 \Rightarrow atts'_1 \quad SG, n \vdash atts_2 \Rightarrow atts'_2}{SG, n \vdash atts_1 atts_2 \Rightarrow atts'_1 atts'_2}$$

There is no unfolding for code gaps since they have already been reduced to template gaps in the flow graph construction. For template gaps we unfold according to the string edges and template edges and check whether the gap may be open:

$$\frac{str \in S(n, g)}{(R, T, S, P), n \vdash \langle [g] \rangle \Rightarrow str}$$

$$\frac{(n, g, m) \in T(n, g) \quad (R, T, S, P), n \vdash t(m) \Rightarrow xml}{(R, T, S, P), n \vdash \langle [g] \rangle \Rightarrow xml}$$

$$\frac{OPEN \in P(n, g)}{(R, T, S, P), n \vdash \langle [g] \rangle \Rightarrow \langle [g] \rangle}$$

For attribute gaps we unfold according to the string edges, and check whether the gap may be open:

$$\frac{str \in S(n, g)}{(R, T, S, P), n \vdash name=[g] \Rightarrow name="str"}$$

$$\frac{OPEN \in P(n, g)}{(R, T, S, P), n \vdash name=[g] \Rightarrow name=[g]}$$

The following function *close* is used on the unfolded templates to plug the empty string into remaining template gaps and remove all attributes with gap values:

$$close(xml) = \begin{cases} \langle name close(atts) \rangle & \\ \quad close(xml') \langle /name \rangle & \text{if } xml = \langle name atts \rangle xml' \langle /name \rangle \\ \epsilon & \text{if } xml = \langle [g] \rangle \\ \quad close(xml_1) close(xml_2) & \text{if } xml = xml_1 xml_2 \\ xml & \text{otherwise} \end{cases}$$

$$close(atts) = \begin{cases} close(atts_1) \ close(atts_2) & \text{if } atts = atts_1 \ atts_2 \\ \epsilon & \text{if } atts = name=[g] \\ atts & \text{otherwise} \end{cases}$$

We now define the language of a summary graph by:

$$\mathcal{L}(SG) = \{close(d) \in XML \mid d \in unfold(SG)\}$$

Compared to the summary graphs in [8] this definition differs in the following ways: First of all, the gap presence map is added. The old algorithm worked under the assumption that all incoming branches to join points in the flow graph would agree on which gaps were open. This was achieved using a simple preliminary program transformation that would convert the “implicit  $\epsilon$ -plugs” of `<bigwig>` [5] into explicit ones using the information from the DynDoc type system [32]. Since JWIG does not inherit this implicit-plug feature from the `<bigwig>` design nor uses a DynDoc-like type system we have added the gap presence map. This map contains the information from the “gap track analysis” in [8], but in addition to finding gaps that *may* be open it also tracks *must* information which we need to verify the use of plug operations later.

Secondly, the present definition is more flexible in that it allows strings to be plugged into template gaps. In [8], template gaps were completely separated from attribute gaps. Thirdly, we generalize the flat string lattice to full regular languages allowing us to potentially capture many string operations.

The summary graph abstraction has evolved through experiments during our previous work on `<bigwig>`. We claim that it is in a finely tuned balance between expressibility and complexity. First, summary graphs are strictly more expressive than e.g. the regular expression types of XDuce [18]. These types can be systematically encoded as summary graphs by using nodes with templates containing only gaps to represent the states of NFAs. Summary graphs are more expressive, since they can similarly remember information about the paths from the roots to given nodes. Another approach would be to use recursive types, however these are even less expressive than regular expression types. On the other hand, any summary graph may be captured as a regular tree language [14]. However, this strictly richer formalism has a high computational complexity compared to summary graphs. Thus, summary graphs are just expressive enough to capture the intricacies of normal control flow in programs and yet sufficiently tractable to allow efficient analysis.

Note that our analysis is *monovariant*, in the sense that each template constant is only represented once. It is possible to perform a more expensive analysis that duplicates summary graph nodes according to some criteria, but we have not yet encountered the need. On the other hand, our analysis is *polyvariant* in XML element constructors, since these are analyzed separately for each occurrence in the templates.

Let  $\mathcal{G}$  be the set of all summary graphs. This set is a lattice where the ordering is defined as one would expect:

$$\begin{aligned} (R_1, T_1, S_1, P_1) \sqsubseteq (R_2, T_2, S_2, P_2) &\Leftrightarrow \\ R_1 \subseteq R_2 \wedge T_1 \subseteq T_2 \wedge \\ \forall n \in N, g \in G : S_1(n, g) \subseteq S_2(n, g) \wedge P_1(n, g) \subseteq P_2(n, g) \end{aligned}$$




Fig. 12. A summary graph whose language is a set of XML documents, each containing a `ul` list with zero or more text items and a `class` attribute. The node on the left is a root, the dashed lines show the gap presence for each gap, and  $L$  denotes some set of strings.

This ordering respects language inclusion: If  $SG_1 \sqsubseteq SG_2$ , then  $\mathcal{L}(SG_1) \subseteq \mathcal{L}(SG_2)$ .

Figure 12 shows an example summary graph consisting of two nodes, a single template edge, and two string edges. The language of this summary graph is the set of XML documents that consist of `ul` elements with a `class="large"` attribute and zero or more `li` items containing some text from the language  $L$ . Note that the gap presence of the `items` gap in the template of the left node contains `OPEN`, so the empty template may be plugged in which corresponds to the case where the list is empty.

### 4.3 Constructing Summary Graphs

At every program point  $\ell$  in the flow graph, each template variable  $x \in X$  is associated a summary graph, as modeled by the *summary graph environment* lattice:

$$SGE = X \rightarrow \mathcal{G}$$

which inherits its structure pointwise from  $\mathcal{G}$ . We compute an element of the lattice for every program point using yet another forward data-flow analysis. Let  $[[xml]]_n$  mean the template constant labeled  $n$ , and let  $tgaps(n)$  and  $agaps(n)$  be the sets of template gap names and attribute gap names, respectively, that occur in the template constant labeled  $n$ . Let  $gaps(n) = tgaps(n) \cup agaps(n)$ . Given an environment lattice element  $\Delta \in SGE$  we define an abstract denotation for template expressions,  $\widehat{\Delta} : EXP_{XML} \rightarrow \mathcal{G}$ :

$$\widehat{\Delta}(exp) = \begin{cases} \Delta(x) & \text{if } exp = x \\ const(gaps(n), n) & \text{if } exp = [[xml]]_n \\ tplug(\Delta(x), g, \Delta(y)) & \text{if } exp = x <[g = y] \text{ and } y \text{ has type XML} \\ splug(\Delta(x), g, string_\ell(y)) & \text{if } exp = x <[g = y] \text{ and } y \text{ has type STRING} \\ (\emptyset, \emptyset, \perp, \perp) & \text{if } exp = \mathbf{null} \end{cases}$$

where  $\perp = \lambda m. \lambda h. \emptyset$  and the auxiliary functions are:

$$const(A, n) = (\{n\}, \emptyset, \perp, \lambda m. \lambda h. \text{if } n = m \wedge h \in A \text{ then } \{\text{OPEN}\} \text{ else } \{\text{CLOSED}\})$$

$$tplug((R_1, T_1, S_1, P_1), g, (R_2, T_2, S_2, P_2)) =$$

$$\begin{aligned}
 & (R_1, \\
 & T_1 \cup T_2 \cup \{(n, g, m) \mid \text{OPEN} \in P_1(n)(g) \wedge m \in R_2\}, \\
 & \lambda m. \lambda h. S_1(m)(h) \cup S_2(m)(h), \\
 & \lambda m. \lambda h. \text{if } g=h \text{ then } P_2(m)(h) \text{ else } \text{merge}(P_1(m)(h), P_2(m)(h)))
 \end{aligned}$$

$$\text{merge}(p_1, p_2) = \text{if } p_1 = \{\text{OPEN}\} \vee p_2 = \{\text{OPEN}\} \text{ then } \{\text{OPEN}\} \text{ else } p_1 \cup p_2$$

$$\begin{aligned}
 \text{splug}((R, T, S, P), g, L) = \\
 & (R, \\
 & T, \\
 & \lambda m. \lambda h. \text{if } g=h \wedge \text{OPEN} \in P(m)(h) \text{ then } S(m)(h) \cup L \text{ else } S(m)(h), \\
 & \lambda m. \lambda h. \text{if } g=h \text{ then } \{\text{CLOSED}\} \text{ else } P(m)(h))
 \end{aligned}$$

For template constants, we look up the set of gaps that appear and construct a simple summary graph with one root and no edges. The *tplug* function models plug operations where the second operand is a template expression. It finds the summary graphs for the two sub-expressions and combines them as follows: The roots are those of the first graph since it represents the outermost template. The template edges become the union of those in the two graphs plus a new edge from each node that may have open gaps of the given name to each root in the second graph. The string edge sets are simply joined without adding new information. For the gaps that are plugged into, we take the gap presence information from the second graph. For the other gaps we use the *merge* function which sets gaps to “definitely open” if they are so in one of the graphs and otherwise takes the least upper bound. The *splug* function models plug operations where the second operand is a string expression. It adds the set of strings obtained by the string analysis for the string expression to the appropriate string edge, and then marks the gap as “definitely closed”. The `null` constant is modeled by the empty set of documents. Attempts to plug or show a null template yield null dereference exceptions at run-time, and we do not wish to perform a specific null analysis.

Having defined the analysis of expressions we can now define transfer functions for the statements. As for the other data-flow analysis, only assignments are interesting. For every XML assignment statement  $x = \text{exp}$ ; the transfer function is defined by:

$$\Delta \mapsto \Delta[x \mapsto \widehat{\Delta}(\text{exp})]$$

and for all other statements the transfer function is the identity function.

By inspecting the *tplug*, *merge*, and *splug* functions it is clear that the transfer function is always monotone. The lattice *SGE* is not finite, but analysis termination is ensured by the following observation: For any program, all summary graph components are finite, except *REG*. However, the string analysis produces only a finite number of regular languages, and we here use at most all possible unions of these. So, only a finite part of *SGE* is ever used.

The result of this analysis is for each program point  $\ell$  a map:

$$\text{summary}_\ell : \text{EXP}_{\text{XML}} \rightarrow \mathcal{G}$$

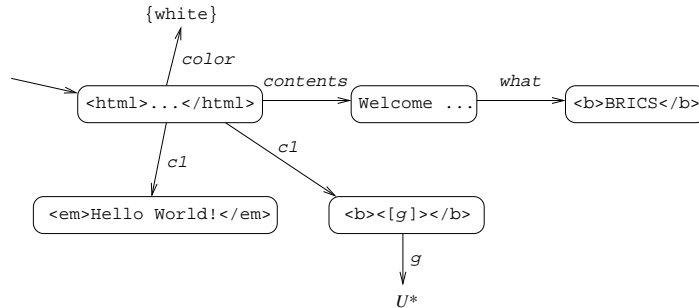
This analysis is conservative like the string analyses, that is, it is sound but not complete: For any execution of the program, any program point  $\ell$ , and any XML

expression  $exp$ , the set of XML documents  $\mathcal{L}(summary_\ell(exp))$  will always contain the value of  $exp$  at  $\ell$ .

The worst-case complexity of this analysis can also be estimated using the formula from Section 3.4. The lattice height is the sum of the heights of the four summary graph components. The node set  $N$  and the gap name set  $G$  both have size  $O(n)$ , again where  $n$  is the size of the original Jwig program. The height of the root node component is thus  $O(n)$ . For each template edge  $(n, g, n')$  which is created during the analysis,  $(n, g)$  determines a specific gap in a specific template in the original Jwig program. Since there can be at most  $O(n)$  of these, we can at most construct  $O(n^2)$  template edges. Similarly, for the string edge and the gap presence map components, all but  $O(n)$  pairs of elements from  $N$  and  $G$  are mapped to a fixed element. For the string analysis, we have argued that the height of the used part of the string lattice is  $O(n)$ , so the string edge component has height  $O(n^2)$ . The co-domain of the gap presence map has constant size, so this component has height  $O(n)$ . In total, the height  $h$  of the summary graph lattice is  $O(n^2)$ . For the same reasons, the sizes of the summary graphs that are constructed are also at most  $O(n^2)$ . All operations on summary graphs are linear in their sizes, so the time  $t$  for computing a summary graph operation is  $O(n^2)$ . Inserting this in the formula gives that the summary graph construction runs in time  $O(n^6)$  in the size of the program. Note that without the flow-graph simplification phase, the formula would have given  $O(n^8)$  instead of  $O(n^6)$ .

#### 4.4 Summary Graphs for the Example

For the Jwig example program from Section 2.5, the following summary graph is generated for the `exit` statement in the `main` method:



Implicitly in this illustration, the gap presence map maps everything to  $\{\text{CLOSED}\}$ , and the string edge map maps to the empty language by default. Because of the simple flow in the example program, the language of this summary graph is precisely the set of XML documents that may occur at run-time. In general, the summary graphs are conservative since they may denote languages that are too large. This means that the subsequent analyses can be sound but not complete.

## 5. PROVIDING STATIC GUARANTEES

The remaining analyses are independent of both the original Jwig program and the flow graphs. All the relevant information is at this stage contained in the inferred summary graphs. This is a modular approach where the “front-end” and



“back-end” analyses may be improved independently of each other. Also, summary graphs provide a good context for giving intuitive error messages.

### 5.1 Plug Analysis

We first validate *plug consistency* of the program, meaning that gaps are always present when subjected to the plug operation and that XML templates are never plugged into attribute gaps. This information is extracted from the summary graph of the template being plugged into.

In the earlier works [32] a similar check was performed directly on the flow graphs. Our new approach has the same precision, even though it relies exclusively on the summary graphs. Furthermore, we no longer require the flow graph to agree on the gap information for all incoming branches in the join points, as mentioned in Section 4.2.

For a specific plug operation  $x \llbracket g = y \rrbracket$  at a program point  $\ell$ , consider the summary graph  $summary_\ell(x) = (R, T, S, P)$  given by the data-flow analysis described in the previous section. To perform the required checks we need to compute a map:

$$GP : (N \rightarrow 2^G) \rightarrow 2^{\{\text{OPEN}, \text{CLOSED}\}}$$

indicating for some gap occurrence function, which is either *tgaps* or *agaps*, whether any  $g$  gaps of the given kind are present (**OPEN**) or absent (**CLOSED**) in the values represented by the summary graph for  $x$ .

Given such a map, we express consistency of the plug operation as follows. Compute  $t = GP(tgaps)$  and  $a = GP(agaps)$ . These values respectively describe the occurrence of template gaps and attribute gaps named  $g$  in the possible values of  $x$  at  $\ell$ . The plug operation is consistent if and only if the following condition is satisfied:

$$\begin{aligned} t = \{\text{OPEN}\} \vee a = \{\text{OPEN}\} & \quad \text{if } y \text{ has type STRING, and} \\ t = \{\text{OPEN}\} \wedge a = \{\text{CLOSED}\} & \quad \text{if } y \text{ has type XML.} \end{aligned}$$

This captures the requirement that string plug operations are allowed on all gaps that are present, while template plug operations only are possible for template gaps.

To calculate the  $GP$  map, we perform a fixed point iteration on the summary graph. We calculate the map:

$$\Gamma : N \times (N \rightarrow 2^G) \rightarrow 2^{\{\text{OPEN}, \text{CLOSED}\}}$$

as the least solution to the constraint system that consists of the following two constraints for each  $n \in N$  and  $M \in \{tgaps, agaps\}$ :

$$\begin{aligned} (g \in M(n) \wedge \text{OPEN} \in P(n)(g)) \vee \\ (\exists h \in tgaps(n) : \exists(n, h, M) \in T : \text{OPEN} \in \Gamma(m, M)) & \Rightarrow \text{OPEN} \in \Gamma(n, M) \end{aligned}$$

$$\begin{aligned} (g \notin M(n) \vee \text{CLOSED} \in P(n)(g)) \wedge \\ (\forall h \in tgaps(n) : \exists(n, h, M) \in T : \text{CLOSED} \in \Gamma(m, M) \\ \vee ((h \neq g \vee M = agaps) \wedge \text{OPEN} \in P(n)(h)) \\ \vee S(n, h) \neq \emptyset) & \Rightarrow \text{CLOSED} \in \Gamma(n, M) \end{aligned}$$

The value of  $\Gamma(n, M)$  contains **OPEN** if there exists an unfolding of the summary graph, starting from  $n$ , which contains an open  $g$  gap of the type specified by

$M$ . Similarly, it contains **CLOSED** if there exists an unfolding which does *not* contain any open  $g$  gaps of the right type. The first constraint specifies under which conditions **OPEN** can be included in  $\Gamma(n, M)$ : Either a potentially open  $g$  gap of the right type is present in the template of the  $n$  node itself, or there is a gap with a template edge to a node where such a gap may occur. The second constraint specifies the conditions for **CLOSED** to be contained in  $\Gamma(n, M)$ : If there are  $g$  gaps of the right type in the template of the  $n$  node itself, then these gaps must be potentially closed, and furthermore, for all template gaps in the template of  $n$ , either something may be plugged in where there is no open  $g$  gap of the right type, or the template gap may still be open and at the same time have a different name or a different type than the ones we are looking for.

Again, the involved lattice is finite and the constraints are monotone, so the least fixed point is well-defined and the iteration terminates. The  $GP$  map is then defined pointwise as the union of  $\Gamma$  applied to the roots of the summary graph:

$$GP(M) = \bigcup_{n \in R} \Gamma(n, M)$$

As mentioned, the summary graphs that are constructed are conservative with respect to the actual values that appear at run-time. However, the plug analysis introduces no new imprecision, that is, this analysis is both sound and complete with respect to the summary graphs: It determines that a given plug operation cannot fail if and only if for every value in  $unfold(summary_\ell(x))$ , the plug operation does not fail. In other words, if the plug analysis detects no errors, it is guaranteed that no `PlugException` will ever be thrown when running the program.

For a specific plug operation, the summary graph can be traversed at most  $2|N|$  times before reaching the fixed point. Each traversal takes time  $O(n^2)$ , where  $n$  is the size of the JWIG program, because the summary graph can have at most  $O(n)$  nodes and  $O(n^2)$  edges. This is done for each of the  $O(n)$  plug operations, so in total, the plug analysis takes time  $O(n^4)$ .

## 5.2 Receive Analysis

We now validate *receive consistency*, meaning that `receive` and `receive[]` operations always succeed. For the single-string variant, `receive`, it must be the case that for all program executions, the last document being shown before the receive operation contained exactly one field of the given name. Also, there must have been at least one show operation between the initiation of the session thread and the receive operation. If these properties are satisfied, it is guaranteed that no `ReceiveException` will ever be thrown when running the program.

The array variant, `receive[]`, always succeeds, so technically, we do not have to analyze those operations. However, we choose to consider it as an error if we are able to detect that for a given `receive[]` operation, there are no possibility of ever receiving other than the empty array. This is to follow the spirit of Java where, for instance, it is considered a compile-time error to specify a cast operation that is guaranteed to fail for all executions.


In case the name of the field is either `submit`, `submit.x`, or `submit.y`, then we know that it comes from a submit button or image. As described in Section 2.3, exactly one value is then always generated. That is, in these cases, both `receive`

and `receive[]` always succeed. For the remainder of this section, we thus assume that the field name is not one among those three.

Given a receive operation, we need to count the number of occurrences of input fields of the given name that may appear in every document sent to the client in an associated `show` operation. For a concrete XHTML/HTML document, this information is defined by Section 17.13.2 in [31]. For a running Jwig program, a conservative approximation of the information can be extracted from the receive edges in the flow graph and the summary graphs of the associated `show` operations.

Compared to the field analysis in `<bigwig>` [32], this situation differs in a number of ways: 1) The present analysis works on summary graphs rather than on flow graphs. 2) In the old analysis, the plug and receive analyses were combined. We separate them into two independent analyses without losing any precision. 3) In `<bigwig>`, a `form` is always inserted automatically around the entire document body. That precludes documents from having other forms for submitting input to other services. As described in Section 2.2, Jwig instead allows multiple forms by identifying those relevant to the session by the absence of an `action` attribute in the `form` element. 4) The notions of tuples and relations in [32] are in Jwig replaced by arrays and `receive[]` operations.

Again, we will define a constraint system for computing the desired information. This information is represented by a value of the following lattice,  $C$ :



The element 0 means that there are always zero occurrences of the field, 1 means that there is always exactly one occurrence, \* means that the number varies depending on the unfolding or that it is greater than one, ◇ represents one or more radio buttons, and ⊥ represents an unknown number. The constraint system applies two special monotone operators on  $C$ :  $\oplus$  for addition and  $\otimes$  for multiplication. These are defined as follows:

$\oplus$	$\perp$	0	◇	1	*
$\perp$	$\perp$	0	◇	1	*
0	0	0	◇	1	*
◇	◇	◇	◇	*	*
1	1	1	*	*	*
*	*	*	*	*	*

$\otimes$	$\perp$	0	◇	1	*
$\perp$	$\perp$	0	$\perp$	$\perp$	$\perp$
0	0	0	0	0	0
◇	$\perp$	0	◇	◇	*
1	$\perp$	0	◇	1	*
*	$\perp$	0	*	*	*

Assume that we are given a summary graph  $(R, T, S, P)$  corresponding to a specific `show` statement. Two special functions are used for extracting information about fields and gaps for an individual node in the summary graph:

$$count : N \rightarrow GFP$$

$$allforms : N \rightarrow 2^{GFP}$$

where  $GFP = (F \rightarrow C) \times (G \rightarrow C)$  shows the number of occurrences of fields and gaps in a specific form. The *allforms* function returns a set of such values, corresponding to the various forms that may appear, and *count* counts disregarding the *form* elements:

$$count(n) = (fcount(n, t(n)), gcount(n, t(n)))$$

$$allforms(n) = \bigcup_{k \in forms(t(n))} \{(fcount(n, k), gcount(n, k))\}$$

where

$$forms(xml) = \begin{cases} \{xml\} & \text{if } xml = \langle \mathbf{form} \ atts \rangle \ xml' \ \langle / \mathbf{form} \rangle \\ & \text{and } atts \text{ does not contain } \mathbf{action} \\ forms(xml') & \text{if } xml = \langle \mathbf{name} \ atts \rangle \ xml' \ \langle / \mathbf{name} \rangle \\ & \text{and } \mathbf{name} \neq \mathbf{form} \\ & \text{or } atts \text{ contains } \mathbf{action} \\ forms(xml_1) \cup forms(xml_2) & \text{if } xml = xml_1 \ xml_2 \\ \emptyset & \text{otherwise} \end{cases}$$

$$gcount(n, xml)(g) = \begin{cases} gcount(n, xml_1)(g) \\ \oplus gcount(n, xml_2)(g) & \text{if } xml = xml_1 \ xml_2 \\ gcount(n, xml')(g) \\ \oplus gcount(n, atts)(g) & \text{if } xml = \langle \mathbf{name} \ atts \rangle \ xml' \ \langle / \mathbf{name} \rangle \\ 1 & \text{if } xml = \langle [g] \rangle \\ & \text{and } P(n, g) = \{\mathbf{CLOSED}\} \\ * & \text{if } xml = \langle [g] \rangle \\ & \text{and } \mathbf{OPEN} \in P(n, g) \\ 0 & \text{otherwise} \end{cases}$$

$$gcount(n, atts)(g) = \begin{cases} gcount(n, atts_1)(g) \\ \oplus gcount(n, atts_2)(g) & \text{if } atts = atts_1 \ atts_2 \\ 1 & \text{if } atts = \langle [g] \rangle \\ & \text{and } P(n, g) = \{\mathbf{CLOSED}\} \\ * & \text{if } atts = \langle [g] \rangle \\ & \text{and } \mathbf{OPEN} \in P(n, g) \\ 0 & \text{otherwise} \end{cases}$$

$$fcount(n, xml)(f) = \begin{cases} fcount(n, xml_1)(f) \\ \oplus fcount(n, xml_2)(f) & \text{if } xml = xml_1 \ xml_2 \\ fcount(n, xml')(f) & \text{if } xml = \langle name \ atts \rangle \ xml' \ \langle /name \rangle \\ & \text{and } name \notin \text{FIELDS} \\ fc(n, atts) & \text{if } xml = \langle name \ atts \rangle \ xml' \ \langle /name \rangle \\ & \text{and } name \in \text{FIELDS} \\ & \text{and } atts \text{ contains } name="f" \\ 0 & \text{otherwise} \end{cases}$$

The *forms* function finds the relevant **form** elements in the given template, *gcount* counts the number of occurrences of a given gap name, and *fcount* counts the number of occurrences of a given field name. Note that the latter two functions need to consider the gap presence map of the summary graph. For the field count we can assume that only valid XHTML is shown because of the show analysis presented in the next section, and we can exploit the restrictions about input field elements described in Section 2.3. The set  $\text{FIELDS} = \{\text{input}, \text{button}, \text{select}, \text{textarea}, \text{object}\}$  contains all names of elements that define input fields. The function  $fc(n, atts)$  counts the number of name–value pairs that may be produced: If *atts* contains `type="radio"`, then it returns  $\diamond$ ; otherwise, if *atts* contains a `type` attribute with value `reset`, `submit`, or `image`, or an attribute with name `disabled` or `declare`, it returns 0; otherwise, if it contains `type="checkbox"` or an attribute named `multiple`, it returns  $*$ , and otherwise it returns 1. In order to detect whether `disabled` or `declare` occur, the gap presence map,  $P(n)$ , needs to be consulted in case of attribute gaps.

With these auxiliary functions in place, we can now define the value  $fp \in C$  representing the number of occurrences of *f* in the possible unfoldings of the summary graph:

$$fp = \bigsqcup_{r \in R} \Phi(r)$$

If for every root, the number of occurrences is always 0, always  $\diamond$ , or always 1, the final result is 0,  $\diamond$ , or 1, respectively; if it sometimes is  $\diamond$  and sometimes 1, the result is 1; and otherwise it is  $*$ . The function  $\Phi$  traverses the nodes in the summary graph, looking for applicable forms:

$$\Phi(n) = \bigsqcup_{(ff, gg) \in \text{allforms}(n)} \text{infields}(n, (ff, gg)) \sqcup \bigsqcup_{(n, h, m) \in T, h \in \text{tgaps}(n)} \Phi(m)$$

The left-hand factor counts the field occurrences for each **form** element that occurs directly in the template of *n*, while the right-hand factor looks at the templates that may be plugged into gaps in *n*.

$$\text{infields}(n, (ff, gg)) = ff(f) \oplus \bigoplus_{h \in G} (gg(h) \otimes \text{inflow}(n, h))$$

$$\text{inflow}(n, h) = \bigsqcup_{(n, h, m) \in T} \text{infields}(m, \text{count}(m))$$

Given a current node *n* and an element  $(ff, gg)$  of *GFP* representing the fields and gaps directly available in a particular form, the *infields* function sums the

field occurrences according to  $ff$  and those that may occur due to plug operations. For the latter part, we iterate through the possible gaps and multiply each count with the gap multiplicity. The *inflow* function follows the template edges and recursively finds the number of field occurrences in the same way as *outflow* but now assuming that we are inside a form.

As usual, we can compute the least fixed point by iteration because the lattice is finite and all operations are monotone. Since the *count* and *allforms* functions never return  $\perp$ , the result,  $fp$ , is always in the set  $\{0, \diamond, 1, *\}$ . The desired properties can now be verified by inspecting that:

$$\begin{aligned} fp &\in \{1, \diamond\} && \text{for } \mathbf{receive} \text{ operations, and} \\ fp &\neq 0 && \text{for } \mathbf{receive[]} \text{ operations.} \end{aligned}$$

for every summary graph computed for some **show** operation that is connected by a receive edge to the receive operation in the flow graph.

Like the plug analysis, this receive analysis is both sound and complete with respect to the summary graphs and the receive edges — assuming that only valid XHTML is ever shown: For a **receive**  $f$  operation, the analysis determines that it cannot fail if and only if for every label  $\ell$  of a **show** node which has an edge to the **receive** node, it is the case that in every XML document in  $\mathcal{L}(\text{summary}_\ell(x))$ , each form without an **action** attribute produces exactly one  $f$  field value. A similar property holds for **receive[]** operations.

For each **receive** and **receive[]** operation, we calculate  $fp$  for every summary graph of an associated **show** operation. Thus,  $fp$  is calculated  $O(n^2)$  times, where  $n$  is the size of the original Jwig program. The auxiliary functions *count* and *allforms* can be pre-computed in time  $O(n)$ . Each argument to *infields* denotes a specific form element in a template constant. Since there are  $O(n)$  template nodes and  $O(n)$  form elements in the program, both  $\Phi$  and *infields* are given at most  $O(n)$  different values as arguments. Since the lattice has constant height, we therefore iterate through the summary graph  $O(n)$  times. Each iteration performs a single traversal of the summary graph which takes time  $O(n^2)$ . In total, the receive analysis runs in time  $O(n^5)$  in the size of the original Jwig program.

### 5.3 Show Analysis

For every **show** statement in the Jwig program, we have computed a summary graph that describes how the XML templates are combined in the program and which XML documents may be shown to the client at that point. This gives us an opportunity to verify that all documents being shown are *valid* with respect to some document type. In particular, we wish to ensure that the documents are valid XHTML 1.0 [30] which is the most commonly used XML language for interactive Web services. XHTML 1.0 is the official XML version of the popular HTML 4.01. It is relatively easy to translate between the two, so in principle our technique works for HTML as well.

Validity of an XML document means that it is well-formed and in addition satisfies some requirements given by a schema for the particular language. The first part, well-formedness, essentially means that the document directly corresponds to a tree structure whose internal nodes are elements by requiring element tags to balance and nest properly. This part comes for free in Jwig, since all XML templates

are syntactically required to be well-formed. The remaining validity requirements specify which attributes a given element may have and which text and subelements that may appear immediately below the element in the XML tree. Such properties are specified using a *schema language*. In XHTML, the requirements are given by a DTD (Document Type Definition) schema plus some extra restrictions that cannot be formalized in the DTD language.

Our validation technique is parameterized by the schema description. Thereby we expect that it will be straightforward to support for instance WML or VoiceXML which are used for alternative interaction methods, in place of XHTML. Rather than using DTD, we apply a novel schema language, *Document Structure Description 2.0* (DSD2) [27]. This schema language is more expressive than DTD, so more validity requirements can be formalized. The expressive power of DSD2 is comparable to that of W3C's XML Schema [36], but DSD2 is significantly simpler, as indicated below.

#### 5.4 The Document Structure Description 2.0 Language

The DSD2 language is designed as a successor to the schema language described in [21; 22]. A DSD2 schema description of an XML language is itself an XML document. A *DSD2 processor* is a tool that takes as input a DSD2 schema and an XML document called the *instance document*. It then inserts default attributes and contents in the instance document according to the schema and checks whether the instance document with defaults inserted is *valid*. This is done by traversing the tree structure of the instance document in a top-down manner. For every element node, default attributes and contents are inserted and it is checked that all requirements about the attributes and contents of the element are satisfied.

The following description of DSD2 is intended to give a brief overview—not to define the language exhaustively. A normative specification document for DSD2 is currently under development [27].

Conceptually, a DSD2 schema consists of a list of *constraints*. A constraint is either a *declaration*, a *requirement*, a *conditional constraint*, or a *default specification*. Furthermore, there are notions of *string normalization*, *keys* and *references*, and *options* which we can ignore here. During the top-down traversal, the processor checks each element in turn. The *current element* is the one currently being checked. This check of an individual element is performed in five steps:

- (1) all applicable constraints are found;
- (2) default attributes and contents are inserted;
- (3) the requirements are checked;
- (4) it is checked that all attributes of the current element are declared; and
- (5) it is checked that the contents matches all contents declarations and that the whole contents are declared.

The following description of the various types of constraints explains these steps in more detail:

*Declarations.* A declaration constraint contains a list of *attribute declarations* and *contents declarations*. An attribute declaration specifies that an attribute with a given name is allowed in the current element provided that the value matches a

given regular expression. A contents declaration is a regular expression over characters and element names that specifies a requirement for the presence, ordering, and number of occurrences of sub-elements and character data. A contents declaration only looks at elements that are mentioned in the regular expression. This subsequence of the contents is called the *projected contents*. If the expression contains any character sub-expressions, all character data in the contents is included in the projected contents. Checking a contents declaration succeeds if the projected contents matches the regular expression. All attributes and contents that have been matched by a declaration are considered to be *declared*.

*Requirements.* A requirement constraint contains boolean expressions that must evaluate to true for the current element. Boolean expressions are built of the usual boolean operators, together with attribute expressions which probe the presence and values of attributes, element expressions which probe the name of the current element, and *parent* and *ancestor* operators which probe whether certain properties are satisfied for the elements above the current element in the instance document tree.

*Conditional constraints.* A conditional constraint contains a list of constraints whose applicability is guarded by a boolean expression. Only when the boolean expression evaluates to true for the current element, the constraints within are processed.

*Defaults.* A default constraint specifies a default value for an attribute or a default contents sequence for an empty element. In case of conflicts, for instance if two default attributes with the same name are applicable, all but the one specified last in the schema are ignored.

For convenience, specifications can be grouped and named for modularity and reuse. Furthermore, the DSD2 schema can restrict the name of the root element: for example, in XHTML, it must be `html`. DSD2 has full support for Namespaces [11]. For XHTML, the namespace `http://www.w3.org/1999/xhtml` is used, and all elements must use the empty namespace prefix.

As an example, the following snippet of the DSD2 description of XHTML 1.0 describes `d1` elements:

```
<if><element name="d1"/>
  <declare>
    <attribute name="compact"><string value="compact"/></attribute>
    <repeat min="1"><union>
      <element name="dt"/><element name="dd"/>
    </union></repeat>
  </declare>
  <constraint ref="ATTRS"/>
</if>
```

These constraints show that a `d1` element may contain a `compact` attribute, provided that its value is `compact`, and that it must contain at least one `dt` or `dd` sub-element. Additionally, `ATTRS`, which is defined elsewhere, describes some additional common attributes that may occur.

The following example (abbreviated with "...") describes `a` elements:

```
<if><element name="a"/>
```



```

<declare>
  <attribute name="name"><stringtype ref="NMTOKEN"/></attribute>
  <attribute name="shape"><stringtype ref="SHAPE"/></attribute>
  ...
  <repeat><union>
    <string/>
    <contenttype ref="PHRASE"/>
    ...
  </union></repeat>
</declare>
<constraint ref="HREFLANG"/>
...
<default name="shape" value="rect"/>
<require>
  <not><ancestor><element name="a"/></ancestor></not>
</require>
</if>

```

This reads: If the current element is named “a”, then the sub-constraints are applicable. First, the attributes `name`, `shape`, etc. are declared. The `stringtype` constructs are references to regular expressions defining the valid attribute values. Then, a contents declaration states that all text is allowed as contents together with some contents expressions defined elsewhere. After that, there are some references to constraint definitions and a default specification for the `shape` attribute. Finally, there is a requirement stating that `a` elements cannot be nested. The latter constraint is an example of a validity requirement that cannot be expressed by DTD or XML Schema.

As a final example, the following requirement can be found in the description of input elements:

```

<require>
  <or>
    <attribute name="type">
      <union><string value="submit"/><string value="reset"/></union>
    </attribute>
    <attribute name="name"/>
  </or>
</require>

```

This states that there must be a `type` attribute with value `submit` or `reset` or a `name` attribute. This is another validity requirement that cannot be expressed concisely in most other schema languages. The whole DSD2 schema for XHTML 1.0 can be found at <http://www.brics.dk/DSD/xhtml1-transitional.dsd>.

## 5.5 Validity Analysis

We show below how the DSD2 processing algorithm explained in the previous section generalizes from concrete XML documents to summary graphs. In fact, the DSD2 language has been designed with summary graph validation in mind. Since the DSD2 language is a generalization of the DTD language, the following algorithm could be adapted to DTD. One benefit of using DSD2 instead of DTD or XML Schema is that it allows us to capture many more errors. Every validity requirement that merely appear as comments in the DTD schema for XHTML can be formalized in DSD2, as exemplified in the previous section. Of course, there still

are syntactic requirements that even DSD2 cannot express. For instance, in tables, the `thead`, `tfoot`, and `tbody` sections must contain the same number of columns, and the `name` attributes of `a` elements must have unique values. Summary graphs clearly do not have the power to count such numbers of occurrences, so we do not attempt to also check these requirements. Still, our approach based on DSD2 captures more syntactic errors than is possible with DTD. Only the uniqueness requirements specified by `ID` and `IDREF` attributes are not checked, but they do not play a central role in the schema for XHTML anyway.

Recall from Section 2.3 that we make a few modification of the documents at run-time just before they are sent to the clients. It is trivial to modify the validation algorithm to take these modifications into account. For instance, rather than requiring one or more entries in all lists, the analysis permits any number since lists with zero entries are always removed anyway.

Given a DSD2 schema and a summary graph  $SG = (R, T, S, P)$  associated to some `show` statement, we must check that every XML document in  $\mathcal{L}(SG)$  is valid according to the schema. The algorithm for validating a summary graph with respect to a DSD2 schema proceeds in a top-down manner mimicking the definition of the unfolding relation in Section 4.2, starting from the root elements in the templates of the summary graph root nodes. In contrast to the analyses described in the previous sections, we will describe this one in a less formal manner because of the many technical details involved. Rather than showing all the complex underlying equation systems or describing the entire algorithm in detailed pseudo-code, we attempt to present an concise overview of its structure.

For each root  $r \in R$ , we perform the following steps:

- (1) Let  $p$  be a pointer to the root element in  $t(r)$ .
- (2) Check that the name of the  $p$  element is valid for a root element.
- (3) Initialize  $M = \emptyset$  and call  $check(r, p, c_0)$ .

If no violations are detected,  $SG$  is valid which implies that all documents in  $\mathcal{L}(SG)$  are valid. The  $check$  procedure is given a summary graph node  $n$ , a pointer  $p$  to an element in  $t(n)$ , and a *context map*  $c$ . It recursively checks validity of the  $p$  element and the sub-trees below it, assuming the specified context. A context map assigns a truth value to every `parent` and `ancestor` expression occurring in the schema. The map  $c_0$  maps everything to *false*.

The set  $M$  is a global *memoization set* which consists of triples that—like the arguments to  $check$ —contain a summary graph node, a pointer to an element in the template of the node, and a context map. This set at all times describes the elements and contexts that already have been checked. We use it to avoid performing redundant computations. For a given schema and summary graph, there are only finitely many possible memoization sets, so termination is guaranteed, even in case of cycles in the summary graph.

The procedure  $check(n, p, c)$  does the following, essentially corresponding to the five steps for checking individual elements in the algorithm from the previous section, but now also considering gaps:

- (1) If  $(n, p, c) \in M$ , then skip all the following steps; otherwise add  $(n, p, c)$  to  $M$  and proceed.

- (2) Compute a new context map  $c'$ . This is done by re-evaluating each **parent** and **ancestor** expression, based on the  $p$  element and the  $c$  map. Note that because of the  $c$  map, this can be done in constant time in the size of the summary graph and the templates.
- (3) Find all constraints in the schema that are applicable for the  $p$  element. This is done by traversing the schema and for each conditional constraint, evaluating its boolean expression. If it evaluates to *true*, it is included. When evaluating boolean expressions, the context map  $c'$  is consulted for evaluating **parent** and **ancestor** sub-expressions. We describe later more details on how the boolean expressions are evaluated.
- (4) Build a collection of defaults that should be inserted. This is done by traversing the applicable defaults. An attribute default is included if the  $p$  element does not have an attribute of that name. A contents default is included if the contents of the  $p$  element is empty. In order to determine whether or not attributes or contents occur, we need to consider the gap presence map,  $P$ , the template edges,  $T$ , and the string edges,  $S$ , in case there are gaps among the attributes or contents of the  $p$  element. Note that we do not actually insert the defaults here but only collect them. When the defaults are collected, we need to update  $c'$  according to the collected defaults since default attributes may affect the truth values of the **parent** and **ancestor** expressions.
- (5) For each applicable requirement, check that its boolean expression evaluates to *true*. If it does not, then the summary graph is not valid.
- (6) For each attribute in the  $p$  element, including the collected default attributes, check that it is declared by the applicable attribute declarations. An attribute is declared by an attribute declaration if the attribute value matches the regular expression of the declaration. In case the attribute is an attribute gap, this amounts to checking inclusion of one regular language of Unicode strings in another. However, it is possible that one attribute declaration specifies that a given attribute may have one set of values and another declaration specifies that the same attribute may also have another set of values. Therefore, in general, we check that all values that are possible according to the string edges,  $S$ , match some declaration. If some attribute is not declared, the summary graph is not valid.
- (7) For each contents declaration, check that the contents of  $p$  element matches the declaration. As previously mentioned, a contents sequence matches a contents declaration if the projected contents is in the language of the regular expression of the declaration. In case there are no gaps in the contents, this is a simple check of string inclusion in a regular language. If there are gaps, the situation is more involved: The template edges from the gaps may lead to templates which at the top-level themselves contain gaps. (The *top-level* of a template is the sequence of elements, characters, and gaps that are not enclosed by other elements.) In turn, this may cause loops of template edges. Therefore, in general, the set of possible contents sequences forms a context-free language, which we represent by a context-free grammar. Without such loops, the language is regular. The problem of deciding inclusion of a context-free language in a regular language is decidable [17], but computationally expensive. For that reason, we

approximate the context-free language by a regular one: Whenever a loop is found in the context-free grammar, we replace it by the regular language  $A^*$  where  $A$  consists of the individual characters and elements occurring in the loop. This allows us to apply a simpler regular language inclusion algorithm. Although loops in summary graph often occur, our experiments show that it is rare that this approximation actually causes any imprecision. In addition to checking that all contents declarations are satisfied, we check that all parts of the contents has been declared, that is, matched by some declaration. If not, the summary graph is not valid. Again, if any declaration contains a character sub-expression, all character data is considered declared.

- (8) Look through the immediate contents of the  $p$  element. For each sub-element  $p'$  that occurs within  $t(n)$ , call  $check(n, p', c')$  recursively. For each template gap  $g$  in the contents, find all each outgoing template edge  $(n, g, n') \in T$  and call  $check(n', p', c')$  recursively for every element  $p'$  occurring at the top-level of  $t(n')$ .

Note that, assuming a fixed schema, steps (1)-(7) can be performed in linear time in the number of attributes in the current element and in the length of its contents, independently of the rest of the  $t(n)$  template and of all the other templates. Because of the memoization, the entire algorithm for whole summary graphs therefore runs in linear time in the size of the templates, which in a sense is optimal.

Also note that in contrast to the algorithm for validating XML documents, this one does not have side-effects on the summary graphs: defaults are *not* inserted in the templates, but the validity checks works as if they were inserted. It would not be possible to explicitly insert the defaults, since each template in general is evaluated in many different contexts during the validity check.

Evaluation of boolean expressions for conditional constraints and requirement constraints is non-trivial because we have to consider all the possible unfoldings of the summary graph. We apply a *four-valued* logic for that purpose. Evaluating a boolean expression results in one of the following values:

- true* – if evaluating the expression on every possible unfolding would result in *true*;
- false* – if they all would result in *false*;
- some* – if some unfolding would result in *true* and others in *false*;
- don't-know* – if the processor is unable to detect whether all, no, or some unfoldings would result in *true*.

All boolean operators extend naturally to these values. The value *don't-know* is for instance produced by the conjunction of *some* and *some*. If the guard of a conditional constraint evaluates to *don't-know*, we terminate with a “don't know” message. However, for our concrete XHTML schema, this can in fact never happen.

Compared with the technique described in [8], we have now moved from an abstract version of DTD to the more powerful DSD2 schema language. Furthermore, by the introduction of four-valued logic for evaluation of boolean expressions, we have repaired a defect that in rare cases caused the old algorithm to fail.

Our algorithm is sound, that is, if it validates a given summary graph it is certain that all unfoldings into concrete XML documents are also valid. Because of the possibility of returning “don't know” and of the approximation of context-free

languages by regular ones, the algorithm is generally not complete. An alternative to our approach of “giving up” when these situations occur would be to branch out on all individual unfoldings and use classical two-valued logic. This shows that the problem is decidable. However, the complexity of the algorithm would then increase significantly, and, for the XHTML schema, false errors can only occur in the rare cases where we actually need to approximate context-free languages by regular ones, as mentioned above.

## 6. IMPLEMENTATION AND EVALUATION

To make experiments for evaluating the JWIG language design and the performance of the program analyses, we have made a prototype implementation. It consists of the following components, roughly corresponding to the structure in Figures 6 and 7:

- a simple desugarer, made with JFlex [23], for translating JWIG programs into Java code;
- the `<bigwig>` low-level run-time system [7] which in its newest version [25] is based on a module for the Apache Web Server [3] and extended with Java support;
- a Java-based run-time system for representing and manipulating XML templates;
- a part of the Soot optimization framework for Java [37; 35], which converts Java bytecode to a more convenient 3-address instruction code language, called Jimple;
- a flow graph package, which operates on the Jimple code generated by Soot and also uses Soot’s CHA implementation;
- a finite-state automaton package with UTF16 Unicode alphabet and support for all standard regular operations [26];
- a summary graph construction package which also performs the string analysis;
- a plug and receive analyzer which performs the checks described in Section 5.1 and 5.2; and
- a DSD2 validity checker with summary graph support.

All parts are written in Java, except the low-level run-time system which is written in C. The Java part of the run-time system amounts to 3000 lines, and the analysis framework is 12,500 lines.

### 6.1 Example: The Memory Game

To give a more complete example of a JWIG service, we present the well-known Memory Game, where the player must match up pairs of cards lying face down. First, the number of pairs is chosen, next the game itself proceeds iteratively, and finally the player is congratulated. A snapshot of the game in progress is seen in Figure 13.

The main session, presented in Figure 14, looks just like a corresponding sequential Java program. The templates being used are presented in Figure 15. The construction of a grid of cards is performed by the `makeCardTable` method presented in Figure 16. The class representing individual cards is seen in Figure 17. In all, the Memory Game is written in 169 lines of JWIG.

By itself, the session concept and the XML templates simplify the program compared to solutions in JSP or Servlets. Furthermore, since the example is analyzed

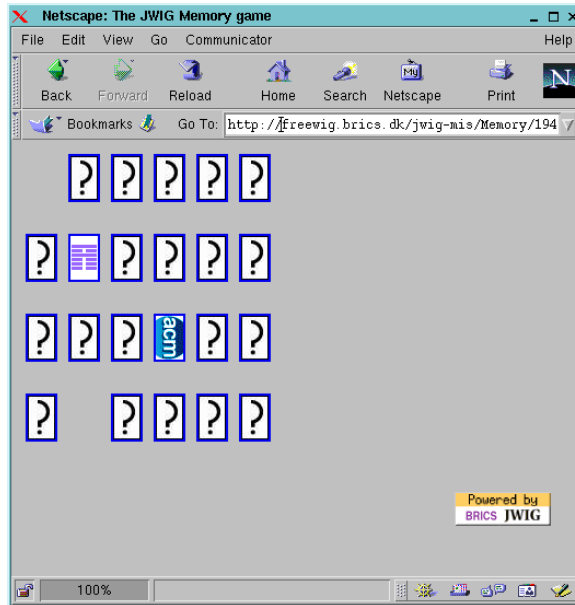


Fig. 13. A snapshot of the Memory Game being played.

without errors, we know that no Jwig exceptions will be thrown while the game is being played. In particular, we are guaranteed that all documents being shown are valid XHTML according to the strict standard imposed by the DSD2 schema.

The Jwig run-time system, which is also used in the `<bigwig>` project, is tailor-made for session-based Web services. Each session thread is associated a unique URL which refers to a file on the server. This file at all times contains the most recent page shown to the client. The session code runs as a JVM thread that lives for the entire duration of the session. In contrast, sessions in Servlet and JSP services run as short-lived threads where the session identity is encoded using cookies or hidden input fields, as described in Section 1.3. This precludes sessions from being bookmarked, such that the client cannot suspend and later resume a session, and the history buffer in the browser typically gets cluttered with references to obsolete pages. In our solution, the session URL functions as an identity of the session, which avoids all these problems. These aspects are described in more detail in [9].

If we introduce an error in the program, by forgetting the `name` attribute in the `input` field in the `welcome` template, then the Jwig analyzer produces the following output:

```
*** Field 'howmany' is never available on line 68
*** Invalid XHTML at line 49
--- element 'input': requirement not satisfied:
<or xmlns="http://www.brics.dk/DSD/2.0/error">
  <attribute name="type">
    <union>
      <string value="submit" />
      <string value="reset" />
    </union>
  </attribute>
</or>
```

```

public class Game extends Session {
    public void main() {
        // ask for number of cards
        int howmany;
        do {
            show wrap <[ body = welcome <[ atmost = images.length ]]);
            howmany = Integer.parseInt(receive howmany);
        } while (howmany < 1 || howmany > images.length);

        // generate random permutation of cards
        Card[] cards = new Card[howmany*2];
        Random random = new Random();
        for (int i = 0 ; i < howmany ; i++) {
            for (int c = 0 ; c < 2 ; c++) {
                int index;
                do {
                    index = random.nextInt(howmany*2);
                } while (cards[index] != null);
                cards[index] = new Card(i);
            }
        }

        // play the game
        int pairsleft = howmany;
        int moves = 0;
        show makeCardTable(cards);
        while (pairsleft > 0) {
            // first card picked
            int firstcard = Integer.parseInt(receive submit);
            cards[firstcard].status = 1;
            show makeCardTable(cards);

            // second card picked
            int secondcard = Integer.parseInt(receive submit);
            cards[secondcard].status = 1;
            moves++;

            // check match
            if (cards[firstcard].value == cards[secondcard].value) {
                cards[firstcard].status = 2;
                cards[secondcard].status = 2;
                if (--pairsleft > 0)
                    show makeCardTable(cards);
            } else {
                show makeCardTable(cards);
                cards[firstcard].status = 0;
                cards[secondcard].status = 0;
            }
        }

        // done, show result
        exit farewell <[ howmany = howmany, moves = moves ];
    }
}

```

Fig. 14. The main session of the Memory Game.

```

private static final XML wrap = [[
  <html>
    <head><title>The JWIG Memory game</title></head>
    <body><form><[body]></form></body>
  </html>
]];

private static final XML welcome = [[
  <h3>Welcome to the JWIG Memory game!</h3>
  <p>How many pairs of cards do you want (from 1 to <[atmost]>)?</p>
  <input type="text" name="howmany"/>
]];

private static final XML farewell = wrap <[ body = [[
  <h3>Thank you for playing this game!</h3>
  <p>You found all <[howmany]> pairs using <[moves]> moves.</p>
]] ];

```

Fig. 15. Templates from the Memory Game.

```

private XML makeCardTable(Card[] cards) {
  XML table = [[ <table><[row]></table> ]];
  for (int y=0; y < (cards.length+COLS-1)/COLS; y++) {
    XML row = [[ <tr><[elem]></tr><[row]> ]];
    for (int x=0; x < COLS; x++) {
      XML elem = [[ <td><[contents]></td><[elem]> ]];
      int index = y*COLS+x;
      if (index < cards.length) {
        elem = elem <[ contents = cards[index].makeCard(index) ];
      }
      row = row <[ elem = elem ];
    }
    table = table <[ row = row ];
  }
  return wrap <[ body = table ];
}

```

Fig. 16. Generating a grid of cards in the Memory Game.

```

</union>
</attribute>
<attribute name="name" />
</or>

```

In the first line, the receive analysis complains that the *howmany* field is never available from the client. The remainder of the error message is from the show analysis, which notices that the `input` element violates the quoted constraint from the XHTML schema. This particular validity error is not caught by DTD validation of the generated document. If the involved element contained gaps, the error message would include a print of all relevant template and string edges and values of the gap presence map, which shows the relevant plug operations. Clearly, such diagnostics are useful for debugging.



```

private class Card {
    public int status;
    public int value;

    public Card(int value) {
        this.status = 0;
        this.value = value;
    }

    public XML makeCard(int index) {
        switch(status) {
            case 0:
                return [[ <input type="image" alt="card"
                           src=[image] name=[index] /> ]]
                    <[ image = back_image, index = index ];
            case 1:
                return [[ <img src=[image] alt=[num] /> ]]
                    <[ image = images[value], num = value ];
            case 2:
                return [[ <img src=[image] alt="" /> ]]
                    <[ image = blank_image ];
            default:
                return null;
        }
    }
}

```

Fig. 17. Representing a card in the Memory Game.

## 6.2 Performance

The Jwig implementation may be evaluated with respect to compile-time, analysis-time, or run-time. The compile-time performance is not an issue, since Jwig programs are simply piped through a Jflex desugarer and compiled using a standard Java compiler. The Jwig runtime performance is not particularly interesting, since we reuse the <bigwig> runtime system and the standard J2SE JVM. The critical component in our system is the extensive collection of static analyses that we perform on the generated class files.

As shown in Figure 7, the static analysis is a combination of many components, which we in the following quantify separately. Our benchmark suite, shown in Figure 18, is a collection of small to medium sized Jwig services, most of which have been converted from corresponding <bigwig> applications [8]. The right-most column shows the total time in seconds for the entire suite of program analyses. For all benchmarks, at most 150 MB memory is used.

The four larger ones are an XML template manager where templates can be uploaded and edited (TempMan), an interactive Web board for on-line discussions (WebBoard), a system for study administration (Bachelor), and a system for management of the JA00 2001 conference (Jaoo).

Figure 19 shows the resources involved in computing the flow graphs on a 1 GHz Pentium III with 1 GB RAM running Linux. For each benchmark we show the time in seconds used by Soot, the time in seconds used by phases 1 through 7

Name	Lines	Templates	Shows/Exits	Total Time
Chat	80	4	3	5.370
Guess	94	8	7	7.147
Calendar	133	6	2	7.029
Memory	167	9	6	9.718
TempMan	238	13	3	7.719
WebBoard	766	32	24	9.769
Bachelor	1078	88	14	115.641
Jaoo	3923	198	9	35.997

Fig. 18. The benchmark services.

Name	Load	Construct	Size Before	Simplify	Size After
Chat	3.205	1.871	238/388	0.232	107/99
Guess	3.242	1.993	286/442	0.152	113/89
Calendar	3.276	2.241	386/672	0.254	124/127
Memory	3.284	2.227	451/765	0.292	143/127
TempMan	3.290	2.604	779/1437	0.883	200/192
WebBoard	3.244	2.660	878/1285	0.775	422/287
Bachelor	3.255	4.303	2278/3676	21.862	1059/914
Jaoo	3.557	5.647	3009/4407	14.045	1406/1008

Fig. 19. Flow graph construction.

Name	Time	Largest Size
Chat	0.103	2/1/5
Guess	0.105	2/1/3
Calendar	0.440	5/9/5
Memory	2.627	7/8/5
TempMan	0.203	11/13/9
WebBoard	1.189	9/11/11
Bachelor	24.673	47/83/24
Jaoo	5.067	33/45/48

Fig. 20. Summary graph construction.

described in Section 3, the size of the resulting flow graph, the time in seconds used by the simplifying Phase 8, and the size of the simplified flow graph. The flow graph sizes are shown as number of nodes and number of flow edges. The loading time is dominated by initialization of Soot. Phases 1 through 7 of the flow graph construction is seen to be linear in the program size. The time for the simplification phase strongly depends on the complexity of the document constructions, which explains the relatively large number for the `Jaoo` service. For all benchmarks, the simplification phase substantially reduces the flow graph size. Furthermore, recall that before the simplification phase, flow edges may have multiple variables, while after simplification, they all have exactly one variable.

Figure 20 quantifies the computation of summary graphs, including the string analysis. For each benchmark we show the total time in seconds and the size of the largest summary graph, in terms of nodes, template edges, and non-trivial string edges. The relatively large numbers for the `Bachelor` example correctly reflects that it constructs complicated documents. Without graph simplification, the total time

Name	Plug	Receive	Show	False Errors
Chat	0.002	0.004	0.953	0
Guess	0.015	0.002	1.638	0
Calendar	0.017	0.000	0.801	0
Memory	0.003	0.002	1.283	0
TempMan	0.004	0.015	0.720	0
WebBoard	0.016	0.003	1.882	0
Bachelor	0.133	0.009	61.406	0
Jaoo	0.059	0.002	7.620	0

Fig. 21. Summary graph analysis.

for the `Memory` example blows up to more than 15 minutes, while the `Jaoo` example was aborted after 90 minutes. We conclude that summary graph construction appears to be inexpensive in practice and that graph simplification is worth the effort.

Figure 21 deals with the subsequent analysis of all the computed summary graphs. For each benchmark we show the total time in seconds for each of the three analyses and the total number of false errors generated by the conservative analyses. In all cases, the times are smaller than for the other phases. Many benchmarks generated errors, but by careful inspection, they were all seen to correctly identify actual XHTML validity errors. Thus, the analysis appears to be precise enough to serve as a real help to the programmer.

We conclude that the Jwig prototype implementation is certainly feasible to use, but that there is room for performance improvements for the implementation.

## 7. PLANS AND IDEAS FOR FUTURE DEVELOPMENT

Our current system can be extended and improved in many different directions which we plan to investigate in future work. These can be divided into language design, program analysis, and implementation issues, and are briefly described in the following.

### 7.1 Language design

So far, the design of Jwig has focused on two topics that are central to the development of interactive Web services: sessions and dynamic construction of Web documents. However, there are many more topics that could benefit from high-level language-based solutions, as shown in [9].

One example is validation of form input. Many Web services apply intricate JavaScript client-side code in the Web documents for checking that input forms are filled in consistently. For instance, it is typical that certain fields must contain numbers or email addresses, or that some fields are optional depending on values entered in other fields. Proper error messages need to be generated when errors are detected such that the clients have the chance to correct them, and extra checks need to be performed on the server. The PowerForms language [6] has been developed to attack the problem of specifying form input validation requirements in a more simple and maintainable way based on regular expressions and boolean logic. It should be straightforward to integrate PowerForms into Jwig.

The current XML cast operation in Jwig is somewhat unsatisfactory for two reasons: 1) if a cast fails due to invalid XHTML, an exception is not thrown immediately since it is not detected until a subsequent show operation; and 2) its expressiveness is limited—for instance, unions of templates cannot be expressed. One solution to this may be to use DSD2 descriptions instead of constant templates in the cast operations. However, to generalize the analyses correspondingly, a technique for transforming a DSD2 description of an XML language into a summary graph is needed. We believe that this is theoretically possible—further investigation will show whether it is also practically feasible.

Another idea is to broaden the view from interactive Web services to whole Web sites comprising many services and documents. The Strudel system [16] has been designed to support generation and maintenance of Web sites according to the design principle that the underlying data, the site structure, and the visual presentation should be separated. A notion of data graphs allows the underlying data to be described, a specialized query language is used for defining the site structure, and an HTML template language that resembles the XML template mechanism in Jwig defines the presentation. We believe that the development of interactive services can be integrated into such a process. For sites that comprise both complex interactive session-based services and more static individual pages, the concepts in the `Service.Session` and `Service.Page` classes could be applied. Jwig could also benefit from a mechanism for specifying dependencies between the pages or sessions and the data, for instance, such that pages are automatically cached and only recomputed when certain variables or databases are modified.

We have shown that our template mechanism is suitable for constructing XHTML documents intended for presentation. If the underlying data of a Web service is represented with XML, as suggested by Strudel, we will need a general mechanism for extracting and transforming XML values. Currently, we only provide the plug operation for combining XML templates—a converse “unplug” operation would be required for deconstructing XML values. Preliminary results suggest that our notion of summary graphs and our analyses generalize to such general XML transformations. XDuce [18] is a related research language designed to make type-safe XML transformations. In XDuce, types are simplified DTDs where we instead use the more powerful DSD2 notation. Furthermore, XDuce is a tiny functional language while Jwig contains the entire Java language. Instead of relying on a type system for ensuring that the various XML values are valid according to some schema definition, we perform data-flow analyses based on summary graphs. Based on these ideas, a current project aims to make Jwig a general and safe XML transformation language.

## 7.2 Program analysis

The experiments indicate that the notion of summary graphs is suitable for modeling the XML template mechanism and that the analysis precision is adequate. However, the preliminary string analysis described in Section 4.1 can be improved. The modular design of the analyses makes it possible to replace this simple string analysis by a more precise one. For example, string concatenation operations could be modeled more precisely by exploiting the fact that regular languages are closed under finite concatenation. Because of loops in the flow graphs, this would in gen-

eral produce context-free languages so a suitable technique for approximating these by regular languages is needed. That essentially amounts to applying widening for ensuring termination. Other operations, such as the substring methods, could also easily be modeled more precisely than with `anystring`. An advanced version of such an analysis would apply flow-sensitivity, such that e.g. `if` statements that branch according to the value of a string variable would be taken into account, and instead of modeling `receive` by `anystring`, the regular languages provided by PowerForms specifications could be applied. A natural extension to these ideas would be to add a “regular expression cast” operator to the JWIG language. As the other cast operations, that would provide a back-door to the analysis which occasionally can be convenient no matter how precise the analysis may be.

The current program analysis is based on the assumption that the medium used for communication with the clients is XHTML. However, since the show analysis is parameterized by a DSD2 description, validity with respect to any XML language describable by DSD2 can be checked instead. Two obvious alternatives are WML, *Wireless Markup Language* [38], which is used for mobile WAP devices with limited network bandwidth and display capabilities, and VoiceXML, *Voice Extensible Markup Language* [4], for audio-based dialogs. Such languages can be described precisely with DSD2. Only the receive analysis requires modification since it needs to identify the forms and fields, or whatever the equivalent notions are in other formalisms.

### 7.3 Implementation

Our current implementation is a prototype developed to experiment with the design and test the performance. This means that there are plenty of ways to improve the performance of the analysis and the run-time system.

We plan to apply the `metafront` syntax macros [10] in a future version to improve the quality of the parsing error messages. This will also allow us to experiment with syntax macros as a means for developing highly domain-specific languages in the context of Java-based interactive Web services.

The application of Soot in the generation of flow graphs is sometimes a bottleneck in the analysis, even though the theoretical complexity of this translation is trivial compared to the other analysis phases. This suggests that we use a tailor-made byte-code to flow-graph converter instead.

Finally, we believe that it is possible to significantly improve the run-time performance for JWIG services by integrating the JWIG run-time system with a Java Enterprise Edition server. For instance, this allows `Service.Page` to become essentially as efficient as JSP code by exploiting that the threads are never suspended by `show` statements. JRockit [2] is a commercial JVM implementation which is tuned for Web servers with high loads. In particular, it supports light-weight threads which will significantly reduce the overhead induced by our session model.

## 8. CONCLUSION

We have defined JWIG as an extension of the Java language with explicit high-level support for two central aspects of interactive Web services: 1) sessions consisting of sequences of client interactions and 2) dynamic construction of Web pages. Compared to other Web service programming languages, these extensions can improve

the structure of the service code. In addition to being convenient during development and maintenance of Web services, this allows us to perform specialized program analyses that check at compile-time whether or not run-time errors may occur due to the construction of Web pages or communication with the clients via input forms. The program analyses are based on a unique notion of *summary graphs* which model the flow of document fragments and text strings through the program. These summary graphs prove to contain exactly the information needed to provide all the desired static guarantees of the program behavior.

This article can be viewed as a case study in program analysis. It contains a total of seven analyses operating on different abstractions of the source program: one for making receive edges during flow graph construction, the reaching definitions analysis in the flow-graph simplification phase, the string analysis, the summary graph construction, and the plug, receive, and show analyses. The whole suite of analyses is modular in the sense that each of them easily can be replaced by a more precise or efficient one, if the need should arise. If, for example, future experience shows that the control-flow information in the flow graphs is too imprecise, one could apply a *variable-type analysis* [35] instead of CHA. Or, if the string analysis should turn out to be inadequate for developing, e.g., WML services, it could be replaced by another. Analysis correctness is given by the correctness of each phase. For instance, the flow graphs conservatively approximate the behavior of the original Jwig programs, the summary graphs conservatively model the template constructions with respect to the flow graphs, and the validity results given by the show analysis are conservative with respect to the summary graphs.

The language extensions permit efficient implementation, and despite the theoretical worst-case complexities of the program analyses, they are sufficiently precise and fast for practical use.

All source code for our Jwig implementation, including API specifications and the DSD2 schema for XHTML 1.0, is available from <http://www.brics.dk/Jwig/>.

## REFERENCES

- Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers - Principles, Techniques, and Tools*. Addison-Wesley, 1986.
- Björn Antonsson, Joakim Dahlstedt, Mattias Joëlson, Marcus Lagergren, Staffan Larsen, David Lindholm, Olof Lindholm, Peter Lönnebring, Mikael Vidstedt, and Calle Wilund. JRocket – the faster server JVM, 2002. <http://www.jrocket.com>.
- Brian Behlendorf et al. The Apache HTTP server project, 2002. <http://httpd.apache.org/>.
- Linda Boyer, Peter Danielsen, Jim Ferrans, Gerald Karam, David Ladd, Bruce Lucas, and Kenneth Rehor. *Voice eXtensible Markup Language, Version 1.0*. W3C, May 2000. W3C Note, <http://www.w3.org/TR/voicexml/>.
- Claus Brabrand. *<bigwig> Version 1.3 – Reference Manual*. BRICS Notes Series NS-00-4, September 2000.
- Claus Brabrand, Anders Møller, Mikkel Ricky, and Michael I. Schwartzbach. PowerForms: Declarative client-side form field validation. *World Wide Web Journal*, 3(4):205–314, 2000.
- Claus Brabrand, Anders Møller, Anders Sandholm, and Michael I. Schwartzbach. A runtime system for interactive Web services. In *Proc. 8th International World Wide Web Conference (WWW8)*. Elsevier, 1999. Also in *Computer Networks* Vol. 31 No. 11-16: 1391-1401, Elsevier.
- Claus Brabrand, Anders Møller, and Michael I. Schwartzbach. Static validation of dynamically generated HTML. In *Proc. Workshop on Program Analysis for Software Tools and Engineering (PASTE 2001)*. ACM, 2001.

- Claus Brabrand, Anders Møller, and Michael I. Schwartzbach. The <bigwig> project. *Transactions on Internet Technology*, 2(1), 2002.
- Claus Brabrand and Michael I. Schwartzbach. The metafront system: Extensible syntax processing. In preparation, 2002.
- Tim Bray, Dave Hollander, and Andrew Layman. *Namespaces in XML*. W3C, January 1999. W3C Recommendation, <http://www.w3.org/TR/REC-xml-names>.
- Tim Bray, Jean Paoli, C. M. Sperberg-McQueen, and Eve Maler. *Extensible Markup Language (XML) 1.0 (Second Edition)*. W3C, October 2000. W3C Recommendation, <http://www.w3.org/TR/REC-xml>.
- David R. Chase, Mark Wegman, and F. Kenneth Zadeck. Analysis of pointers and structures. In *Proc. Conference on Programming Language Design and Implementation (PLDI'90)*, pages 296–310. ACM, June 1990.
- H. Comon, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree automata techniques and applications. <http://www.grappa.univ-lille3.fr/tata>, 1999.
- Jeffrey Dean, David Grove, and Craig Chambers. Optimization of object-oriented programs using static class hierarchy analysis. In *Proc. European Conference on Object-Oriented Programming (ECOOP'95)*, volume 952 of *LNCS*, pages 77–101. Springer, August 1995.
- Mary F. Fernandez, Dan Suciu, and Igor Tatarinov. Declarative specification of data-intensive Web sites. In *Proc. 2nd Conference on Domain-Specific Languages (DSL '99)*. USENIX, October 1999.
- John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979.
- Haruo Hosoya and Benjamin C. Pierce. XDuce: A typed XML processing language. In *Proc. 3rd International Workshop on the Web and Databases (WebDB 2000)*, volume 1997 of *LNCS*, pages 226–244. Springer, May 2000.
- Jason Hunter and Brett McLaughlin. JDOM, 2001. <http://jdom.org/>.
- John B. Kam and Jeffrey D. Ullman. Monotone data flow analysis frameworks. *Acta Informatica*, 7:305–317, 1977.
- Nils Klarlund, Anders Møller, and Michael I. Schwartzbach. Document Structure Description 1.0, 2000. BRICS Notes Series NS-00-7. <http://www.brics.dk/DSD/>.
- Nils Klarlund, Anders Møller, and Michael I. Schwartzbach. The DSD schema language. *Automated Software Engineering*, 2002. To appear. Earlier version appears in ACM Proc. 3rd Workshop on Formal Methods in Software Practice (FMSP 2000).
- Gerwin Klein. JFlex - the fast scanner generator for Java, 2001. <http://www.jflex.de/>.
- David A. Ladd and J. Christopher Ramming. Programming the web: An application-oriented language for hypermedia services. In *4th Intl. World Wide Web Conference (WWW4)*. W3C, 1995.
- Anders Møller. The <bigwig> runtime system, 2001. <http://www.brics.dk/bigwig/runwig/>.
- Anders Møller. `dk.brics.automaton` - finite-state automata and regular expressions for Java, 2001. <http://www.brics.dk/automaton/>.
- Anders Møller. Document Structure Description 2.0. In preparation, 2002.
- Flemming Nielson, Hanne Riis Nielson, and Chris Hankin. *Principles of Program Analysis*. Springer, 1999.
- Gerald Oskoboiny. HTML Validation Service, 2001. <http://validator.w3.org/>.
- Steven Pemberton et al. *XHTML 1.0: The Extensible HyperText Markup Language*. W3C, January 2000. W3C Recommendation, <http://www.w3.org/TR/xhtml1>.
- Dave Raggett, Arnaud Le Hors, and Ian Jacobs. *HTML 4.01 Specification*. W3C, December 1999. W3C Recommendation, <http://www.w3.org/TR/html4>.
- Anders Sandholm and Michael I. Schwartzbach. A type system for dynamic Web documents. In *Proc. Principles of Programming Languages (POPL'00)*. ACM, 2000.
- Sun Microsystems. Java Servlet Specification, Version 2.3, 2001. <http://java.sun.com/products/servlet>.

Sun Microsystems. JavaServer Pages Specification, Version 1.2, 2001. <http://java.sun.com/products/jsp>.

Vijay Sundaresan, Laurie J. Hendren, Chrislain Razafimahefa, Raja Vallee-Rai, Patrick Lam, Etienne Gagnon, and Charles Godin. Practical virtual method call resolution for Java. In *Proc. Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2000)*, pages 264–280. ACM, 2000.

Henry S. Thompson, David Beech, Murray Maloney, and Noah Mendelsohn. *XML Schema Part 1: Structures*. W3C, May 2001. <http://www.w3.org/TR/xmlschema-1/>.

Raja Vallee-Rai, Laurie Hendren, Vijay Sundaresan, Patrick Lam, Etienne Gagnon, and Phong Co. Soot - a Java optimization framework. In *Proc. IBM Centre for Advanced Studies Conference (CASCON'99)*, pages 125–135. IBM, 1999.

WAP Forum. *Wireless Markup Language, Version 2.0*. Wireless Application Protocol Forum, September 2001. <http://www.wapforum.org/>.

All <bigwig> papers are available from <http://www.brics.dk/bigwig/>.



## Recent BRICS Report Series Publications

- RS-02-11 Aske Simon Christensen, Anders Møller, and Michael I. Schwartzbach. *Extending Java for High-Level Web Service Construction*. March 2002.
- RS-02-10 Ulrich Kohlenbach. *Uniform Asymptotic Regularity for Mann Iterates*. March 2002. 17 pp.
- RS-02-9 Anna Östlin and Rasmus Pagh. *One-Probe Search*. February 2002. 17 pp.
- RS-02-8 Ronald Cramer and Serge Fehr. *Optimal Black-Box Secret Sharing over Arbitrary Abelian Groups*. February 2002.
- RS-02-7 Anna Ingólfssdóttir, Anders Lyhne Christensen, Jens Alsted Hansen, Jacob Johnsen, John Knudsen, and Jacob Illum Rasmussen. *A Formalization of Linkage Analysis*. February 2002. vi+109 pp.
- RS-02-6 Luca Aceto, Zoltán Ésik, and Anna Ingólfssdóttir. *Equational Axioms for Probabilistic Bisimilarity (Preliminary Report)*. February 2002. 22 pp.
- RS-02-5 Federico Crazzolaro and Glynn Winskel. *Composing Strand Spaces*. February 2002. 30 pp.
- RS-02-4 Olivier Danvy and Lasse R. Nielsen. *Syntactic Theories in Practice*. January 2002. 34 pp. This revised report supersedes the earlier BRICS report RS-01-31.
- RS-02-3 Olivier Danvy and Lasse R. Nielsen. *On One-Pass CPS Transformations*. January 2002. 18 pp.
- RS-02-2 Lasse R. Nielsen. *A Simple Correctness Proof of the Direct-Style Transformation*. January 2002.
- RS-02-1 Claus Brabrand, Anders Møller, and Michael I. Schwartzbach. *The <bigwig> Project*. January 2002. 36 pp. This revised report supersedes the earlier BRICS report RS-00-42.
- RS-01-55 Daniel Damian and Olivier Danvy. *A Simple CPS Transformation of Control-Flow Information*. December 2001.