

An Embedded Domain-Specific Language for Type-Safe Server-Side Web Scripting

PETER THIEMANN

Universität Freiburg, Germany

WASH/CGI is an embedded domain-specific language for server-side Web scripting. Due to its reliance on the strongly typed, purely functional programming language Haskell as a host language, it is highly flexible and—at the same time—it provides extensive guarantees due to its pervasive use of type information.

WASH/CGI can be structured into a number of sublanguages addressing different aspects of the application. The *document sublanguage* provides tools for the generation of parameterized XHTML documents and forms. Its typing guarantees that almost all generated documents are valid XHTML documents. The *session sublanguage* provides a session abstraction with a transparent notion of session state and allows the composition of documents and Web forms to entire interactive scripts. Both are integrated with the *widget sublanguage* which describes the communication (parameter passing) between client and server. It imposes a simple type discipline on the parameters that guarantees that forms posted by the client are always understood by the server. That is, the server never asks for data not submitted by the client and the data submitted by the client has the type requested by the server. In addition, parameters are received in their typed internal representation, not as strings. Finally, the *persistence sublanguage* deals with managing shared state on the server side as well as individual state on the client side. It presents shared state as an abstract data type, where the script can control whether it wants to observe mutations due to concurrently executing scripts. It guarantees that states from different interaction threads cannot be confused.

Categories and Subject Descriptors: D.1.1 [**Programming Techniques**]: Applicative (Functional) Programming; D.2.2 [**Software Engineering**]: Design Tools and Techniques—*Software libraries*; D.3.2 [**Programming Languages**]: Language Classifications—*Specialized application languages*; H.3.4 [**Information Storage and Retrieval**]: Systems and Software; I.7.2 [**Document and Text Processing**]: Document Preparation

General Terms: Design, Languages

Additional Key Words and Phrases: Interactive Web services, Web programming

This work was done during a stay the Oregon Graduate Institute, Portland, OR.

This work was supported by VolkswagenStiftung.

Author's address: Institut für Informatik, Albert-Ludwigs-Universität Freiburg, Georges-Köhler-Allee 079, D-79110 Freiburg, Germany; email: thiemann@acm.org.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2005 ACM 1533-5399/05/0200-0001 \$5.00

1. INTRODUCTION

Since its inception in 1991, the Web has dramatically changed its face. What started as a hypertext system for distributing scientific documents has changed into a ubiquitous global network for exchanging all kinds of data. Along with the content, the presentation of information in the Web has changed. Originally, Web servers acted like file servers that produced stored documents on demand. Nowadays, most Web pages are computed on the fly according to user profiles, language and image preferences, or the latest news. Not only are those pages generated dynamically, they also support all sorts of interaction, from chatting, blogging,¹ and browsing catalogs to buying groceries and booking entire travel arrangements online.

Clearly, the implementation of all those interactive services requires substantial effort for developing software that runs on the Web servers and/or associated application servers. Despite the fact that it is deceptively easy to program dynamic Web pages (using CGI [CGI 1.0 1999], for instance), it turns out that programming reliable interactive Web services is very hard. This is due to a number of reasons.

First, the protocol underlying all Web services is HTTP, the Hypertext Transfer Protocol [Fielding et al. 1999]. Much of HTTP's success comes from its simple one-shot request-response model of communication. However, most Web services require notions of session or state management to implement interactive transactions reliably. The lack of a session concept in HTTP places a burden on the programmer who has to implement sessions of for an interactive Web service.

Second, the original “script-centric” approach to programming a Web service is derived from the one-shot model. Each request starts a “script” on the server that generates the answer page and terminates. It is a nontrivial software engineering problem to get a number of scripts to cooperate even on a single server because their interdependencies are recorded in unchecked string references, that is, attributes containing URLs.

Third, all communication between browser and server is in the form of name-value pairs where the names are defined in one program (the script generating the currently displayed form) and used in another program (the script processing the submitted form). The requirement that defined and used names match gives rise to a nontrivial consistency problem among a set of programs. And, again, the problem is caused by the use of strings for external naming.

Fourth, the values communicated to the server are strings. For that reason each submitted value must be checked and erroneous submissions must be rejected before the actual processing can commence. To make the application reliable, a large fraction of application code must be dedicated to validate the input values [Scott and Sharp 2002].

Fifth, XHTML does not provide any means to structure forms into small interaction components. Hence, all form-related code is monolithic in the sense

¹Short for Web log, a blog is a Web page that serves as a publicly accessible personal journal for an individual [BLOG 2003].

that even a form with many unrelated components (as, for example, on a typical portal site) must be generated and processed by one chunk of code.

Sixth, the dynamically generated documents should be guaranteed to adhere to the XHTML standard. This can be tricky, in particular, if documents are generated from highly parameterized parts.

Seventh, the management of persistent state on the server is a difficult problem due to the nature of the browser interface that permits the use of the back button and cloning of windows.

Last but not least, there is a plethora of further considerations for programmers and maintainers of Web services. Among them is the whole area of security, the question of synchronizing accesses to shared resources on the server (database transactions, in particular) among different session threads, as well as issues pertaining to managing and evolving Web services.

Hence, there is a growing demand for languages and tools that support this kind of programming task. WASH/CGI addresses the first seven problems listed above by providing sublanguages for document generation, session management, widget composition, and state management. Despite the very specific problems of the Web programming domain, it is not necessary to design and implement a language from scratch. Instead, it turns out that a little language (a library) embedded in the purely functional programming language Haskell can solve the first seven problems. The key to make this work is the consistent use of abstract data types, type-safe interfaces, and compositionality in the design of the sublanguages.

Specifically, WASH/CGI offers the following contributions to Web programming.

- (1) WASH/CGI has a built-in notion of a session with transparently handled session state.
- (2) The session sublanguage allows the construction of complicated interactions from Web forms and other documents. Many interactions may be specified within one program and there is no need to assure the consistency of a set of programs.
- (3) Input values from a Web form are communicated via lexical binding. Since an external name is not required, this eliminates a common source of errors.
- (4) All input values are received by the script in their typed internal representation. There is no need to explicitly convert strings to values in a script. This conversion is performed and validated once and for all by the WASH/CGI implementation. Errors are flagged automatically, without further programming. Furthermore, there is an extension mechanism that allows programmers to define their own datatypes, validations, and conversions.
- (5) The widget sublanguage provides the means for specifying interaction fragments, where an arbitrary part of a form may be specified *together* with the code processing it. These interaction fragments may be composed and reused even across applications.

- (6) All documents generated by a WASH/CGI script are guaranteed to be well-formed (due to an internal tree representation) and to adhere to the XHTML standard to a large degree.
- (7) WASH/CGI provides abstract data types for persistent server-side and client-side state. Their interface is engineered to detect inconsistencies due to concurrent updates of the state information.

Meanwhile, a number of languages provide support for Web programming in one form or another (see the discussion in Section 10). However, contributions (3), (4), (5), and (7) are unique to WASH/CGI.

Since many readers will not be familiar with the Haskell language, Section 2 provides a brief overview of its main features. The code examples in the running text are designed to be understandable with just this introduction. After these preliminaries, Section 3 introduces the document sublanguage. It first discusses the basic interface and then enriches the interface to guarantee quasi-validity of generated documents. Section 4 first reviews the requirements on implementing sessions in Web applications and discusses some implementation schemes. Then it presents the basic WASH/CGI operators for sessions. Section 5 describes the widget sublanguage, which provides the building blocks for constructing typed interactive services. Various notions of state make sense for Web applications. Besides the session state that is an implicit part of every WASH/CGI application, Section 6 introduces the persistence sublanguage which supports two notions of persistent state, client-side and server-side state. Section 7 puts all of the above to work and develops a complete example. All implementation related discussion is gathered in Section 8. This section first explains the implementation ideas in general terms and concentrates the discussion of host language specifics to one subsection. Full appreciation of this subsection requires deeper knowledge of Haskell than can be obtained from Section 2. We include it to demonstrate the feasibility of the whole scheme and the simplicity and elegance of its implementation. The rest of the article performs a detailed assessment of the design decisions taken (Section 9), discusses related work (Section 10), and finally concludes.

2. PRELIMINARIES

This section is designed to convey just enough background on the Haskell programming language so that readers can understand the examples in the text and that they can code simple examples by themselves. The Haskell language report [Haskell98 1998] provides a complete reference.

Identifiers come in two flavors, alphanumeric and symbolic, and in two characteristics, for values and for data constructors. Symbolic identifiers like `+` are used like infix operators. Data constructors are distinguished by an initial capital letter or an initial colon “`:`” if they are symbolic. Comments start with two dashes--and stop at the end of the line.

2.1 Types and Their Values

Haskell is a strongly typed language with parametric polymorphism, overloading, and type inference. For each expression, the compiler computes a most general type and rejects programs that might cause type conflicts at runtime.

Basic Haskell types used in this article are

- `Int`, fixed-length machine integers with constants written in the usual way and equipped with the usual infix operations,
- `Bool`, truth values with elements `True` and `False`,
- `Char`, Unicode characters with constants written `'x'`, `'y'`, `'z'`, ... ,
- `String`, lists of Unicode characters, written in the familiar string notation `"abc"`, and
- `()`, the *unit type* with the single element `()`.

Haskell provides type constructors to build new types from existing ones. There are a number of standard constructors available. If t , t_1 , and t_2 are types, then so are

- $[t]$, the type of lists with elements of type t , a list is written as either $[e_1, e_2, \dots]$ or $e_1 : e_2 : \dots []$, where $:$ is the infix list construction operator and $[]$ is the empty list,
- `Maybe t` , the type of optional values of type t , where each element is either `Nothing` or `Just e` for some e of type t ,
- $t_1 \rightarrow t_2 \rightarrow \dots \rightarrow t$, the type of functions that map values of type t_1, t_2, \dots to a result of type t ,
- `IO t` , the type of I/O actions that yield results of type t .

Further types and type constructors may be defined by the programmer.

Haskell supports parametric polymorphism. Hence, types may contain type variables which can be substituted for by any type. For example, the function `length` that determines the length of a list has type `[a] -> Int`. In that type, `a` is a type variable indicating that `length` yields a result of type `Int` no matter what type is substituted for `a`, that is, no matter what the type of the elements of the list is.

2.2 Syntax of Declarations and Expressions

The declaration

$$f :: t$$

specifies a type signature for f , as in a C prototype or a Java interface. It declares that f has type t . This declaration is optional as the compiler can usually compute t from the definition of f . A definition for f has the form

$$f = e$$

or—if f is a function with formal parameters x_1, x_2, \dots, x_n —

$$f \ x_1 \ x_2 \ \dots \ x_n = e$$

In either case, e is an expression formed from

- variables x_1, x_2, \dots ;
- constants;
- infix operators applied to expressions, for example, $e_1 + e_2$;
- functions applied to expressions, for example, $f e_1 e_2 \dots e_n$;
- let $x_1 = e_1 \dots$ in e , definition of values x_1, \dots for use in e (and e_1, \dots);
- conditional expressions if e then e_1 else e_2 ;
- case expressions, which are used to decompose, for example, lists

```
length xs =
  case xs of
    []      -> 0
    x : xs -> 1 + length xs
```

and optional values

```
fromMaybe x opt =
  case opt of
    Nothing -> x          -- returns x if opt == Nothing
    Just y  -> y          -- returns y if opt == Just y
```

A unique syntactic feature of Haskell is the *offside rule*. It specifies scope by indentation and is used with `let`, `do`, `case`, and other expressions. Without specifying it formally, the idea is to vertically line up items in the same scope. For example, the start of the second alternative `Just y -> ...` in the above case expressions is indicated by indenting it to just below the respective first alternative.

2.3 Type Classes and Overloading

Haskell includes a notion of parametric overloading by restricting type variables to range over a proper subset of all types. These sets are called *type classes*. They are specified inductively and are associated with a set of overloaded member functions. For example, in

```
(==) :: (Eq a) => a -> a -> Bool
```

the type variable `a` is restricted by prepending the constraint `Eq a`. This declaration specifies the type of the (infix) equality function `==`. The definition of `==` for a particular type `t -> t -> Bool` (i.e, substituting `t` for the type variable `a` in the type of `==`) is provided via a special declaration syntax, the *instance declaration*. The same declaration makes a type `t` a member of the type class `Eq`.

The type class `Eq` is part of the Haskell standard. Further standard classes are `Read` and `Show`.

```
read :: (Read a) => String -> a
show :: (Show a) => a -> String
```

The class `Read` collects types `a` that have a parser function and the class `Show` collects types `a` that have an unparser function. The intention is that each type `t` in class `Read` has a member function `read :: String -> t` to convert a string into a value of type `t`. Similarly, the member function `show :: t -> String` converts a value into a string.

2.4 IO

I/O operations and other side effects cannot be performed at arbitrary places in a program because Haskell is a pure functional language where evaluation can be performed in arbitrary order. To deal with I/O, there is an abstract datatype `IO a` for *I/O actions* that yield values of type `a`. For example, the action that reads a character from standard input is `getChar` and its type is `IO Char`. There are operators to create trivial actions (`return x` is an action that does not perform I/O but just returns the value `x`) and to combine actions by putting them in sequence (`action >>= nextAction` is a combined action that first performs `action` and then passes the returned value to function `nextAction` that computes the following action depending on the returned value). I/O actions are passive (like a data structure) until they are processed at the toplevel of the program. At that point, they are executed in the sequence specified by the `>>=` operator.

The `do` notation is another syntax for composing actions:

```
do x <- action1
   y <- action2
   ...
   finalAction
```

This expression builds an action that first performs `action1` and binds its result to `x`. Next it performs `action2` (which may depend on `x`) and binds its result to `y`. And so on until `finalAction`, which may depend on all previously computed results. The binding part `x <-` may also be omitted. In this case, the result is ignored.

For example, the following program

```
main =
  do pwds <- readFile "/etc/passwd"
     putStrLn "Your password file:"
     putStrLn pwds
```

reads the contents of the file `/etc/passwd` into the string `pwds`, outputs the line `Your password file:`, and then goes on to print the file's contents. The results of the `putStrLn` actions are ignored.

The operations `return` and `>>=` as well as the `do` notation may be used for other kinds of actions. In fact, all these operations are overloaded using the standard type class `Monad`.

3. THE DOCUMENT SUBLANGUAGE

The document sublanguage supports the creation of XHTML documents [XHTML 1.0 2000]. As such, it has to provide the means to create XHTML document nodes, that is, element, attribute, text, and comment nodes. WASH/CGI does so by defining an element construction function for each element (it is named in the same way as the element it constructs), as well as one construction function for each other kind of node, attribute, text, and comment [Thiemann 2002b]. In addition, the WASH distribution includes a preprocessor that enables the direct use of XML markup syntax in the program (see Thiemann [2003b]), however, the article sticks to using the construction functions.

Each element construction function takes as an argument a sequence of child elements, text nodes, and attributes (arbitrarily mixed) and returns a singleton sequence that contains just the newly constructed element. Similarly, the function `attr` takes an attribute name *name* and its value *value* (both strings) and constructs a singleton sequence with the attribute assignment *attr=value*.² Also, the function `text` constructs a singleton sequence with a text node from a string.

The notion of a sequence of document nodes is a unique feature of the document sublanguage. Besides the element constructors, it provides operators for constructing empty sequences and for concatenating sequences.

3.1 Composing Documents

For reasons explained below in Section 5, sequences are constructed using actions as explained for I/O in Section 2.4. Hence, the operators for composing actions are also used to concatenate sequences of document nodes: the `do` notation and the binary infix operator `>>=`. There are two more alternatives for concatenation, `##` and `>>`, which will be explained below in Section 5. The operator `empty` constructs the empty sequence of document nodes. For example, the code fragment

```
do p (text "This is my first CGI program!")
  p (do text "My hobbies are"
      ul (do li (text "swimming")
              li (text "music")
              li (text "skiing"))))
```

creates a document with two paragraphs. The first one contains the text node *This is my first CGI program!* and the second one contains some text and an embedded itemized list. The XHTML output generated from this code looks as follows (pretty-printed for readability):

```
<p>This is my first CGI program!</p>
<p>My hobbies are
  <ul>
```

²The preferred notation `@name` for accessing the attribute *name* in XML processing languages cannot be adopted because the symbol `@` is reserved in Haskell. The closest possible approximation is the notation `(@@ name)value` with mandatory parentheses.

```

    <li>swimming</li>
    <li>music</li>
    <li>skiing</li>
  </ul>
</p>

```

As an example for generating an attribute, consider the code fragment

```

p (do attr "style" "color: red"
    text "This text is red")

```

which creates the following XHTML element that is rendered as a paragraph with red text:

```

<p style="color: red">This text is red</p>

```

A parameterized document is a Haskell function that constructs a sequence of document nodes. For example, WASH/CGI provides a standard document wrapper, `standardPage`, which takes two arguments, a title string `ttl` and a sequence of document nodes `nodes`, and returns a proper XHTML page. It is implemented as follows:

```

standardPage ttl nodes =
  html (do head (title (text ttl))
           body (do h1 (text ttl)
                   nodes))

```

This small example already demonstrates that our notation for documents scales easily to arbitrary parameterization. For this kind of document template to work satisfactorily, it is crucial that the `nodes` parameter can construct a *sequence* of document nodes, not just a single node. With WASH/CGI, `nodes` may provide attributes for the enclosing body element as well as an arbitrary number of element and text nodes which are to become children of `body`.

3.2 Valid Documents

A frequently recurring problem with generated XHTML documents is their validity, or rather lack thereof. A valid document is well-formed and adheres to its DTD, in this case the XHTML DTD [XHTML 1.0 2000]. Well-formedness just means that opening and closing tags are properly nested. Beyond that, the DTD places restrictions on which child nodes (attributes, text, other elements) may be used with which elements. The child elements and text are also subject to ordering constraints imposed by regular expressions. While it is easy to check the validity of a static document using one of the existing checkers (e.g., Validator [2002]), it is not easy to guarantee that a program generating XHTML will only produce valid documents.

Since the actions of the document sublanguage construct an internal tree representation of the document without exposing its string rendition to the programmer, well-formedness of the generated documents comes for free. To guarantee validity, it is possible to impose all constraints defined in a DTD,

by exploiting Haskell's type class-based overloading mechanism [Thiemann 2002b]. In practice, it turns out to be too restrictive to program fully DTD-compliant document generators in that way. For that reason, the document sublanguage approximates validity with *quasi validity*. Quasi-validity enforces the parent-child restrictions from the DTD, but it ignores constraints on the sequence and multiplicity of child nodes. See Section 9 for further discussion.

WASH/CGI enforces quasi validity through the types of the construction functions for document nodes:

```
text    :: (Monad m, AdmitChildCDATA e)
        => String          -> WithHTML e m ()
ul      :: (Monad m, AdmitChildUL e)
        => WithHTML UL m a -> WithHTML e m a.
```

The result type of each function has the form `WithHTML e m a` where `WithHTML` is an abstract type of sequences of document nodes or (more precisely) of document construction actions. Such an action can produce a result of type `a` and it can incorporate another kind of action `m` (e.g., `m` might be `IO` so that I/O can be performed during the construction of a document).

The remaining type parameter `e` of the `WithHTML` data type determines the name of the parent element to which a particular node will be attached. To apply the `text` node constructor, the parent node must admit a child of type `CDATA` as indicated by `AdmitChildCDATA e`. Similarly, the `ul` function can only be used inside a parent that admits a `ul` child element. In addition, its own children will all have a `UL` element as a parent, hence the type `WithHTML UL m a` for the argument of `ul`. An analogous mechanism may be employed for checking attribute occurrences.

WASH provides, for each kind of document node, a type and a type class. For example, the type `CDATA` models a text node and the type class `AdmitChildCDATA` models the set of nodes that take a text node as a child, the type `UL` models an element of name `ul` and the type class `AdmitChildUL` models the set of nodes that take a `ul` element as a child, and so on. The values of these types are not interesting, the types are only present to make information about the document available to Haskell's type checker.

The use of types for document nodes scales nicely to parameterized documents. For example, the `standardPage` function from Section 3.1 has type

```
standardPage :: (Monad m, AdmitChildHTML context) =>
    String -> WithHTML BODY m a -> WithHTML context m a
standardPage ttl nodes =
    html (do head (title (text ttl))
            body (h1 (text ttl) >> nodes)).
```

The type says that the `ttl` argument must be a string, that `nodes` must be a sequence of nodes admissible as children to a `body` element, and that the result must be used in a context that admits an `html` element. In this case, the type accurately reflects the requirements imposed by the XHTML DTD.

While the types involved are not always perspicuous, the programmer only gets exposed to them if the program may generate invalid documents. In that

case, a type error message like `No instance for (AdmitChtmlUL A)` arising from use of `'ul'` at... will indicate that the code attempts to generate a document with an `ul` element nested within an `a` element. Furthermore, it is usually unnecessary to provide type annotations like the one above for `standardPage` because Haskell's type inference algorithm will compute it at compile-time.

4. THE SESSION SUBLANGUAGE

To proceed from the simple one-shot request-response model dictated by HTTP to a more session-oriented programming model, we first discuss what comprises a session for an interactive Web service. Once we have laid out the requirements and the implementation options for sessions, we explain the design of WASH/CGI's session sublanguage.

4.1 Sessions on the Web

For Web scripting purposes, a session is an alternating sequence of Web forms and form submissions starting with a form submission:

```
POST url0; FORM f1; POST url1; FORM f2; POST url2;...
```

Each `POST` stands for an HTTP request sent by a client and each `FORM` stands for the corresponding response. Each `urli` is a URL [Berners-Lee et al. 1998] that addresses a particular resource on a server. For simplicity, we consider data that accompanies `POST` requests part of the `urli`. We also abstract from the implementation issue that one conceptual `FORM` response might result in the creation of new browser windows and that user input might be required in more than one browser window before a `POST` request is generated. Two conditions make such a sequence a session:

- (1) Each `FORM` response of the server is determined by the prefix of the sequence up to the response.
- (2) For each request `POST urli`, the `urli` is drawn from a set of URLs determined by `FORM fi`, with `url0` drawn from a set of external entry points.

The concept of a session implies the existence of a *session state* that is shared between the client and the server. Unfortunately, this is at odds with the underlying transport protocol of the Web, HTTP [Fielding et al. 1999]. Due to its stateless nature, it is simple to implement, but it leaves the burden of implementing sessions to the application programmer. To understand the issues involved in implementing sessions, we examine the computation of the `FORM` response from its request more closely.

Whenever the Web server receives a request for a particular kind of URL,³ it does not simply deliver the contents of the file addressed by the URL, but rather executes the file as an external program and passes the parameters of the request to that external program. A standard convention for passing these parameters is the Common Gateway Interface (CGI) [CGI 1.0 1999], hence the

³For example, referring to a file in a directory which is script-enabled or a file with suffix `.cgi`, depending on the configuration.

name “CGI scripts” for these programs. The CGI script computes the response FORM, returns it to the Web server, and terminates. Usually, the response is an XHTML document that contains XHTML *form* elements, so that the interaction can continue from that response.

To implement sessions under these conditions requires a CGI script to determine the current session state before it can do anything else.⁴ The usual technique to recover the session state is to store a token in each FORM so that the token is passed along with each POST *url* operation continuing the session. Such a token may be implemented in a number of well-known ways.

- (1) With URL rewriting, the token gets incorporated directly in the URL, either as a path component or via the parameter mechanism.
- (2) Using cookies [Kristol 2001], the token gets stored on the client side and it is passed along with each request for a qualifying URL.
- (3) Using hidden input fields [XHTML 1.0 2000], the token gets stored in the response document itself (without being displayed) and is passed to the server as part of the request, along with other parameters entered by the user.

In each case, the application maintains a mapping from tokens to current session states to recover the session state from a token. Quite often, the token only identifies the session but not the position inside the session.

Unfortunately, the token-based approach leads to a number of problems which are due to the way that users can navigate the Web [Graunke and Krishnamurthi 2002].

- (1) Users can hit the back button to go back in the history of the browser and resume the interaction from some earlier point.
- (2) Users can clone the session at any point and continue with each clone in different ways.
- (3) Users can store the current form in a file and may resume the interaction several months later. (Bookmarking)

Exploitation of the navigation features conflicts with the linear session model proposed at the start of this section. A tree structure with each FORM a potential branching point would provide a better model. Using the back button corresponds to pruning the current session by returning to an ancestor, cloning corresponds to pursuing more than one branch of the session at the same time, and bookmarking corresponds to electing an arbitrary node as a root node (it might also be considered as delayed cloning).

If the session involves any kind of dependency on server state (e.g., access to external databases), then a deviation from the linear progression may cause the user’s view of the session (the browser window) to be out of sync with the server’s state. Since cloning yields two apparently independent views of the session, at least one of them must be out of sync with the server’s state. Things are aggravated by the fact that the server can only observe the

⁴The compile-time part of the session state may be provided via the name of the script, but remembering run-time values requires other techniques.

effects of the navigation actions after the fact, but not the navigation action itself.

The consistency problems with browser navigation have lead researchers to two opposite points of view. One camp [Brabrand et al. 2002] considers browser navigation harmful and disallows it altogether. The other camp [Graunke and Krishnamurthi 2002] considers browser navigation as a new paradigm for graphical user interfaces.

Since browser navigation is an integral part of Web-based user interfaces, it seems better to find ways of supporting it than to abolish it right away. The discussion above indicates that using one token per session and one associated snapshot of the session state on the server is a bad idea. Using the back button will bring the server's state out of sync with what the user is actually looking at. Cloning is even worse because it mixes up the states from two or more interaction sequences.

Clearly, a new token must be assigned for each interaction step (i.e., for each generated form). However, this creates another problem because the server will have to store very many token-state associations. Particularly in the presence of bookmarking, there is no way for the server to figure out which of these tokens are still required and which can be discarded. Quite often, timeouts are used to address this problem, but they do not really solve the problem.

Hence, we are pursuing another option that relies on the observation that hidden fields may store large chunks of data whereas URLs and cookies are restricted in size. The idea is not to store *any* state information on the server but rather store the session information in the token itself. The server then puts the token into a hidden field of the response page returned to the client.⁵

That way, the server need not maintain a token-state association table (hence, there are no problems with discarding entries of such a table), restoring the state amounts to reading the token, and the state can be regenerated at each interaction step. In consequence, our approach works with arbitrary combinations of going back, cloning, and bookmarking.

Programmers often specialize the compile-time part of the session state by implementing each interaction step in a different CGI script. Keeping this set of programs in sync requires

- that each script run for the request `POST urli` only generates forms `FORM fi+1` that are understood by `urli+1` and
- that all scripts agree on the session state that is passed along and the way it is represented.

This consistency requirement is a nontrivial software engineering problem and it is the cause of annoying errors in many existing Web applications [Graunke et al. 2003]. The problems are not tied to CGI scripts but recur for other scripting techniques like Servlets and Java Server Pages.

⁵See Section 8.1.1 for a discussion of the security issues.

4.2 Sessions in WASH/CGI

WASH/CGI addresses the problems discussed in the previous section by

- enabling the programmer to implement the entire application in one program and
- providing a transparent notion of session state to the programmer.

In fact, the programmer's perspective on the Web application is similar to writing a GUI program. The difference is that the GUI is specified in XHTML and that the structure of the interaction is limited due to the underlying request-response protocol.

Sessions are constructed from documents by using two operators.

```
ask  :: WithHTML XHTML_DOCUMENT CGI a -> CGI ()
tell :: (CGIOutput a) => a -> CGI ()
```

Intuitively, `ask` takes an XHTML document and sends it to the browser. The document may contain a form with attached actions to continue the session. The `tell` operator takes any value that can be output in an HTTP response, converts it to the appropriate format, and sends it to the browser. The use of `tell` concludes the session.

A more technical description of the operators requires an explanation of the types. A value of type `CGI a` is a *CGI action* that computes a result of type `a`. CGI actions are the basic building blocks for sessions, they comprise the session sublanguage. CGI actions can be composed (i.e., put in sequence) using the operator `>>=` and the `do` notation.

A value of type `WithHTML context CGI a` is an *HTML action* that creates a sequence of document nodes. In fact, the XHTML construction operations introduced in Section 3 all have types analogous to `p`.

```
p :: (Monad m, AdmitChildP context) =>
    WithHTML P m a -> WithHTML context m a
```

The context `XHTML_DOCUMENT` in the type of `ask` indicates that the context expects a single `html` element to construct an XHTML document.

The `(Monad m)` constraint indicates that `m a` is the type of an *action returning a result of type a*. In the type `m a`, the type variable `m` ranges over *type operators* (constructors) whereas `a` ranges over types, as usual. That is, the implementation of `p` will work for all kinds of actions that we substitute for `m`. In many cases, `m` will be `CGI`, but sometimes other kinds of actions, like `IO`, are used.

Coming back to `ask` and `tell`, `ask` is a function that takes an HTML action and yields a CGI action. Conceptually, this CGI action constructs the document and then presents it to the browser.

`tell` is an overloaded function (as indicated by the `(CGIOutput a)` constraint of the type) that maps any value that can be converted to an HTTP response to a CGI action. The action delivers this response to the browser. For example, the response type `a` can be an `HTML Element`, a `String` (which is shown as `text/plain` document), a `FileReference` (which specifies a content type and a file name to be sent to the browser), a `StatusCode` (as defined in the HTTP

standard [Fielding et al. 1999]), or a `Location` (which specifies a redirection URL).

The `ask` operation is different from `tell` in that the XHTML document created by its HTML action argument may contain CGI actions that lead to further interaction steps. In contrast, a response returned by `tell` concludes an interaction sequence.

The typical script is specified as a single CGI action. Since Haskell only accepts an IO action as a main program, there is a function `run` that maps a CGI action to an IO action. Conversely, there is a function `io` that embeds an IO action into a CGI action.

```
run  :: CGI () -> IO ()
io   :: (Read a, Show a) => IO a -> CGI a.
```

The constraint `(Read a, Show a)` in the type of `io` restricts the values that may be returned from IO actions embedded in CGI actions.⁶

At this point, we can write our first script.

```
main =
  run (ask (standardPage "Hello World!" myinfo))

myinfo =
  do p (text "This is my first CGI program!")
    p (do text "My hobbies are"
          ul (do li (text "swimming")
                  li (text "music")
                  li (text "skiing")))).
```

The argument of `run` is a CGI action that displays the result of `standardPage` applied to the title "Hello World!" and the document `myinfo`. Since the latter document does not specify further CGI actions, the interaction is finished after the document is returned.

A typical use of `tell` would be to deliver file contents for downloading, indicate a redirection, or issue an error message. Here are some sample uses:

```
-- deliver a sound file "beback.wav"
tell
  FileReference { fileReferenceName = "beback.wav"
                 , fileReferenceContentType = "audio/x-wav"
                 , fileReferenceExternalName = "I'll be back"
                 }

-- relocate to a mail window
tell (Location ("mailto:" ++ email ++
               "?subject=" ++ subject ++
               "&body=" ++ message))
```

⁶This is further discussed in Section 8.

```
-- issue an error message
tell (Status 404 "Not Found" Nothing).
```

The first example applies `tell` to a record with three fields of type `String` (with field names `fileReferenceName` and so on). In the second example, the `++` operator performs string concatenation and `email`, `subject`, and `message` are supposed to be variables of type `String`.

5. THE WIDGET SUBLANGUAGE

In this section, we introduce the means to create interactive applications. In the first step, we attach simple continuation actions to a Web form. In the further steps, we introduce input widgets and explain how their values are passed to callback actions.

5.1 Continuation Buttons

Many operations of the widget sublanguages are wrappers for creating input elements with suitable attributes. The declaration

```
type HTMLField context a =
  WithHTML INPUT CGI () -> WithHTML context CGI a
```

defines a type synonym `HTMLField` that captures the common type for these wrapper operations.

A value of type `HTMLField context a` expects a sequence of attributes for the `INPUT` field (of type `WithHTML INPUT CGI ()`) and returns a one-element sequence with the input element.⁷

Creating a button and attaching a CGI action to it relies on the operation

```
submit0 :: (AdmitChildINPUT context) =>
  CGI () -> HTMLField context ().
```

Using a CGI action `action` of type `CGI ()`, `submit0` action constructs an input element with attribute `type="submit"`. In the browser, this element is rendered as a button and clicking the button submits the form to the server. This submission will trigger the action on the server. According to the XHTML specification, an input element is only legal inside of a form element. The latter should be created using the `makeForm` operation:

```
makeForm :: (AdmitChildFORM context) =>
  WithHTML FORM CGI a -> WithHTML context CGI ().
```

This operation is not called `form` because it does more than just creating a form element: It also fills in the attributes `action` and `method` and stores the current session token as a hidden field in the form.⁸

To extend the example, we might display the hobbies on a separate page. See Figure 1 for a rendering of the code below.

⁷Technically, the `context` type should be restricted to contexts admitting an input element, but this restriction is not expressible with a type synonym.

⁸See Section 8.1 for more information what is actually stored in this field.



Fig. 1. Rendering of “Hello World!”.

```

main =
  run page1

page1 =
  ask (standardPage "Hello World!" (makeForm myinfo1))

myinfo1 =
  do p (text "This is my second CGI program!")
      submit0 page2 (attr "value" "Click for my hobbies")

page2 =
  ask (standardPage "My hobbies are" (makeForm myinfo2))

myinfo2 =
  ul (do li (text "swimming")
          li (text "music")
          li (text "skiing")).
    
```

The main novelty is in the line

```
submit0 page2 (attr "value" "Click for my hobbies")
```

It specifies a submission button which executes the CGI action `page2` when clicked. The button displays the text specified by the `value` attribute.

The above combination of `ask`, `standardPage`, and `makeForm` is so common that it has become part of the standard library:

```

standardQuery :: String -> WithHTML BODY CGI a -> CGI ()
standardQuery ttl nodes =
  ask (standardPage ttl (makeForm nodes)).
    
```

5.2 Callbacks with Parameters

Input widgets are HTML actions that actually compute something. In addition to creating a suitable XHTML element with the correct set of attributes, a widget constructor also returns an *input handle*. It is via the input handle that the code can later extract the value entered into the widget. For example, the widget constructor

```

textInputField :: (AdmitChildINPUT context) =>
  HTMLField context (InputField String INVALID)
    
```



Fig. 2. Rendering of parameterized “Hello World!”.

creates a textual input field and returns its handle of type `InputField String INVALID`. Upon creation, this handle is *invalid* (as indicated by the `INVALID` in the type) because it has not received a value, yet. In contrast to other approaches, the programmer is never required to name the input widget, thus eliminating a common source of errors.⁹ Instead, the handle is bound to a variable and passed along to those callback actions that wish to inspect the value.

Passing an input handle to a callback action requires the use of the `submit1` operation. Its type is

```
submit1 :: (AdmitChildINPUT context) =>
  InputField a INVALID
  -> (InputField a VALID -> CGI ())
  -> HTMLField context ()
```

That is, in `submit1` handle callback, `submit1` takes an invalid input handle, `handle`, gets the value for the handle and turns it internally into a valid input handle (indicated by `VALID` in the type of callback’s argument), and then passes the latter as a parameter to `callback`. The result is again an `HTMLField`, which takes attributes for the button and constructs its XHTML representation, an input with attribute `type="submit"`. Inside the callback action, the operation `value` extracts the value from a valid handle:

```
value :: InputField a VALID -> a.
```

To illustrate parameter passing, we extend the example further by asking for the Web surfer’s name before showing the hobbies. Figure 2 shows the browser’s rendition of the code.

```
main =
  run page1

page1 =
  standardQuery "Hello World!" $
  do p (text "This is my third CGI program!")
    inf <- p (do text "Enter your name "
      textInputField (attr "size" "10"))
    submit1 inf page2 (attr "value" "Click for my hobbies")
```

⁹In those cases where the name of the input field is externally mandated (as with ECML [Eastlake, 3rd and Goldstein 2001]), the name attribute can be used explicitly.

```

page2 inf =
  standardQuery "My hobbies are" $
  do p (text "Hi, " ## text (value inf) ## text "!")
      ul (do li (text "swimming")
              li (text "music")
              li (text "skiing")).

```

The \$ infix operator used in the above code stands for function application. Using the infix operator (instead of just juxtaposition) avoids the need for putting the two do blocks in parentheses.

At this point, we also have to explain how the values computed by HTML actions are propagated. The notation

```

do ...aaa...
  x <- htmlAction
  ...bbb...
  finalAction

```

extracts the value computed by `htmlAction` and binds it to `x`. The variable `x` is visible inside the remainder of the action, that is, in `...bbb...` and `finalAction`. The result computed by the entire `do` expression is the result of `finalAction`.

Each XHTML constructor (e.g., `p`, `ul`, `li`, ...) returns the value computed by its argument action, as evidenced by their type, for example,

```

p :: (AdmitChildP context) =>
    WithHTML P CGI a -> WithHTML context CGI a.

```

The operators `##` and `>>` differ in the way that they propagate the results of their operands. They can be easily explained in terms of the `do` notation:

```

act1 ## act2 =
  do x <- act1      -- result is returned
     act2           -- result is ignored
  return x

```

```

act1 >> act2 =
  do act1           -- result is ignored
     act2           -- result is returned

```

Applying the above rules, in

```

do p (text "This is my third CGI program!")
  inf <- p (do text "Enter your name "
               textInputField (attr "size" "10"))
  submit1 inf page2 (attr "value" "Click for my hobbies"),

```

the result of the `do` block is the result of `textInputField ...`, that is, the input handle for the field. This handle is then passed through by the `p` constructor so that the variable `inf` is bound to the input handle.

5.3 Typed Input Fields

The `textInputField` of the previous subsection was only good for input of strings. However, there is a more general operation

```
inputField :: (AdmitChildINPUT context, Reason a, Read a) =>
             HTMLField context (InputField a INVALID)
```

that creates a *typed input field* for any type `a`, provided this type belongs to type classes `Reason` and `Read`. The constraint `Reason a` implies that the type `a` has an operation `reason :: a -> String` which provides an explanation of the input syntax for a value of type `a`. The constraint `Read a` implies that type `a` has an operation `read :: String -> a` which converts a string into a value of type `a`. The latter class `Read` is a standard type class of Haskell [Haskell98 1998]. Of course, the explanation generated by `reason` should match the syntax parsed by `read`.

Clearly, there is no guarantee that a Web surfer enters a parsable value into a type field. For that reason, the `submit1` function *validates* each handle before it passes it on to the callback action. A handle is validated only if its input is parsable to a value of the desired type. Hence, a valid handle guarantees the presence of a value. If `submit1` cannot validate a handle passed to it, it redisplay the last form with the erroneous input marked.

As an example, consider the code below which reads an integer and displays it again.

```
main =
  run page1

page1 =
  standardQuery "Numeric Input" $
  do inf <- p (do text "Enter a number "
                 inputField (attr "size" "10"))
             submit inf page2 (attr "value" "Process")

page2 inf =
  let n :: Int
      n = value inf
  in
  standardQuery "Your Number" $
  p (text "Your number was " ## text (show n) ## text "!").
```

In `page1`, the essential change is the use of `inputField` instead of `textInputField`. In `page2`, the `n :: Int` in the `let` declares that `n` is an integer and then extracts the integer value from the input handle `inf`. The function `show` converts (in this example) an integer to a string.

Without any further programming, the program behaves as shown in Figure 3. The left-hand snapshot shows the browser window after typing in an incorrect input, but before clicking the “Process” button. The right-hand snapshot shows the window after clicking the button. The erroneous input field

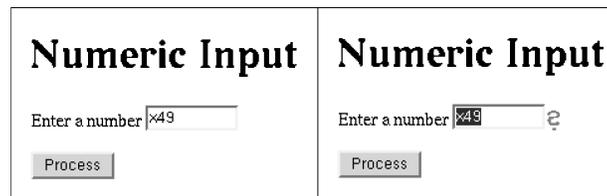


Fig. 3. Capturing errors with typed input fields.

is marked with a red question mark and its contents are selected. In addition, the reason phrase appears in the status line of the browser as soon as the pointer is above the input field.

With this approach it is easy to create a new input field with customized syntax. It amounts to defining a new type and providing a parser and a reason function for it. Doing so requires some skill in Haskell programming. The following example defines a type `NonEmpty` for nonempty strings and then provides a parser (instance `Read NonEmpty`) and a description of the input syntax (instance `Reason NonEmpty`).

```
newtype NonEmpty = NonEmpty { unNonEmpty :: String }
-- NonEmpty str constructs a NonEmpty value from a string
-- unNonEmpty n extracts the string from a NonEmpty value

instance Read NonEmpty where
  readsPrec i =
    do str <- many1 anyChar      -- regular expression: .+
       return (NonEmpty str)    -- construct value of type NonEmpty

instance Reason NonEmpty where
  reason _ = "non empty string".
```

To avoid an extended discussion of Haskell, we will not comment on it further at this point. However, the code snippet should make it clear that defining a new input field type is a simple programming task.

Using the new `NonEmpty` type is as simple as using numeric input in the example above:

```
page2 inf =
  let nonemptystring = unNonEmpty (value inf) in
  -- nonemptystring is guaranteed to be a non-empty string.
```

5.4 Nontextual Input Fields

The remaining input widgets (buttons, check boxes, selection boxes, images, etc) are all be expressed in terms of typed input handles in WASH/CGI. For example, check boxes yield handles of type `InputField Bool INVALID`, buttons can be made to return any type a provided values of type `a` have a string representation, selection boxes can return any value whatsoever, and images yield a handle of type `InputField (Int, Int) INVALID`, that is, they return a pair of

integers. Details about them may be found in the WASH/CGI documentation [Thiemann 2003a].

5.5 Combining Input Fields

Up to now, the examples have only dealt with callback actions that take a single handle as a parameter. To pass multiple handles to a callback requires a more general type than the one given for `submit1`. The function `submit` is this upwards-compatible generalization for `submit1` and its type is

```
submit :: (AdmitChildINPUT context, InputHandle h) =>
         h INVALID -> (h VALID -> CGI ()) -> HTMLField context ().
```

The type `h INVALID` generalizes the former `InputField a INVALID`. The type class `InputHandle` restricts the possible instantiations of `h`.¹⁰ Creating a pair of handles requires the use of a special pairing data constructor,

```
F2 :: a x -> b x -> F2 a b x.
```

Furthermore, there is a rule (the instance declaration for class `InputHandle` and type `F2`) that states that `F2 h1 h2` belongs to `InputHandle` if both `h1` and `h2` are `InputHandles`, too. Technically, the type `F2 h1 h2` is a partially applied (type) function that distributes its third, `x`, argument to its first and second argument. For convenience, WASH/CGI defines types similar to `F2` for other tupling operators and for lists.

As an example, the following code may be used to query for a name and a password. The check function enforces that both name and password fields are non-empty. The function `passwordInputField` works just like `inputField` but generates an input element with `type="password"` so that the characters entered are not echoed.

```
login =
  standardQuery "Login" $ table $
    do nameF <- tr (td (text "Name ") >>
                   td (inputField (attr "size" "8")))
       passF <- tr (td (text "Password ") >>
                   td (passwordInputField (attr "size" "8")))
       tr (td (submit (F2 nameF passF) check (attr "value" "LOGIN"))))

check (F2 nameF passF) =
  let name = unNonEmpty (value nameF)
      pass = unNonEmpty (value passF)
  in
  standardQuery "Authenticate" ...
```

¹⁰The variable `h` ranges over *type constructors* and the type class `InputHandle` specifies a set of *type constructors*, not plain types. This is a standard feature of Haskell [Haskell98 1998].

| | | |
|-----------------------|----------------------|---------|
| <input type="radio"/> | Pay by Credit Card | |
| | <input type="text"/> | Card No |
| | <input type="text"/> | Expires |
| <input type="radio"/> | Pay by Bank Transfer | |
| | <input type="text"/> | Acct No |
| | <input type="text"/> | Routing |
| Submit Query | | |

Fig. 4. Rendition of the payment form.

5.6 Virtual Input Fields

In some cases, it is useful to have input fields that do not have a screen rendition. An example for such a field arises with the use of the standard radio input type. An input element of type `radio` gives rise to a radio button. In HTML, all radio buttons of the same name form an implicit radio group, that is, the browser makes sure that at most one of the radio buttons of this group is pressed.

WASH/CGI makes this concept explicit by providing a function

```
radioGroup :: Read a =>
    WithHTML context CGI (RadioGroup a INVALID).
```

Its use fixes the type of the value returned by any radio button that belongs to this group to be `a`. The button constructor `radioButton` takes as parameters the returned handle of type `RadioGroup a INVALID` to attach buttons to the group and a value of type `a` to be returned when the button is activated.

```
radioButton :: (AdmitChildINPUT context, Show a) =>
    RadioGroup a INVALID -> a -> HTMLField context ()
```

This design makes sure that

- radio buttons are not grouped by accident as they always have to be created inside a specific radio group,
- all buttons of a group return a value of the same type, and
- error processing can be done once and for all.

The next section has an example of using radio buttons.

5.7 Decision Trees

The `submit` function explained above requires that each callback validates one fixed subset of the input fields present in the form. For some forms, this requirement is too restrictive. For example, consider a typical form fragment from the payment processing part of a shopping application rendered in Figure 4:

```
data ModeOfPayment = PayCredit | PayTransfer
-- declares an enumerated data type ModeOfPayment
-- with two symbolic values PayCredit and PayTransfer
```

```

paymentOptions =
  do rg <- radioButton rg PayCredit empty
     tr (do td (radioButton rg PayCredit empty)
           td (text "Pay by Credit Card"))
     ccnrF <- tr ((td empty >> td (inputField (attr "size" "16")))
                ## td (text "Card No"))
     ccexF <- tr ((td empty >> td (inputField (attr "size" "5")))
                ## td (text "Expires"))
     tr (do td (radioButton rg PayTransfer empty)
           td (text "Pay by Bank Transfer"))
     acctF <- tr ((td empty >> td (inputField (attr "size" "10")))
                ## td (text "Acct No"))
     routF <- tr ((td empty >> td (inputField (attr "size" "8")))
                ## td (text "Routing"))
     tr (td (submit (F5 rg ccnrF ccexF acctF routF) payCGI empty))

payCGI (F5 rg ccnrF ccexF acctF routF) =
  let ccnr = unCreditCardNumber (value ccnrF)
      expMonth = cceMonth (value ccexF)
      acct = unAllDigits (value acctF)
      rout = unAllDigits (value routF)
  in
  standardQuery "Payment Confirmation"
  (text "Thanks for shopping at TinyShop.Com!")

```

It contains a group of two radio buttons and four textual input fields. The variable `rg` is the handle to the value representing the selected radio button and the variables ending in `F` are handles to the values of the other input fields. The last line submits the handles of the two radio buttons, `rg`, the credit card number, `ccnrF`, its expiration date, `ccexF`, the account number, `acctF`, and its routing number, `routF`, to the callback `payCGI`. The form is fully typed and set up to reject unreasonable inputs without any further programming due to the way the handles are accessed in `payCGI`. For example, the credit card number and the account number are both submitted to the Luhn check,¹¹ the expiration date must have the form as printed on the card, and so on. Before submitting the form, a customer would select one mode of payment by clicking one of the radio buttons and then fill in the two fields associated to it. However, the validation built into `submit` tries to validate *all five* input handles, thus insisting that both credit card data *and* the full account information are present!

To handle such cases requires introducing an additional abstract data type, a decision tree of type `DTree context`, where the `context` argument is used for validity checking as explained in Section 3.2. A decision tree is either a leaf that contains an unconditional CGI action or it is a node that contains a group of input handles (of type `h INVALID`) and a function that maps their validated

¹¹A standard checksum algorithm for credit card numbers and bank account numbers, see for example [LuhnCheck 1999].

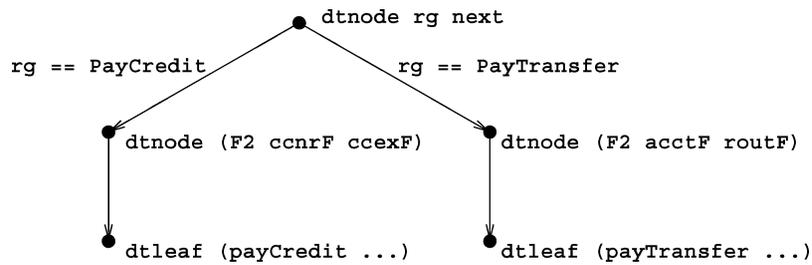


Fig. 5. Decision tree for the payment example.

versions to the next decision tree. Hence, decision trees have two constructor functions:

```

dtleaf :: CGI () -> DTree context
dtnode :: InputHandle h =>
        h INVALID -> (h VALID -> DTree context) -> DTree context
    
```

The first, `dtleaf`, creates a trivial decision tree from a CGI action. The `dtnode` constructor takes an invalid input handle, validates it, and (upon successful validation) passes it to the second argument that builds the rest of the decision tree. The function `submitx` transforms a decision tree into a submit button that can be put into a Web form:

```

submitx :: (AdmitChildINPUT context) =>
        DTree context -> HTMLField context ()
    
```

It turns out that the original submission operators are expressible with the decision tree-based operators. Submitting through the trivial decision tree, `submitx (dtleaf action)`, is equivalent to `submit0 action` and submitting through a one-stage decision tree, `submitx (dtnode handles (dtleaf . callback))`, is equivalent to `submit handles callback`.¹²

The decision tree required in the example first dispatches on the payment mode extracted from the handle `rg` for the radio buttons and then validates either the credit card fields or the account information. The leaves of the tree just invoke the respective code for processing the payment, `payCredit` or `payTransfer`, where the computed price is assumed to be present in variable `price`. Figure 5 shows the tree constructed in the code below.

```

let paymentOptions = dtnode rg next
    next paymodeF =
        case value paymodeF of
            PayCredit   -> dtnode (F2 ccnrF ccexF)
                               (dtleaf . payCredit price)
            PayTransfer -> dtnode (F2 acctF routF)
                               (dtleaf . payTransfer price)
in tr (td (submitx paymentOptions empty))
    
```

¹²The infix dot operation, `.`, stands for function composition.

```

payCredit price (F2 ccnrF ccexF) =
  let ccnr = unCreditCardNumber (value ccnrF)
      expMonth = cceMonth (value ccexF)
  in standardQuery "Confirm Credit Payment" $
    p $ text "Thanks for shopping with us!"

```

The `payTransfer` function is similar.

Further constraints on input values might be worth checking. For example, two password fields should contain the same entry, a list of values should be entered in ascending order, etc. Such constraints fall outside the realm of decision trees and must be enforced by the application.

5.8 Summary

The examples above demonstrate that the widget sublanguage provides the means of specifying and combining input elements in almost arbitrary ways. Fragments of interactions can be specified and implemented separately, and later be combined on one Web page. Even in this combination, all interactions are still fully type safe and every submitted data is checked for validity without programmer intervention.

All validation techniques discussed may be implemented in two locations, on the server and on the client. While the server-side implementation is mandatory (and part of the WASH implementation), a client-side implementation is also feasible by generating suitable JavaScript programs [Netscape Corporation 2003]. Client-side validation is desirable because it reduces latency by avoiding communication with the server.

Beyond the facilities discussed, it is also possible to compose new widgets from existing ones. However, the techniques required to do so are out of the scope of this paper.

6. THE PERSISTENCE SUBLANGUAGE

WASH/CGI provides two abstractions for persistent state beyond the builtin notion of session state: server state and client state. The main difference between these two is their scope. While the server state is globally visible to all programs running on the server, the client state is only visible to scripts started from one particular client. Hence, the global state acts as a database with which scripts may interact while the client state keeps information particular to one client (i.e., cookies [Kristol 2001]).

Compared to the usual string-based interfaces to server state and client state, the WASH/CGI interface is fully type-safe, even across separate programs. That is, any attempt to store data of one type and retrieve it as a value of another type will fail.

6.1 Server-Side State

Server-side state is defined as a parameterized, abstract data type `T a` with the following signature.

```

init :: (Types a, Read a, Show a) => String -> a -> CGI (T a)
get  :: (Types a, Read a)         => T a      -> CGI a

```

```

set  :: (Types a, Read a, Show a) => T a -> a    -> CGI (Maybe (T a))
add  :: (Types a, Read a, Show a) => T [a] -> a    -> CGI (T [a])
current :: (Types a, Read a)      => T a        -> CGI (Maybe (T a))

```

A value of type $T\ a$ is a handle to a *persistent entity* (i.e., a value in the server-side state) of type a . These handles are the only means of accessing the state from the script. The idea is that each handle uniquely identifies a particular *value* of the persistent entity. Hence, getting a handle may be considered as getting a snapshot of the persistent entity. We call a handle *current* as long as its snapshot still reflects the current server-side state. In the implementation, currency of a handle is checked using version numbers. This design addresses two issues.

First, handles help keep the session state small: Since a handle identifies a value, the session state need not contain the value which may be very large.

Second, a handle allows the script to explicitly notice changes to the value by concurrently executing scripts. The script writer can choose to either continue to use the value denoted by the handle (which is sufficient as long as no `set` operation is attempted) or to observe the new value by discarding the handle and obtaining a current handle. The latter choice places the obligation on the programmer to check if the new value is still consistent with the assumptions of the program.

The constraints placed on the types of the interface are of a technical nature. The constraint `Types a` indicates that there must be a term representation for the type a and the constraints `Read a` and `Show a` indicate that there must be conversion functions to and from strings for values of type a .

The operation `init name initialValue` attempts to create a new persistent entity with name `name` and initialized with `initialValue`. If it succeeds in creating a new entity, it returns a handle pointing to `initialValue`. If it cannot create a new entity because there is already a persistent entity of the same name and type, then the returned handle points to its current value.

The operation `get handle` returns the value of `handle`. This value may not be current because the server-side state may have evolved further after creation of the handle.

The operation `set handle newValue` attempts to overwrite the value of the persistent entity pointed to by `handle` with `newValue`. If `handle` is current, then it succeeds and returns a handle to the new value. Otherwise, it fails. In the latter case, it is up to the programmer to further deal with the situation.

The operation `add handle newValue` deals with the special case where the persistent value is a set, represented by a list where the order of elements does not matter. In that case, adding `newValue` to the set always makes sense, regardless whether the handle is current. Hence, `add handle newValue` never fails and the value of the handle returned from the `add` operation is guaranteed to contain `newValue`.

Finally, the operation `current handle` returns a current handle to the persistent entity pointed to by `handle`. This operation is potentially dangerous because it explicitly allows to observe changes to the persistent entity made

by concurrently running scripts. After using `current`, it is the obligation of the programmer to check that the value identified by the handle is still the intended one.¹³

The problem in implementing this interface is that the server might be forced to keep all different values of a persistent entity until the end of time. However, unlike session state, users will expect that server-side state evolves. Hence, the server can reasonably garbage collect values that have not been used recently and discard them. A script that tries to get a handle that has been discarded receives an exception which can be translated into a user-level error message.

Here is a snippet from an example program that uses a persistent entity to store a list of scores and player names for a game.

```
gameover myScore =
  do initialHandle <- init "hiscores" []
     currentHandle <- add initialHandle myScore
     hiScores      <- get currentHandle
     standardQuery "Hall of Fame" (ul (mapM_ showItem hiScores))

showItem (score, name) =
  li (do text name
        text ": "
        text (show score)).
```

In the second line, the variable `initialHandle` becomes a handle to the current value of the persistent entity named `hiscores`, which is initialized to the empty list `[]`. The next line adds the new score to the list of scores as explained above. The fourth line gets a list of scores that is guaranteed to contain `myScore`. The sixth line generates a page containing an unordered list where the individual list items are generated by applying the function `showItem` to each element of `hiScores` (that's what `mapM_` does). The function `showItem` takes a pair consisting of a string name and an integer score and generates a list item from it.

6.2 Client-Side State

The interface for accessing persistent client-side state is very similar to the one for server-side state, but there are subtle differences in their behavior. The operations are also given in terms of a parameterized abstract data type $T\ a$.¹⁴ They provide an abstraction layer on top of the implementation in terms of cookies [Kristol 2001].

```
check  :: (Read a, Show a, Types a) => String    -> CGI (Maybe (T a))
init   :: (Read a, Show a, Types a) => String -> a -> CGI (T a)
get    :: (Read a, Show a, Types a) => T a      -> CGI (Maybe a)
set    :: (Read a, Show a, Types a) => T a -> a   -> CGI (Maybe (T a))
```

¹³The semantics of persistent entities and the `current` operation in particular has been misinterpreted by Graunke and others [Graunke et al. 2003] who claim that the WASH/CGI implementation of server-side state is faulty.

¹⁴A program that uses both, client-side and server-side state, must keep them apart using Haskell's module system.

```

current :: (Read a, Show a, Types a) => T a      -> CGI (Maybe (T a))
delete  :: (Types a)                  => T a      -> CGI ()

```

Although a `handle` still denotes a snapshot of the client-side state, the value corresponding to this snapshot may no longer be available by the time the script accesses the `handle`. This can happen if the user runs several scripts dealing with the same client-side value in different browser windows at the same time.

For that reason, `get handle` may fail! This is reflected in its returning a value of type `Maybe a`: if successful, it returns `Just theValue`, otherwise it returns `Nothing`. In the latter case, the script must obtain a new `current handle` to continue.

The operations, `init`, `set`, and `current`, work as for server-side state. In addition, there are operations `check`, to check for the existence of a cookie and return a `handle` on success, and `delete`, for removing a cookie. There is currently no way for a program to check if setting the cookie on a client is successful. If the client does not return the cookie with the next request an error message will be posted to the browser saying that cookies need to be enabled to run the script. Beyond the operations shown here, there are further operations to manipulate the expiration time of client-side state.

WASH/CGI does not impose any restriction on where in the code the client-side state may be read or written. This is in contrast to other Web programming libraries that expose details of the cookie implementation by requiring that all cookie related computations must be completed before any part of the response page is sent to the browser. This restriction is due to the fact that cookies are part of the HTTP headers, which must be flushed before the contents of the page are sent.

7. A COMPLETE EXAMPLE

In this section, we give the code for a complete example, a guess-number game inspired by code from various tutorials on Servlets [Sun Microsystems 2003] as well as from the original paper on MAWL [Ladd and Ramming 1995]. In the game, the script first generates a secret number. Then the user tries to guess the number guided by the answers “too small” or “too large”. Guessing continues until the number has been found. To make the example more interesting, we are using a cookie to store the secret number (although session state would be sufficient) and we have added a high score list to demonstrate the use of persistent entities. The full code comprises 83 lines. The script is part of a suite of demo programs on the WASH/CGI web page [WASH 2001].

In the code, the cookie (client-side state) module is accessed via name `C`, whereas the persistence (server-side state) module is named `P`. Hence, the operations as well as the types from either module are prefixed with either “`C.`” or “`P.`”. For instance, `P.T` is the type of a server-side state handle whereas `C.T` is the type of a client-side state handle.

We start with code to initialize the high score list to the empty list of scores. The name of this persistent entity is `GuessNumber`. A score is pair of the number of guesses used and the name of the player.

```

type NumGuesses = Int
type PlayerName = String
type Score = (NumGuesses, PlayerName)

```

```

highScoreStore :: CGI (P.T [Score])
highScoreStore = P.init "GuessNumber" [].

```

This code just defines `highScoreStore` as a CGI action that initializes a handle to the persistent store. The action is only executed when incorporated into the main action.

```

type SecretNumber = Int
setNumber :: SecretNumber -> CGI (C.T (NumGuesses, SecretNumber))
setNumber secret = C.create "theNumber" (0, secret).

```

The operation `setNumber` defines the initializer for the cookie. The cookie contains a pair of the current number of guesses and the secret number.

```

main :: IO ()
main = run mainCGI

```

```

mainCGI :: CGI ()
mainCGI =
  standardQuery "Guess a number"
    (do submit0 playTheGame (attr "value" "Play the game")
        submit0 admin (attr "value" "Check scores"))).

```

The first two lines are standard for starting a WASH/CGI program. The action `mainCGI` implements the start of the interaction. It presents the player with a choice to either run the program or check the high scores, which are both activated using a parameterless button (`submit0`).

Of the two choices the `admin` function is the simpler one.

```

admin =
  do highScoreList <- highScoreStore
     highScores <- P.get highScoreList
     standardQuery "GuessNumber - High Scores" $ do
       table $ do
         attr "border" "border"
         tr (th (text "Name") ## th (text "# Guesses"))
         mapM_ showScore (sort highScores)

```

```

showScore (guesses, name) =
  tr (td (text name) ## td (text (show guesses)))

```

`admin` first initializes a handle `highScoreList` to the high score store and reads its contents into `highScores`. Then it presents a high score table with two columns, one for the player's name and one for the number of guesses. The function `mapM_` applies function `showScore` to each element of the sorted list

of high scores and pastes the resulting HTML actions together. The function `showScore` translates a score into one row of the enclosing HTML table.

The other choice in the menu is the `playTheGame` function.

```
message = text "I've thought of a number between 1 and 100."

playTheGame =
  do aNumber <- io (randomRIO (1,100))
    numHandle <- setNumber aNumber
    play numHandle message.
```

It first chooses a random number between 1 and 100 (inclusive) using the predefined I/O action `randomRIO`. Then it stores the number (and the number of guesses so far) in a cookie and falls through to function `play`.

The function `play` is parameterized over the handle to the cookie `aHandle` and the message `aMessage` to display. From the main menu it is called with text "I've thought of a number ...". This message is not restricted to be a text, it may contain any document nodes that are appropriate at this place. The actual code fragment dealing with the interaction is placed in a paragraph after the message.

```
play aHandle aMessage =
  standardQuery "Guess a number" $
    do aMessage
      p (do text "Make a guess "
          guessF <- inputField (attr "size" "3")
          submit guessF (processGuess aHandle) empty).
```

`play` is straightforward. It displays the message, asks for a guess, extracts a handle to the guess from the input field and passes it to `submit`. The callback function `processGuess` performs the actual work.

```
processGuess aHandle guessF =
  let aGuess = value guessF in
  do mValue <- C.get aHandle
    case mValue of
      Just (nGuesses, aNumber) ->
        let nGuesses' = nGuesses + 1 in
        if aNumber == aGuess
        then do C.delete aHandle
              youGotIt nGuesses' aNumber
        else do Just aHandle' <- C.set aHandle (nGuesses', aNumber)
              play aHandle' (do text "Your guess "
                              text (show aGuess)
                              text " was too "
                              if aGuess < aNumber
                              then (b (text "small."))
                              else (b (text "large."))))
```

```

Nothing ->
  standardQuery "Don't do that!" $ do
    text "You are trying to outwit me by playing with "
    text "the back button and by cloning windows!"
    submit0 (return ()) (fieldVALUE "Restart").

```

Given the handle to the cookie and the next guess, `processGuess` first tries to read the current value of the cookie. If this fails for some reason (case `Nothing`), then it displays an appropriate error message. Otherwise, it increments the number of guesses and checks if the guess was correct. If it was, then the cookie is deleted and control is passed to `youGotIt` to display a winning message. Otherwise, the cookie is updated to the new number of guesses and `play` is invoked recursively with a message (some text nodes and a `b` element) announcing the result of comparing the two numbers.

```

youGotIt nGuesses aNumber =
  standardQuery "You got it!" $
  do text "CONGRATULATIONS!"
     br empty
     text ("It took you " ++ show nGuesses ++ " tries to find out.")
     br empty
     text "Enter your name for the hall of fame "
     nameF <- textInputField (attr "size" "20")
     submit nameF (addToHighScore nGuesses) empty.

```

There is not much to say about `youGotIt`. It displays a parameterized HTML page and then falls through to `addToHighScore` as soon as a name is entered in the input field.

```

addToHighScore nGuesses nameF =
  let name = value nameF in
  if name == "" then restart else
  do highScoreList <- highScoreStore
     P.add highScoreList (nGuesses, name)
     restart.

```

The function first obtains the contents of the input field (unrestricted text). If the field is non-empty, then it gets a handle to the persistent `highScoreList` and adds the new score to it. In either case, the interaction wraps back to the start using the predefined `restart` action.

8. IMPLEMENTATION

The description of the implementation comes in two parts. The first sections describe the techniques in general terms and the last section gives implementation details specific to Haskell. This last section is not essential for understanding the techniques used. However, it conveys that the techniques map to Haskell code in a natural way.

8.1 Sessions

In Section 4, we stated that each form sent to the browser contains a hidden token for the session state. Clearly, sheer size of the data does not permit to store a memory snapshot of the application in the form.

At this point, Haskell's purely functional nature becomes essential. Since all functions definable in Haskell are free of side effects, they deliver the same result whenever they are applied to the same arguments. I/O is achieved through an abstract data type `IO a` of I/O actions. Manipulation of these actions is still free of side effects, actions are only executed when they are returned from the main function. Interactive programs are possible because later I/O actions may depend on the results of earlier I/O actions.

Since all external influences to a Haskell program are mediated through I/O actions, it is possible to replay a Haskell program by keeping a log of the results of the first run of each I/O action. For this to work, the program must be denied direct access to I/O by interposing another abstract data type, in our case `CGI a`, that encapsulates the replay mechanism.

The replay mechanism maintains the log in the form of two lists, a *recorded list* and a *replay list*. The replay list is examined for the results of operations that have already been done. It is consumed during replay. The recorded list keeps all previous results and is extended with new results, once the replay list is depleted. To initiate a replay, both lists are initialized with the saved recorded list from a former run.

Concretely, here is how the operations `io` and `ask` of Section 4.2 work. The CGI action `io ioAction` first checks the replay list for a saved result. If the result is available, it is removed from the replay list and passed to the next CGI action (without executing `ioAction`). If the replay list is depleted, then `ioAction` is executed, its result is appended to the recorded list, and also passed to the next CGI action.

The CGI action `ask document` is slightly more complicated because its results are not immediately available. Instead, they must be extracted from the next request received from the browser, in another execution of the same program. Like `io`, `ask` first checks the replay list for a saved result. If the result is available, it is removed from the replay list and passed to the next CGI action, which is specified by `document` (details follow in Section 8.2). If the replay list is depleted, then `document` is converted to an HTML form and passed to the server as an HTTP response. At this point, the program terminates. However, the returned HTML page has a hidden input field which contains the current value of the recorded list, that is, all responses to `io` and `ask` actions leading up to that form. When the script receives the answer from the form, the request from the browser contains

- (1) the last value of the recorded list and
- (2) the input parameters for the last form.

From these two items, the script constructs a new log by appending the input parameters to the saved value of the recorded list. The resulting log is used to

| action | <i>current</i> | <i>next(answer)</i> | |
|---------------|----------------|---------------------|-----------------------------|
| replay list | answer:rest | rest | answer from log; |
| recorded list | prms | prms | <i>current</i> not executed |
| replay list | [] | [] | answer from executing |
| recorded list | prms | prms++[answer] | <i>current</i> action |

Fig. 6. The replay mechanism.

initialize both logs, the replay list as well as the recorded list. Then the script restarts from the first CGI action.

Figure 6 illustrates the mechanism by showing the state of the logs before executing the *current* action and before executing the *next* action that depends on the result *answer* of *current*.¹⁵ The case where the replay list contains the answer works the same for *io* and for *ask*. In the second case, where the replay list is depleted, *ask* does *not* reach the *next (answer)* state directly. Instead, upon receipt of *answer* the whole program is replayed starting from the *first* action on the log `prms++[answer]`:

| action | <i>first</i> | ... | <i>current</i> | <i>next(answer)</i> |
|---------------|----------------|-----|----------------|---------------------|
| replay list | prms++[answer] | ... | [answer] | [] |
| recorded list | prms++[answer] | ... | prms++[answer] | prms++[answer] |

To get the replay mechanism initialized, the first invocation of the program sets both lists to the empty list.

8.1.1 Security. Storing the session state in the form sent back to the browser may have security implications if it reveals internal server information. In addition, the session state often contains sensitive user input, like passwords. WASH/CGI addresses both concerns by providing a one-time pad encryption for the session state. This way, the required degree of security is entirely in the hands of the system administrator, who has to prime the encryption with data from an appropriate random source.

8.1.2 Limiting the Size of the Log. Another concern is that the size of the log grows without bounds during a long interaction. However, quite often interactions can be structured so that there is a main trunk of interaction, like returning to a main menu and starting subinteractions from the menu. There are two ways of dealing with this case.

First, the programmer may start a subinteraction in a new browser window by using the `target="_blank"` attribute for the submission button that activates the subinteraction. The subinteraction proceeds in the new window and terminates by presenting a page without buttons. After closing the window the interaction continues at the main menu in the previous window. It is also possible to deal with both windows in an interleaved fashion. The point of this structuring is that the subinteractions never get incorporated in the main interaction trunk. Consequently, the main log does not grow unless some action updates the menu window.

¹⁵The infix operators `:` and `++` are the cons and append operations on lists.

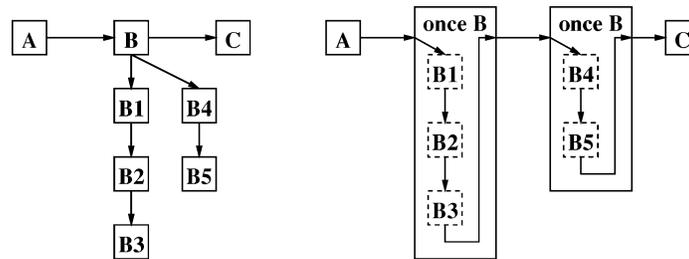


Fig. 7. Spawning interactions vs. once.

Alternatively, the programmer may use the operator `once` to achieve a similar effect without having to worry about opening new windows. The type of `once` is

```
once :: (Read a, Show a) => CGI a -> CGI a
```

Its parameter is a CGI action that is once run to completion, involving any number of nested `io`, `ask`, and `once` operations. When the action returns its result, the log entries for the nested operations are replaced by a single log entry for the result. When `once` is replayed later on, it simply grabs the result from the log without passing control to the nested operations.

Figure 7 illustrates the difference between the two approaches. The left picture shows the interaction tree resulting from spawning subinteractions from a menu node **B**. Each node of the tree corresponds to an entry in the log. Hence, at point **C**, the log has length three. The right picture shows a similar interaction using the `once` operator. The dashed boxes correspond to log entries that are removed by the `once` operator after that part of the interaction is complete. Hence, the log is temporarily extended (to length four at **B3** and **B5**) but then shrunk back again, so that the final log length at **C** is four.

8.2 Widgets

Earlier, in Section 5.3, we have seen that widgets are generated by functions of the following type:

```
type HTMLField context a =
  WithHTML INPUT CGI () -> WithHTML context CGI a
inputField :: (AdmitChildINPUT context, Reason a, Read a) =>
  HTMLField context (InputField a INVALID)
```

To understand how `inputField` works, we have to look a little bit inside the type `WithHTML INPUT CGI ()`. This abstract type threads the current state of the script throughout the whole construction of a Web form. The current state of the script contains all the information typically available in CGI scripts, as well as the replay and recorded lists. In addition, the `ask` form operation determines if there is already an answer to form in the log and passes this answer to the construction of form.

Hence, the widget operations work in two modes, depending on whether the input from the form is present. The first time a widget constructor (e.g., `inputField`) is executed, the input is not yet present so that the constructor

just generates its HTML representation. If the input for the form is present, then each widget constructor checks that the answer provides acceptable input for it and stores the input in its handle (after converting it from a string to a value of the expected type). This conversion may also fail, but this fact is only recorded in the handle data structure.

To tell individual instances of the widgets apart, they are automatically numbered and named. This naming is also managed using the state threaded through by `WithHTML context CGI ()`.

A submit button is slightly more complicated because it also specifies a callback action, as evident from the type of its constructor:

```
submit :: (AdmitChildINPUT context, InputHandle h) =>
         h INVALID -> (h VALID -> CGI ()) -> HTMLField context ()
```

Before submission of the form, each use of `submit handles callback` just generates XHTML like any other widget. The code detects this state by checking that no input for this form is available. If input to the form is present, then each submit button checks if it has been clicked (by checking if an input for its name is present). This check will be successful for at most one button. The selected button then tries to validate all input handles passed to it. If all of them are valid, then the constructor of the button composes a CGI action by applying `callback` to the validated input handles and stores the resulting action in the CGI state. Later on, when the construction and examination of the page is complete, the `ask` operation checks if an action has been stored in the CGI state and will invoke it.

If some of the handles do not validate, the constructor of the button includes JavaScript code in the XHTML page that puts error indicators on the erroneous input fields. Since it does not store an action in the CGI state, the enclosing `ask` will just display the XHTML page with the error indicators.

With this approach, a callback function *never* receives invalid data and it is guaranteed that only those input fields required for a particular callback are validated.

8.3 Persistence

The main stepping stone for handling state effectively in WASH/CGI is the design of the type-safe interfaces shown in Section 6. Once the interfaces are in place, the actual implementation is relatively straightforward. A handle to either kind of state is represented by a pair of the name of the entity and a version number.

Server-side state is a type-indexed mapping from version numbers to values. Getting the value from a handle always returns the value corresponding to the version number in the handle. Consequently, the result of a `get` operation need not be stored in the log. Setting the value through a handle fails if the current version number on the server does not match the version number of the handle.

Client-side state is similar, but the mapping is only defined for the current version number. Hence, the results of all operations must be stored in the log.

A naive implementation might assign version numbers sequentially by just incrementing the version number at each set or add operation. However, it is safer to use random numbers as version numbers. To see this, consider a scenario where some house keeping operation removes the entire history of server-side state. Using sequence numbers that start from 0 again, there might be a pending transaction stored on some browser that still refers to an old sequence number *before* the house keeping took place. Continuing that transaction will lead to inconsistencies. Hence, it is important that each version number is used at most once. A good random number generator can guarantee that with high probability [Gutmann 1998].¹⁶

8.4 Haskell Specifics

The internal representation of an XHTML document relies on two (abstract) data types `ELEMENT_` and `ATTR_`. A document tree is built using the interface

```
element_ :: String -> [ATTR_] -> [ELEMENT_] -> ELEMENT_
attr_    :: String -> String -> ATTR_
add_     :: ELEMENT_ -> ELEMENT_ -> ELEMENT_
add_attr_ :: ELEMENT_ -> ATTR_ -> ELEMENT_.
```

The functions `element_` and `attr_` are the obvious constructor functions. The function `add_ parent child` adds a child element and the function `add_attr_ parent attr` adds an attribute to a parent element. However, this interface is not accessible in a WASH/CGI program, it is only used to build the next level of abstraction.

The datatype that encapsulates generation of XHTML, `WithHTML c m a`, is a monad provided that `m` is a monad. That is, `WithHTML c` is a state monad transformer [Liang et al. 1995]. The state manipulated by `WithHTML c m` is an HTML element and the variable `c` is a phantom type variable that does not occur in the constructors of the type:

```
newtype WithHTML c m a =
  WithHTML (ELEMENT_ -> m (a, ELEMENT_)).
```

The element passed as an argument is the currently open element, that is, in terms of the generated document it corresponds to the closest enclosing opening HTML tag. Each construction function for an element, attribute, or text node adds the node into this context.

The constructor for building a `<title>` element serves as an example.

```
title ma =
  WithHTML (\enclosingElem ->
    do (a, titleElem) <- runWithHTML ma (element_ "title" [] [])
       return (a, add_ enclosingElem titleElem)).
```

¹⁶A reviewer suggests to use UUIDs which provide globally unique resource names without requiring a centralized authority [ISO (International Organization for Standardization) 1996]. A UUID consists of the hardware address of a network adapter, a time stamp, and a random number. Its use would effectively decrease the probability of a clash between version numbers from one server to another server to zero, but it would not have significant advantages in a setup with one server.

```

io :: (Read a, Show a) => IO a -> CGI a
io ioAction =
  CGI (\ cgistate ->
    case inparm cgistate of
      PAR_RESULT str : rest ->
        case reads str of
          (result, ""):_ ->
            return (result, cgistate { inparm = rest })
          _ ->
            reportError "Result unreadable"
              (text "Cannot read " ## text (show str)) cgistate
    [] ->
      do result <- ioAction
        return (result,
          cgistate { outparm = PAR_RESULT (show result) :
                    outparm cgistate })
    _ ->
      reportError "Out of sync" empty cgistate)

```

Fig. 8. Implementation of `io`.

The parameter `enclosingElem` contains the enclosing element. `runWithHTML ma` installs its argument, an empty `<title>` element, as the new enclosing element and inserts subelements according to action `ma`. This action takes place in the underlying monad, hence the result is a pair from the result `a` of running `ma` and the transformed `<title>` element. To complete the constructor function, the result `a` is paired with the enclosing element after adding the transformed `<title>` element to it and then returned. For a user of a `WithHTML c m` monad, the internal management of the element data is completely invisible.

The abstract datatype `CGI a` that encapsulates a CGI action is a monad. More precisely, it is a state monad lifted over the IO monad:

```

newtype CGI a =
  CGI (CGIState -> IO (a, CGIState)).

```

The `CGIState` is a record type that contains the replay and recorded lists, information derived from the CGI environment, information specific to the page currently generated, the current encoder for the state token (see Section 8.1.1), and data structures for dealing with cookies (see Section 6.2).

Figure 8 shows the code for the `io` operation (see Section 4). It serves as an example to see the log machinery in action. `CGIState` is a record type with fields `inparm` (for the replay list) and `outparm` (for the recorded list). The recorded list is stored in reverse order to make appending new results at the end a constant-time operation. Each element of the log can be an answer to an I/O operation of the form `PAR_RESULT str`, where `str` is the value encoded as a string, or an association list of CGI parameters (an answer to a Web form which is decoded by `ask`). The case expression examines the replay list. If its first element is the result of an I/O operation, then its string representation is converted back to a data value using `read` and returned as the value of the computation. The thus consumed answer is removed from the log as it is passed along. If the replay list is empty, then the code executes the `ioAction`, returns its result and extends

the *recorded list* with the textual encoding of the result.¹⁷ If the replay list starts with another kind of answer, then somebody has been tampering with the state tokens in the Web form or there is a version conflict between the code that has created the Web form and the presently running code. Either case is captured and displayed as an error.

9. ASSESSMENT

Since the creation of WASH/CGI in 2001, a number of applications have been developed with it. The applications range from demo programs of up to 50 lines through mini-applications (generic time table software and a mini shop, both in less than 200 lines) to applications like conference submission software (900 lines), a mail reader (550 lines relying on a 900-line email processing library and a 2000-line mail store management system), and a submission system for exercises (670 lines). The library source is roughly 3700 lines of Haskell code. It was first released on the Web on 25 April 2002. There were 90 downloads between July and September 2003.

The conference submission software was used for two workshops, once for the entire cycle with approximately 40 users (submitters, reviewers, and PC members) and once with 6 PC members just for collecting reviews and discussion. The only real problem that arose with this application was one submission where the submitter had pasted an entire postscript file into one of the text fields of the form instead of using the upload button. This triggered a bug in the read function of the standard library which caused the replay log for this session to become unreadable. No data was lost, however, the submitter resubmitted a number of times, due to the error message.

The exercise submission software was used by 136 students submitting exercises for a software engineering lecture and three administrators checking submissions. The user interface comprises 11 screens and the administrator interface contributes another 12 screens, for a total of 23 screens. The administrator interface includes some fairly complicated tabular screens, one with three buttons per user and another with four buttons and two input fields per submitted exercise. There have been 2177 sessions during the semester with a total of 232 MB of responses generated. Unfortunately, the Web server's log does not contain the size of the requests, so there is no way of calculating the total traffic. However, the size of the internal database was roughly 21 MB at the end of the semester. There was no major problem with the software despite the fact that it was evolved during the semester to keep up with changing requirements. The total development time was roughly 30 hours with the initial version up and running inside of five hours.

Regarding performance, the conference submission software becomes sluggish once it has stored about a hundred reviews. However, the sluggishness is due to the simpleminded implementation of the review database, it is not caused by the internal replay log of the WASH system. This is confirmed by

¹⁷This code relies on the assumption that `read (show x) == x`. This is recommended practice in the Haskell report [Haskell98 1998] and is indeed the case for the builtin types and the automatically derived instances of the Read and Show classes.

the exercise submission software which implements the database differently. In front of the browser, its latency is no different than that of other comparable Web services.

We have successfully integrated the WASH/CGI library with a Web server written in Haskell by Simon Marlow [Marlow 2002]. In this server, a WASH/CGI application can run like a servlet in its own thread. The integration clearly indicates that the ideas of WASH/CGI are not limited to CGI applications, but that they are applicable to Web programming in general.

We have taken every effort to keep the design as simple and intuitive as possible. One deliberate guideline was *not* to create a new language but to adhere to the Haskell98 standard [Haskell98 1998]. Although it is tempting to rely on one of the many implemented extensions to the Haskell type system [Thiemann 2002a], there were satisfactory solutions to the document validation problem (see Section 3.2), the treatment of submission buttons with decision trees (see Section 5.7), and a problem with the lifetime of input handles completely within the Haskell98 language. We will continue to work on the preprocessor mentioned at the beginning of Section 3 because it can reduce the slope of the learning curve for programmers who are not fluent in Haskell.

The document sublanguage guarantees well-formedness of all generated documents, but only quasi validity as detailed in Section 3.2. This design is motivated by the following considerations:

- Quasi validity guarantees validity for about 95 percent of all HTML elements. Out of 89 elements, it is fully adequate for 68 (which just restrict children to zero or more occurrences of certain elements) and mostly adequate for an additional 16 elements (which restrict to zero or one occurrences or one or many occurrences). Only the remaining 5 elements restrict the order of the child elements (HMTL, HEAD, TABLE, FIELDSET, MAP).
- A checker for full validity (including attribute occurrences) can be implemented in the Haskell type system [Thiemann 2002b]. However, programming in the presence of this checker gets very awkward and limiting, as described in that work.
- The checker for full validity requires some (well-known) extensions to the Haskell type system.

How to support the Web browser's back and clone buttons is a religious issue among Web programmers. While some welcome its additional power and flexibility and regard it as a new paradigm for structuring user interfaces [Graunke and Krishnamurthi 2002], others regard it as intrinsically adverse to the concept of an interactive transaction [Brabrand et al. 1999]. The implementation suggested in the present work admits the use of the back button as well as cloning and bookmarking (see Section 4). This decision leaves the choice of the underlying interaction model to the application programmer. If navigation in the style of a Web browser makes sense for an application, then WASH/CGI can support it. If the application requires a strictly sequential transaction model, then that can also be enforced using client-side state (Section 6.2). However, the WASH/CGI programming model is *not* tied to the implementation. An

alternative implementation that rules out the use of the back and clone buttons would be more efficient and at the same time it would essentially remove all restrictions currently placed on IO.

Another issue that is automatically addressed by our implementation of session state (Section 4) is scalability. Since all state information is present in the state token delivered to the browser, it does not matter which server processes the next submission. Hence, WASH/CGI scripts can be easily run on a farm of Web servers. The only caveat here lies in the accesses to server-side state, which must be shared among these Web servers.¹⁸

10. RELATED WORK

Meijer's CGI library [Meijer 2000] implements a low-level facility for accessing the input to a CGI script and for creating its output. While it supports a tree-based representation of documents, it does not encapsulate its manipulation. In addition, it still preserves the CGI-style script-centric programming style, it does not have a notion of session.

Hughes [2000] has devised the powerful concept of arrows, a generalization of monads. His motivating application is the design of a CGI library that implements sessions. Indeed, the functionality of his library was the major source of inspiration for our work. Our work indicates that monads are sufficient to implement sessions (Hughes also realized that (J. Hughes, 2000, private communication)). Furthermore, WASH/CGI extends the functionality offered by the arrow CGI-library with a novel representation of HTML and typed compositional forms. Also, the callback-style of programming advocated here is not encouraged by the arrow library.

Hanus's library [Hanus 2001] for server-side scripting in the functional-logic language Curry comes close to the functionality that we offer. In particular, its design inspired our investigation of a callback-style programming model. While his library uses logical variables to identify input fields in HTML forms, we are able to make do with a purely functional approach. Our approach only relies on the concept of a monad, which is fundamental for a real-world functional programmer.

Another line of approaches to Web programming using functional languages relies on the presence of first-class continuations in the programming language Scheme [Queinnec 2000]. The main idea is to use a continuation to take a snapshot of the state of the script's thread in the server after sending the form to the browser. The server then stores the continuation in an association list under a previously generated random key, which is noted in the form, too. When the form is submitted later on, the server extracts the key, retrieves the continuation, and resumes the suspended thread. Conceptually, this is similar to the log that we are using. Technically, there are two important differences. First, WASH has to reconstruct the current state from the log when the next user request arrives. Using a continuation of a server thread, this reconstruction is immediate. However, stale continuations may fill up the server's memory, so there is no clear advantage. Second, WASH relies only on CGI whereas the

¹⁸The present implementation does not support this.

continuation approach requires either a continually running threaded server that preserves the association list containing the continuations or a means of serializing the association list. Third, the continuation approach is only viable in an environment that supports first-class continuations (like the Scheme language), whereas WASH relies solely on typing and abstraction.

The tight coupling of scripts to the server in the continuation-based approach is leveled off by a transformational approach due to Graunke et al. [2001]. Their approach transforms arbitrary interactive programs into CGI scripts. The transformation consists of three steps, a transformation to continuation-passing style, lambda lifting, and defunctionalization. The resulting program consists of a number of dispatch functions (arising from defunctionalization) and code pieces implementing interaction fragments. Compared to the WASH approach, they reconstruct the state from a marshalled continuation closure whereas WASH replays the interaction. Furthermore, they implement mutable state using cookies stored on the client side. WASH also provides an interface for mutable references (the `Cookie` module) which is implemented in terms of cookies. However, the interface is slightly different to the standard Haskell interface for mutable references because access to a cookie-based reference may fail due to a mismatch of version numbers (see Section 8.3). WASH's programming interface allows to detect such an error and recover from it, whereas Graunke et al propose to terminate the script with an exception.

Bigwig [Brabrand et al. 2002] is a system for writing Web applications. It provides a number of domain specific customizable languages for composing dynamic documents, specifying interactions, accessing databases, etc. It compiles these languages into a combination of standard Web technologies, like HTML, CGI, applets, and JavaScript. Like our library, Bigwig provides a session abstraction. However, the Bigwig notion is more restrictive in that sessions may neither be backtracked nor forked. Each Bigwig session has a notion of a current state, which cannot be subverted. Moreover, the implementation of sessions is different and relies on a special runtime system that keeps one session thread alive for the entire duration of the session [Brabrand et al. 1999]. Bigwig includes a sophisticated static analysis that guarantees that all generated output is a valid HTML document. The same analysis also finds types for forms. In contrast, WASH/CGI relies entirely on Haskell's type system so that all checking is left to the Haskell compiler. While WASH only approximates validity, it performs a full type check for forms since (typed) field values are directly passed to (typed) callback-actions.

Bigwig has been further developed into JWIG [Christensen et al. 2003]. JWIG is an extension of Java with features of Bigwig, in particular the notion of sessions, the guarantee about valid XHTML output, and the guarantee that fields demanded by the program are present in the input submitted by the browser. JWIG adds a number of new statements and a special XML data type to the Java language. Beyond Bigwig, JWIG has an integrated notion of shared server-side state, but relies on the Java primitives for concurrency control. It does not have a special abstraction for server-side state that catches inconsistencies due to the (mis-) use of browser features as explained in Section 6.1.

It is implemented by translation into Java. The translator performs extensive static program analysis to make the above guarantees. Since the analysis only covers the Jwig classes, it uses a number of conservative assumptions about library classes and the rest of the application. In contrast, since WASH/CGI relies solely on Haskell typing, no such conservative assumptions are required and all library and application code can be used in a WASH script without restriction.

MAWL [Ladd and Ramming 1995; Atkinson et al. 1997] is a domain specific language for specifying form-based interactions. It was the first language to offer a typing of forms against the code that received its input from the form. It provides a subset of Bigwig’s functionality, for example, the facilities for document templates are much more limited.

Guide [Levy 1998] is a rule-based language for CGI programming. It supports a simple notion of document templates, similar in style to that offered by MAWL. It provides only rudimentary control structure: it sequentially processes a list of predicates and takes the action (i.e, it displays a HTML page) associated to the first succeeding predicate. It supports the concept of a session, a limited form of concurrency control, as well as session-wide and global variables. However, it neither supports typed input, nor composition of input fields, nor facilities to ensure that the variables used in a script match the variables defined in a form.

In previous work [Thiemann 2000, 2002b], we have investigated the use of Haskell for generating valid HTML and XML documents, where the type system enforces adherence to the DTD. WASH/CGI contains a simplified variant of that work as explained in Section 3.2.

The LAML system [Nørmark 2002] is an approach to authoring XHTML documents based on the Scheme programming language. LAML provides construction functions for document nodes (called mirror functions), similar in style to our construction functions. It performs dynamic checks at runtime to guarantee that the generated output adheres to a DTD, but it cannot give static guarantees. LAML relies on Scheme’s functional abstraction to provide parameterized documents. WASH relies on Haskell’s functional abstraction but can give static guarantees due to its reliance on the type system. Finally, while there are various extensions for authoring documents that span multiple URLs, there is neither a session concept nor a concept for accessing persistent state built into LAML. Hence, it appears that LAML is primarily targeted for conveniently producing static Web pages.

Java Server Pages [Peligrí-Llopart and Cable 1999] and Servlets [Sun Microsystems 2003] are popular approaches to implementing Web services. They also provide a notion of session persistence and encapsulate its implementation. Unfortunately, they do not guarantee type safety and they do not provide the advanced support for generating HTML and forms as we do. More fundamentally, both approaches still support the page-centric or script-centric approaches to Web programming, whereas WASH/CGI supports a session-centric view of the application. Similar argumentation applies to other frameworks like ASP and ASP.NET [Homer and Sussman 2002], PHP [PHP 2003], Python Server Pages [PSP 1999], ColdFusion, and Zope.

11. CONCLUSION

WASH/CGI provides a high-level declarative approach to programming Web services. It is based on a session-centric programming model that hides away the details of the underlying protocols. This approach makes programming a Web service similar to programming an HTML-based GUI application.

WASH/CGI provides extensive guarantees ranging from well-formedness and quasi validity for all generated documents through type-safe input fields to type-safe server-side and client-side state. These guarantees are intrinsic in using WASH/CGI, no application code is required to achieve them. They are obtained by consistent use of abstract data types, type-safe interfaces, and compositionality.

Due to the host language Haskell, each aspect of WASH/CGI can be abstracted and parameterized. For example, a response page may be built from any number of fragments where each fragment comes with its own functionality. Again, the integration of these fragments does not require coding beyond assembling the fragments into one response page. The high potential for parameterization is unique to WASH/CGI.

While the efficiency of the implementation of server-side state and of the replay mechanism could certainly be improved, for example, by compiling WASH/CGI or by abolishing the use of the back button and cloning in the browser, we feel that its current state presents a good checkpoint. The fundamental ideas are proven to work, the implementation is stable, and many convenience features are integrated that alleviate the construction of Web services.

The current release of the WASH system is available for downloading at <http://www.informatik.uni-freiburg.de/~thiemann/WASH/>.

ACKNOWLEDGMENTS

Many thanks are due to the anonymous reviewers whose extensive comments helped to improve the presentation of this work.

REFERENCES

- ATKINSON, D., BALL, T., BENEDIKT, M., BRUNS, G., COX, K., MATAGA, P., AND REHOR, K. 1997. Experience with a domain specific language for form-based services. In *Proceedings of the Conference on Domain-Specific Languages (DSL)* (Santa Barbara, Calif.). USENIX.
- BERNERS-LEE, T., FIELDING, R. T., AND MASINTER, L. 1998. Uniform resource identifiers (URI): Generic syntax. <http://www.faqs.org/ftp/rfc/rfc2396.txt>.
- BLOG 2003. blog. <http://www.webopedia.com/TERM/B/blog.html>.
- BRABRAND, C., MØLLER, A., SANDHOLM, A., AND SCHWARTZBACH, M. I. 1999. A runtime system for interactive Web services. *Comput. Netw.* 31, 11–16, 1391–1401.
- BRABRAND, C., MØLLER, A., AND SCHWARTZBACH, M. 2002. The <bigwig> Project. *ACM Trans. Internet Tech.* 2, 2, 79–114.
- CGI 1.0 1999. CGI: Common gateway interface. <http://www.w3.org/CGI/>.
- CHRISTENSEN, A. S., MØLLER, A., AND SCHWARTZBACH, M. I. 2003. Extending Java for high-level Web service construction. *ACM Trans. Prog. Lang. Syst.* 25, 6, 814–875.
- EASTLAKE, 3RD, D. E., AND GOLDSTEIN, T. 2001. ECML v1.1: Field specifications for E-commerce. <http://www.faqs.org/rfcs/rfc3106.html>.
- FIELDING, R., GETTYS, J., MOGUL, J., FRYSTYK, H., MASINTER, L., LEACH, P., AND BERNERS-LEE, T. 1999. Hypertext transfer protocol. <http://www.faqs.org/rfcs/rfc2616.html>.

- GRAUNKE, P., FINDLER, R. B., KRISHNAMURTHI, S., AND FELLEISEN, M. 2001. Automatically restructuring programs for the Web. In *Proceedings of ASE-2001: The 16th IEEE International Conference on Automated Software Engineering* (San Diego, Calif.). IEEE Computer Society Press, Los Alamitos, Calif., pp. 211–222.
- GRAUNKE, P. T., FINDLER, R. B., KRISHNAMURTHI, S., AND FELLEISEN, M. 2003. Modeling Web interactions. In *Proceedings of the 12th European Symposium on Programming* (Warsaw, Poland). Lecture Notes in Computer Science. Springer-Verlag, New York.
- GRAUNKE, P. T. AND KRISHNAMURTHI, S. 2002. Advanced control flows for flexible graphical user interfaces: or, growing GUIs on trees or, bookmarking GUIs. In *Proceedings of the 24th International Conference on Software Engineering (ICSE-02)* (New York, NY.). ACM, New York, pp. 277–296.
- GUTMANN, P. 1998. Software generation of practically strong random numbers. In *Proceedings of the 7th USENIX Security Symposium* (San Antonio, Tex.). USENIX.
- HANUS, M. 2001. High-level server side Web scripting in Curry. In *Practical Aspects of Declarative Languages, Proceedings of the Third International Workshop, PADL01*. Lecture Notes in Computer Science (Las Vegas, Nev.). Springer-Verlag, New York.
- HASKELL98. 1998. Haskell 98, a non-strict, purely functional language. <http://www.haskell.org/definition>.
- HOMER, A. AND SUSSMAN, D. 2002. *Professional ASP.NET 1.0* (2 ed.). Wrox Press Inc.
- HUGHES, J. 2000. Generalising monads to arrows. *Sci. Comput. Prog.* 37, 67–111.
- ISO (INTERNATIONAL ORGANIZATION FOR STANDARDIZATION). 1996. Information technology-Open systems interconnection-Remote procedure call (RPC). ISO/IEC 11578:1996.
- KRISTOL, D. M. 2001. HTTP Cookies: Standard, privacy, and politics. *ACM Trans. Internet Tech.* 1, 2 (Nov.), 151–198.
- LADD, D. A. AND RAMMING, J. C. 1995. Programming the Web: An application-oriented language for hypermedia service programming. World Wide Web Consortium. pp. 567–586.
- LEVY, M. R. 1998. Web programming in Guide. *Softw.-Pract. Exper.* 28, 15 (Dec.), 1581–1603.
- LIANG, S., HUDAK, P., AND JONES, M. 1995. Monad transformers and modular interpreters. In *Proceedings of the 1995 ACM SIGPLAN Symposium on Principles of Programming Language* (San Francisco, Calif.). ACM, New York, pp. 333–343.
- LUHNCHECK 1999. Luhn formula. http://www.webopedia.com/TERM/L/Luhn_formula.html.
- MARLOW, S. 2002. Developing a high-performance Web server in Concurrent Haskell. *J. Funct. Prog.* 12, 4&5 (July), 359–374.
- MEIJER, E. 2000. Server-side Web scripting with Haskell. *J. Funct. Prog.* 10, 1 (Jan.), 1–18.
- NETSCAPE CORPORATION 2003. JavaScript central. <http://devedge.netscape.com/central/javascript/>.
- NØRMARK, K. 2002. Programmatic WWW authoring using Scheme and LAML. In *Web Engineering Track of the 11th International World Wide Web Conferene—WWW2002*, page 296.
- PELIGRI-LLOPART, E. AND CABLE, L. 1999. Java Server Pages Specification. <http://java.sun.com/products/jsp/index.html>.
- PHP. 2003. PHP: Hypertext processor. <http://www.php.net/>.
- PSP. 1999. Python server pages. <http://www.ciobriefings.com/psp/>.
- QUEINNEC, C. 2000. The influence of browsers on evaluators or, continuations to program Web servers. In *Proc. International Conference on Functional Programming 2000* (Montreal, Ont., Canada). P. Wadler Ed. ACM, New York, pp. 23–33.
- SCOTT, D. AND SHARP, R. 2002. Abstracting application-level Web security. In *Proceedings of 11th ACM International World Wide Web Conference* (Hawaii). ACM, New York.
- SUN MICROSYSTEMS 2003. Java Servlet Specification, Version 2.4. <http://java.sun.com/products/servlet/>.
- THIEMANN, P. 2000. Modeling HTML in Haskell. In *Practical Aspects of Declarative Languages, Proceedings of the Second International Workshop, PADL'00* (Boston, Mass.). Lecture Notes in Computer Science, vol. 1753. Springer-Verlag, Los Alamitos, Calif., pp. 263–277.
- THIEMANN, P. 2002a. Programmable type systems for domain specific languages. In *Selected Papers of Workshop on Functional and Logic Programming, WFLP 2002*, Volume 76 of ENTCS, (Grado, Italy). M. Comini and M. Falaschi, Eds. Elsevier Science Publishers, Amsterdam, The Netherlands.

- THIEMANN, P. 2002b. A typed representation for HTML and XML in Haskell. *J. Funct. Prog.* 12, 4-5 (July), 435–468.
- THIEMANN, P. 2003a. WASH/CGI reference manual. <http://www.informatik.uni-freiburg.de/~thiemann/haskell/WASH>. To appear.
- THIEMANN, P. 2003b. XML templates and caching in WASH. In *Proceedings of the 2003 ACM SIGPLAN Haskell Workshop* (Uppsala, Sweden). J. Jeuring, Ed. ACM, New York, pp. 19–26.
- VALIDATOR. 2002. Markup validation service. <http://validator.w3.org/>.
- WASH. 2001. Web authoring system in Haskell (WASH). <http://www.informatik.uni-freiburg.de/~thiemann/haskell/WASH>.
- XHTML 1.0. 2000. XHTML 1.0: The extensible hypertext markup language. <http://www.w3.org/TR/xhtml1>.

Received February 2003; accepted June 2003