

# Prosper: A Framework for Extending Prolog Applications with a Web Interface

Levente Hunyadi  
hunyadi@users.sourceforge.net

Budapest University of Technology and Economics  
Department of Computer Science and Information Theory  
1117 Budapest, Magyar Tudósok körútja 2., Hungary  
Phone: +36 1 463-2585 Fax: +36 1 463-3147

Keywords: web integration; application development framework

**Abstract.** Clear separation of presentation and code-behind, declarative use of visual control elements and a supportive background framework to automate recurring tasks are fundamental to rapid web application development. In this paper, a framework is presented that facilitates extending Prolog applications with a web front-end. The framework relies on Prolog to the greatest possible extent, supports code re-use, and integrates easily into existing web server solutions. An architecture comprising two layers is elaborated. The lower layer, attached directly yet flexibly via FastCGI to the web server, handles issues of threading, basic session and context management, while the upper layer produces output based on XML source documents, which can be thought of as the Prolog counterparts of ASP.NET and Java server pages. The framework has been developed with SWI-Prolog and contains low-level communication code in C.

## 1 Introduction

The idea of a comprehensive web framework was brought forth by the development of a web application to assist in teaching declarative programming and help in course-related administration. In this scenario, the ideal framework should not only address the challenges of the web environment but should also remain effective from a declarative programmer's view. While existing web development frameworks give excellent support for structured display of data in web pages, they fail to provide a straightforward way to embed Prolog calls such as those required for evaluating programming assignments.

The proposed framework, named *PROlog Server Pages Extensible aRchitecture* (or *Prosper* in short), realizes a web environment for the direct disposal of Prolog programs. By connecting directly to the web server, it eliminates multiple levels of indirection and reduces the speed penalty incurred per request. More importantly, it supports designing complex web pages in existing HTML editors independently from actual Prolog application code. Pages are augmented with special markup that carries extra information for Prosper. These special elements

can incorporate elaborate visual presentation logic into web pages driven by values calculated by Prolog predicates in application code. This allows seamless creation of an interface to Prolog solutions without much coding or re-structuring of the original program for the sake of a web interface.

Prosper is implemented mainly in SWI-Prolog and partially in C. SWI-Prolog is compliant to part one of Prolog ISO standard and has comprehensive support for multi-threading. ISO-compliance caters for portability while multi-threading helps harness the potential in parallel execution. Network communication interfaces have been written in C to ensure maximum performance. The full source code of Prosper is available at SourceForge.net [2].

The paper is structured as follows. Section 2 gives a brief introduction to web technologies underlying the proposed framework. Section 3 elaborates on design trade-offs, inspects related work and analyzes possible approaches to create a Prolog web framework with special attention to the chosen approach. In Section 4, the architecture of the proposed framework is laid out. Section 5 traces the way a request produces a reply in Prosper by means of an example. Section 6 gives implementation details and compares how Prosper fares in terms of speed w.r.t. various other setups. Section 7, with which this paper concludes, outlines possible ways of extension and future work.

Throughout the paper, a sound knowledge of Prolog with basics on SGML-languages (especially XML [5] and (X)HTML [3]) and some experience in developing web applications with ASP.NET [11] and/or Java [4] is assumed. One should consult the indicated sources for further details on these technologies.

## 2 Background

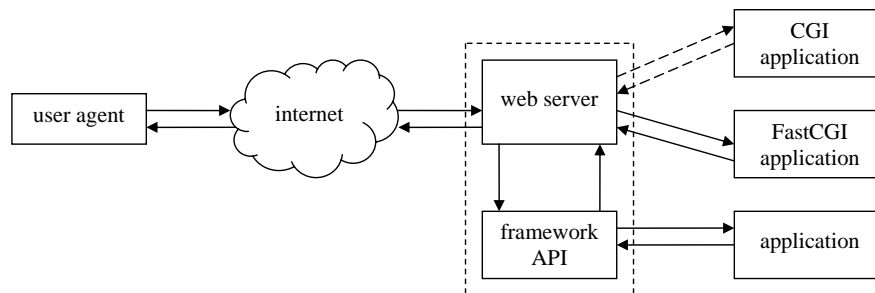
Essentially, the web operates in a request-and-reply manner according to the Hyper-Text Transfer Protocol. First, the client formulates a *request* querying a document. The request is received by the server, which looks up the requested document in its file system and returns it to the client in *reply* [8]. In the case of *dynamic content generation*, the request received by the server does not correspond to a file system entry, rather it is forwarded to an external application that outputs the reply based on the request *context*. Web application development frameworks are inserted into the chain either in place of the server (e.g. Java web solutions) or between server and external application (e.g. the ASP.NET framework), and expose a programmer-friendly view of the web environment to the application developer.

Web frameworks taking the place of the server require a thorough implementation to provide general web service functionality (e.g. serve static content) with sufficient security. For this end, it is often desirable to use a trusted web server behind which applications are placed rather than using a separate endpoint for each application. In this scenario, frameworks are often connected to servers by means of server APIs (application programming interfaces). Here, the application is loaded as a module of the server and the server forwards requests

matching some criteria (URL pattern or extension) to the application instead of processing them itself. This is called *strong coupling*.

Common Gateway Interface [1] describes a protocol that provides *loose coupling*. In order to process matching requests, the server invokes a non-integrated external application with the given request context (query parameters, user settings, etc.) and returns the output it produces to the client. Loose coupling separates the processes involved, which therefore minimally affect each other, increasing flexibility. In addition, fatal errors in the application do not endanger the server. Nonetheless, repetitive invocation of an external program can take up valuable resources by successive process initializations (re-opening database connections, etc.). FastCGI [6] is a persistent version of the CGI protocol that allows applications to remain active after serving a request (thereby maintaining database connections, etc.) yet preserving their independence (i.e. no modification of the server is required and the application works with web servers of multiple vendors).

The request and response chains and the relationships of the various application types are shown in Figure 1.



**Fig. 1.** The web service chain. The dashed rectangle indicates the boundary of the process space while the dashed arrows refer to non-persistent connection.

Applications that process many simultaneous requests have to be *multi-threaded* so that processing a request does not keep others waiting. Hence, each request is assigned a separate newly initialized thread. However, on high demand this can lead to the so-called *thrashing*, where threads with already assigned jobs lose computing resources to dozens of rapidly launched new threads, eventually leading to no thread performing useful task. Thus, applications often make use of the *worker thread model*. In this model, a constant number of threads execute concurrently. Jobs are assigned threads from a pool, which return to the pool after the job is complete. This allows fast processing of simultaneous requests with elimination of thread startup costs and stability upon high demand.

Many web development platforms make use of the *model-view paradigm*. In this paradigm, *application logic* (what the program does) and visual *presentation* (how the results are displayed) are strictly separated and logic does not reference presentation. This allows the two to be created more independently and caters for easier maintenance of both. While not every web development framework makes it compulsory, the pattern can be considered fairly wide-spread.

### 3 Possible Approaches and Related Work

In general, two approaches exist as to extending a Prolog application with a web interface.

(a) Generating web content directly from a Prolog program, possibly with the help of general-purpose libraries, is a straightforward approach. The PiLLOW library [7], available in many Prolog implementations, is a notable representative of this approach. As exemplified by the library, the close relationship of Prolog terms and the hierarchical structure of HTML easily lends itself to composing the web page in the form of terms, which are then transformed to and output as plain text on demand. By means of uninstantiated variables in the term representation, simple templates can be created.

Nevertheless, a Prolog term representation is inherently not visual and integrates poorly into existing web authoring tools. Moreover, the approach does not promote clear separation of application logic and presentation, seriously violating the model-view paradigm and leading to more difficult maintenance. Also, a stand-alone Prolog server replying to requests on a dedicated port is often assumed, which is hard to incorporate into a complex environment with an existing web server. However, a library such as PiLLOW can take off the majority of recurring tasks from the programmer's shoulder and can contribute greatly to web application development, especially in simple scenarios. Commonly aided tasks include parsing HTTP GET and POST parameters, generating forms, HTTP cookies and session maintenance.

(b) Embedding pieces of Prolog in the presentation layer (Figure 2) is another natural approach, which can be thought of as the “inside out” version of the previous one, motivated by various successful server-side technologies such as PHP. Here, web pages are composed as (X)HTML rather than as Prolog terms, and Prolog calls are inserted in the text by means of special escape sequences. The helper library parses the page into a predicate consisting of a series of write/1 statements and the equivalents of the embedded Prolog calls. Many projects that take this approach exist in the Prolog domain, [10] and [12] are two such examples.

```
<?, member(number=N, Get),  
    forall( ( between(1, N, X), factorial(X, Y) ), ?>  
<li>The factorial of <?= X ?> is <?= Y ?>.</li>  
<? ), ?>
```

**Fig. 2.** An excerpt from an HTML server page composed with embedded Prolog escape sequences. The snippet lists all factorials from 1 to N. N is specified as a query string parameter.

Albeit simple, this approach is generally insufficient for larger projects as it is weakly structured. Apparently, even repetitively displaying a block of text as in Figure 2 produces code that is difficult to comprehend. More complex nesting is even harder to implement only by means of skipping in and out of escaped blocks.

Clearly, escape sequences lead to interleaved application logic and presentation, and are hence extremely difficult to maintain or extend.

Another variant of approach (b) is composing web pages in an external framework, such as JSP or ASP.NET, and embedding foreign language calls to Prolog. PrologBeans for Java and PrologBeans.NET for the .NET platform [15], both available as SICStus extensions, are representatives of this variant. Here, all web-related issues are handled by an external framework, which provides optimized solutions to general patterns in web authoring and offers rapid application development. In order to call Prolog predicates, however, wrapper objects, written in the native language of the framework, are required that marshal calls to the Prolog engine. In fact, from a design perspective, the approach is comprised of two parts, so-called *stubs*. The wrapper object constitutes the first stub, while its Prolog counterpart the other. The stubs maintain a TCP or piped connection to each other through which Prolog call parameters and results are transmitted, usually as a stream of characters.

While practical in harnessing the benefits of a web development framework, this approach undoubtedly requires experience in programming both Prolog and the external encapsulating language. From a performance point of view, stubs introduce a further level of indirection into the web service chain and often lead to inefficient operation because the Prolog and the foreign language execution model are vastly different. Lastly, debugging Prolog embedded in foreign code is substantially harder, which can greatly increase development time.

Prosper offers a balanced mix of the two main approaches. It is a variant of approach (a) in the sense that the majority of request processing and content generation is performed in Prolog or Prolog-integrated libraries. Only Prolog programming experience is required and development is eased through improved debugging. On the other hand, it is closer to approach (b) in the sense that it adopts the model-view paradigm of rapid application development frameworks by splitting web applications into an application logic and a presentation layer.

```
<html logic-module="factorial">
  <h1>Factorial example</h1>
  <psp:assign var="E" expr="{atom_number(http_get(number))}">
    <psp:for-all function="between(1, E)" iterator="N">
      <li><psp:insert function="factorial(N)" /></li>
    </psp:for-all>
  </psp:assign>
</html>
```

---

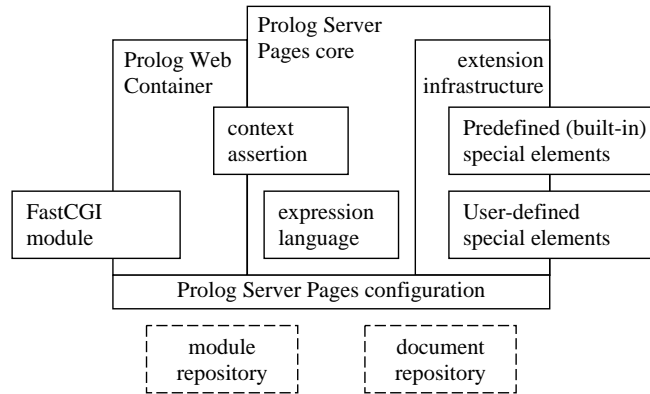
```
:- module(factorial, [factorial/2]).
factorial(Number, Factorial) :- ...
```

**Fig. 3.** The Prosper example document `factorial.xhtml` (above) and the Prolog module `factorial.pl` associated with it (below). Some XHTML elements (e.g. `ul`) have been omitted and full namespaces are not shown for brevity.

Application logic is coded as a regular Prolog module, while presentation is an (X)HTML document with some elements carrying extra information for Prosper to realize visual transformation rules. Figure 3 shows a web page that lists all factorials up to N, functionally equivalent to the web page generated by the snippet in Figure 2. Despite its verbosity, the presentation layer is not interleaved with application logic and retains its structure as a regular XHTML document. Roughly speaking, Prosper can be viewed as an extension of PiLLoW with a more robust visual front-end. Section 4 elaborates on the design of the proposed framework.

## 4 Architectural overview

From a design perspective, Prosper can be decomposed into two major layers (Figure 4). The lower layer, *Prolog Web Container*, maintains a direct persistent connection to the web server through the FastCGI protocol. The FastCGI module transmits data to and from the Prolog framework. In addition to communicating with the web server, Prolog Web Container parses headers and payload associated with HTTP requests into Prolog terms and generates them for replies, maintains a worker thread pool and assigns jobs to threads. The primary task of the container is to isolate the communication protocol and provide a natural view of request, session and profile data for the programmer. In accordance, the container provides similar facilities as other Prolog libraries in use, PiLLoW in particular, i.e. reversing content encoding, parsing query strings, etc.



**Fig. 4.** The architecture of the proposed framework.

*Prolog Server Pages*, built on top of the container, defines an XML-based document model. The conventional XML document model is extended with *special elements* belonging to a dedicated namespace each of which realizes a *transformation rule*. A transformation rule describes how the (visual) content of an element is transformed based on attributes, and the local and global *context* of the given element. Local context corresponds to variables instantiated in server

documents, while global context refers to request context as extracted by Prolog Web Container and exposed as Prolog predicates by the *context assertion* module. In assigning values to local variables, Prosper offers the so-called *expression language*. Expression language can be seen as an extension to the `is/2` predicate to include basic atom manipulation, request context variables and user-defined functions.<sup>1</sup>

Prosper includes a predefined set of special elements implementing the most common transformation rules such as conditionals and iteration constructs. However, the set of transformation rules is not restricted. Relying on the *extension infrastructure*, the user may create new modules that contain hook predicates registered for steps associated with reply generation. Modules correspond to XML namespaces and exported hook predicate names to element names in server page documents. In fact, it is via hook predicates that the predefined transformation rules are realized in the framework, which means – in the extreme case – that they can also be redefined. Special elements and their implementor hook predicates are declared in a *configuration file*. The configuration file also holds connection settings to the web server and parallel execution parameters required by Prolog Web Container.

Apart from the visual part of Prolog Server Pages, the *logic modules* give real power to the architecture. While independent from Prolog Server Pages documents, they provide the code-behind that encapsulates true application logic as conventional Prolog modules. Server pages can reference code-behind in a variety of ways: assign server page variables based on application logic, test for the satisfiability of predicates (goals) and formulate conditions using the return value of functions, thereby affecting visual layout.

Prolog modules constituting application code reside in a dedicated directory, the so-called *module repository*. Similarly, Prosper maintains a *document repository*, which is the default location to search for server pages.

## 5 Generating a reply

In order to get a deeper insight into the internals of the framework, in this section we will trace how a request dynamically produces a reply in Prosper. As an example, let us suppose that the user has entered a URL into his browser's location bar that corresponds to a web page which lists all factorials up to 3 (e.g. `http://prosper.iit.bme.hu/factorial.xhtml?number=3`).

Once received by the web server, based on configuration settings, the server detects that this HTTP request is to be forwarded to Prosper for reply generation. It dispatches a FastCGI request, which is intercepted by one of the idle Prolog Web Container worker threads.<sup>2</sup> The thread extracts the context asso-

---

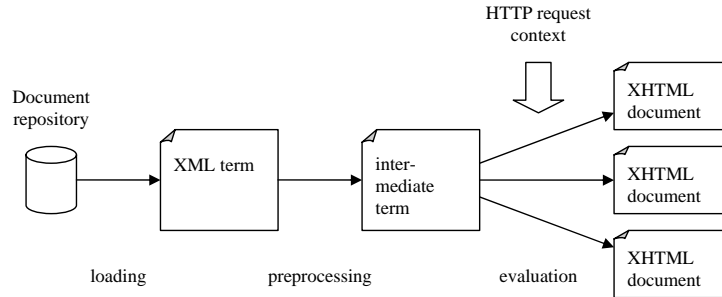
<sup>1</sup> In this paper, a Prolog *function* corresponds to a predicate all of whose arguments are strictly inbound except for the last, which is strictly outbound, and which should be unified with a ground term and is interpreted as the return value of the function.

This corresponds to the Mercury [9] definition of function.

<sup>2</sup> See predicate `worker/1` in module `prosper_server` [2].

ciated with the request as Prolog terms. The context typically includes query parameters in the URL (typically for HTTP GET requests), HTML form data passed as payload (typically for HTTP POST requests) and the session identifier. The Prolog representation of the context is handed over to Prolog Server Pages. In our example, the request context only contains GET parameters, represented by the list `[number='3']`.

First, Prolog Server Pages *loads* the document associated with the URL. The loaded document is *preprocessed* into a so-called intermediate term (IT) representation. Context is then *asserted* into a dedicated module and the document is *evaluated*. Evaluation ends with generating *output*, which is returned by Prolog Web Container to the web server as a FastCGI reply. Transformation rules are associated with both the preprocessing and the evaluation phase. The crucial difference is that in the preprocessing phase, no external context information is available, while evaluation-time transformation is context-sensitive (Figure 5). The aforementioned steps are elaborated below.



**Fig. 5.** The stages of reply generation in the case of a non-cached server page document.

*Loading a Prosper document.* The role of the loading phase<sup>3</sup> is to fetch a referenced document from disk and construct its Prolog XML term representation, similar to the one used by the PiLLOW library [7].

Whenever an HTTP request corresponds to a server page that has not been loaded, Prosper looks for the page in the document repository. Let us suppose that the URL entered by the imaginary user does not correspond to a loaded document. Therefore, the document `factorial.xhtml` is loaded and parsed into a Prolog XML term representation as seen in Figure 6. This representation mainly consists of nested `element/3` terms, where the arguments represent:

1. the name of the XML element after namespace resolution;
2. a list of attributes associated with the element;
3. a list of nested XML nodes as `element/3` terms for XML elements or atoms for character data.

*Preprocessing phase.* The goal of the preprocessing phase,<sup>4</sup> the next link in the service chain, is to validate the loaded document. Preprocessing ensures that

<sup>3</sup> Implemented in `import_page/3` in module `prosper_core` [2].

<sup>4</sup> Implemented in `markup_to_term/6` in module `prosper_core` [2].



```

element(html, ['logic-module'=factorial], [
  element(h1, [ ], ['Factorial example']),
  element(php:assign, [var='E', expr='{atom_number(http_get(number))}'], [
    element(php:for-all, [function='between(1, E)', iterator='N'], [
      element(li, [ ], [
        element(php:insert, [function='factorial(N)'], [ ])
      ])
    ])
  ])
])

```

**Fig. 6.** The XML term representation of the example document in Figure 3.

special elements referenced by the document exist, they are used correctly in terms of syntax, and that the logic module associated with the document is loaded and compiled.

As previously mentioned, special elements correspond to transformation rules. What the transformation rule exactly does depends on the attributes associated with the element and its context. In our example, `php:assign`, `php:for-all` and `php:insert` are special elements, assuming the namespace `php` is registered with Prosper.<sup>5</sup> The `php:assign` special element can have a `var` attribute, which specifies the name of the variable to introduce in the scope of the element. Similarly, `php:insert` is used with the attribute `function` in the example to insert a return value but could also be used in conjunction with `expr` to insert the value of an expression.

However, in the preprocessing phase no context information associated with the HTTP request is available; it has not yet been asserted. In spite of this, verifying attributes, parsing atoms into terms, etc. are already possible. These operations are performed by *preprocessing-time hooks* for each special element. A hook predicate interprets element attributes and/or contents and has the following signature (*elementName* denotes the name of the special element without the namespace):

```
elementName (+VarTypes, +Attrs, +Contents, -Terms)
```

Here, `Attrs` is a list of `Name=Value` pairs, which consists of attributes that parameterize the element. `Contents` is a list of inner elements in XML term representation. `Terms` is the single output argument of the predicate, which is the IT representation (preprocessed form) of the element and is commonly bound to a single-element list of the following form:<sup>6</sup>

```
[ extension(ModuleName:Predicate, ContentTerms) ].
```

<sup>5</sup> For conciseness, namespaces are not written out as full URLs, even though in the actual implementation, they are used in that manner.

<sup>6</sup> In fact, `Terms` is a list of atoms, `element/3` and `extension/2` terms. However, only `extension/2` terms are subject to evaluation in a later phase thus `Terms` is usually a list with a single `extension/2` element.

In this term, `ContentTerms` has similar semantics as `Terms` in the enclosing element: it is the IT representation of the enclosed child elements. This suggests a recursive way of operation. Indeed, albeit not compulsory, most special elements compute their own IT representation based on that of their descendants. `Predicate` corresponds to a Prolog predicate, which (augmented with some additional arguments) will be called in the evaluation phase to generate output. In other words, `Predicate` is the evaluation-time transformation rule associated with the special element parameterized with `Attrs`. For instance, a different `Predicate` is associated with a `psp:assign` element if it assigns a variable based on an expression than if based on a function call. In fact, arguments present in `Attrs` as an association list are converted into positional arguments with appropriate conversions where necessary (e.g. atoms converted to Prolog goals).

Hook predicates should never fail but should signal malformed syntax (such as an unrecognized attribute) by throwing an exception.<sup>7</sup> This guarantees that the document is syntactically well-formed at the end of the preprocessing phase.

```

element(html, [ ], [
  element(h1, [ ], ['Factorial example']),
  extension(assign_expression('E', EL), [
    extension(for_all(factorial:between(1, E)-['E'=E], 'N'), [
      element(li, [ ], [
        extension(insert_function(factorial:factorial(N)-['N'=N]), [ ])
      ])
    ])
  ])
])

```

**Fig. 7.** The intermediate term representation of the example document. *EL* denotes the execution plan of the expression language term and is omitted for conciseness.

To better comprehend the preprocessing phase, we compare the XML term representation of our example document in Figure 6 with its preprocessed version in Figure 7. In the case of the root element `html`, the two representations are identical, except for the attribute `logic-module`. This attribute binds a conventional Prolog module (the code-behind) to the Prosper document. Any predicates that occur in the document are auto-qualified with the name of this module during the preprocessing phase.

The first notable difference is `psp:assign`, which has been converted into an `extension/2` term. `assign_expression` is the name of a predicate that computes an expression language (EL) term and assigns its value to a variable. The scope of the variable is the contents of the `psp:assign` element. The function `http_get` in the EL term returns the string value of a query string variable, while `atom_number`, as its name suggests, converts its operand to a Prolog number. Just as documents, EL expressions are preprocessed, yielding an *execution plan*, which

<sup>7</sup> Failure prevents extracting the context of the actual error.

is not shown in Figure 7. The execution plan is a compound term that contains (1) the uninstantiated variables in the expression and (2) the module-qualified names of the Prolog functions to call to compute the result.

The representation of the special element `psp:for-all` has also changed substantially. The atom in its attribute called `function` has been converted into a real Prolog term augmented with a list of uninstantiated variables in it. `for_all(Function-Insts, Variable)` is a predicate that instantiates variables in `Function` and calls it, returning results in a local variable. Subsequent solutions are obtained through backtracking. Note the number of arguments to `between/3` (the third, output argument is absent) and the auto-qualification.

As we have already seen, a unique code-behind file can be associated with a Prosper document through the `logic-module` attribute. The exact location of the file is either directly specified in the `logic-module` attribute as an absolute path, or it may be a relative path, in which case it is searched for w.r.t. the module repository.

For the sake of higher performance, Prosper caches preprocessed documents. If a document is available in the cache, the loading and preprocessing phases are skipped.

*Request context assertion.* Context information available in Prosper documents and logic modules is loaded in the request context assertion phase. The primary goal of this phase is to expose HTTP request context (such as request parameters and session variables) to EL functions and logic module predicates in a natural manner without having to propagate an extra argument encapsulating the context. Predicates in the module `psp` store context information by means of *thread-local* blackboard primitives [13]. Whenever a mutable (session) value is modified while the request is served (e.g. a session variable is assigned to), changes are recorded in a dynamic fact database at the end of the subsequent evaluation phase. Hence, no particular thread is associated with any session and any worker thread may serve any request. Worker threads load current values from the dynamic fact database into blackboard primitives before the evaluation phase and record new values when evaluation ends.

Logic modules have access to context by calling predicates exported by the module `psp`. For instance, the `http_get/2` and `session/2` predicates retrieve the value of a GET and a session variable, respectively, and the predicate `session_set/2` assigns a value to a session variable.<sup>8</sup> For maximum conformance to the Prolog execution model, they all support backtracking, i.e. assignments to session variables are undone upon failure in a logic module predicate.<sup>9</sup>

*Evaluation phase.* In the last major phase, *evaluation*,<sup>10</sup> the preprocessed document is transformed w.r.t. the available request context. By the end of the

---

<sup>8</sup> For a full list, see exported predicates in module `psp` [2].

<sup>9</sup> SWI-Prolog provides backtrackable destructive assignment on blackboard primitives.

<sup>10</sup> Implemented in `term_to_elements/3` in module `prosper_extensions` [2].

evaluation phase, the document has been transformed into a term representation the string equivalent of which is ready to be sent back directly to the client as response.

From a declarative point of view, each IT element represents an (evaluation-time) transformation rule, influenced by (1) term contents, (2) asserted HTTP request- and session-related data that is globally accessible in the entire document and (3) local variables assigned by outer special elements (i.e. that encapsulate the element to which the rule corresponds).

In the case of `element/3` terms, the (recursive) transformation rule is trivial: transformation rules are applied to each child element with the same context as valid for the parent element and the evaluated form of the parent element comprises of the results of these transformation rules. For `extension/2` terms, recall that the first argument corresponds to a hook predicate assembled in the preprocessing phase: this is what represents the transformation rule. From a procedural point of view, in fact, the IT representation is traversed top-down, at each depth invoking hook predicates or the trivial transformation rule, where hook predicates may introduce new local variables before processing the children of the term they correspond to.

Local variables are means to store and reuse calculated data within server page documents. In contrast to globally available data (loaded into the thread-local module `psp` in the context assertion phase), they are accessed as Prolog variables rather than predicates and they are confined to the server page document in which they are introduced and may not be directly used in code-behind files. More precisely, the scope of local variables is always restricted to the descendants of the element in which they are assigned and are hidden by variables of the same name. Server page local variables have similar semantics as Prolog or XSLT variables in the sense that they can be assigned only once. Contrary to Prolog, however, variables cannot remain uninstantiated and are unified immediately in the element in which they are introduced.

Figure 8 shows the evaluated version of the preprocessed document in Figure 7. The result should not be surprising. For the elements `html`, `h1` and `li`, the trivial transformation rule has been applied and they are intact except for their recursively processed contents. The IT equivalents of special elements `psp:assign`, `psp:for-all` and `psp:insert` are absent from the output but their effect is apparent. The local variable `E`, which is introduced by `assign_expression`, has been used to instantiate unbound variables in the function `between(1, E)`, and the iteration variable `N` of `for_all` has been used multiple times to call the function `factorial(N)`. `N` behaves as expected, taking a different value for each loop of the iteration.

From the perspective of the framework, local variables are in fact `Name=Value` members in an association list. The association list is initially empty for the root element but may be extended with further members by any transformation rule, in which case the recursively processed descendant elements see the extended list. In our example, the evaluation-time transformation rule associated with the

```

element(html, [ ], [
  element(h1, [ ], ['Factorial example']),
    element(li, [ ], ['1']),
    element(li, [ ], ['2']),
    element(li, [ ], ['6'])
  ])
])

```

**Fig. 8.** The evaluated form of the example document.

`psp:assign` element prepends the variable `E` to the name-value list, while the rule related to `psp:for-all` does so with `N`.

## 6 Implementation and Evaluation

Prosper is implemented mainly in SWI-Prolog and partially in C. The most notable SWI-specific extra services utilized by the framework are XML document parsing and generation, blackboard primitives, multi-threading and basic thread communication.

The framework comprises of the following major components:

1. The *server module* implements Prolog Web Container.
2. The *core module* manages the lifecycle of a Prosper page. In particular, it imports pages on demand, initiates context assertion, page preprocessing and evaluation, and outputs error documents.
3. The *context module* asserts and retracts thread-local data via blackboard primitives to expose request, session and profile values, all of which are manipulated through dedicated predicates of the module `psp`.
4. The *extension module* contains predicates essential to special element implementors. It includes helper predicates to aid XML attribute parsing and the predicates `element_to_terms/3` and `term_to_elements/3`, which realize page preprocessing and evaluation, respectively. The latter two predicates are called by transformation rule hooks to recursively process child elements.
5. The *built-in elements module* contains the predefined set of special elements, including simple and compound conditionals, iteration constructs, variable assignment and insertion.
6. The *expression language module* is responsible for expression language execution plan generation and expression evaluation.
7. The *FastCGI foreign language module*, written in C, implements the FastCGI protocol.

While primarily designed to increase designer and programmer performance, the proposed architecture is comparable to other Prolog-based technologies in terms of speed. In a loopback scenario (i.e. server and client were running on the same machine), different configurations were polled by HTTP requests with GET parameters. All configurations parsed the query string, computed a simple

arithmetic expression based on query parameters, and displayed results in a web page. CGI and FastCGI-based applications (Prosper inclusive) connected to Apache/2.0.54, .NET applications ran on the built-in web server provided with Visual Studio 2005 (see Table 1). Benchmarking was performed by ApacheBench 2.0.41 on an AMD Athlon64 3000+ running Microsoft Windows XP Professional SP2.

**Table 1.** Comparative performance of various frameworks. The table shows cumulative response times in seconds for 1000 requests with 2 concurrent threads. In test cases *small* and *large*, responses of about sizes 1kB and 50kB were requested with few embedded Prolog calls. In test case *Prolog-intensive*, the architectures had to call about 50 Prolog predicates in application code-behind to produce a result of about 3kB in size. n.a. indicates that there is no overhead of a Prolog call-intensive setup or it is not meaningful for the test case.

Development tool	Application model	small	large	intensive
SICStus Prolog 3.12.5	CGI, saved state	165.78	225.33	n.a.
SWI-Prolog 5.5.33	CGI	39.47	60.17	n.a.
PrologBeans.NET	ASPX	6.297	7.781	91.91
Prosper (PWC + PSP)	multi-threaded FCGI	2.688	8.719	5.828
Prosper (PWC only)	multi-threaded FCGI	1.938	6.953	n.a.
SWI-Prolog 5.6.27	standalone server	1.266	6.313	n.a.
static html content		0.875	1.406	n.a.

Three cases are of special interest. The standalone multi-threaded HTTP server shipped with SWI-Prolog can serve as the basis for comparing the performance of Prolog-based frameworks. It provides convenience tools for HTTP reply generation but intermixes presentation and application logic. The difference in speed between Prosper with Prolog Web Container and SWI's standalone server gives an estimate of the cost of using an intermediary FastCGI transmission. The extra overhead of Prosper with the Prolog Server Pages document model shows the relative cost of having a separate presentation and application logic layer.

## 7 Summary, Perspectives for Future Work

In this paper, a novel framework facilitating the development of web-oriented Prolog applications has been presented. With a persistent multi-threaded architecture, an XML-based document model and a set of reusable transformation rules, it provides an efficient yet convenient way to create web applications in Prolog. Code changes required in existing Prolog modules for the sake of web presentation are minimal and web pages constituting the presentation layer can be composed with a declarative way of thinking in any arbitrary XML editor. Presentation and application logic are clearly separated, thus application logic

can be debugged and maintained independently. Lastly, the framework integrates well in existing web server scenarios and is open for extension.

However, the proposed solution is not yet comprehensive. Often, web applications rely on external databases and Prolog implementations scarcely provide a natural, transparent and homogeneous way of accessing data in these backing stores. The architecture should be extended to provide adequate support for databases both in terms of flexibility and performance. XML-based requests and web services are also possible ways of extension, which can greatly contribute to the versatility of the framework.

## Acknowledgements

The author wishes to acknowledge Péter Szabó for the original idea of this framework and for pinpointing crucial development issues as well as the insightful comments of Péter Szeredi on the drafts of this paper.

## References

1. *The CGI specification*, <http://hoo.hoo.ncsa.uiuc.edu/cgi/interface.html>
2. *Project page of Prosper*, <http://sourceforge.net/projects/prospear>, module Prosper in CVS `pserver:anonymous@prospear.cvs.sourceforge.net:/cvsroot/prospear`
3. *XHTML 1.0 The Extensible HyperText Markup Language* (Second Edition), January 26 2000 (revised 1 August 2002), <http://www.w3.org/TR/xhtml1/>
4. Eric ARMSTRONG et al. *The J2EE 1.4 Tutorial* (For Sun Java System Application Server Platform Edition 8.1 2005Q2 UR2), Sun Microsystems, June 7, 2005
5. Tim BRAY et al. *Extensible Markup Language (XML) 1.0* (Fourth Edition), August 16 2006, <http://www.w3.org/TR/2006/REC-xml-20060816/>
6. Mark R. BROWN, *FastCGI specification*, Document Version: 1.0, Open Market, Inc., April 29, 1996
7. Daniel CABEZA and Manuel HERMENEGILDO, *The PiLLOW Web Programming Library*, Reference Manual, The CLIP Group, School of Computer Science, Technical University of Madrid, January 5, 2001, <http://www.clip.dia.fi.upm.es/Software/pillow/pillow.html>
8. R. FIELDING et al., *Hypertext Transfer Protocol – HTTP/1.1*, RFC 2616, Network Working Group, The Internet Society, June 1999
9. Fergus HENDERSON et al., *The Mercury Language Reference Manual*, Version 0.12.2, The University of Melbourne, 2006
10. Benjamin JOHNSTON, *Prolog Server Pages*, <http://www.benjaminjohnston.com.au/template.prolog?t=psp>
11. Dan HURWITZ and Jesse LIBERTY, *Programming ASP.NET*, 3rd Edition, O’Reilly, October 2005
12. Mauro Di NUZZO, *Prolog Server Pages: A server-side scripting language based on Prolog*, Version 0.2, April 2006 <http://www.prologonlinereference.org/psp.psp>
13. Jan WIELEMAKER, *Native preemptive threads in SWI-Prolog*, Social Science Informatics (SWI), University of Amsterdam
14. Jan WIELEMAKER, Z. HUANG, Lourens VAN DER MEIJ, *SWI-Prolog and the Web*, <http://hcs.science.uva.nl/projects/SWI-Prolog/articles/TPLP-plweb.pdf>
15. *PrologBeans and PrologBeans.NET for SICStus Prolog*, <http://www.sics.se/sicstus/docs/latest/html/sicstus/PrologBeans.html>