

# Links: Web Programming Without Tiers\*

Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop

University of Edinburgh

**Abstract.** Links is a programming language for web applications that generates code for all three tiers of a web application from a single source, compiling into JavaScript to run on the client and into SQL to run on the database. Links supports rich clients running in what has been dubbed ‘*Ajax*’ style, and supports concurrent processes with statically-typed message passing. Links is *scalable* in the sense that session state is preserved in the client rather than the server, in contrast to other approaches such as Java Servlets or PLT Scheme. Client-side concurrency in JavaScript and transfer of computation between client and server are both supported by translation into continuation-passing style.

## 1 Introduction

A typical web system is organized in three tiers, each running on a separate computer (see Figure ??). Logic on the middle-tier server generates pages to send to a front-end browser and queries to send to a back-end database. The programmer must master a myriad of languages: the logic is written in a mixture of Java, Perl, PHP, and Python; the pages described in HTML, XML, and JavaScript; and the queries are written in SQL or XQuery. There is no easy way to ensure that data interfaces between the languages match up — that a form in HTML or a query in SQL produces data of a type that the logic in Java expects. This is called the *impedance mismatch* problem. The problem is exacerbated because code for the browser or database is often generated at runtime, making web applications error prone and difficult to debug.

Links eliminates impedance mismatch by providing a single language for all three tiers. In the current version, Links translates into JavaScript to run on the browser and SQL to run on the database. The server component is written in O’Caml; it consists of a static analysis phase (including Hindley-Milner typechecking), a translator to JavaScript and SQL, and an interpreter for the Links code that remains on the server. All run-time code generation is performed by the compiler in a principled manner, rather than by the programmer using techniques such as string processing; principled generation of this code makes applications less error prone and easier to debug, as well as making it now possible to perform type checking for information passed between the three tiers. The programmer still needs to think about which code runs in which location, especially regarding security of data, but now in a context where the fundamental properties of well-formed code and well-typed data are guaranteed.

Increasingly, web applications designers are migrating work into the browser. “Rich client” systems, such as Google Mail and Google Maps, use a new style of interaction

---

\* Submitted to FMCO ‘06




Fig. 1. Three-tier model

dubbed “*Ajax*” [?]. Client-side Links compiles into JavaScript, a functional language widely available on most browsers. JavaScript is notoriously variable across platforms, so we have designed the compiler to target a common subset that is widely available. It would be easy to extend Links to support additional target languages, such as Flash or Java. However, we expect the popularity of Ajax will mean that standard and reliable versions of JavaScript will become available over the next few years.

Links is a strict, typed, functional language. It incorporates ideas proven in other functional languages, including:

- database query optimization, as found in Kleisli and elsewhere,
- continuations for web interaction, as found in PLT Scheme and elsewhere, and
- concurrency with message passing, as found in Erlang and elsewhere,

All three of these features work better with immutable values rather than mutable objects. In Links, side effects play a limited (though important) role, being used for updates to the database and display, and communication between concurrent processes. Types ensure consistency between forms in the browser, logic in the server, and queries on the database—and between senders and receivers in a concurrent program.

Links programs are *scalable* in the sense that session state is preserved in the client rather than the server. Many commercial web tools (like J2EE) and most research web tools (including current releases of PLT Scheme [?] and Mozart QHTML [?]) are not scalable in this sense.

In Links, all server state is serialized and passed to the client, then restored to the server when required. This *resumption passing style* extends the *continuation passing style* commonly used in PLT Scheme and elsewhere. Links functions are labelled as to whether they are intended to execute on the client or the server; functions running on the client may invoke those on the server, and vice-versa.

*Database programming.* Queries are written in the Links notation and compiled into SQL, a technique pioneered by Kleisli [?,?] and now used in LINQ [?].

*Web interaction.* The notion that a programming language could provide support for web interaction first appears in the programming language MAWL [?]. The notion of *continuation* from functional programming has been particularly fruitful, being applied by a number of researchers to improve interaction with a web client, including Quiennec

[?], Graham [?] (in a commercial system sold to Yahoo and widely used for building web stores), Felleisen and others [?,?], and Thiemann [?].

*Concurrency.* Links supports concurrent programming in the client, using “share nothing” concurrency where the only way processes can exchange data is by message passing, as pioneered in Erlang [?] and Mozart [?].

*XML programming.* Links provides convenient syntax for constructing XML data, similar to that provided by XQuery [?]. The current version does not support regular expression types, as found in XDuce and other languages, but we may add them in a future version. Regular expression types were given a low priority, because they are already well understood from previous research [?].

*Other languages.* Other languages for web programming include Xtatic [?], Scala [?], Mozart [?], SML.NET [?], F# [?], C $\omega$  (based on Polyphonic C# [?] and Xen [?]), HOP [?] and Ocsigen [?]. These languages have many overlaps with Links, as they are also inspired by the functional programming community.

However, none of these languages shares Links’ objective of generating code for all three tiers of a web application from a single source — scripts for the front-end client, logic for the middle-tier server, and queries for the back-end database. We expect that providing a single, unified language as an alternative to the current multiplicity of languages will be a principal attraction of Links.

*This paper.* We introduce Links by describing three examples in Section ?? . Section ?? sketches our features for concurrency and client-server interaction. Section ?? gives an SQL-compatible subset of Links and details how it is compiled into SQL, while Section ?? describes “mailbox typing” for message-passing. Section ?? discusses some shortcomings of the current implementation and Section ?? concludes.

## 2 Links by example

This section introduces Links by a series of examples. The reader is encouraged to try these examples online at

<http://groups.inf.ed.ac.uk/links/examples/>

We begin with an example to introduce the basic functionality of Links, and then present two further examples that demonstrate additional capabilities: a draggable list, and a progress bar.

### 2.1 Dictionary suggest

The Dictionary Suggest application presents a text box, into which the user may type a word. As the user types, the application displays a list of words that could complete the one being typed. A dictionary database is searched, and the first ten words beginning with the given prefix are presented (see Figure ??). In addition, the user can add, update

## Dictionary suggest

### Search for definitions

fun

### Click a definition to edit it

**fun** Sport, merriment, frolicsome amusement.

**funambulate** To walk or to dance on a rope.

**funambulation** Ropedancing.

**funambulatory** Narrow, like the walk of a ropedancer.

**funambulatory** Performing like a ropedancer.

**funambulist** A ropewalker or ropedancer.

**funambulo** Alt. of Funambulus

**funambulus** A ropewalker or ropedancer.

Word: function

Meaning: A public or social ceremony or gathering; a festivity or entertainment, esp. one somewhat formal.

Update Cancel

Delete

**function** A quantity so connected with another quantity, that if any alteration be made in the latter there will be a consequent alteration in the former. Each quantity is said to be a function of the other. Thus, the circumference of a circle is a function of the diameter. If  $x$  be a symbol to which different numerical values can be assigned, such expressions as  $x^2$ ,  $3x$ ,  $\text{Log } x$ , and  $\text{Sin } x$ , are all functions of  $x$ .

### New definition

Word:

Meaning:

Add

**Fig. 2.** Dictionary Suggest screenshot. An entry for “function” appearing in the word list has just been clicked, so its entry has expanded into an editable form.

and delete a definition by clicking on it. To add a new definition the user fills in and submits the form at the bottom of the page by clicking ‘Add’.

To update an existing definition the user clicks on one of the suggestions, which then expands into a form containing the existing definition; then the user edits the form and submits it by clicking ‘Update’. This form also includes buttons for cancelling the update (which restores the original suggestion) and deleting the definition.

This application is of interest because it must perform a database lookup and update the display at every keystroke. Applications such as Google Suggest [?] have a similar structure.

The Links version is based on an ASP.NET version, available online [?], using the same data. It extends the ASP.NET version by allowing the definitions to be added, updated and deleted. The dictionary contains 99,320 entries. For comparison, the ASP.NET version responds to a keystroke in about 0.1s (due to the difficulty of instrumenting third-party code, measurements were made with a stopwatch. The average of a dozen measurements, allowing for a similarly measured reaction time, was 0.1s). In the Links

```

var defsTable =
  table "definitions" with
    (id:String, word:String, meaning:String)
  where id readonly from database "dictionary";

fun newDef(def) server { insert defsTable values [def] }
fun updateDef(def) server {
  update (var d <-- defsTable) where (d.id == def.id)
  set (word=def.word, meaning=def.meaning)
}
fun deleteDef(id) server {
  delete (var def <-- defsTable) where (def.id == id)
}

fun completions(s) server {
  if (s == "") [] else {
    take(10, for (var def <-- defsTable)
      where (def.word ~ /s.*/)) orderby (def.word)
    [def])
  }
}

fun suggest(s) client {
  replaceChildren(format(completions(s)),
    getNodeById("suggestions"))
}

fun editDef(def) client {
  redraw(
    <form l:onsubmit="{
      var def = (id=def.id, word=w, meaning=m); updateDef(def);
      redraw(formatDef(def), def.id)}" method="POST">
    <table>
      <tr><td>Word:</td><td>
        <input l:name="w" value="{def.word}"/></td></tr>
      <tr><td>Meaning:</td><td>
        <textarea l:name="m" rows="5" cols="80">
          {stringToXml(def.meaning)}</textarea></td></tr>
    </table>
    <button type="submit">Update</button>
    <button l:onclick="{redraw(formatDef(def), def.id)}">
      Cancel</button>
    <button l:onclick="{deleteDef(def.id); redraw([], def.id)}"
      style="position:absolute; right:0px">Delete</button>
    </form>,
    def.id)
}

```

Fig. 3. Dictionary suggest in Links (1)

```

fun redraw(xml, defId) client {
  replaceChildren(xml, getNodeById("def:"+defId))
}

fun formatDef(def) client {
  <span l:onclick="{editDef(def)}">
    <b>{stringToXml(def.word)}</b>
    {stringToXml(def.meaning)}<br/>
  </span>
}

fun format(defs) client {
  <#>
    <h3>Click a definition to edit it</h3>
    for (var def <- defs)
      <span class="def" id="def:def.id">{formatDef(def)}</span>
  </#>
}

fun addForm(handler) client {
  <form l:onsubmit="{handler!NewDef((word=w, meaning=m))}">
    <table>
      <tr><td>Word:</td><td>
        <input type="text" l:name="w"/></td></tr>
      <tr><td>Meaning:</td><td>
        <textarea l:name="m" rows="5" cols="80"/></td></tr>
      <tr><td><button type="submit">Add</button></td></tr>
    </table>
  </form>
}

var handler = spawn {
  fun receiver(s) {
    receive {
      case Suggest(s) -> suggest(s); receiver(s)
      case NewDef(def) ->
        newDef(def);
        replaceChildren(addForm(self()), getNodeById("add"));
        suggest(s); receiver(s)
    }
  }
  receiver("")
};

```

**Fig. 4.** Dictionary suggest in Links (2)

```

<html>
<head>
  <style>.def {{ color:blue }}</style>
  <title>Dictionary suggest</title>
</head>
<body>
  <h1>Dictionary suggest</h1>
  <h3>Search for definitions</h3>
  <form l:onkeyup="{handler!Suggest(s)}">
    <input type="text" l:name="s" autocomplete="off"/>
  </form>
  <div id="suggestions"/>
  <h3>New definition</h3>
  <div id="add">{addForm(handler)}</div>
</body>
</html>

```

**Fig. 5.** Dictionary suggest in Links (3)

version, over 36 trials with various prefixes, total response time measured on average 649ms with a standard deviation of 199ms; subtracting the time spent performing the database query in each trial, the average time taken was 297ms with a standard deviation of 54ms. Given that no effort has been spent trying to optimize the Links system, this seems to indicate acceptable performance at this stage.

The code for the application is shown in Figures ??-??; following is a short walk-through of the code. On each keystroke, a *Suggest* message containing the current contents of the text field is sent to the *handler* process. The *handler* process passes the text content to the function *suggest*. This function calls *completions*, which executes on the server, to find the first ten words with the given prefix, and *format* (executing on the client) to format the list returned. Doing the server interaction in a separate *handler* process allows the user interaction to remain responsive, even while looking up suggestions.

The rest of the code is concerned with modifying the database. A form for adding definitions is created by the function *addForm*. Clicking ‘Add’ sends a *NewDef* message to the *handler* process containing a new definition. The *handler* process calls *newDef* to add the definition, then resets the form and updates the list of suggestions (in case the new definition appears in the current list of suggestions).

Clicking on a definition invokes the function *editDef*, which calls the function *redraw* in order to replace the definition with a form for editing it. Clicking ‘Cancel’ reverses this operation. Clicking ‘Update’ or ‘Delete’ performs the corresponding modification to the definition by calling *updateDef* or *deleteDef* on the server, and then updates the list of suggestions by calling the function *redraw* on the client.

Having sketched the basic structure of the example, we now describe several of the key features of Links, illustrating them by our example code. The features we cover

are syntax, types, XML, regular expressions, interaction, list comprehensions, database access and update, concurrency, and partitioning the program onto client and server.

The core of Links is a fairly standard functional programming language with Hindley-Milner type inference. One missing feature is exception handling, which we plan to add in a future version.

*Syntax.* The syntax of Links resembles that of JavaScript. This decision was made not because we are fond of this syntax, but because we believe it will be familiar to our target audience. Low-level syntactic details become a matter of habit that can be hard to change: we conjecture that using the familiar  $f(x, y)$  in place of the unfamiliar  $f\ x\ y$  will significantly lower the barrier to entry of functional programming.

One difference from JavaScript syntax is that we do not use the keyword **return**, which is too heavy to support a functional style. Instead, we indicate return values subtly (perhaps too subtly), by omitting the trailing semicolon. The type checker indicates an error if a semicolon appears after an expression that returns any value other than the unit value  $()$ .

*Types.* Links uses Hindley-Milner type inference with row variables [?]. As basic types Links supports integers, floats, characters, booleans, lists, functions, records, and variants.

- A list is written  $[e_1, \dots, e_k]$ , and a list type is written  $[A]$ .
- A lambda abstraction is written **fun**  $(x_1, \dots, x_k) \{e\}$ , and a function type is written  $(A_1, \dots, A_k) \rightarrow B$ .
- A record is written  $(f_1=e_1, \dots, f_k=e_k)$ , and a record type is written  $(f_1:A_1, \dots, f_k:A_k \mid r)$ , where  $r$  is an optional row variable. Field names for records begin with a lower-case letter.
- A variant is written  $F_i(e_i)$  and a variant type is written  $[ \mid F_1:A_1, \dots, F_k:A_k \mid r \mid ]$ . Field names for variants begin with an upper-case letter.

Strings are simply lists of characters, and tuples are records with natural number labels. Apart from the table declaration in Figure ??, none of the examples in the paper explicitly mention types. This is partly because type inference renders type annotations unnecessary and partly to save space. Nevertheless, it should be stressed that all of the examples are type-checked statically, and static typing is an essential part of Links.

Links currently does not support any form of overloading; we expect to support overloading in future using a simplified form of type classes. As with regular expression types, this was left to the future because it is well understood from other research efforts. In particular, the WASH and iData systems make effective use of type classes to support generic libraries [?,?].

*XML.* Links includes special syntax for constructing and manipulating XML data. XML data is written in ordinary XML notation, using curly braces to indicate embedded code. Embedded code may occur either in attributes or in the body of an element. The Links notation is similar to that used in XQuery, and has similar advantages. In



particular, it is easy to paste XML boilerplate into Links code. The parser begins parsing XML when a `<` is immediately followed by a legal tag name; a space must always follow `<` when it is used as a comparison operator; legal XML does not permit a space after the `<` that opens a tag. Links also supports syntactic sugar `<#> . . . </#>` for specifying an XML forest literal as in the function *format* in Figure ??.

The client maintains a data structure representing the current document to display, called the Document Object Model, or DOM for short. Often this structure is represented in some form of HTML, such as XHTML, the dialect of XML corresponding to HTML. Links provides library functions to access and modify the DOM, based on the similar operations specified by the W3C DOM standard. Links supports two types for manipulating XML: *DomNode* is a mutable reference (similar to a *ref* type in ML), while *Xml* is an immutable list of trees. There is an operation that converts the former to the latter by making a deep copy of the tree rooted at the node, returning it in a singleton list. We expect eventually to support regular expression types for XML that refine each of these two types, and to support a notation like XPath for manipulating trees, but as these points are well understood (but a lot of work to implement) they are not a current priority.

*Regular expressions.* Matching a string against a regular expression is written  $e \sim /r/$  where  $r$  is a regular expression. Curly braces may be used to interpolate a string into a regular expression, so for example  $/\{s\} . */$  matches any string that begins with the value bound to the variable  $s$ .

*Interaction.* The Links code specifies display of an XML document, the crucial part of which is the following:

```
<form l:onkeyup="{handler!Suggest(s)}">
  <input type="text" l:name="s" autocomplete="off"/>
</form>
<div id="suggestions"/>
```

The `l:name` attribute specifies that the string typed into the field should be bound to a string variable  $s$ .

The attributes `l:name` and `l:onkeyup` are special. The attribute `l:onkeyup` is followed by Links code in curly braces that is *not* immediately evaluated, but is evaluated whenever a key is released while typing in the form. (Normally, including curly braces in XML attributes, as elsewhere in XML, causes the Links code inside the braces to be evaluated and its value to be spliced into the XML.) The `l:name` attribute on an input field must contain a Links variable name, and that variable is bound to the contents of the input field.

The attributes that Links treats specially are `l:name` and all the attributes connected with events: `l:onchange`, `l:onsubmit`, `l:onkeyup`, `l:onmousemove`, and so on. These attributes are prefixed with `l:`, using the usual XML notation for namespaces; in effect, `l` denotes a special Links namespace.

The scope of variables bound by `l:name` is the Links code that appears in the attributes connected with events. Links static checking ensures that a static error is raised if a name is used outside of its scope; this guarantees that the names mentioned

on the form connect properly to the names referred to by the application logic. In this, Links resembles MAWL and Jwig, and differs from PLT Scheme or PHP.

Our experience has shown that this style of interaction does not necessarily scale well, and it may be preferable to use a library of higher-order forms as developed in WASH or iData. We return to this point in Section ??.

*List comprehensions.* The Links code calls the function *suggest* each time a key is released. This in turn calls *completions* to find the first ten completions of the prefix, and *format* to format the results as HTML for display. Both *completions* and *format* use **for** loops, in the former case also involving **where** and **orderby** clauses. These constructs correspond to what is written as a *list comprehension* in languages such as Haskell or Python. Each comprehension is equivalent to an ordinary expression using standard library functions, we give three examples, where the comprehension is on the left and its translation is on the right.

```
for (var x <- e1)          concatMap(fun(x) {e2}, e1)
  e2

for (var x <- e1)          concatMap(
  where (e2)              fun(x) {if (e2) e3 else []},
  e3                      e1)

for (var x <- e1)          concatMap(
  orderby (e2)           fun(x) {e3},
  e3                     orderBy(fun(x) {e2}, e1))
```

Here *concatMap*(*f*, *xs*) applies function *f* to each element in list *xs* and concatenates the results, and *orderBy*(*f*, *xs*) sorts the list *xs* so that it is in ascending order after *f* is applied to each element. The **orderby** notation is conceptually easier for the user (since there is no need to repeat the bound variables) and technically easier for the compiler (because it is closer to the SQL notation that we compile into, as discussed in the next section; indeed, currently **orderby** clauses only work reliably in code that compiles into SQL, because the absence of overloading means we have not yet implemented comparison on arbitrary types in the client or server).

*Database.* Links provides facilities to query and update SQL databases, where database tables are viewed as lists of records. We hope to provide similar facilities for XQuery databases in the future.

Links provides a **database** expression that denotes a connection to a database, and a **table** expression that denotes a table within a database. A database is specified by name (and optional configuration data, which is usually read from a configuration file), and a table is specified by the table name, the type signature of a row in the table, and the database containing the table.

The type of a table is distinct from the type of its list of records, since tables (unlike lists) may be updated. The coercion operation *asList* takes a table into the corresponding list, and **for** (**var** *x* ← *e1*) *e2* (with a long arrow) is equivalent to **for** (**var** *x* ← *asList* (*e1*)) *e2* (with an ordinary arrow).

In the following example, there is a table of words, where each row contains three string fields: the word itself, its type (noun, verb, and so on), and its meaning (definition). Typically, one expresses queries using the Links constructs such as **for**, **where**, and **orderby**, and functions on lists such as *take* and *drop*. The Links compiler is designed to convert these into appropriate SQL queries over the relevant tables whenever possible. For example, the expression

```
take(10, for (var def <-- defsTable)
      where (def.word ~ /s.*/)) orderby (def.word)
[def])
```

compiles into the equivalent SQL statement

```
SELECT def.meaning AS meaning, def.word AS word
FROM definitions AS def
WHERE def.word LIKE '{s}%'
ORDER BY def.word ASC
LIMIT 10 OFFSET 0
```

This form of compilation was pioneered in systems such as Kleisli [?,?], and is also central to Microsoft's new Language Integrated Query (LINQ) for .NET [?]. A related approach, based on abstract interpretation rather than program transformation, is described by Wiederman and Cook [?].

Note that the match operator `~` on regular expressions in Links is in this case compiled into the `LIKE` operator of SQL, and that the call to *take* in Links is compiled into `LIMIT` and `OFFSET` clauses in SQL. At run time, the phrase `{s}` in the SQL is replaced by the string contained in the variable *s*, including escaping of special characters where necessary. Links also contains statements to update, delete from, and insert into a database table, closely modeled on the same statements in SQL.

The `LIMIT` and `OFFSET` clauses are not part of the SQL standard, but are available in PostgreSQL, MySQL, and SQLite, the three targets currently supported by Links. For non-standard features such as this, Links can generate different code for different targets; for instance, it generates different variants of `INSERT` for PostgreSQL and MySQL.

Section ?? presents a subset of Links that is guaranteed to compile into SQL queries.

*Concurrency.* Links allows one to spawn new processes. The expression `spawn {e}` creates a new process and returns a fresh process identifier, the primitive `self()` returns the identifier of the current process, the command `e1 ! e2` sends message *e2* to the process identified by *e1*, and the expression

```
receive {
  case p1 -> e1
  ...
  case pn -> en
}
```

extracts the next message sent to the current process (or waits if there is no message), and executes the first case where the pattern matches the message. (Unlike Erlang, it is an error if no case matches, which fits better with our discipline of static typing.)

By convention, the messages sent to a process belong to a variant type. Event handlers are expected to execute quickly, so we follow a convention that each handler either sends a message or spawns a fresh process. For instance, in the Dictionary Suggest

application, each event handler sends a message to a separate handler process that is responsible for finding and displaying the completions of the current prefix. This puts any delay resulting from the database lookup into a separate process, so any keystroke will be echoed immediately. Furthermore, the messages received by the handler process are processed sequentially, ensuring that the updates to the DOM happen in the correct order.

As well as any new processes spawned by the user, there is one distinguished process: the *main* process. The top-level program and all event handlers are run in the main process. Having all event handlers run in a single process allows us to guarantee that events are processed in the order in which they are received.

*Client-server.* The keyword **server** in the definition of *completions* causes invocations of that function to be evaluated on the server rather than the client, which is required because the database cannot be accessed directly from the client.

When the Dictionary Suggest application is invoked, the server does nothing except to transmit the Links code (compiled into JavaScript) to the client. If the user presses and releases a key then the *suggest* function will continue to run on the client as its definition is annotated with the keyword **client**. The client runs autonomously until *completions* is called, at which point the *XMLHttpRequest* primitive of JavaScript is used to transmit the arguments to the server and creates a new process on the server to evaluate the call. No server resources are required until *completions* is called. This contrasts with Java Servlets, which keep a process running on the server at all times, or with PLT Scheme, which keeps a continuation cached on the server.

Location annotations **server** and **client** are only allowed on top-level function definitions. If a top-level function definition does not have a location annotation then it will be compiled for both the client and the server. A location annotation should be read as “this function *must* be run in the specified location”. The implementation strategy for client-server communication is discussed further in Section ??.

## 2.2 Draggable list

The draggable list application demonstrates the use of the concurrency primitives in Links to manage the state of an interactive GUI component. It displays an itemized list on the screen. The user may click on any item in the list and drag it to another location in the list (see Figure ??).

The code for the draggable list is shown in Figure ?. The example exploits Links’ ability to run concurrent processes in the client. Each draggable list is monitored by a separate process.

For each significant event on an item in the list (mouse up, mouse down, mouse out) there is a bit of code that sends a message to the right process, indicating the event. The process itself is coded as two mutually recursive functions, corresponding to two states. The process starts in the waiting state; when the mouse is depressed it changes to the dragging state; and when the mouse is released it reverts to the waiting state. When the mouse is moved over an adjacent item while in the dragging state, the dragged item and the adjacent item are swapped.




Fig. 6. Draggable list: before and after dragging

Both functions take, as parameter, the DOM `id` of the draggable-list element, thus they know what part of the DOM they control. The `dragging` function takes an additional parameter, designating the particular item that is being dragged. Both functions are written in tail recursive style, each one calling either itself (to remain in that state) or the other (to change state). This style of coding state for a process was taken from Erlang.

In this simple application, the state of the list can be recovered just by examining the text items in the list. In a more sophisticated application, one might wish to add another parameter representing the abstract contents of the list, which might be distinct from the items that appear in the display.

Observe that the code allows multiple lists to coexist independently, each monitored by its own process.

### 2.3 Progress bar

The progress bar application demonstrates symmetric client and server calls. (Note that unlike standard Ajax frameworks, in Links client code may call server code and vice-versa, equally easily.) Here some computation is performed on the server, and the progress of this computation is demonstrated with a progress bar (see Figure ??). When the computation is completed, the answer is displayed (see Figure ??).

The code for the progress bar application is shown in Figure ?. In this case, the computation performed on the server is uninteresting, simply counting up to the number typed by the user. In a real application the actual computation chosen might be more interesting.

Periodically, the function `computation` (running on the server) invokes the function `showProgress` (running on the client) to update the display to indicate progress. When this happens, the state of the computation on the server is pickled and transmitted to the client, exploiting continuation-passing style. The client is passed just the name of a function to call, its arguments, and an object representing the server-side continuation to be invoked when the client call is finished. Note that no state whatsoever is maintained on the server when computation is moved to the client. The implementation is discussed in more detail in Section ??.

```

fun waiting(id) {
  receive {
    case MouseDown(elem) ->
      if (isElementNode(elem)
          && (parentNode(elem) == getNodeById(id)))
        dragging(id,elem)
      else waiting(id)
    case MouseUp -> waiting(id)
    case MouseOut(newElem) -> waiting(id)
  }
}

fun dragging(id,elem) {
  receive {
    case MouseUp -> waiting(id)
    case MouseDown(elem) ->
      if (isElementNode(elem)
          && (parentNode(elem) == getNodeById(id)))
        dragging(id,elem)
    case MouseOut(toElem) ->
      if (isElementNode(toElem)
          && (parentNode(elem) == getNodeById(id)))
        {swapNodes(elem,toElem); dragging(id,elem)}
      else dragging(id,elem)
  }
}

fun draggableList (id,items) client {
  var dragger = spawn { waiting(id) };
  <ul id="{id}"
    l:onmouseup="{dragger!MouseUp}"
    l:onmouseuppage="{dragger!MouseUp}"
    l:onmousedown="{dragger!MouseDown(getTarget(event))}"
    l:onmouseout="{dragger!MouseOut(getToElement(event))}">
    { for (var item <- items) <li>{ item }</li> }
  </ul>
}


<html><body>
<h1>Draggable lists</h1>
<h2>Great Bears</h2>
{
  draggableList("bears",
    ["Pooh", "Paddington", "Rupert", "Edward"])
}
</body></html>

```

Fig. 7. Draggable lists in Links



**Fig. 8.** Progress bar



**Fig. 9.** Progress bar displaying the final answer

One advantage of this design is that if the client quits at some point during the computation (either by surfing to another page, terminating the browser, or taking a hammer to the hardware), then no further work will be required of the server.

On the other hand, the middle of an intensive computation may not be the best time to pickle the computation and ship it elsewhere. A more sophisticated design would asynchronously notify the client while continuing to run on the server, terminating the computation if the client does not respond to such pings after a reasonable interval. This is not possible currently, because the design of Links deliberately limits the ways in which the server can communicate with the client, in order to guarantee that long-term session state is maintained on the client rather than the server. This is not appropriate in all circumstances, and future work will need to consider a more general design.

### 3 Client-server computing

A Links program can be seen as a distributed program that executes across two locations: a client (browser) and a server. The programmer can optionally annotate a function definition to indicate where it should run. As mentioned before, code on the client can invoke server code, and vice-versa.

This symmetry is implemented on top of the asymmetric mechanisms offered by web browsers and servers. Current standards only permit the browser to make a direct request to the server. The server can return a value to the browser when done, but there is no provision for the server to invoke a function on the browser directly—so implementing our symmetric calls requires some craft.

Our implementation is also *scalable*, in the sense that session state, when it is captured automatically, is preserved in the client, thus requiring no server resources except when the server is actively working. This is significant since server resources are at a premium in the web environment.

We achieve this by using a variation of continuation-passing style, which we call *resumption-passing style*. It is now common on the web to use continuations to “invert

```

fun compute(count, total) server {
  if (count < total) {
    showProgress(count, total);
    compute(count+1, total)
  } else "done counting to " ++ intToString(total)
}

fun showProgress(count, total) client {
  var percent =
    100.0 *. intToFloat(count) /. intToFloat(total);
  replaceNode(
    <div id="bar"
      style="width:floatToString(percent)%;
        background-color: black">|</div>,
    getNodeById("bar")
  )
}

fun showAnswer(answer) client {
  domReplaceNode(
    <div id="bar">{stringToXml(answer)}</div>
    getNodeById("bar")
  );
}

<html>
  <body id="body">
    <form l:onsubmit=
      "{showAnswer(compute(0, stringToInt(n)))}">
      <input type="text" l:name="n"/>
      <input type="submit"/>
    </form>
    <div id="bar"/>
  </body>
</html>

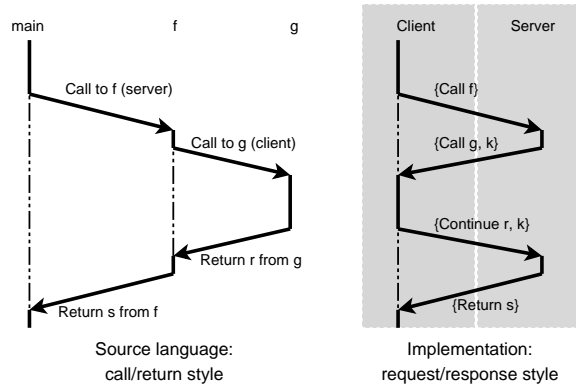
```

**Fig. 10.** Progress bar in Links

back the inversion of control” [?], permitting a server process to retain control after sending a *form* to the client. But Links permits a server to retain control after invoking an *arbitrary function* on the client.

Figure ?? shows how the call/return style of programming offered by Links differs from the standard request/response style, and how we use request/response style to emulate the call/return style. The left-hand diagram shows a sequence of calls between functions annotated with “client” and “server.” The solid line indicates the active thread of control as it descends into these calls, while the dashed line indicates a stack frame which is waiting for a function to return. In this example, `main` is a client function





**Fig. 11.** Semantic behaviour of client/server annotations

which calls a server function  $f$  which in turn calls a client function  $g$ . The semantics of this call and return pattern are familiar to every programmer.

The right-hand diagram shows the same series of calls as they actually occur at runtime in our implementation. The dashed line here indicates that some local storage is being used as part of this computation. During the time when  $g$  has been invoked but has not yet returned a value, the server stores nothing locally, even though the language provides an illusion that  $f$  is “still waiting” on the server. All of the server’s state is encapsulated in the value  $k$ , which it passed to the client with its call.

To accomplish this, the program is compiled to two targets, one for each location, and begins running at the client. In this compilation step, server-annotated code is replaced on the client by a remote procedure call to the server. This RPC call makes an HTTP request indicating the server function and its arguments. The client-side thread is effectively suspended while the server function executes.

Likewise, on the server side, a client function is replaced by a stub. This time, however, the stub will not make a direct call to the client, since in general web browsers are not addressable by outside agents. However, since the server is always working on behalf of a client request, we have a channel on which to communicate. So the stub simply gives an HTTP response indicating the server-to-client call, along with a representation of the server’s continuation, to be resumed when the server-to-client call is complete. Upon returning in this way, all of the state of the Links program is present at the client, so the server need not store anything more.

When the client has completed the server-to-client call, it initiates a new request to the server, passing the result and the server continuation. The server resumes its computation by applying the continuation.

### 3.1 Client-side Concurrency

A Links program is concurrent: many threads can run simultaneously within the client or within the server, communicating only through message queues.<sup>1</sup> We implement concurrency in the client, even though JavaScript does not offer any concurrency primitives. To accomplish this, we use the following techniques:

- compiling to continuation-passing style,
- inserting explicit context-switch instructions,
- placing any server calls asynchronously (with *XMLHttpRequest*),
- eliminating the stack between context switches using either *setTimeout* or a trampoline.

Producing JavaScript code in CPS allows us to reify each process’s control state, but since JavaScript does not implement tail-call elimination, there is a risk of overflowing the stack.

To manage the stack and to schedule threads, the compiler wraps every function application and every continuation application in calls to a special “yield” function which does the dirty work. Most calls to *\_yield* are simply no-ops, but after every *\_yieldGranularity* calls, it will collapse the JavaScript stack and may schedule a new thread to run.

A quick aside about the concurrency semantics: Normally, threads can pre-empt one another at any time. However, inside an event handler, we disallow pre-emption—this allows events to be handled in the order they arrive, for otherwise one event handler might be pre-empted by a later event even before the earlier handler noted the event’s occurrence! In order to prevent event handlers from blocking concurrency, we recommend that they be short-lived (ideally they just send a message or spawn a new process).

With this in mind, we use two techniques for managing the stack, one of which also allows a waiting thread to run; we call these the “timeout technique” and the “exception trampoline.” The latter does not transfer control to other threads and is used during synchronous execution, as with event handlers.

To see these techniques, refer to the implementation of *\_yield* for function application<sup>2</sup> in Figure ?? . In detail, the two techniques work as follows:

- In the timeout technique, the thread relinquishes control by returning completely to the top level (that is, to the browser), but only after calling the built-in *setTimeout* function, which adds a given function to a pool of thunks, each of which will be invoked no sooner than the given interval. Since the JavaScript code is all in continuation-passing style, returning from any function returns from the whole call stack, thus emptying it—but the timeout thunk has captured the continuation of that thread, so nothing is lost. The thunk we use simply manages the “process id” global and then calls  $f(a, k)$  to continue the thread.

<sup>1</sup> Note that presently there is no message-based communication between client and server

<sup>2</sup> The “yield” function for continuation application is the same except it takes as arguments continuation  $k$  and value  $a$ , and the function applications  $f(a, k)$  are replaced with continuation applications  $k(a)$ .

Besides giving other timeout thunks the chance to run, returning to the top level allows any event handlers or any *XMLHttpRequest* callbacks to run.

The *\_sched\_pause* is a lower bound on how long to wait before allowing this callback to run. Ideally this value should be zero, so there would be no delay between the moment a thread yields and the moment it is eligible to run again. But for some browsers a *\_sched\_pause* of zero runs the callback immediately, rather than putting it in the pool, thus blocking other threads.

- In the “exception trampoline” method, an exception is thrown containing the current continuation. Throwing the exception unwinds the stack; the trampoline then catches the exception and invokes the continuation. The trampoline code is shown in Figure ??.

Unsurprisingly, preliminary tests indicate that the output of the CPS-translating compiler is slower than comparable direct-style code. The performance penalty appears to be significant but not disastrous: the total run-time of compute-intensive programs we measured was within an order of magnitude of equivalent hand-written JavaScript versions.

```
function _yield(f, a, k) {
  ++_yieldCount;
  if ((_yieldCount % _yieldGranularity) == 0) {
    if (!_handlingEvent) {
      var current_pid = _current_pid;
      setTimeout((function() {
        _current_pid = current_pid;
        f(a, k)}),
        _sched_pause);
    }
    else {
      throw new _Continuation(function () { f(a,k) });
    }
  }
  else {
    return f(a,k);
  }
}
```

**Fig. 12.** The yield function

Client-to-server calls are implemented using the asynchronous mode of the function *XMLHttpRequest*, to which we pass a callback function that invokes the continuation of the calling process. Using the asynchronous mode allows threads to run via *setTimeout* or triggering user-event handlers.

```

function _Continuation(v) { this.v = v }

function _wrapEventHandler(handler) {
  return function(event) {
    var active_pid = _current_pid;
    _current_pid = _mainPid;
    _handlingEvent = true;
    var cont = function () { handler(event) }
    for (;;) {
      try {
        cont();
        _handlingEvent = false;
        _current_pid = active_pid;
        return;
      }
      catch (e) {
        if (e instanceof _Continuation) {
          cont = e.v;
          continue;
        }
        else {
          _handlingEvent = false;
          _current_pid = active_pid;
          throw e;
        }
      }
    }
  }
}

```

Fig. 13. The event handler trampoline

### 3.2 Related work: continuations for web programming

The idea of using continuations as a language-level technique to structure web programs has been discussed in the literature [?,?] and used in several web frameworks (such as Seaside, Borges, PLT Scheme, RIFE and Jetty) and applications, such as ViaWeb's e-commerce application [?], which became Yahoo! Stores. Most these take advantage of a call/cc primitive in the source language, although some implement a continuation-like object in other ways. Each of these systems creates URLs which correspond to continuation objects. The most common technique for establishing this correspondence is to store the continuation in memory, indexed by a unique identifier which is included as part of the URL.

Relying on in-memory tables makes such a system vulnerable to system failures and difficult to distribute across multiple machines. Whether stored in memory or in a database, the continuations can require a lot of storage. Since URLs can live long in bookmarks, emails, and other media, it is impossible to predict how long a continuation should be retained. Most of the above frameworks destroy a continuation after a set period of time. Even with a modest lifetime, the storage cost can be quite high, as each request may generate many continuations and there may be many simultaneous users.

Our approach differs from these implementations by following the approach described by the PLT Scheme team [?]. Rather than storing a continuation object at the server and referring to it by ID, we serialize continuations (or at least, their "data" portions), embedding them completely in URLs and hidden form variables. The resulting representation includes only the code-pointers and free-variable bindings that are needed in the future of the computation.

We have not yet addressed the issue of security. Security is vital. Our current implementation completely exposes the state of the application to the client. A better system needs to cryptographically encode any sensitive data when it is present on the client.

## 4 Query compilation

A subset of Links expressions compiles precisely to a subset of SQL expressions. Figure ?? gives the grammar of the SQL-compilable subset of Links expressions and Figure ?? gives the grammar for the subset of SQL to which we compile.

All terms are identified up to  $\alpha$ -conversion (that is, up to renaming of bound variables, field names and table aliases). We allow free variables in queries, though they are not strictly allowed in SQL: values for these variables will be supplied either at compile time, during query rewriting, or else at runtime. We use vector notation  $\vec{v}$  to mean  $v_1, \dots, v_k$ , where  $1 \leq k$  and abuse vector notation in the obvious way in order to denote tuples, record patterns and lists of tables. For uniformity, we write an empty 'from' clause as **from**  $\bullet$ . In standard SQL the clause would be omitted altogether. We assume a single database  $db$ , and write **table**  $t$  **with**  $(f_1, \dots, f_k)$  for  $asList(\mathbf{table} \ t \ \mathbf{where} \ (f_1:A_1, \dots, f_k:A_k) \ \mathbf{from} \ db)$ .

Any expression  $e$  can be transformed into an equivalent SQL query by repeated application of the rewrite rules given in Figure ?. To give the rewriting process sufficient freedom, we extend the non-terminal  $s$  from Figure ?? to give ourselves a working grammar.

(expressions)	$e ::= take(n, e) \mid drop(n, e) \mid s$
(simple expressions)	$s ::= \mathbf{for} (pat \leftarrow s) s$ $\quad \mid \mathbf{let} x = b \mathbf{in} s$ $\quad \mid \mathbf{where} (b) s$ $\quad \mid \mathbf{table} t \mathbf{with} \bar{f}$ $\quad \mid [(\bar{b})]$
(basic expressions)	$b ::= b_1 \mathit{op} b_2$ $\quad \mid \mathbf{not} b$ $\quad \mid x$ $\quad \mid lit$ $\quad \mid z.f$
(patterns)	$pat ::= z \mid (\bar{f}=\bar{x})$
(operators)	$op ::= \mathbf{like} \mid > \mid = \mid < \mid <> \mid \mathbf{and} \mid \mathbf{or}$
(literal values)	$lit ::= \mathbf{true} \mid \mathbf{false} \mid string\text{-}literal \mid n$
(finite integers)	$i, m, n$
(field names)	$f, g$
(variables)	$x, y$
(record variables)	$z$
(table names)	$t$

**Fig. 14.** Grammar of an SQL-compilable subset of Links.

(queries)	$q ::= \mathbf{select} Q \mathbf{limit} \mathit{ninf} \mathbf{offset} n$
(query bodies)	$Q ::= cols \mathbf{from} tables \mathbf{where} c$
(column lists)	$cols ::= \bar{c}$
(table lists)	$tables ::= \bar{t} \mathbf{as} \bar{a} \mid \bullet$
(SQL expressions)	$c, d ::= c \mathit{op} d$ $\quad \mid \mathbf{not} c$ $\quad \mid x$ $\quad \mid lit$ $\quad \mid a.f$
(integers)	$\mathit{ninf} ::= n \mid \infty$
(table aliases)	$a$

**Fig. 15.** The SQL grammar that results from query compilation.

```

TABLE
  table  $t$  with  $\bar{f}$ 
   $\rightarrow$  ( $a$  is fresh)
  select  $a.\bar{f}$  from  $t$  as  $a$  where true limit  $\infty$  offset 0
LET
  let  $x = b$  in  $s \rightarrow s[b/x]$ 
TUPLE
  [ $\bar{b}$ ]  $\rightarrow$  select  $\bar{b}$  from  $\bullet$  where true limit  $\infty$  offset 0
JOIN
  for ( $(\bar{f}=\bar{x}) <-$ 
    select  $\bar{c}_1$  from  $\bar{t}_1$  as  $\bar{a}_1$  where  $d_1$  limit  $\infty$  offset 0)
    select  $\bar{c}_2$  from  $\bar{t}_2$  as  $\bar{a}_2$  where  $d_2$  limit  $\infty$  offset 0
   $\rightarrow$  ( $\bar{a}_1$  and  $\bar{a}_2$  are disjoint)
  select  $\bar{c}_2[\bar{c}_1/\bar{x}]$  from  $\bar{t}_1$  as  $\bar{a}_1, \bar{t}_2$  as  $\bar{a}_2$ 
    where  $d_1$  and  $d_2[\bar{c}_1/\bar{x}]$  limit  $\infty$  offset 0
WHERE
  where ( $b$ ) (select  $\bar{c}$  from  $\bar{t}$  as  $\bar{a}$  where  $d$  limit  $m$  offset  $n$ )
   $\rightarrow$ 
  select  $\bar{c}$  from  $\bar{t}$  as  $\bar{a}$  where  $d$  and  $b$  limit  $m$  offset  $n$ 
RECORD
  for ( $z <-$  select  $\bar{c}$  from  $\bar{t}$  as  $\bar{a}$  where  $d$  limit  $m$  offset  $n$ )  $s$ 
   $\rightarrow$  ( $\bar{x}$  not free in  $s$ )
  for ( $(\bar{f}=\bar{x}) <-$  select  $\bar{c}$  from  $\bar{t}$  as  $\bar{a}$ 
    where  $d$  limit  $m$  offset  $n$ )  $s[\bar{x}/z.\bar{f}]$ 
TAKE
  take( $i, \text{select } Q \text{ limit } m \text{ offset } n$ )
   $\rightarrow$ 
  select  $Q$  limit  $\min(m, i)$  offset  $n$ 
DROP
  drop( $i, \text{select } Q \text{ limit } m \text{ offset } n$ )
   $\rightarrow$ 
  select  $Q$  limit  $m-i$  offset  $n+i$ 
DROP $_{\infty}$ 
  drop ( $i, \text{select } Q \text{ limit } \infty \text{ offset } n$ )
   $\rightarrow$ 
  select  $Q$  limit  $\infty$  offset  $n+i$ 

```

Fig. 16. Database rewrite rules

$$s ::= \dots \mid \varrho$$

We write  $s[b/x]$  for the substitution of  $b$  for  $x$  in  $s$ . Again we abuse the vector notation in the obvious way in order to denote simultaneous pointwise substitution. Note that  $min$  is a meta language (i.e. compiler) operation that returns the minimum of two integers, and is not part of the target language.

**Proposition 1.** *The database rewrite rules are strongly normalising and confluent.*

*Proof.* Define the size  $|e|$  of an expression  $e$  as follows.

$$\begin{aligned} |take(n, s)| &= 1 + |s| \\ |drop(n, s)| &= 1 + |s| \\ |\mathbf{for} (x \leftarrow s1) s2| &= 2 + |s1| + |s2| \\ |\mathbf{for} ((\bar{f} = \bar{x}) \leftarrow s1) s2| &= 1 + |s1| + |s2| \\ |\mathbf{let} x = b \mathbf{in} s| &= 1 + |s| \\ |\mathbf{where} (b) s| &= 1 + |s| \\ |\mathbf{table} t| &= 1 \\ |[ (b) ]| &= 1 \end{aligned}$$

Each rewrite rule strictly reduces the size of an expression, thus the rules are strongly normalising. The rewrite rules are orthogonal and hence weakly confluent. Confluence follows as a direct consequence of weak confluence and strong normalisation.

## 5 Statically typed message passing

Concurrency in Links is based on processes that send messages to mailboxes. Each process can read from a single mailbox which it owns. The message passing model is inspired by Erlang [?], but unlike in Erlang, mailboxes in Links are statically typed. To achieve this, we add a special mailbox type to the typing context, and annotate function types with the type of their mailbox.

We here present a tiny core calculus to capture the essence of process and message typing in Links. There is nothing particularly difficult here, but it does capture the essence of how our type inference works. The simple description of our core has already proven useful to researchers at MSR Cambridge working with Links.

We let  $A, B, C, D$  range over types,  $s, t, u$  range over terms, and  $x$  range over variables. A type is either the empty type  $0$ ; unit type  $1$ ; a process type  $P(A)$ , for a process that accepts messages of type  $A$ ; or a function type  $A \rightarrow^C B$ , for a function with argument of type  $A$ , result of type  $B$ , and that may accept messages of type  $C$  when evaluated. Links also supports record types and variant types (using row typing), but as those are standard we don't describe them here. Typically, a process will receive messages belonging to a variant type, and we use the empty variant type (which is isomorphic to the empty type  $0$ ) to assert that a function does not receive messages.

A typing judgement has the form  $\Gamma; C \vdash t : A$ , where  $\Gamma$  is a typing context pairing variables with types,  $C$  is the type of messages that may be accepted during execution of the term,  $t$  is the term, and  $A$  is the type of the term.



The typing rules are shown in Figure ???. The rules for variables and the constant of unit type are standard. The type of a lambda abstraction is labelled with the type  $C$  of messages that may be received by its body, where the context of the abstraction may accept messages of an unrelated type  $D$ ; while in an application, the function must receive messages of the same type  $C$  as the context in which it is applied.

If  $t$  is a term of unit type that receives messages of type  $C$ , then `spawn  $t$`  returns a fresh process identifier of type  $P(C)$ , where the context of the spawn may receive messages of an unrelated type  $D$ . The term `self` returns the process identifier of the current process, and so has type  $P(C)$  when it appears in a context that accepts messages of type  $C$ . If  $s$  is a term yielding a process identifier of type  $P(D)$  and  $t$  is a term of type  $D$ , then `send  $s$   $t$`  is a term of unit type that sends the message  $t$  to the process identified by  $s$ ; it evaluates in a context that may receive messages of an unrelated type  $C$ . Finally, the term `receive` returns a value of type  $C$  when it is executed in a context that may receive messages of type  $C$ .

Links syntax differs slightly from the core calculus. The type  $A \rightarrow^C B$  is written  $A -\{C\} \rightarrow B$ , or just  $A \rightarrow B$  when  $C$  is unspecified. The core expression `self` is written as a function call `self()`, the core expression `send  $t$   $u$`  is written  $t ! u$ , and the Links expression

```

receive {
  case  $p1 \rightarrow e1$ 
  ...
  case  $pn \rightarrow en$ 
}

```

corresponds to a switch on the value returned by the receive operator. (Recall that by convention the type of messages sent to a process, i.e. its mailbox type, is a variant type.)

We give a formal semantics for the typed message-passing calculus via a translation into  $\lambda(\text{fut})$ , Niehren et al’s concurrent extension of call-by-value  $\lambda$ -calculus [?]. For our purposes we need not consider all of the details of  $\lambda(\text{fut})$ . We just describe the image of the translation and then rely on previous results [?]. The image of the translation is simply-typed call-by-value  $\lambda$ -calculus extended with the constants of Figure ???.

The expression `thread ( $\lambda x.u$ )` spawns a new process binding its *future* to the variable  $x$  in  $u$  and also returning this future. A future can be thought of as placeholder for a value that is set to contain a real value once that value becomes known. Futures are central to  $\lambda(\text{fut})$ , but not to our translation.

The other constants are used for manipulating asynchronous *channels*, which can be used for implementing mailboxes. The constants `newChannel`, `put` and `get` are used for creating a channel, writing to a channel, and reading from a channel. They are not built in to  $\lambda(\text{fut})$ , but are easily implemented in terms of  $\lambda(\text{fut})$ -primitives, when  $\lambda(\text{fut})$  is extended with product types. The type of channels is  $\text{Chan}(A)$ <sup>3</sup>.

We write  $\Gamma \vdash^{\text{fut}} u : A$  for the  $\lambda(\text{fut})$  typing judgement “ $u$  has type  $A$  in context  $\Gamma$ .” The translation on terms and types is given in Figure ???. At the top level a fresh channel corresponding to the mailbox of the main process is introduced. Thereafter, the transla-

<sup>3</sup> In fact  $\text{Chan}(A)$  is just the type  $(A \rightarrow \text{unit}) \times (\text{unit} \rightarrow A)$  (i.e. a pair of put and get functions), but that is not relevant to our presentation.

$$\begin{array}{c}
\overline{\Gamma; x : A; C \vdash x : A} \quad \overline{\Gamma; C \vdash () : \mathbf{1}} \\
\frac{\Gamma; x : A; C \vdash u : B}{\Gamma; D \vdash \lambda x. u : A \rightarrow^C B} \quad \frac{\Gamma; C \vdash s : A \rightarrow^C B \quad \Gamma; C \vdash t : A}{\Gamma; C \vdash st : B} \\
\frac{\Gamma; C \vdash t : \mathbf{1}}{\Gamma; D \vdash \text{spawn } t : P(C)} \quad \overline{\Gamma; C \vdash \text{self} : P(C)} \\
\frac{\Gamma; C \vdash s : P(D) \quad \Gamma; C \vdash t : D}{\Gamma; C \vdash \text{send } s t : \mathbf{1}} \quad \overline{\Gamma; C \vdash \text{receive} : C}
\end{array}$$

**Fig. 17.** Statically typed message passing

$$\begin{array}{l}
\text{thread} : (A \rightarrow A) \rightarrow A \\
\text{newChannel} : \text{unit} \rightarrow \text{Chan}(A) \\
\text{put} : \text{Chan}(A) \rightarrow A \rightarrow \text{unit} \\
\text{get} : \text{Chan}(A) \rightarrow \text{unit} \rightarrow A
\end{array}$$

**Fig. 18.** Constants for the  $\lambda(\text{fut})$  translation

$$\begin{array}{l}
\llbracket u \rrbracket = \text{let } ch \text{ be newChannel } () \text{ in } \llbracket u \rrbracket^{ch} \\
\llbracket () \rrbracket^{ch} = () \quad \llbracket \text{spawn } u \rrbracket^{ch} = \text{let } ch' \text{ be newChannel } () \text{ in } \\
\quad \text{let } \_ \text{ be thread } (\lambda \_. \llbracket u \rrbracket^{ch'}) \text{ in } ch' \\
\llbracket x \rrbracket^{ch} = x \quad \llbracket \text{send } s t \rrbracket^{ch} = \text{put } \llbracket s \rrbracket^{ch} \llbracket t \rrbracket^{ch} \\
\llbracket \lambda x. u \rrbracket^{ch} = \lambda x. \lambda ch'. \llbracket u \rrbracket^{ch'} \quad \llbracket \text{receive} \rrbracket^{ch} = \text{get } ch \\
\llbracket st \rrbracket^{ch} = \llbracket s \rrbracket^{ch} \llbracket t \rrbracket^{ch} \quad \llbracket \text{self} \rrbracket^{ch} = ch \\
\llbracket \mathbf{1} \rrbracket = \mathbf{1} \quad \llbracket P(A) \rrbracket = \text{Chan}(\llbracket A \rrbracket) \quad \llbracket A \rightarrow^C B \rrbracket = \llbracket A \rrbracket \rightarrow \llbracket P(C) \rrbracket \rightarrow \llbracket B \rrbracket
\end{array}$$

**Fig. 19.** Translation into  $\lambda(\text{fut})$

tion on terms is annotated with the current channel. Unit and variables are just mapped to themselves. The current mailbox is threaded through functions and applications. The spawn, send and receive constructs are mapped to thread, put and get, and self is simply mapped to the channel that corresponds to the current mailbox. The type translation maps unit to unit, process types to channel types, and function types to function types with an extra channel argument.

It is straightforward to prove by induction on typing derivations that the transformation respects typing judgements.

**Proposition 2.** *The transformation  $\llbracket \cdot \rrbracket$  respects typing judgements.*

$$\Gamma; C \vdash u : A \quad \text{iff} \quad \llbracket \Gamma \rrbracket \vdash^{\text{fut}} \text{let } ch^{\llbracket C \rrbracket} \text{ be } \text{newChannel}() \text{ in } \llbracket u \rrbracket^{ch} : A$$

By Proposition ?? and standard results for  $\lambda(\text{fut})$  [?], we can be sure that programs that are well-typed in our message-passing calculus “do not go wrong”.

## 6 Issues

Links is still at an early stage of development. There are a number of shortcomings in the current design, which we plan to address. We mention some of them here.

*Form abstractions.* A flaw in the design of forms is that Links does not support abstraction over form components. As described in section ??, form interaction in Links is specified using the attributes `l:name`, which binds the value of a form field to a Links variable, and `l:onsubmit`, which contains Links code to be executed when a form is submitted. Variables bound by `l:name` are in scope within the `l:onsubmit` attribute of the lexically enclosing form. This design has several shortcomings. Firstly, it is not possible to abstract over input elements, since all input elements must be lexically contained within the form to which they belong. Secondly, it is not possible to vary the number of input elements within a form, since each input element is associated with a Links variable, and the number of variables is fixed at compile time. Thirdly, form components are not composable: it is not possible to compose input elements to create new form components which have the same interface as the language-supplied primitives, since `l:name` is only attachable to `input XML` literals.

To illustrate these problems, say we wish to construct a *date* component which we will use to allow the user to enter a number of dates into a form.

```
var date =
  <#>
  month: <input l:name="month"/>
  day:   <input l:name="day"/>
</#>
```

We would then like to use instances of *date* in a form elsewhere in the program:

```
<form l:onsubmit="{e}">
  Arrival: {date}
  Departure: {date}
</form>
```

Unfortunately, Links provides no way to refer within the `l:onsubmit` attribute to the values entered into the date fields, since the input elements are not lexically enclosed within the form. The only way to write the program is to textually include the date code within the form, then refer to the individual input elements of each month and day within *e*. This violates abstraction in two ways: the component must be defined (twice!) at the point of use, and the individual elements of the component cannot be concealed from the client.

If we try to construct a form with a number of fields determined at runtime then we run into further difficulties:

```
fun (n) {
  <form l:onsubmit="{e}">{
    for (var i <- range(n))
      Field {i}: <input l:name="x"/>
  }</form>
}
```

This code is also problematic: there is no way to refer to the different input fields generated in the loop, since all have the same name! Such constructs are consequently disallowed, and Links provides no way to construct forms with dynamically-varying numbers of fields.

The `l:name` mechanism provides a limited (but clean) way to refer to DOM nodes that are `<input>` elements. A related problem is that there is currently no other way of statically referring to a DOM node, which means we have to resort to calling the function `getNodeById` as illustrated in the draggable lists example.

Other systems also suffer from these problems to a greater or lesser degree. In PLT Scheme [?], composing new forms from old or making forms with a varying number of fields is possible, but not straightforward. In WASH [?] and iData [?], building forms that return new types requires declaring instances for those types in particular type classes, and in WASH a special type constructor needs to be used for forms with a varying number of fields. It was surprising for us to realize that composability of forms has received relatively little attention.

We are currently working on a design for light-weight modular form components, which uses a uniform framework to support form abstraction, varying numbers of form fields, and a mechanism that allows DOM nodes to be statically bound.

*Synchronous messages.* As far as we are aware, our mailbox typing scheme for Links is novel, though the translation into  $\lambda(\text{fut})$  shows that it can be simulated using typed concurrent languages with support for typed channels.

Links message passing is asynchronous, but often one wants synchronous message passing, where one sends a message and waits for a result. An example where synchronous message passing is needed is the extension of the draggable lists example to support reading the content of draggable lists.<sup>4</sup>

Although Links does not directly support synchronous message passing, it can be simulated in the same way as in Erlang. In order to send a synchronous message from process *P* to process *Q* the following sequence of events must occur:

---

<sup>4</sup> See <http://groups.inf.ed.ac.uk/links/examples/>.

- $P$  sends a message  $(M, P)$  to  $Q$ ,
- $P$  blocks waiting for a reply  $Reply(N)$ ,
- on receiving the message  $(M, P)$  the process  $Q$  sends a message  $Reply(N)$  to  $P$ ,
- $P$  receives the reply message  $Reply(N)$ .

However, Erlang is untyped; in Links, the simulation of synchronous messages has the unfortunate side-effect of polluting the type of process  $P$ 's mailbox, to include a different constructor for each distinct type of reply message associated with synchronous messages sent by  $P$ .

Type pollution of this kind makes it hard to structure programs cleanly. We are considering moving to a different model of concurrency that makes it easier to support typed synchronous message passing, such as the Join calculus [?], or even removing explicit concurrency from the language altogether and replacing it with something like functional reactive programming [?].

*Database abstraction.* As we showed in Section ??, we can guarantee to compile a useful fragment of Links into SQL. However, this fragment still lacks many important features. We do not yet deal with aggregate functions such as `count`, `sum`, and `average`, and we only support `order by` attached to a single loop iterator (rather than a nest of loop iterators). One can express the equivalent of a `group by` construct as a nested loop, but we do not yet translate code into `group by` when appropriate, nor yet support any other form of nested queries.

We also do not yet provide adequate support for abstracting over a query. For instance, suppose we try to abstract over the condition in the Dictionary Suggest example. We can amend the `completions` function to take a condition encoded as a function from unit to boolean in place of the prefix.

```

fun completions(cond) server {
  if (s == "") [] else {
    take(10, for (var def <-- defsTable)
              where (cond()) orderby (def.word)
              [def])
  }
}

```

Only once the condition is known is it possible to compile this to efficient SQL. The current compiler produces spectacularly inefficient SQL, which returns the entire dictionary, leaving the Links interpreter to then filter according to the condition and then take the first 10 elements of the filtered list. To produce adequately efficient code, it is necessary to inline calls to `completions` with the given predicate. But our compiler does not yet perform inlining, and in general inlining can be difficult in the presence of separate compilation.

Some of the problems outlined above are easy. For instance, we should be able to support aggregates in the same style that we support `take` and `drop`; and it is straightforward to extend our techniques to support `orderby` on nested loops. However, other problems are more challenging; in particular, it is not immediately clear how to extend the fragment described in Section ?? to support abstraction over predicates usefully.

Microsoft's Linq addresses all these issues, but does so by (a) requiring the user to write code that closely resembles SQL and (b) a subtle use of the type system to dis-

tinguish between a procedure that returns a value and a procedure that returns code that generates that value. It remains an open research problem to determine to what extent one can support efficient compilation into SQL without introducing separate syntax and semantics for queries as is done in Linq.

## 7 Conclusion

We have demonstrated that it is possible to design a single language in which one can program all three tiers of a web application. This eliminates the usual impedance mismatch and other problems connected with writing an application distributed across multiple languages and computers. We have shown that we can support Ajax-style interaction, programming applications such as Dictionary Suggest and draggable lists. The language compiles to JavaScript to run in the client, and SQL to run on the database.

Ultimately, we would like to grow a user community for Links. This depends on a number of factors, many of which are not well understood. One important factor seems to be the availability of good libraries. Good interfaces to other systems may help to rapidly develop useful functionality. A vexing question here is how to access facilities of languages that take a mostly imperative view of the world without vitiating the mostly functional approach taken by Links.

Already, researchers at Microsoft Cambridge and the University of Maryland have begun projects to add security features to Links; and one paper has been published that uses a modified version of Links to enhance security [?].

To make Links truly successful will require bringing together researchers and developers from many localities, many perspectives, and many communities. We would like to join our efforts with those of others — let us know if you are interested!

## References

1. D. L. Atkins, T. Ball, G. Bruns and K. C. Cox. Mawl: A domain-specific language for form-based services. *Software Engineering*, 25(3):334–346, 1999.
2. J. Armstrong. Concurrency oriented programming in Erlang. Invited talk, FFG 2003.
3. V. Balat. Ocsigen: typing web interaction with objective Caml *Proceedings of the 2006 workshop on ML*, Portland, Oregon, September 2006.
4. N. Benton, L. Cardelli, C. Fournet. Modern concurrency abstractions for C#. *TOPLAS*, 26(5), 2004.
5. N. Benton, A. Kennedy, and C. Russo. Adventures in interoperability: the SML.NET experience. *PPDP*, 2004.
6. G. Bierman, E. Meijer, W. Schulte. Programming with rectangles, triangles, and circles. *XML Conference*, 2003.
7. P. Buneman, S. Naqvi, V. Tannen, and L. Wong. Principles of programming with complex objects and collection types. *TCS*, 149(1), 1995.
8. R. Burstall, D. MacQueen, and D. Sannella. Hope: An experimental applicative language. *Lisp Conference*, 1980.
9. S. El-Ansary, D. Grolaux, P. Van Roy, M. Rafea. Overcoming the multiplicity of languages and technologies for web-based development. *Mozart/Oz Conference*, LNCS 3389, 2005.
10. C. Fournet, G. Gonthier. The Join Calculus: a language for distributed mobile programming. *Applied Semantics*, LNCS 2395, 2002.

11. J. Garret. Ajax: a new approach to web applications. 2005.
12. P. Graham. "Method for client-server communications through a minimal interface." United States Patent no. 6,205,469. March 20, 2001.
13. P. Graham. Beating the averages. 2001.
14. P. Graunke, R. B. Findler, S. Krishnamurthi, M. Felleisen. Automatically restructuring programs for the web. *ASE*, 2001.
15. P. Graunke, S. Krishnamurthi, S. van der Hoeven, M. Felleisen. Programming the web with high-level programming languages. *ESOP*, 2001.
16. V. Gapeyev, M. Levin, B. Pierce, A. Schmitt. The Xtatic experience. *PLAN-X*, 2005.
17. [labs.google.com/suggest](http://labs.google.com/suggest)
18. Microsoft Corporation. *DLinq: .NET Language Integrated Query for Relational Data* September 2005
19. A. Møller, M. Schwartzbach. The design space of type checkers for XML transformation languages. *ICDT, LNCS 3363*, 2005.
20. J. Niehren, J. Schwinghammer, G. Smolka. A Concurrent Lambda Calculus with Futures. *TCS*, 364(3), 2006.
21. G. Narra. ObjectGraph Dictionary.  
<http://www.objectgraph.com/dictionary/how.html>
22. M. Odersky *et al.* An overview of the Scala programming language. Technical report, EPFL Lausanne, 2004.
23. R. Plasmeijer and P. Achten. iData For The World Wide Web: Programming Interconnected Web Forms. *FLOPS*, 2006.
24. F. Pottier and D. Rémy. The essence of ML type inference. In B. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 10, pages 389–489. MIT Press, 2005.
25. C. Queinnec. Continuations to program web servers. *ICFP*, 2000.
26. Christian Queinnec *Inverting back the inversion of control or, continuations versus page-centric programming*, SIGPLAN Not., 2003
27. J. Reynolds. *Definitional interpreters for higher-order programming languages*. ACM '72: Proceedings of the ACM annual conference, 1972.
28. Ruby on Rails.  
<http://www.rubyonrails.org/>
29. M. Serrano, E. Gallesio, and F. Loitsch. HOP, a language for programming the Web 2.0. *Proceedings of the First Dynamic Languages Symposium*, Portland, Oregon, October 2006.
30. D. Syme. F# web page.  
[research.microsoft.com/projects/ilx/fsharp.aspx](http://research.microsoft.com/projects/ilx/fsharp.aspx)
31. P. Thiemann. WASH/CGI: server-side web scripting with sessions and typed, compositional forms. *PADL*, 2002.
32. J. Trevor, N. Swamy and M. Hicks. Defeating Script Injection Attacks with Browser-Enforced Embedded Policies. *World Wide Web*. May 2007
33. P. Van Roy. Convergence in language design: a case of lightning striking four times in the same place. *FLOPS*, 2006.
34. Z. Wan and P. Hudak. Functional reactive programming from first principles. In *PLDI '00: Proceedings of the ACM SIGPLAN 2000 conference on Programming language design and implementation*, Vancouver, British Columbia, Canada, 2000.
35. B. Wiederman and W. Cook. Extracting queries by static analysis of transparent persistence. *POPL*, 2007.
36. L. Wong. Kleisli, a functional query system. *JFP*, 10(1), 2000.
37. XML Query and XSL Working Groups. XQuery 1.0: An XML Query Language, W3C Working Draft, 2005.