Distributed WWW Programming using (Ciao-)Prolog and the PiLLoW library*

DANIEL CABEZA and MANUEL HERMENEGILDO

CLIP Group
http://www.clip.dia.fi.upm.es – http://www.cliplab.org
Facultad de Informática
Universidad Politécnica de Madrid (UPM)
28660-Boadilla del Monte, Madrid, Spain
(*e-mail:* {dcabeza,herme}@fi.upm.es)

Abstract

We discuss from a practical point of view a number of issues involved in writing distributed Internet and WWW applications using LP/CLP systems. We describe PiLLoW, a *public-domain* Internet and WWW programming library for LP/CLP systems that we have designed in order to simplify the process of writing such applications. PiLLoW provides facilities for accessing documents and code on the WWW; parsing, manipulating and generating HTML and XML structured documents and data; producing HTML forms; writing form handlers and CGI-scripts; and processing HTML/XML templates. An important contribution of PiLLoW is to model HTML/XML code (and, thus, the content of WWW pages) as terms. The PiLLoW library has been developed in the context of the Ciao Prolog system, but it has been adapted to a number of popular LP/CLP systems, supporting most of its functionality. We also describe the use of concurrency and a high-level model of client-server interaction, Ciao Prolog's *active modules*, in the context of WWW programming. We propose a solution for client-side downloading and execution of Prolog code, using generic browsers. Finally, we also provide an overview of related work on the topic.

Keywords: WWW, HTML, XML, CGI, HTTP, Distributed Execution, (Constraint) Logic Programming.

1 Introduction

The wide diffusion of the Internet and the popularity of the "World Wide Web" (Berners-Lee *et al.*, 1994) –WWW– protocols are effectively providing a novel platform that facilitates the development of new classes of portable and user-friendly distributed applications. Good support for network connectivity and the protocols and communication architectures of this novel platform are obviously requirements for any programming tool to be useful in this arena. However, this alone may not be enough. It seems natural that significant parts of network applications will require symbolic and numeric capabilities which are not necessarily related with distribution. Important such capabilities are, for example, high-level symbolic information processing, dealing with combinatorial problems, and natural language processing in general. Logic Programming (LP) (Kowalski, 1974; Colmerauer, 1975) and Constraint Logic Programming (CLP) systems (Jaffar & Lassez, 1987; Van Hentenryck, 1989; Colmerauer, 1990; M. Dincbas & Van Hentenryck, 1990; ECRC, 1993) have been shown particularly successful at tackling these issues (see, for example, the

---

[0] This paper is an expanded and improved version of (Cabeza & Hermenegildo, 1996a; Cabeza & Hermenegildo, 1996b; Cabeza & Hermenegildo, 1997; Cabeza *et al.*, 1996).

proceedings of recent conferences on the "Practical Applications of Prolog" and "Practical Applications of Constraint Technology"). It seems natural to study how LP/CLP technology fares in developing applications which have to operate over the Internet.

In fact, Prolog, its concurrent and constraint based extensions, and logic programming languages in general have many characteristics which appear to set them particularly well placed for making an impact on the development of practical networked applications, ranging from the simple to the quite sophisticated. Notably, LP/CLP systems share many characteristics with other recently proposed network programming tools, such as Java, including dynamic memory management, well-behaved structure and pointer manipulation, robustness, and compilation to architecture-independent bytecode. Furthermore, and unlike the scripting or application languages currently being proposed (e.g., shell scripts, Perl, Java, etc.), LP/CLP systems offer a quite unique set of additional features including dynamic databases, search facilities, grammars, sophisticated meta-programming, and well understood semantics.

In addition, most LP/CLP systems also already offer some kind of low level support for remote communication using Internet protocols. This generally involves providing a *sockets* (ports) interface whereby it is possible to make remote data connections via the Internet's native protocol, TCP/IP. A few systems support higher-level communication layers on top of this interface including linda-style blackboards (e.g., SICStus Prolog (Carlsson, 1988), Ciao (Carro & Hermenegildo, 1999; Cabeza & Hermenegildo, 1995; Hermenegildo & CLIP Group, 1994; Hermenegildo *et al.*, 1995a; Hermenegildo *et al.*, 1999a; Hermenegildo *et al.*, 1999b; Bueno *et al.*, 1997), BinProlog/μ^2-Prolog (Tarau, 1996; Bosschere, 1989), etc.) or shared variable-based communication (e.g., KL1 (Chikayama *et al.*, 1994), AKL (Janson & Haridi, 1991), Oz (Smolka, 1994), Ciao (Hermenegildo *et al.*, 1995b; Cabeza & Hermenegildo, 1995), etc.). In some cases, this functionality is provided via libraries, building on top of the basic TCP/IP primitives. This is the case, for example, of the SICStus and Ciao distributed linda-style interfaces. In fact, as we have shown in previous work, shared-variable based communication can also be implemented in conventional systems via library predicates, by using attributed variables (Hermenegildo *et al.*, 1995b; Cabeza & Hermenegildo, 1995). In addition to these communication primitives, several systems offer concurrency and even higher-level abstractions (distributed objects, mobile code, ...) which are very useful for developing general-purpose distributed applications.

Our concrete interest here is WWW applications. These applications generally use specific high-level protocols (such as HTTP or FTP), data formats (such as HTML or XML), and application architectures (e.g., the CGI interface) which are different from, e.g., the shared-variable or linda-based protocols typically used in other types of distributed applications. In this paper we study how good support for these WWW-related protocols, data formats, and architectures can be provided for LP/CLP systems, building on the widely available interfaces to the basic TCP/IP protocols. Our aim is to discuss from a practical point of view a number of the new issues involved in writing WWW applications using LP/CLP systems, as well as the architecture of some typical solutions. In the process, we will describe P*i*LL*o*W ("Programming in Logic Languages on the Web"), a public domain Internet/WWW programming library for LP/CLP systems which, we argue, significantly simplifies the process of writing such applications. P*i*LL*o*W provides facilities for generating HTML/XML structured documents by handling them as Herbrand terms, producing HTML *forms*, writing form handlers, processing HTML/XML templates, accessing and parsing WWW documents, etc. We also describe the architecture of some

relatively sophisticated application classes, using a high-level model of client-server interaction, *active modules* (Cabeza & Hermenegildo, 1995). Finally we describe an architecture for automatic LP/CLP code downloading for local execution, using just the library and generic browsers.

Apart from the tutorial value of the paper, we present a number of technical contributions which include the idea of representing HTML and XML code (and structured text in general) as Prolog terms, the use of the logical variable in such terms leading to a model of an "HTML template" (a pair comprising a term with free variables and a dictionary associating names to those variables), the notion of "active logic modules" and its application to solving efficiency issues in CGI interaction in a very simple way, the idea of "Prolog scripts" and its application to CGIs, and the identification of a number of features that should be added to existing systems in order to facilitate the programming of WWW applications – mainly concurrency.

The argument throughout the paper is that, with only very small limitations in functionality (which disappear when concurrency is added, as in systems such as BinProlog/μ^2-Prolog, AKL, Oz, KL1, and Ciao Prolog), it is possible to add an extremely useful Internet/WWW programming layer to any LP/CLP system without making any significant changes in the implementation. We argue that this layer can simplify the generation of applications in LP/CLP systems including active WWW pages, search tools, content analyzers, indexers, software demonstrators, collaborative work systems, MUDs and MOOs, code distributors, etc.

The purpose of the paper is also to serve as a tutorial, containing sufficient information for developing relatively complex WWW applications in Prolog and other LP and CLP languages using the P*i*LL*o*W library. The P*i*LL*o*W library has been developed in the context of the Ciao Prolog system, but it has been adapted to a number of other popular LP/CLP systems, supporting most of its functionality. The Ciao Prolog system and the P*i*LL*o*W library can be freely downloaded from `http://www.clip.dia.fi.upm.es` and `http://www.cliplab.org`.

2 Writing basic cgi-bin applications

The simplest way of writing WWW applications is through the use of the "Common Gateway Interface" (CGI). A CGI executable is a standard executable file but such that the HTTP server (the program that responds to HTTP requests in a machine which serves a WWW site) can tell that it in fact contains a program that is to be run, rather than a document text that is to be sent to the client (the browser) as usual. The file can be distinguished by belonging to a special directory, commonly named `cgi-bin`, or by a special filename ending, such as `.cgi`. This is normally set during configuration of the HTTP server. The basic idea behind the CGI interface is illustrated in Figure 1. When the user selects an address of a CGI executable in a document, such as `http://www.xxx.yyy/cgi-bin/hello_world` (or perhaps `http://www.xxx.yyy/foo/hello_world.cgi`) the browser issues a standard document request (1). The HTTP server, recognizing that it is a CGI executable rather than a document, starts the executable (2), and during such execution stores the output of the executable in a buffer (3). Upon termination of the executable, the contents of the buffer (which should be in a format that the browser can handle, such as HTML) are returned to the browser as if a normal page with that content had been accessed (4).
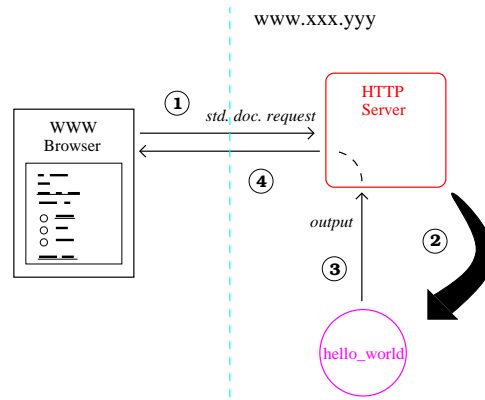
Fig. 1. The CGI interface

The following is an example of how a very simple such executable can be written in an LP/CLP language. The source might be as follows:[1]

```
main :-
    write('Content-type: text/html'), nl, nl,
    write('<HTML>'),
    write('Hello world.'),
    write('</HTML>').
```

And the actual executable could be generated as usual, for example in the Ciao system, using the standalone compiler, by writing at a UNIX shell "`ciaoc -o hello_world.cgi hello_world`".[2] The executable then has to be placed in an appropriate place (accessible via an HTTP address by a browser) and have the right permissions for being executed by the server (for example, in some systems this means being executable by the user "`nobody`").

In systems which make executables through saved states (which usually have the disadvantage of their generally large size), at the system prompt one could create the executable by writing something like:

```
?:- compile('hello_world.pl'),
     save('hello_world.cgi'), main.
```

3 LP/CLP Scripts for CGI Applications

CGI executables are often small- to medium-sized programs that perform relatively simple tasks. This, added to the slow speed of the network connection in comparison with that of executing a program (which makes program execution speed less important) has made scripting languages (such as shell scripts or Perl) very popular for writing these programs. The popularity is due to the fact that no compilation is necessary (extensive string handling capabilities also play an important role in the case of Perl), and thus changes and updates to the program imply only editing the source file.

Logic languages are, a priori, excellent candidates to be used as scripting languages.[3] However, the relative complication in making executables (needing in some

---

[1] Note that in the examples presented, and in order to shorten them, the HTML code may be slightly simplified, and as a result of this it may not be completely standard-conforming. However, the examples can be used as is with all popular browsers.

[2] It often convenient to use options (such as `ciaoc`'s `-s` or `-S`) which will generate a standalone executable which is independent of any libraries.

[3] For example, the built-in grammars and databases greatly simplify many typical script-based applications.
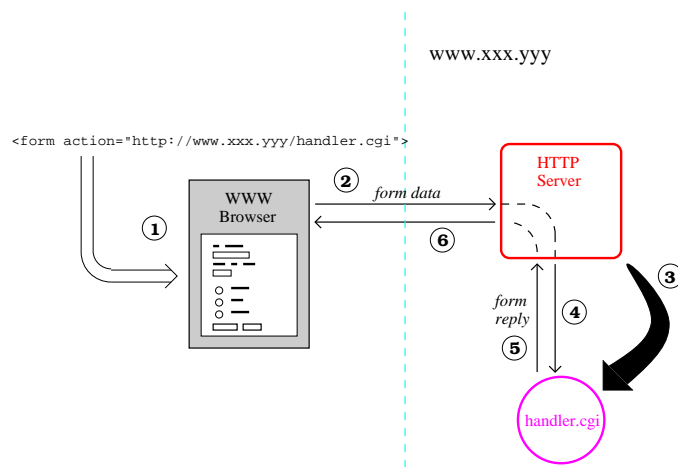
Fig. 2. The Forms interface

systems to start the top-level, compile or consult the file, and make a saved state) and the often large size of the resulting executables may deter CGI application programmers. It appears convenient to provide a means for LP/CLP programs to be executable as scripts, even if with reduced performance.

It is generally relatively easy to support scripts with the same functionality in most LP/CLP systems. In Ciao, the program `ciao-shell` –which has also been adapted to SICStus (Hermenegildo, 1996)– accomplishes this task, by first loading the file given to it as the first argument (but skipping the first lines and avoiding loading messages) and then starting execution at `main/1` (the argument provides the list of command line options). Then, for example, in a Unix system, the following program can be run directly as a script without any need for compilation:

```
#!/usr/local/bin/ciao-shell

main(_) :-
    write('Content-type: text/html'), nl, nl,
    write('<HTML>'),
    write('Hello world.'),
    write('</HTML>').
```

Note that in some UNIX versions either the program `ciao-shell` must be included in the `/etc/shells` listing or the first line should be replaced by these two:

```
#!/bin/sh
exec ciao-shell $0 "$@"
```

The execution of Prolog scripts may be optimized in some systems. For example, in Ciao, the first time a script is run it is also compiled and its bytecode is saved to a file. At subsequent times, if the script has not changed, the object code is retrieved from that file, avoiding compilation or interpretation overhead.

4 Form Handling in HTTP

So far we have shown CGI executables which produce output, but this output is not a function of input coming from the request, which is obviously of limited interest. CGI executables become most useful when combined with HTML forms. HTML forms are HTML documents (or parts of HTML documents) which include special

fields such as text areas, menus, radio buttons, etc. which allow providing input to CGI executables. The steps involved in the handling of the input contained in a form are illustrated in Figure 2. When a document containing a form is accessed via a form-capable browser (Mosaic, Netscape, Lynx, etc.), the browser displays the input fields, buttons, menus, etc. indicated in the document, and *locally* allows the user to perform input by modifying such fields. However, this input is not ultimately handled by the browser. Instead, it will be sent to a "handler" CGI program, which can be anywhere on the net, and whose address must be given in the form itself (1). Forms generally have a "submit" button such that, when pressed, the input provided through the menus, text areas, etc. is sent by the browser to the HTTP server corresponding to the handler (2). Two methods for sending this input exist: "GET" and "POST". In the meantime, the sending browser waits for a response from that program, which should come in the form of a new HTML document. The handler program is invoked in much the same way as a cgi-bin application (3), except that the information from the form is supplied to the handler (in different ways depending on the system, the method of invocation, and the content type) (4). This information is encoded in a predefined format, which relates each piece of information to the corresponding field in the form, by means of a keyword associated with each field. The handler then identifies the information corresponding to each field in the original form, processes it, and then responds by writing an HTML document to its standard output (5), which is forwarded by the server to the waiting browser when the handler terminates (6). An important point to be noted is that, as with simple cgi-bin applications, the handler is started and should terminate for each transaction. The reader is referred, for example, to (Grobe & Naseer, 1998) for a more complete introduction to CGI scripts and HTML forms.

5 Writing Form Handlers with P*i*LL*o*W

The only complication in writing form handlers compared to writing simple CGI applications is the need to capture and parse the form data. As we said before, this data can be provided in several ways, depending on the system and the method used to invoke the form, and is encoded with escape sequences. It is relatively easy to write a Prolog program to parse such input (using, for example, definite clause grammars –DCGs). The P*i*LL*o*W library provides some predicates which do this and simplify the whole task, hiding the low-level protocol behind. The principal predicates provided include:

get_form_input(`Dic`) Translates input from the form (with either the POST or GET methods, and even with CONTENT_TYPE multipart/form-data) to a dictionary *Dic* of *attribute*=*value* pairs. It translates empty *value*s (which indicate only the presence of an attribute) to the atom '$empty', values with more than one line (from text areas or files) to a list of lines as strings, the rest to atoms or numbers (using name/2). This is implemented using DCG parsers.

get_form_value(`Dic,Var,Val`) Gets value *Val* for attribute *Var* in dictionary *Dic*. Does not fail: value is '' if not found (this simplifies merging form producers and form handlers, see later).

form_empty_value(`V`) Useful to check that a value *V* from a text area is empty (filters spaces, newlines, linefeeds, etc.).

form_default(`Val,Default,NewVal`) Useful when a form is only partially filled (and also in the first invocation of a combined form handler/producer – see Section 7). If the value of *Val* is empty then *NewVal=Default*, else *NewVal=Val*.

my_url(`URL`) Returns in *URL* the Uniform Resource Locator (WWW address) of this cgi executable.

`form_request_method(`*Method*`)` Returns in *Method* the method of invocation of the form handler ("GET" or "POST").

For example, suppose we want to make a handler which implements a database of telephone numbers and is queried by a form including a single entry field with name `person_name`. The handler might be coded as follows:

```
#!/usr/local/bin/ciao-shell

:- include(library(pillow)).

main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
    write('Content-type: text/html'), nl, nl,
    write('<HTML><TITLE>Telephone database</TITLE>'), nl,
    write('<IMG SRC="phone.gif">'),
    write('<H2>Telephone database</H2><HR>'),
    write_info(Name),
    write('</HTML>').

write_info(Name) :-
    form_empty_value(Name) ->
        write('You have to provide a name.')
  ; phone(Name, Phone) ->
        write('Telephone number of <B>'),
        write(Name),
        write('</B>: '),
        write(Phone)
  ; write('No telephone number available for <B>'),
    write(Name),
    write('</B>.').

phone(daniel, '336-7448').
phone(manuel, '336-7435').
phone(sacha,  '543-5316').
```

The code above is quite simple. On the other hand, the interspersion throughout the text of calls to `write` with HTML markup inside makes the code somewhat inelegant. Also, there is no separation between computation and input/output, as is normally desirable. It would be much preferable to have an encoding of HTML code as Prolog terms, which could then be manipulated easily in a more elegant way, and a predicate to translate such terms to HTML for output. This functionality, provided by the P*i*LL*o*W library, is presented in the next section.

6 Handling HTML as Prolog Terms

Since LP/CLP systems perform symbolic processing using Herbrand terms, it seems natural to be able to handle HTML code directly as terms. Then, such structures only need to be translated by appropriate predicates to HTML code when they need to be output. In general, this relationship between HTML code and Prolog terms allows viewing any WWW page as a Herbrand term. The predicates which provide this functionality in P*i*LL*o*W are:

`output_html(`*F*`)` Accepts in *F* an HTML term (or a list of HTML terms) and
sends to the standard output the text which is the rendering of the term(s) in
HTML format.

`html2terms(`*Chars,Terms*`)` (also, `xml2terms/2`) Relates a list of HTML
(resp. XML) terms and a list of ASCII characters which are the rendering of
the terms in HTML format. This predicate is reversible (but it normalizes in
the reverse direction –see later). `output_html/2` uses this predicate to transform
HTML terms in characters. Again, this is implemented via DCG parsing.

In an *HTML term* certain atoms and structures represent special functionality
at the HTML level. An HTML term can be recursively a list of HTML terms. The
following are legal HTML terms:

```
hello
[hello, world]
["This is an ", em('HTML'), " term"]
```

When converting HTML terms to characters, `html2terms/2` translates special
structures into the corresponding format in HTML, applying itself recursively to
their arguments. Strings are always left unchanged. HTML terms may contain logic
variables, provided they are instantiated before the term is translated or output.
This allows creating documents piecemeal, back-patching of references in docu-
ments, etc.

In the following sections we list the meaning of the principal Prolog structures
that represent special functionality at the HTML level. Only special atoms are
translated, the rest are assumed to be normal text and will be passed through to
the HTML document.

6.1 General Structures

Basically, HTML has two kinds of components: HTML *elements* and HTML *en-
vironments*. An HTML element has the form "`<NAME Attributes >`" were `NAME`
is the name of the element and *Attributes* is a (possibly empty) sequence of at-
tributes, each of them being either an attribute name or an attribute assignment
as `name="Value"`.

An HTML environment has the form "`<NAME Attributes > Text </NAME>`"
were `NAME` is the name of the environment an *Attributes* has the same form as
before.

The general Prolog structures that represent these two HTML constructions are:

Name `$`*Atts* ('`$/2`' is defined as an infix, binary operator.) Represents an HTML
element of name *Name* and attributes *Atts*, were *Atts* is a (possibly empty) list
of attributes, each of them being either an atom or a structure `name=value`. For
example, the term

```
img$[src='images/map.gif',alt="A map",ismap]
```

is translated into the HTML source

```
<img src="images/map.gif" alt="A map" ismap>
```

Note that HTML is not case-sensitive, so we can use lower-case atoms.

name(*Text*) (A term with functor *name*/1 and argument *Text*) Represents an
HTML environment of name *name* and included text *Text*. For example, the
term

```
address('clip@dia.fi.upm.es')
```

is translated into the HTML source

```
<address>clip@dia.fi.upm.es</address>
```

`name(Atts,Text)` (This is a term with functor *name/2* and arguments *Atts* and *Text*) Represents an HTML environment of name *name*, attributes *Atts* and included text *Text*. For example, the term

```
a([href='http://www.clip.dia.fi.upm.es/'],"Clip home")
```

represents the HTML source

```
<a href="http://www.clip.dia.fi.upm.es/">Clip home</a>
```

`env(Name,Atts,Text)` Equivalent to *Name*(*Atts*,*Text*).

`begin(Name,Atts)` It translates to the start of an HTML environment of name *Name* and attributes *Atts*. There exists also a `begin(Name)` structure. Useful, in conjunction with the next structure, when including in a document output generated by an existing piece of code (e.g. *Name* = `pre`). Its use is otherwise discouraged.

`end(Name)` Translates to the end of an HTML environment of name *Name*.

Now we can rewrite the previous example as follows (note how the use of the logic variable `Response` allows injecting the result of the call to `response/1` into the output term, using unification):

```
#!/usr/local/bin/ciao-shell

:- include(library(pillow)).

main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
    response(Name,Response),
    output_html([
        'Content-type: text/html\n\n',
        html([title('Telephone database'),
              img$[src='phone.gif'],
              h2('Telephone database'),
              hr$[],
              Response])]). %% Using the logic variable.

response(Name, Response) :-
    form_empty_value(Name) ->
        Response = 'You have to provide a name.'
  ; phone(Name, Phone) ->
        Response = ['Telephone number of ',b(Name),': ',Phone]
  ; Response = ['No telephone number available for ',b(Name),'.'].

phone(daniel, '336-7448').
phone(manuel, '336-7435').
phone(sacha,  '543-5316').
```

Any HTML construction can be represented with these structures (except comments and declarations, which could be included as atoms or strings), but the P*i*LL*o*W library provides additional, specific structures to simplify HTML creation.

6.2 Specific Structures

In this section we will list some special structures for HTML which P*i*LL*o*W understands. While in many cases using the general structures (with the native HTML

names) is probably good practice, using specific structures such as these can sometimes be convenient. Also, some of these structures have special functionality (e.g., `prolog_term/1`). A predicate `html_expansion/2` is provided which allows defining new structures (tables, layers, etc). Specific structures include (the reader is referred to the P*i*LL*o*W manual for a full listing):

`start` Used at the beginning of a document (translates to `<html>`).

`end` Used at the end of a document (translates to `</html>`).

`--` Produces a horizontal rule (translates to `<hr>`).

`\\` Produces a line break (translates to `<br>`).

`$` Produces a paragraph break (translates to `<p>`).

`comment(`*Comment*`)` Used to insert an HTML comment (translates to `<!-- `*Comment*` -->`).

`declare(`*Decl*`)` Used to insert an HTML declaration – seldom used (translates to `<!`*Decl*`>`).

`image(`*Addr*`)` Used to include an image of address (URL) *Addr* (translates to an `<img>` element).

`image(`*Addr*`,`*Atts*`)` As above with the list of attributes *Atts*.

`ref(`*Addr*`,`*Text*`)` Produces a hypertext link, *Addr* is the URL of the referenced resource, *Text* is the text of the reference (translates to `<a href="`*Addr*`">`*Text*`</a>`).

`label(`*Label*`,`*Text*`)` Labels *Text* as a target destination with label *Label* (translates to `<a name="`*Label*`">`*Text*`</a>`).

`heading(`*N*`,`*Text*`)` Produces a heading of level *N* ($1 \leq N \leq 6$), *Text* is the text to be used as heading – useful when one wants a heading level relative to another heading (translates to a `<h`*N*`>` environment).

`itemize(`*Items*`)` Produces a list of bulleted items, *Items* is a list of corresponding HTML terms (translates to a `<ul>` environment).

`enumerate(`*Items*`)` Produces a list of numbered items, *Items* is a list of corresponding HTML terms (translates to an `<ol>` environment).

`description(`*Defs*`)` Produces a list of defined items, *Defs* is a list whose elements are definitions, each of them being a Prolog sequence (composed by `','`/2 operators). The last element of the sequence is the definition, the other (if any) are the defined terms (translates to an `<dl>` environment).

`nice_itemize(`*Img*`,`*Items*`)` Produces a list of bulleted items, using the image *Img* as bullet. The predicate `icon_address/2` provides a colored bullet.

`preformatted(`*Text*`)` Used to include preformatted text, *Text* is a list of HTML terms, each element of the list being a line of the resulting document (translates to a `<pre>` environment).

`verbatim(`*Text*`)` Used to include text verbatim, special HTML characters (`<`,`>`,`&`,`"`) are translated into its quoted HTML equivalent.

`prolog_term(`*Term*`)` Includes any prolog term *Term*, represented in functional notation. Variables are output as `_`.

`nl` Used to include a newline in the HTML source (just to improve human readability).

`entity(`*Name*`)` Includes the entity of name *Name* (ISO-8859-1 special character).

`cgi_reply` This is not HTML, rather, the CGI protocol requires this content descriptor to be used by CGI executables (including form handlers) when replying (translates to "`Content-type: text/html`").

`pr` Includes in the page a graphical logo with the message "Developed using the P*i*LL*o*W Web programming library", which points to the manual and library source.

With these additional structures, we can rewrite the previous example as follows
(note that in this example the use of `heading/2` or `h2/1` is equally suitable):

```
#!/usr/local/bin/ciao-shell

:- include(library(pillow)).

main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
    response(Name,Response),
    output_html([
        cgi_reply,
        start,
        title('Telephone database'),
        image('phone.gif'),
        heading(2,'Telephone database'),
        --,
        Response,
        end]).

response(Name, Response) :-
    form_empty_value(Name) ->
        Response = 'You have to provide a name.'
  ; phone(Name, Phone) ->
        Response = ['Telephone number of ',b(Name),': ',Phone]
  ; Response = ['No telephone number available for ',b(Name),'.'].

phone(daniel, '336-7448').
phone(manuel, '336-7435').
phone(sacha,  '543-5316').
```

We have not included above the specific structures for creating forms. They are
included and explained in the following section.

6.3 Specific Structures for Forms

In this section we explain the structures which represent the various elements related
to forms:

`start_form(`*Addr*`[,`*Atts*`])` Specifies the beginning of a form. *Addr* is
the address (URL) of the program that will handle the form, and
Atts other attributes of the form, as the method used to invoke it.
If *Atts* is not present the method defaults to POST. (Translates to
`<form action="`*Addr*`" `*Atts* `>`.)

`start_form` Specifies the beginning of a form without assigning address to the
handler, so that the form handler will be the cgi-bin executable producing the
form.

`end_form` Specifies the end of a form (translates to `</form>`).

`checkbox(`*Name*`,`*State*`)` Specifies an input of type `checkbox` with name *Name*,
State=on if the checkbox is initially checked (translates to an `<input>` element).

`radio(`*Name*`,`*Value*`,`*Selected*`)` Specifies an input of type `radio` with name *Name*
(several radio buttons which are interlocked must share their name), *Value* is the
the value returned by the button, if *Selected*=*Value* the button is initially checked
(translates to an `<input>` element).

input(*Type*,*Atts*) Specifies an input of type *Type* with a list of attributes *Atts*.
Possible values of *Type* are text, hidden, submit, reset, . . . (translates to an
<input> element).

textinput(*Name*,*Atts*,*Text*) Specifies an input text area of name *Name*. *Text*
provides the default text to be shown in the area, *Atts* a list of attributes (trans-
lates to a <textarea> environment).

option(*Name*,*Val*,*Options*) Specifies a simple option selector of name *Name*,
Options is the list of available options and *Val* is the initial selected option
(if *Val* is not in *Options* the first item is selected) (translates to a <select>
environment).

menu(*Name*,*Atts*,*Items*) Specifies a menu of name *Name*, list of attributes *Atts*
and list of options *Items*. The elements of the list *Items* are marked with the
prefix operator '$' to indicate that they are selected (translates to a <select>
environment).

For example, in order to generate a form suitable for sending input to the previ-
ously described phone database handler one could execute the following goal:

```
output_html([
    start,
    title('Telephone database'),
    heading(2,'Telephone database'),
    $,
    start_form('http://www.clip.dia.fi.upm.es/cgi-bin/phone_db.pl'),
    'Click here, enter name of clip member, and press Return:',
    \\,
    input(text,[name=person_name,size=20]),
    end_form,
    end]).
```

Of course, one could have also simply written directly the resulting HTML doc-
ument:

```
<html>
<title>Telephone database</title>
<h2>Telephone database</h2>
<p>
<form method="POST"
 action="http://www.clip.dia.fi.upm.es/cgi-bin/phone_db.pl">
Click here, enter name of clip member, and press Return:
<br>
<input type="text" name="person_name" size="20">
</form>
</html>
```

7 Merging the Form Producer and the Handler

An interesting practice when producing HTML forms and handlers is to merge the
operation of the form producer and the handler into the same program. The idea is
to produce a generalized handler which receives the form input, parses it, computes
the answer, and produces a new document which contains the answer to the input,
as well as a new form. A special case must be made for the first invocation, in
which the input would be empty, and then only the form should be generated. The

following is an example which merges the producer and the handler for the phones database:[4]

```
#!/usr/local/bin/ciao-shell

:- include(library(pillow)).

main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
    response(Name,Response),
    output_html([
        cgi_reply,
        start,
        title('Telephone database'),
        image('phone.gif'),
        heading(2,'Telephone database'),
        --,
        Response,
        start_form,
        'Click here, enter name of clip member, and press Return:',
        \\,
        input(text,[name=person_name,size=20]),
        end_form,
        end]).

response(Name, Response) :-
    form_empty_value(Name) ->
        Response = []
  ; phone(Name, Phone) ->
        Response = ['Telephone number of ',b(Name),': ',Phone,$]
  ; Response = ['No telephone number available for ',b(Name),'.',$].

phone(daniel, '336-7448').
phone(manuel, '336-7435').
phone(sacha,  '543-5316').
```

This combination of the form producer and the handler allows producing applications that give the impression of being interactive, even if each step involves starting and running the handler to completion. Note that forms can contain fields which are not displayed and are passed as input to the next invocation of the handler. This allows passing state from one invocation of the handler to the next one.

Finally, a note about testing and debugging CGI scripts: this is unfortunately not as straightforward as it could be. Useful techniques include carefully checking permissions, looking at the data logs of the server, replacing predicates such as `get_form_...` with versions that print what is really being received, etc.

8 Templates

A problem in the previous programs is that the layout of the output page is not easily configurable –it is hard-coded in the source and can only be changed by

---

[4] Notice that when only one text field exists in a form, the form can be submitted by simply pressing "Return" inside the text field.

modifying the program. This is something that a normal user (or even an expert programmer if the size of the program is large) may not want to do. In order to address this, P*i*LL*o*W provides a facility for reading in "HTML templates" (also XML templates), and converting them into a term format in which it is very natural to manipulate them. An HTML template is a file which contains standard HTML code, but in which "slots" can be defined and given an identifier by means of a special tag. These slots represent parts of the HTML code in which other HTML code can be inserted. Once the HTML template is read by P*i*LL*o*W, such slots appear as free logic variables in the corresponding P*i*LL*o*W terms. In this way, the user can define a layout with an HTML editor of choice, taking care of marking the "left out" parts with given names. These parts will then be filled appropriately by the program. The functionality associated with parsing such terms is encapsulated in the following predicate:

html_template(*Chars*, *Terms*, *Dict*) Parses the string *Chars* as the contents of an HTML template and unifies *Terms* with the list of HTML terms comprised in the template, substituting occurrences of the special tag <V>*name*</V> with prolog variables. *Dict* is instantiated to the dictionary of such substitutions, as a list of *name*=*Variable* pairs.

In the following example a template file called TlfDB.html is assumed to hold the formatting of the output page, defining an HTML variable called "response" which will be substituted by the response of the CGI program. Note that the predicate file_to_string/2 (defined in Ciao library file_utils) reads a file and returns in its second argument the contents of the file as a list of character codes. Note also that calling html_template/3 with the third argument instantiated to [response = Response] has the effect of instantiating the "slot" in HTML_terms to the contents of Response (this makes use of the fact that there is only one slot on the template; normally, a call to member/2 is used to locate the appropriate *name*=*Variable* pair).

```
#!/usr/local/bin/ciao-shell

:- include(library(pillow)).
:- use_module(library(file_utils)).

main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
    response(Name,Response),
    file_to_string('TlfDB.html', Contents),
    html_template(Contents, HTML_terms, [response = Response]),
    output_html([cgi_reply|HTML_terms]).

response(Name, Response) :-
    form_empty_value(Name) ->
        Response = []
  ; phone(Name, Phone) ->
        Response = ['Telephone number of ',b(Name),': ',Phone,$]
  ; Response = ['No telephone number available for ',b(Name),'.',$].

phone(daniel, '336-7448').
phone(manuel, '336-7435').
phone(sacha,  '543-5316').
```

An example of the contents of the template file could be:

```
<HTML><HEAD><TITLE>Telephone database</TITLE></HEAD>
<BODY background="bg.gif">
<IMG src="phone.gif">
<H2>Telephone database</H2>
<HR>
<V>response</V>
<FORM method="POST">
Click here, enter name of clip member, and press Return:<BR>
<INPUT type="text" name="person_name" size="20"></FORM>
</BODY>
</HTML>
```

9 Accessing WWW documents

The facilities presented in the previous sections allow generating HTML documents, including forms, and handling the input coming from forms. In many applications such as search tools, content analyzers, etc., it is also desirable to be able to access documents on the Internet. Such access is generally accomplished through protocols such as FTP and HTTP which are built on top of TCP/IP. In LP/CLP systems which have TCP/IP connectivity (i.e., a sockets/ports interface) the required protocols can be easily coded in the source language using such facilities and DCG parsers. At present, only the HTTP protocol is supported by PiLLoW. As with HTML code, the library uses an internal representation of Uniform Resource Locators (URLs), to be able to manipulate them easily, and provides predicates which translate between the internal representation and the textual form. The facilities provided by PiLLoW for accessing WWW documents include the following predicates:

url_info(*URL,Info*) Translates a URL *URL* to an internal structure *Info* which details its various components and vice-versa. For now non-HTTP URLs make the predicate fail. E.g.
url_info('http://www.foo.com/bar/scooby.txt',Info)
gives Info = http('www.foo.com',80,"/bar/scooby.txt"),
url_info(URL, http('www.foo.com',2000,"/bar/scooby.txt")
gives URL = "http://www.foo.com:2000/bar/scooby.txt" (a string).

url_info_relative(*URL,BaseInfo,Info*) Translates a relative URL *URL* which appears in the HTML page referred to by *BaseInfo* (given as an url_info structure) to a complete url_info structure *Info*. Absolute URLs are translated as with the previous predicate. E.g.
url_info_relative("/guu/intro.html",
http('www.foo.com',80,"/bar/scoob.html"), Info)
gives Info = http('www.foo.com',80,"/guu/intro.html")
url_info_relative("dadu.html",
http('www.foo.com',80,"/bar/scoob.html"), Info)
gives Info = http('www.foo.com',80,"/bar/dadu.html").

url_query(*Dic,Args*) Translates a list of *attribute=value* pairs *Dic* (in the same form as the dictionary returned by get_form_input/1) to a string *Args* for appending to a URL pointing to a form handler.

fetch_url(*URL,Request,Response*) Fetches a document from the Internet. *URL* is the Uniform Resource Locator of the document, given as a url_info structure. *Request* is a list of options which specify the parameters of the request, *Response* is a list which includes the parameters of the response. The request parameters available include:

head To specify that we are only interested in the header.

`timeout(`*`Time`*`)` *Time* specifies the maximum period of time (in seconds) to wait for a response. The predicate fails on timeout.

`if_modified_since(`*`Date`*`)` Get   document only if newer than *Date*. An example of a structure that represents a date is `date('Tuesday',15,'January',1985,'06:14:02')`.

`user_agent(`*`Name`*`)` Provide a user-agent field.

`authorization(`*`Scheme`*`,`*`Params`*`)` Provides an authentication field when accessing restricted sites.

`name``(`*`Param`*`)` Any other functor translates to a field of the same name (e.g. `from('user@machine')`).

The parameters which can be returned in the response list include (see the HTTP/1.0 definition for more information):

`content(`*`Content`*`)` Returns in *Content* the actual document text, as a list of characters.

`status(`*`Type`*`,`*`Code`*`,`*`Phrase`*`)` Gives the status of the response. *Type* can be any of `informational`, `success`, `redirection`, `request_error`, `server_error` or `extension_code`, *Code* is the status code and *Phrase* is a textual explanation of the status.

`pragma(`*`Data`*`)` Miscellaneous data.

`message_date(`*`Date`*`)` The time at which the message was sent.

`location(`*`URL`*`)` The document has moved to this URL.

`http_server(`*`Server`*`)` Identifies the server responding.

`allow(`*`methods`*`)` List of methods allowed by the server.

`last_modified(`*`Date`*`)` Date/time at which the sender believes the resource was last modified.

`expires(`*`Date`*`)` Date/time after which the entity should be considered stale.

`content_type(`*`Type`*`,`*`Subtype`*`,`*`Params`*`)` Returns the MIME type/subtype of the document.

`content_encoding(`*`Type`*`)` Encoding of the document (if any).

`content_length(`*`Length`*`)` *Length* is the size of the document, in bytes.

`authenticate(`*`Challenges`*`)` Request for authentication.

`html2terms(`*`Chars`*`,`*`Terms`*`)` We have already explained how this predicate transforms HTML terms to HTML format. Used the other way around it can parse HTML code, for example retrieved by `fetch_url`. The resulting list of HTML terms *Terms* is normalized: it contains only `comment/1`, `declare/1`, `env/3` and `$/2` structures.

For example, a simple fetch of a document can be done as follows:

```
url_info('http://www.foo.com',UI), fetch_url(UI,[],R),
member(content(C),R), html2terms(C, HTML_Terms).
```

Note that if an error occurs (the document does not exist or has moved, for example) this will simply fail. The following call retrieves a document if it has been modified since October 6, 1999:

```
fetch_url(http('www.foo.com',80,"/doc.html"),
    [if_modified_since('Wednesday',6,'October',1999,'00:00:00')],R).
```

This last one retrieves the header of a document (with a timeout of 10 seconds) to get its last modified date:

```
fetch_url(http('www.foo.com',80,"/last_news.html"),[head,timeout(10)],R),
member(last_modified(Date),R).
```

The following is a simple application illustrating the use of `fetch_url` and `html2terms`. The example defines `check_links(`*URL*`,`*BadLinks*`)`. The predicate fetches the HTML document pointed to by *URL* and scours it to check for links which produce errors when followed. The list *BadLinks* contains all the bad links found, stored as compound terms of the form: `badlink(Link,Error)` where `Link` is the problematic link and `Error` is the error explanation given by the server.

```
check_links(URL,BadLinks) :-
        url_info(URL,URLInfo),
        fetch_url(URLInfo,[],Response),
        member(content_type(text,html,_),Response),
        member(content(Content),Response),
        html2terms(Content,Terms),
        check_source_links(Terms,URLInfo,[],BadLinks).

check_source_links([],_,BL,BL).
check_source_links([E|Es],BaseURL,BL0,BL) :-
        check_source_links1(E,BaseURL,BL0,BL1),
        check_source_links(Es,BaseURL,BL1,BL).

check_source_links1(env(a,AnchorAtts,_),BaseURL,BL0,BL) :-
        member((href=URL),AnchorAtts), !,
        check_link(URL,BaseURL,BL0,BL).
check_source_links1(env(_Name,_Atts,Env_html),BaseURL,BL0,BL) :- !,
        check_source_links(Env_html,BaseURL,BL0,BL).
check_source_links1(_,_,BL,BL).

check_link(URL,BaseURL,BL0,BL) :-
        url_info_relative(URL,BaseURL,URLInfo), !,
        fetch_url_status(URLInfo,Status,Phrase),
        ( Status \== success ->
          name(P,Phrase),
          name(U,URL),
          BL = [badlink(U,P)|BL0]
        ; BL = BL0
        ).
check_link(_,_,BL,BL).

fetch_url_status(URL,Status,Phrase) :-
        fetch_url(URL,[head,timeout(20)],Response), !,
        member(status(Status,_,Phrase),Response).
fetch_url_status(_,timeout,"Timeout").
```

10 Providing Code Through the WWW

A facility which can be easily built on top of the primitives presented so far is that of "remote WWW modules," i.e., program modules which reside on the net at a particular HTTP address in the same way that normal program modules reside in a particular location in the local file system. This allows for example always fetching the most recent version of a given library (e.g., P*i*LL*o*W) when a program is compiled. For example, the form handler of Section 6.1, if rewritten as

```
#!/usr/local/bin/ciao-shell
```
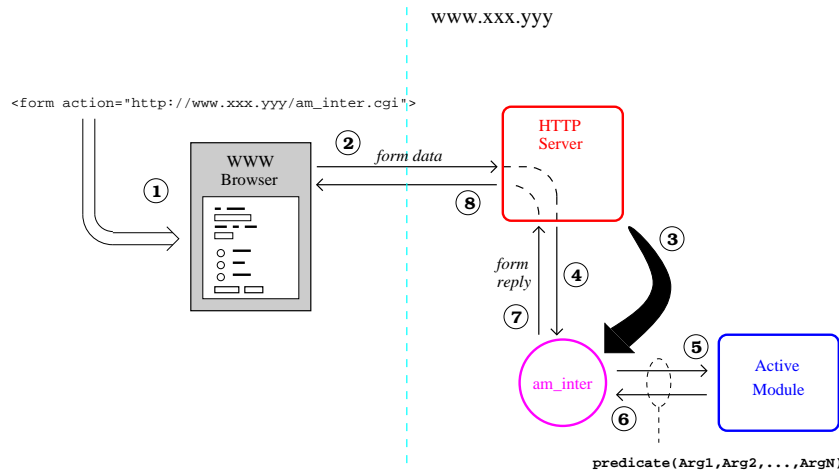
www.xxx.yyy

Fig. 3. The Forms Interface using Active Modules

```
:- use_module('http://www.clip.dia.fi.upm.es/lib/pillow.pl').

main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
...
```

would load the current version of the library each time it is executed. This general-
ized module declaration is just syntactic sugar, using `expand_term`, for a document
fetch, using `fetch_url`, followed by a standard `use_module` declaration. It is ob-
viously interesting to combine this facility with caching strategies. An interesting
(and straightforward to implement) additional feature is to fetch remote byte-code
(as generally done by `use_module`), if available, but this is only possible if the two
systems use the same byte-code (this can normally be checked easily in the byte-
code itself). Also, it may be interesting to combine this type of code downloading
with WWW document accesses, so that code is downloaded automatically when a
particular document is fetched. This issue is addressed in Section 12. Finally, there
are obvious security issues related to downloading code in general, which can be
addressed with standard techniques such as security signatures.

11 A High-Level Model of Client-Server Interaction: Active Modules

Despite its power, the cgi-bin interface also has some shortcomings. The most se-
rious is perhaps the fact that the handler is started and expected to terminate for
each interaction. This has two disadvantages. First, no state is preserved from one
query to the next. However, as mentioned before, this can be fixed by passing the
state through the form (using *hidden* fields), by saving it in a temporary file at the
server side, by using "cookies", etc. Second, and more importantly, starting and
stopping the application may be inefficient. For example, if the idea is to query
a large database or a natural language understanding system, it may take a long
time to start and stop the system. In order to avoid this we propose an alternative
architecture for cgi-bin applications (a similar idea, although not based on the idea
of active modules, has been proposed independently by Ken Bowen (Bowen, 1996)).

The basic idea is illustrated in Figure 3. The operation is identical to that of
standard form handlers, as illustrated in Figure 2, up to step 3. In this step, the
handler started is not the application itself, but rather an interface to the actual

application, which is running continuously and thus contains state. Thus, only the interface is started and stopped with every transaction. The interface simply passes the form input received from the server (4) to the running application (5) and then forwards the output from the application (6) to the server before terminating, while the application itself continues running. Both the interface and the application can be written in LP/CLP, using the predicates presented. The interface can be a simple script, while the application itself will be typically compiled.

An interesting issue is that of communication between interface and application. This can of course be done through sockets. However, as a cleaner and much simpler alternative, the concept of active modules (Cabeza & Hermenegildo, 1995) can be used to advantage in this application. An active module (or an active object, if modularity is implemented via objects) is an ordinary module to which computational resources are attached (for example, a process on a UNIX machine), and which resides at a given (socket) address on the network.[5] Compiling an active module produces an executable which, when running, acts as a server for a number of relations, which are the predicates exported by the module. The relations exported by the active module can be accessed by any program on the network by simply "loading" the module and thus importing such "remote relations." The idea is that the process of loading an active module does not involve transferring any code, but rather setting up things so that calls in the local module are executed as remote procedure calls to the active module, possibly over the network. Except for compiling it in a special way, an active module is identical from the programmer point of view to an ordinary module. Also, a program using an active module imports it and uses it in the same way as any other module, except that it uses "use_active_module" rather than "use_module" (see below). Also, an active module has an address (network address) which must be known in order to use it. The address can be announced by the active module when it is started via a file or a name server (which would be itself another active module with a fixed address).

We now present the constructs related to active modules in Ciao:

:- use_active_module(*Module,Predicates*) A declaration used to import the predicates in the list *Predicates* from the active module *Module*. From this point on, the code should be written as if a standard use_module/2 declaration had been used. The declaration needs the following predicate to be accessible from the module.

module_address(*Module,Address*) This predicate must return in *Address* the address of *Module*, for any active module imported in the code. There are a number of standard libraries defining versions of this predicate.

save_addr_actmod(*Address*) This predicate should define a way to publish *Address*, to be used in active modules (the name of the active module is taken as the name of the current executable). There are a number of standard libraries defining versions of this predicate, which are in correspondence with the libraries which define versions of the previous predicate.

make_actmod(*ModuleFile,PublishModule*) Makes an active module executable from the module residing in *ModuleFile*, using address publish module of name *PublishModule*. When the executable is run (for example, at the operating system level by "*Module* &"), a socket is created and the hook predicate save_addr_actmod/1 mentioned above (which is supposed to be defined in

---

[5] It is also possible to provide active modules via a WWW address. However, we find it more straightforward to simply use socket addresses. In any case, this is generally hidden inside the access method and can be thus made transparent to the user.

PublishModule) is called in order to export the active module address as required. Then, a standard driver is run to attend network requests for the module exported predicates. Note that the code of *ModuleFile* does not need to be written in any special way.

This scheme is very flexible, allowing to completely configure the way active modules are located. This is accomplished by writing a pair of libraries, one defining the way an active module address is published, and a second defining the way the address of a given active module is found. For example, the Ciao standard libraries include as an example an implementation (libraries `filebased_publish` and `filebased_locate`) which uses a directory accessible by all the involved machines (via NFS) to store the addresses of the active modules, and the `module_address/2` predicate examines this directory to find the required data. Other solutions provided as examples include posting the address at a WWW address (`webbased_publish` / `webbased_locate`), and an implementation of a name server, that is, another active module (this one with a known, fixed address) that records the addresses of active modules and supplies this data to the modules that import it, serving as a contact agency between servers and clients.

From the implementation point of view, active modules are essentially daemons: Prolog executables which are started as independent processes at the operating system level. In the Ciao system library, communication with active modules is implemented using sockets (thus, the address of an active module is a UNIX socket in a machine). Requests to execute goals in the module are sent through the socket by remote programs. When such a request arrives, the process running the active module takes it and executes it, returning through the socket the computed results. These results are then taken by the remote processes.

Thus, when the compiler finds a `use_active_module` declaration, it defines the imported predicates as remote calls to the active module. For example, if the predicate *P* is imported from the active module *M*, the predicate would be defined as

```
P :- module_address(M,A), remote_call(A,P)
```

Compiling the following code as an active module, by writing at the Ciao toplevel "`make_actmod(phone_db, 'actmods/filebased_publish')`" (or, using the standalone compiler, by executing "`ciaoc -a 'actmods/filebased_publish' phone_db`"), creates an executable `phone_db` which, when started as a process (for example, by typing "`phone_db &`" at a UNIX shell prompt) saves its address (i.e., that of its socket) in file `phone_db.addr` and waits for queries from any module which "imports" this module (it also provides a predicate to dynamically add information to the database):

```
:- module(phone_db,[response/2,add_phone/2]).

response(Name, Response) :-
    form_empty_value(Name) ->
        Response = 'You have to provide a name.'
  ; phone(Name, Phone) ->
        Response = ['Telephone number of ',b(Name),': ',Phone]
  ; Response = ['No telephone number available for ',b(Name),'.'].

add_phone(Name, Phone) :-
    assert(phone(Name, Phone)).

:- dynamic phone/2.
phone(daniel, '336-7448').
```

```
phone(manuel, '336-7435').
phone(sacha,  '543-5316').
```

The following simple script can be used as a cgi-bin executable which will be the active module interface for the previous active module. When started, it will process the form input, issue a call to `response/2` (which will be automatically handled by the `phone_db` active module), and produce a new form before terminating. It will locate the address of the `phone_db` active module via the `module_address/2` predicate defined in library 'actmods/filebased_locate'.

```
#!/usr/local/bin/ciao-shell

:- use_active_module(phone_db,[response/2]).
:- use_module(library('actmods/filebased_locate')).
:- include(library(pillow)).


main(_) :-
    get_form_input(Input),
    get_form_value(Input,person_name,Name),
    response(Name,Response),
    output_html([
        cgi_reply,
        start,
        title('Telephone database'),
        image('phone.gif'),
        heading(2,'Telephone database'),
        --,
        Response,
        $,
        start_form,
        'Click here, enter name of clip member, and press Return:',
        \\,
        input(text,[name=person_name,size=20]),
        end_form,
        end]).
```

There are many enhancements to this simple schema which, for brevity, are only sketched here. One is to add concurrency to the active module (or whatever means of handling the client-server interaction is being used), in order to handle queries from different clients concurrently. This is easy to do in systems that support concurrency natively, such as Ciao, BinProlog/μ^2-Prolog, AKL, Oz, and KL1. We feel that Ciao can offer advantages in this area because it offers compatibility with Prolog and CLP systems while at the same time efficiently supporting concurrent execution of clause goals via local or distributed threads (Carro & Hermenegildo, 1999). Such goals can communicate at different levels of abstraction: sockets/ports, the shared fact database (similarly to a blackboard), or shared variables. BinProlog/μ^2-Prolog also supports threads, with somewhat different communication mechanisms (Tarau, 1996; Bosschere, 1989). Finally, as shown in (Szeredi *et al.*, 1996), it is also possible to exploit the concurrency present in or-parallel Prolog systems such as Aurora for implementing a multitasking server.

It is also interesting to set up things so that a single active module can handle different forms. This can be done even dynamically (i.e., the capabilities of the active

```
<form action="http://localhost/wpl_questioner.cgi">
```
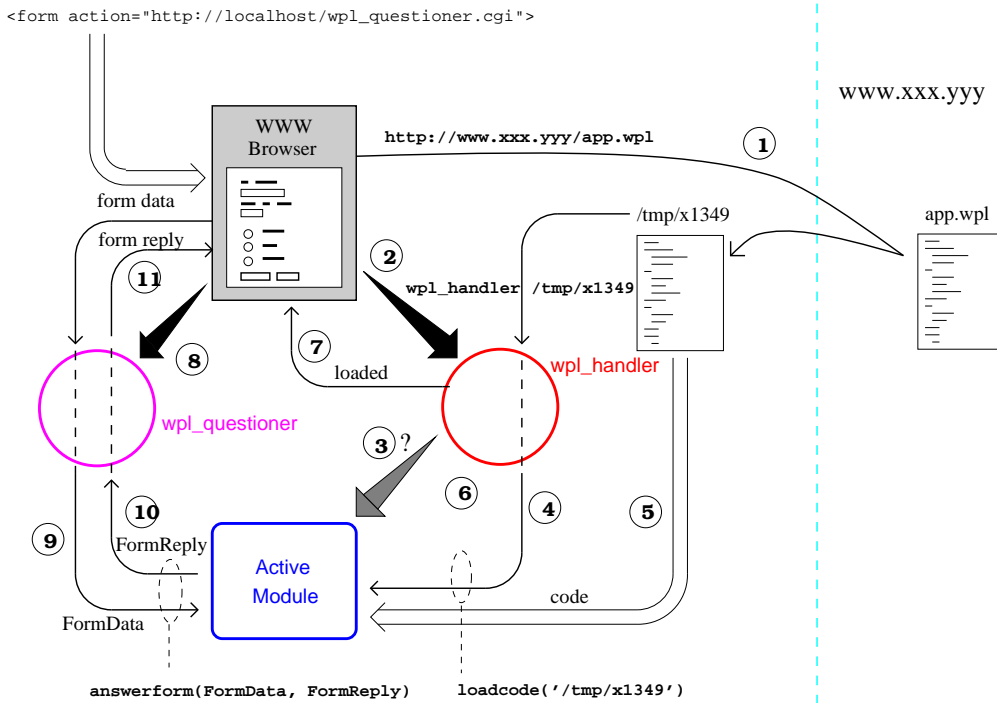

Fig. 4. Automatic code downloading architecture

module are augmented on the fly, being able to handle a new form), by designating a directory in which code to be loaded by the active module would be put, the active module consulting the directory periodically to increase its functionalities. Finally, another important issue that has not been addressed is that of providing security, i.e., ensuring that only allowed clients connect to the active module. As in the case of remote code downloading, standard forms of authentication based on codes can be used.

12 Automatic Code Downloading and Local Execution

In this section we describe an architecture which, using only the facilities we have presented in previous sections, allows the downloading and local execution of Prolog (or other LP/CLP) code by accessing a WWW address, without requiring a special browser. This is a complementary approach to giving WWW access to an active module in the sense that it provides code which will be executed in the client machine (à la Java). More concretely, the functionality that we desire is that by simply clicking on a WWW pointer, and transparently for the user, remote Prolog code is automatically downloaded in such a way that it can be queried via forms and all the processing is done locally.

To allow this, the HTTP server on the server machine is configured to give a specific `mime.type` (for example `application/x-prolog`) to the files which will hold WWW-downloadable Prolog code (for example those with a special suffix, like `.wpl`). On the other side, the browser is configured to start the `wpl_handler` helper application when receiving data of type `application/x-prolog`. This `wpl_handler` application is the interface to a Prolog engine which will execute the WWW downloaded code, acting as an active module. We now sketch the procedure (see figure 4):

1. The form that will be used to query the downloaded code (and which we assume already loaded on the browser) contains a link which points to a

WWW-aware Prolog code file. Clicking on this link produces the download as explained below. Note that for browsers that can handle `multipart/mixed` mime types (such as most modern browsers), the form and the code file could alternatively be combined in the same document. However, for brevity, we will only describe the case when they are separate. The handler for the form is specified as the local cgi-bin executable `wpl_questioner.cgi`.

2. As the server of the file tells the browser that this page is of type `application/x-prolog`, the browser starts a `wpl_handler` and passes the file to it (in this example by saving the file in a temporal directory and passing its name).

3. The `wpl_handler` process checks whether a Prolog engine is currently running for this browser and, if necessary, starts one. This Prolog engine is configured as an active module.

4. Then, through a call to a predicate of the active module "`loadcode(File)`" the handler asks the active module to read the code.

5. The active module reads the code and compiles it.

6. `wpl_handler` waits for the active module to complete the compilation and writes a "done" message to the browser.

7. The browser receives the "done" message.

8. Now, when the "submit" button in the form is pressed, and following the standard procedure for forms, the browser starts a `wpl_questioner` process, sending it the form data.

9. The `wpl_questioner` process gets this form data, translates it to a dictionary *FormData* and passes it to the active module through a call to its exported predicate `answerform(FormData,FormReply)`.

10. The active module processes this request, and returns in *FormReply* a WWW page (as a term) which contains the answer to it (and possibly a new form).

11. The `wpl_questioner` process translates *FormReply* to raw HTML and gives it back to the browser, dying afterwards. Subsequent queries to the active module can be accomplished either by going back to the previous page (using the "back" button present in many browsers) or, if the answer page contains a new query form, by using it. In any case, the procedure continues at 8.

The net effect of the approach is that by simply clicking on a WWW pointer, remote Prolog code is automatically downloaded to a local Prolog engine. Queries posed via the form are answered locally by the Prolog engine.

There are obvious security issues that need to be taken care of in this architecture. Again, standard authentication techniques can be used. However, since source code is being passed around, it is comparatively easy to verify that no dangerous predicates (for example, perhaps those that can access files) are executed. Note again that it is also possible to download bytecode, since this is supported by most current LP/CLP systems, using a similar approach.

13 Related Work

Previous general purpose work on WWW programming using computational logic systems includes, to the best of our knowledge, the publicly available `html.pl` library (Cabeza & Hermenegildo, 1996a) and manual, and the LogicWeb system (Loke & Davison, 1996) (the P*i*LL*o*W library was also described previously in (Cabeza *et al.*, 1996)). The `html.pl` library was built by D. Cabeza and M. Hermenegildo, using input from L. Naish's forms code for NU-Prolog and M. Hermenegildo and F. Bueno's experiments building a WWW interface to the CHAT-80 (Warren & Pereira, 1982) program. It was released as a publicly available WWW

library for LP/CLP systems and announced, among other places, in the Internet `comp.lang.prolog` newsgroup (Cabeza & Hermenegildo, 1996b). The library has since been ported to a large number of systems and adapted by several Prolog vendors, as well as used by different programmers in various institutions. In particular, Ken Bowen has ported the library to ALS Prolog and extended it to provide group processing of forms and an alternative to our use of active modules (Bowen, 1996). The present work is essentially a significant extension of the `html.pl` library.

The main other previous body of work related to general-purpose interfacing of logic programming and the WWW that we have knowledge of is the LogicWeb (Loke & Davison, 1996) system, by S.W. Loke and A. Davison. The aim of LogicWeb is to use logic programming to extend the concept of WWW pages, incorporating in them programmable behavior and state. In this, it shares goals with Java. It also offers rich primitives for accessing code in remote pages and module structuring. The aims of LogicWeb are different from those of `html.pl`/P*i*LL*o*W. LogicWeb is presented as a system itself, and its implementation is done through a tight integration with the Mosaic browser, making use of special features of this browser. In contrast, `html.pl`/P*i*LL*o*W is a general purpose library, meant to be used by a general computational logic systems and is browser-independent. `html.pl`/P*i*LL*o*W offers a wide range of functionalities, such as syntax conversion between HTML and logic terms, access predicates for WWW pages, predicates for handling forms, etc., which are generally at a somewhat lower level of abstraction than those of LogicWeb. We believe that using P*i*LL*o*W and the ideas sketched in this paper it is possible to add the quite interesting functionality offered by LogicWeb to standard LP and CLP systems. We have shown some examples including access to passive remote code (modules with an `ftp` or `http` address) from programs and automatic remote code access and querying using standard browsers and forms. In addition, we have discussed active remote code, where the functionality, rather than the code itself, is exported.

More recently, a larger body of work on the topic was presented at the workshop held on the topic of Logic Programming and the Internet at the 1996 Joint International Conference and Symposium on Logic Programming (where also a previous version of this paper was presented). The work presented in (S.W. Loke & Sterling, 1996) is based on LogicWeb, and aims to provide distributed lightweight databases on the WWW. As with the basic LogicWeb system, we believe that the P*i*LL*o*W library can be used to implement in other systems the interesting ideas proposed therein. As briefly mentioned before, the work in (Szeredi *et al.*, 1996) proposes an architecture similar to that of our active modules in order to handle form requests. In this solution the handling multiple requests is performed by using or-parallelism. While we feel that and-parallelism (as in &-Prolog's or Ciao's threads) is more natural for modeling this kind of concurrency, the ideas proposed are quite interesting. The ECLiPSe HTTP-library (Ph. Bonnet & Thomsen, 1996), aimed at implementing INTERNET agents, offers functionality that is in part similar to that of the Ciao `html.pl`/P*i*LL*o*W libraries, including facilities that are similar to our active modules. The approach is different, however, in several respects. The ECLiPSe library implements special HTTP servers and clients. In contrast, P*i*LL*o*W uses standard HTTP servers and interfaces. Using special purpose servers may be interesting because the approach possibly allows greater functionality. On the other hand this approach in general requires either the substitution of the standard server on a given machine or setting the special server at a different socket address from the standard one. The ECLiPSe library also contains functionality that is related to our active modules, although the interface provided is at a lower level. Finally, other papers describing very interesting WWW applications are being presented

regularly, which underline the suitability of computational logic systems for the task. We believe that the Ciao P*i*LL*o*W library can contribute to making it even easier to develop such applications in the future.

Additional work on the topic of Logic Programming and the Internet can be found in the proceedings of the workshop sponsored by the Compulog-Net research network. The reader is referred to the tutorials and papers presented in these two workshops for more information on a number of applications, other libraries, and topics such as interfacing and compilation from computational logic systems to Java. Examples of Prolog systems interfaced with Java are BinProlog (see `http://clement.info.umoncton.ca/BinProlog`), Ciao (Bueno *et al.*, 1997), and others (Calejo, 1999). Experimental Prolog to Java compilers have been built both in academia (see for example jProlog at `http://www.cs.kuleuven.ac.be/~bmd/PrologInJava/`) and Commercially (see for example the IF Prolog tools `http://www.ifcomputer.com`). This approach is quite attractive, although the results cannot compete in performance with conventional Prolog compilers (it is open for research whether improvements in Java performance or improved Prolog-to-Java compilation technology can bridge the gap). Other commercial work on the topic of interfacing Prolog and the WWW (in addition to that done on the ALS system mentioned above) include the Amzi! Prolog WebLS System (`http://www.amzi.com/share.htm`) and the LPA PrologWeb System (`http://www.lpa.co.uk`).

Recent work using P*i*LL*o*W includes the "Web Integrator"(Davulcu *et al.*, 1999) –a webbase system that integrates data from various Web sources, and allows users to query these Web sources as if they were a single database– and `WebDB` (Cabeza & Hermenegildo, 1998) –a WWW-based database management interface. Also, within the RadioWeb project (Partners, 1997), we have developed (in collaboration with the group of M. Codish at Ben Gurion University) a constraint-based language for describing WWW page layout and style rules and an engine which, by interpreting these rules, can generate WWW sites which dynamically adapt to parameters such as user characteristics (Cederberg & CLIP Group, 1999).

Additional applications developed with the P*i*LL*o*W library can be accessed from the P*i*LL*o*W WWW site (see later). A page with pointers to the proceedings of the previously mentioned workshops, as well as other information (including technical reports and tutorial) regarding the topic of Logic Programming, Constraint Programming, and the Internet is maintained at `http://www.clip.dia.fi.upm.es/lpnet/`.

14 Conclusions and Future Work

We have discussed from a practical point of view a number of issues involved in writing Internet and WWW applications using LP/CLP systems. In doing so, we have described P*i*LL*o*W, an Internet/WWW programming library for LP/CLP systems. P*i*LL*o*W provides facilities for generating HTML/XML structured documents, producing HTML forms, writing form handlers, processing HTML/XML templates, accessing and parsing WWW documents, and accessing code posted at HTTP addresses. We have also described the architecture of some application classes, including automatic code downloading, using a high-level model of client-server interaction, *active modules*. Finally we have also described an architecture for automatic LP/CLP code downloading for local execution, using generic browsers. We believe that the Ciao P*i*LL*o*W library can ease substantially the process of developing WWW applications using computational logic systems.

We have recently developed several extensions to the library (for example, for

setting and getting "cookies"), and sample applications which make extensive use of concurrency (on those LP/CLP systems that support it) to overlap network requests. We have also developed a complementary library for interfacing Prolog with the Virtual Reality Modeling Language VRML (Smedbäck *et al.*, 1999).

In addition to being included as part of the Ciao system, the P*i*LL*o*W library is provided as a standard, standalone public domain library for SICStus Prolog and other Prolog and CLP systems, supporting most of its functionality. Please contact the authors or consult our WWW site `http://www.clip.dia.fi.upm.es` and the P*i*LL*o*W page at `http://www.clip.dia.fi.upm.es/Software/pillow/-pillow.html` for download details and an up-to-date online version of the P*i*LL*o*W manual. The Ciao Prolog system is also freely available from `http://www.clip.dia.fi.upm.es` and `http://www.ciaoprolog.org`.

Acknowledgments

References

Berners-Lee, T., Cailliau, R., Luotonen, A., Nielsen, H.F., & Secret, A. (1994). The World-Wide Web. *Communications of the ACM*, **37**(8), 76–82.

Bosschere, K. De. (1989). Multi–Prolog, Another Approach for Parallelizing Prolog. *Pages 443–448 of: Proceedings of parallel computing.* Elsevier, North Holland.

Bowen, K. 1996 (March). *Personal communication.* Available from `http://www.als.com/als/html_pl.html`.

Bueno, F., Cabeza, D., Carro, M., Hermenegildo, M., López-García, P., & Puebla, G. 1997 (August). *The Ciao Prolog System. Reference Manual.* The Ciao System Documentation Series–TR CLIP3/97.1. School of Computer Science, Technical University of Madrid (UPM).

Cabeza, D., & Hermenegildo, M. (1995). Distributed Concurrent Constraint Execution in the Ciao System. *Proc. of the 1995 compulog-net workshop on parallelism and implementation technologies.* Utrecht, NL: U. Utrecht / T.U. Madrid. Available from `http://www.clip.dia.fi.upm.es/`.

Cabeza, D., & Hermenegildo, M. 1996a (March). *html.pl: An HTML package for (C)LP systems.* Spain. Available from `http://www.clip.dia.fi.upm.es/miscdocs/`.

Cabeza, D., & Hermenegildo, M. 1996b (February). *LP/CLP HTML and WWW interface publicly available.* Posting in `comp.lang.prolog`. Available from `http://www.clip.dia.fi.upm.es/`.

Cabeza, D., & Hermenegildo, M. 1997 (April). WWW Programming using Computational Logic Systems (and the PILLOW/Ciao Library). *Proceedings of the workshop on logic programming and the WWW at WWW6.*

Cabeza, D., & Hermenegildo, M. 1998 (June). *A WWW Database Management Interface for Prolog.* Technical Report CLIP2/98.0. School of Computer Science, Technical University of Madrid (UPM), Facultad Informática UPM, 28660-Boadilla del Monte, Madrid-Spain.

Cabeza, D., Hermenegildo, M., & Varma, S. 1996 (September). The P*i*LL*o*W/Ciao Library for INTERNET/WWW Programming using Computational Logic Systems. *Proceedings of the 1st workshop on logic programming tools for internet applications.* Available from `http://clement.info.umoncton.ca/~lpnet`.

Calejo, M. (1999). Java+prolog: A land of opportunities. *Pages 1–2 of: Proceedings of the first international conference on the practical application of constranint technologies and logic programming.* The Practical Application Company Ltd. Also available at `http://dev.servisoft.pt/interprolog/paclp99/default.htm`.

Carlsson, M. 1988 (February). *Sicstus Prolog User's Manual.* Po Box 1263, S-16313 Spanga, Sweden.

Carro, M., & Hermenegildo, M. (1999). Concurrency in Prolog Using Threads and a Shared Database. *Pages 320–334 of: 1999 international conference on logic programming.* MIT Press, Cambridge, MA, U.S.A.

Cederberg, Per, & CLIP Group, The. 1999 (June). *Flexible Layout and Styling – The LaSt Language.* Technical Report D2.2.M3 CLIP 3/99.0. RADIOWEB Project.

Chikayama, T., Fujise, T., & Sekita, D. (1994). A Portable and Efficient Implementation of KL1. Tick, Evan (ed), *Proc. of the 1994 icot/nsf workshop on parallel and concurrent programming.* U. of Oregon.

Colmerauer, A. (1975). *Les gramaire de metamorphose.* Tech. rept. Univ. D'aix-Marseille, Groupe De Ia.

Colmerauer, A. (1990). An Introduction to Prolog III. *Communications of the ACM,* **28**(4), 412–418.

Davulcu, Hasan, Freire, Juliana, Kifer, Michael, & Ramakrishnan, I.V. 1999 (June). A layered architecture for querying dynamic web content. *ACM SIGMOD international conference on management of data.* URL: `http://www.acm.org/sigmod/sigmod99/eproceedings/`.

ECRC. (1993). *Eclipse user's guide.* European Computer Research Center.

Grobe, M., & Naseer, Hasan. 1998 (July). *An instantaneous introduction to CGI scripts and HTML forms.* Available from `http://www.cc.ukans.edu/~acs/docs/other/forms-intro.shtml`.

Hermenegildo, M. 1996 (April). *Writing "Shell Scripts" in SICStus Prolog.* Posting in `comp.lang.prolog`. Available from `http://www.clip.dia.fi.upm.es/`.

Hermenegildo, M., & CLIP Group, The. (1994). Some Methodological Issues in the Design of Ciao - A Generic, Parallel, Concurrent Constraint System. *Pages 123–133 of: Principles and practice of constraint programming.* LNCS, no. 874. Springer-Verlag.

Hermenegildo, M., Bueno, F., García de la Banda, M., & Puebla, G. 1995a (December). The Ciao Multi-Dialect Compiler and System: An Experimentation Workbench for Future (C)LP Systems. *Proceedings of the ilps'95 workshop on visions for the future of logic programming.* Available from `http://www.clip.dia.fi.upm.es/`.

Hermenegildo, M., Cabeza, D., & Carro, M. (1995b). Using Attributed Variables in the Implementation of Concurrent and Parallel Logic Programming Systems. *Pages 631–645 of: Proc. of the twelfth international conference on logic programming.* MIT Press.

Hermenegildo, M., Bueno, F., Cabeza, D., Carro, M., M. García de la Banda, López-García, P., & Puebla, G. (1999a). The Ciao Multi-Dialect Compiler and System: An Experimentation Workbench for Future (C)LP Systems. *Pages 65–85 of: Parallelism and Implementation of Logic and Constraint Logic Programming.* Commack, NY, USA: Nova Science.

Hermenegildo, M., Puebla, G., & Bueno, F. (1999b). Using Global Analysis, Partial Specifications, and an Extensible Assertion Language for Program Validation and Debugging. *Pages 161–192 of:* Apt, K. R., Marek, V., Truszczynski, M., & Warren, D. S. (eds), *The Logic Programming Paradigm: a 25–Year Perspective.* Springer-Verlag.

Jaffar, Joxan, & Lassez, Jean-Louis. (1987). Constraint Logic Programming. *Pages 111–119 of: ACM symposium on principles of programming languages.* ACM.

Janson, S., & Haridi, S. (1991). Programming Paradigms of the Andorra Kernel Language. *Pages 167–183 of: 1991 international logic programming symposium.* MIT Press.

Kowalski, R. A. (1974). Predicate Logic as a Programming Language. *Pages 569–574 of: Proceedings ifips.*

Loke, S.W., & Davison, A. (1996). Logic programming with the World-Wide Web. *Pages 235–245 of: 7th. ACM Conference on Hypertext.* ACM Press. Available from `http://www.cs.unc.edu/~barman/HT96/P14/lpwww.html`.

M. Dincbas, H. Simonis, & Van Hentenryck, P. (1990). Solving Large Combinatorial Problems in Logic Programming. *Journal of logic programming,* **8**(1 & 2), 72–93.

Partners, The RADIOWEB Project. 1997 (July). *RADIOWEB EP25562: Automatic Generation of Web Sites for the Radio Brodcasting Industry – Project Description / Technical Annex*. Technical Report. RADIOWEB Project.

Ph. Bonnet, S. Bressan, L. Leth, & Thomsen, B. 1996 (September). Towards ECLiPSe Agents on the INTERNET. *Proceedings of the 1st workshop on logic programming tools for internet applications*. Available from `http://clement.info.umoncton.ca/~lpnet/lpnet2.html`.

Smedbäck, G., Carro, M., & Hermenegildo, M. (1999). Interfacing Prolog and VRML and its Application to Constraint Visualization. *Pages 453–471 of: The practical application of constraint technologies and logic programming*. The Practical Application Company.

Smolka, G. 1994 (November). *The Definition of Kernel Oz*. DFKI Oz documentation series. German Research Center for Artificial Intelligence (DFKI).

S.W. Loke, A. Davison, & Sterling, L. 1996 (September). Lightweight Deductive Databases on the World-Wide Web. *Proceedings of the 1st workshop on logic programming tools for internet applications*. Available from `http://clement.info.umoncton.ca/~lpnet/lpnet10.html`.

Szeredi, Péter, Molnár, Katalin, & Scott, Rob. 1996 (September). Serving Multiple HTML Clients from a Prolog Application. *Proceedings of the 1st workshop on logic programming tools for internet applications*. Available from `http://clement.info.umoncton.ca/~lpnet/lpnet9.html`.

Tarau, P. 1996 (April). *Binprolog 5.00*. Posting in `comp.lang.prolog`. Available from `http://clement.info.umoncton.ca/~tarau`.

Van Hentenryck, P. (1989). *Constraint Satisfaction in Logic Programming*. MIT Press.

Warren, D.H.D., & Pereira, F. C. N. (1982). An Efficient, Easily Adaptable System For Interpreting Natural Language Queries. *American journal of computational linguistics*, **8**(3-4), 110–122.