

Rethinking Web interaction

Vincent Balat

Labratoire *Preuves, Programmes et Systèmes*,
CNRS – Université Paris Diderot vincent.balat@pps.jussieu.fr,
<http://www.pps.jussieu.fr/~balat>

Abstract. Web sites are evolving into ever more complex distributed applications. But current Web programming tools are not fully adapted to this evolution, and force programmers to worry about too many inessential details. We want to define an alternative programming style better fitted to that kind of applications. To do that, we propose an analysis of Web interaction in order to break it down into very elementary notions, based on semantic criteria instead of technological ones. This allows defining a common vernacular language to describe the concepts of current Web programming tools, but also some other new concepts. This results in a significant gain of expressiveness. The understanding and separation of these notions also makes it possible to get strong static guarantees, that can help a lot during the development of complex applications.

1 Introduction

Nowadays, Web sites behave more and more like real applications, with a high-level of interactivity on both the server and client sides. For this reason, they deserve well-designed programming tools, with features like high-level code structuring and static typing. These tools must take into account the specificities of that kind of application. One of these specificities is the division of the interface into *pages*, connected to each other by links. These pages are usually associated to *URLs* which one can bookmark. It is also possible to turn back to one page using the *back button*. This makes the dynamics of the interface completely different from a regular application. Another specificity is that this kind of applications is highly dependent on standards as they will be executed on various platforms.

Web programming covers a wide range of fields, from database to networking. The ambition of this paper is not to address them all, nor to deal with the full generality of *service oriented computing*. We concentrate on what we will call *Web interaction*; that is, the interaction between a user and a Web application, through a browser interface. We place ourselves in the context of writing such an application, that communicates with one or several servers, and with the ability to make part of the computation in the browser.

We want to make a first step towards a formalization of this interaction with a twofold goal. First we want to increase the expressiveness of Web frameworks.

Second, we want to improve the reliability of Web applications by using well defined concepts and static validation of the code.

The concepts we present here have been implemented in a complete Web programming framework for the Ocsigen Web server, called Eliom [23].

1.1 A common vernacular language for describing Web interaction

Web development is highly constrained by technologies. First it relies on the HTTP protocol, which is non-connected and stateless. Then, Web applications must be executable in various browsers that implement more or less accurately common standards and recommendations.

We want to detach ourselves as much as possible from these constraints and think about how we would like to program Web interaction. Obviously this is also a question of taste. But rather than proposing yet another programming model from scratch, we start by analyzing common Web programming practices, in order to understand the notions they use. Then we decompose them in very elementary notions that can be used to describe the features of most of Web programming tools, from PHP to JSP or Microsoft.NET Web Forms, etc. We will observe that current frameworks impose too many artificial restrictions. Ideally, we would like to give a generic language, flexible enough to describe all possible behaviours, without imposing any artificial restriction due to one technology.

We place ourselves at a semantic level rather than at a technical one. Moving away from technical details will allow to increase the *expressiveness* of Web programming frameworks. In the domain of programming languages, high-level concepts have been introduced over the years, for example genericity, inductive types, late binding, closures ... They make easier the implementation of some complex behaviours. We want to do the same for the Web. For example the notion of “*sending a cookie*” benefits from being abstracted to a more semantic notion like “*opening a session*” (which is already often the case today). Also it is not really important for the programmer to know how URLs are formed. What matters is the service we want to speak about (and optionally the parameters we want to send it).

This abstraction from technology allows two things:

- First it increases the expressiveness of the language by introducing specific concepts closer to the behaviours we want to describe (and irrespective of the way they are implemented). From a practical point of view, this allows to implement complex behaviours in very few lines of code.
- Having well-designed dedicated concepts also allows to avoid wrong behaviours. We forbid unsafe technical possibilities either by making them inexpressible, or by static checking.

1.2 Improving the reliability of Web applications

As Web sites are currently evolving very quickly into complex distributed applications, the use of strongly and statically typed programming languages for the

Web becomes more and more helpful. Using scripting languages was acceptable when there was very little dynamic behaviour in Web pages, but current Web sites written with such languages are proving to be very difficult to evolve and maintain. Some frameworks are counterbalancing their weaknesses by doing a lot of automatic code generation (for example [28]). But this does not really improve the *safety* of programs. In the current state of knowledge, we are able to do much better, and Web programming must benefit from this.

Static validation of pages One example where static typing revolutionizes Web programming concerns the validation of pages. Respecting W3C recommendations is a necessity to ensure portability and accessibility of Web sites. The novelty is that there now exist typing systems sophisticated enough to statically ensure a page's validity! [10, 3, 13] We do not mean checking the validity of pages once generated, but really to be sure that the program that builds your XML data will always generate something valid, even in the most particular cases.

For example, even if a programmer has checked all the pages of his site in a validator, is he sure that the HTML table he creates dynamically will never be empty (which is forbidden)? What if for some reason there is no data? He must be very conscientious to think about all these cases. And it is most likely that the evolutions of the program will break the validity of pages. In most cases, problems are discovered much later, by users ...

In lots of cases, such errors will even make the generated output unusable, for example for XML data intended to be processed automatically. And the best means to be sure that this situation will never happen is to use a typing system that will prevent you from putting the service on-line if there is the slightest risk for something wrong to be generated ...

For people not accustomed to such strong typing systems, this may seem to impose too much of a constraint to programmers. And indeed, it increases a bit the initial implementation time (by forcing you to take into account all cases). But it also saves such a huge amount of debugging time, that the use of such typing systems really deserves to be generalized. For now, these typing systems for XML are used in very few cases of Web services, and we are not aware of any Web programming framework (except Eliom, see section 6.1) that uses them. Our experience is that it is not difficult to use once you get used to the main rules of XHTML grammar, if you have clear enough error messages.

Validity of Web interaction Static checking and abstraction of concepts can also benefit in many other ways to Web programming, and especially to Web interaction. Here are a few examples:

- In a link, do the types (and names) of parameters match the types expected by the service it points to?
- Does a form match the service it points to?
- Do we have broken links?

It is not so difficult to have these guarantees, even if almost no Web programming framework are doing so now. All what is needed is a programming language *expressive* enough (in the sense we explained above).

Improving the ergonomics of Web sites Lots of Web developers are doing implementation errors resulting in reduced ease of use (wrong use of sessions or GET and POST parameters, etc.). Take as example a famous real estate Web site that allows to browse through the results of a search; but if you set a bookmark on one of the result pages, you never go back to the same page, because the URL does not refer to the advertisement itself, but to the rank in the search ... We will see that a good understanding of concepts can avoid such common errors.

1.3 Overview of the paper

Sections 2 and 3 are devoted to the definition of our vernacular language for describing the services provided by a Web application. Section 2 explains the advantage of using an abstract notion of service instead of traditional page-based programming and string URLs. Section 3 presents a new service identification and selection method. It shows how powerful this notion of service can be made, by separating it into several kinds. This results in a very new programming style for Web interaction.

2 Abstracting services

As explained above, we want to formalize Web interaction, that is, the behaviour of a Web application in reaction to the actions of the user. What happens when somebody clicks on a link or submits a form? A click often means that the user is requesting a new document: for example a new page that will replace the current one (or one part of it). But it can also cause some actions to take place on the server or the client. Let's try to enumerate the different kinds of reactions. A click (or a key strike) from the user may have the following main effects:

1. Modifying the application interface. That is, changing the page displayed by the browser (or one part of the page), or opening a new window or tab with a new page,
2. Changing the URL displayed by the browser (protocol, server name, path, parameters, etc.),
3. Doing some other action, like the modification of a state (for example changing some database values),
4. Sending hidden data (like form data, or files),
5. Getting some data to be saved on the user's hard disk.

Two important things to notice are that each of these items is optional, and may either involve a distant server, or be processed locally (by the browser).

This decomposition is important, as a formalization of Web interaction should not omit any of these items in order not to restrict the freedom of the programmer. Note that all these items are described semantically, not technically.

2.1 The role of URLs

The item “Changing the URL” above is a really significant one and is one key to understand the behaviour of Web applications. This section is devoted to the understanding of that notion. URLs are entry points to the Web site. Changing the URL semantically means: giving the possibility to the user to turn back to this point of interaction later, for example through bookmarks.

Note that, unlike many Web sites, a good practice is to keep the URL as readable as possible, because it is an information visible to users that may be typed manually.

Forgetting technical details about URLs The syntax of URLs is described by the Internet standard STD 66 and RFC 3986 and is summarized (a bit simplified) here:

`scheme://user:pwd@host:port/path?query#fragment`

The path traditionally describes a file in the tree structure of a file system. But this view is too restrictive. Actually, the path describes the hierarchical part of the URL. This is a way to divide a Web site into several sections and subsections.

The query string syntax is commonly organized as a sequence of ‘*key=value*’ pairs separated by a semicolon or an ampersand, e.g.

`key1=value1&key2=value2&key3=value3.`

This is the part of the URL that is not hierarchical.

To a first approximation, the path corresponds to the service to be executed, and the query to parameters for this service. But Web frameworks are sometimes taking a part of the path as parameters. On the contrary, part of the query, or even of the host, may be used to determine the service you want to call. This will be discussed later in more detail.

Note that the *fragment* part of the URL only concerns the browser and is not sent to the server.

The item “Changing the URL” is then to be decomposed semantically into these sub-tasks:

1. changing the protocol to use,
2. changing the server (and port) to which the request must be made,
3. choosing a hierarchical position (path) in the Web site structure, and specifying non hierarchical information (query) about the page,
4. and optionally: telling who you are (credentials) and the fragment of the page you want to display.

URL change and service calls There are two methods to send form data using a browser: either in the URL (GET method) or in the body of the HTTP request (POST method). Even if they are technical variants of the same concept (a function call), their semantics are very different with respect to Web interaction. Having parameters in the URL allows to turn back to the same document

later, whereas putting them in the request allows to send one-shot data to a service (for example because they will cause an action to occur).

We propose to focus on this semantical difference rather than on the way it is implemented. Instead of speaking about POST or GET parameters, we prefer the orthogonal notions of *service calls* and *URL change*. It is particularly important to forget the technical details if we want to keep the symmetry between server and client side services. Calling a local (javascript for example) function is similar to sending POST data to a server, if it does not explicitly change the URL displayed by the browser.

Semantically speaking, in modern Web programming tools, changing the URL has no relation with *calling a service*. It is possible to call a service without changing the URL (because it is a local service, or because the call uses POST parameters). On the contrary, changing the URL may be done without calling a service. You want to change the URL for only one reason: give the user a new entry point to the Web site, to which he can come back when he wants to ask the same service once again, for example by saving it in a bookmark.

2.2 Services as first class values

One of the main principles on which is based our work is: “*consider services as first class values*”, exactly as functional languages consider functions as first class values¹. That is: we want to manipulate services as abstract data (that can for example be given as parameter to a function). This has several advantages, among which:

- The programmer does not need to build the syntax of URLs himself. Thus it is really easy to switch between a local service and a distant one.
- All the information about the service is taken automatically from the data structure representing the service, including the path to which the service is attached and parameter names. This has a very great consequence: if you change the URL of one of your services, even the name of one of its parameters, you do not need to change any link or form towards this service, as they are all built automatically. This means that you’ll never get broken links or wrong parameter names (at least for internal services, i.e. services belonging to your Web site)!

Some recent frameworks already have an abstraction of the notion of service (for example [28, 15, 5]), but do not take the full benefit of it. Our notion of service must be powerful enough to take into account all the possibilities described above, but without relying on their technical implementation.

A service is some function taking parameters and returning some data, with possibly some side effects (remote function calls). The server is a provider of *services*. Client side function calls can also be seen as calls to certain services. The place where services take place is not so significant. This allows to consider

¹ Note that we are using the generic term *service* and not *Web service*, which has been given a precise meaning by many other works.

a Web site with two versions of some services, one on server side, the other on client side, depending on the availability of some resources (network connection, or browser plug-ins for example).

The language must provide some way to define these services, either using a specific keyword or just through a function call.

Once we have this notion, we can completely forget the old “page-based” view of the Web where one URL was supposed to be associated to one file on the hard disk. Thus it is possible to gain a lot of freedom in the organization and modularity of your code, and also, as we will see later, in the way services are associated to URLs. One of the goals of next section is precisely to discuss service identification and selection, that is, how services are chosen by the server from the hierarchical and non-hierarchical parts of the URL, and hidden parameters.

3 A zoology of services

3.1 Values returned by services

A first classification of services may be made according to the results they send. In almost all Web programming tools, services send HTML data, written as a string of characters. But as we’ve seen before, it is much more interesting to build the output as a tree to enable static type checking. To keep full generality, we will consider that a service constructs a result of any type, that is then sent, possibly after some kind of serialization, to the browser which requested it. It is important to give to the programmer the choice of the kind of service he wants.

A reflection on return types of services will provide once again a gain of expressiveness. Besides plain text or typed XHTML trees, a service may create for example a redirection. One can also consider using a service to send a file. It is also important to give the possibility to services to choose themselves what they want to send. For example, some service may send a file if it exists, or an HTML page with an error message on the other case.

We also introduce a new kind of output called *actions*. Basically sending an action means “no output at all”. But the service may perform some action as side effect, like modifying a database, or connecting a user (opening a new session). From a technical point of view, actions implemented server side are usually sending a 204 (No content) HTTP status code. Client side actions are just procedures. We will see some examples of use of actions and how to refine the concept in section 3.4.

3.2 Dynamic services, or continuation-based Web programming

Dynamic services Modern Web frameworks propose various solutions to get rid of the lack of flexibility induced by a one-to-one mapping between one URL and one service. But almost none of them take the full benefit of this, and especially of one very powerful consequence: the possibility to dynamically create new services!

For example, if you want to add a feature for your Web site, or even if you want occasionally to create a service depending on previous interaction with one user. For example, if one user wants to book a plane ticket to go from Auckland to Poznan, the system will look in a database for available planes and dynamically create the services corresponding to booking each of them. Then it displays the list of tickets, with, on each of them, a link towards one of these dynamic services. Thus you'll be sure that the user will book the ticket he expects, even if he duplicates his browser window or uses the back button. This behaviour is really simple to implement with dynamic services and rather tricky with traditional Web programming. Witness the huge number of Web sites which do not implement this correctly.

If you want to implement such behaviour without dynamic services, you'll need to save somewhere all the data the service depends on. One possibility is to put all this data in the link, as parameters, or in hidden form data. Another possibility, for example if the amount of data is prohibitive, is to save it on the server (for example in a database table) and send only the key in the link.

With dynamic service creation, all this contextual data is recorded automatically in the *environment* of the function implementing the service². This function is created dynamically according to some dynamic data (recording the past of the interaction with the user). It requires a functional language to be implemented easily.

Continuations This feature is equivalent to what is known as *continuation-based Web programming*. This technique was first described independently by Christian Queinnec, John Hughes and Paul Graham [25, 27, 26, 11, 14]. There are basically two ways to implement dynamic services. The first consists of viewing the sending of a Web page by the server as a function call, (a question asked of the user) that will return for example the values of a form or a link pressed. The problem is that you never know in advance to which question the user is answering (because he may have pressed the back button of his browser or he may have duplicated the page). Christian Queinnec solves this problem by using the Scheme control operator `call/cc` that allows to name the current point of execution of the program (called *continuation*) and to go back to this continuation when needed.

The second solution is the one we propose. It is symmetric to the first one, as it consists in viewing a click on a link or a form as a remote function call. Each link or form of the page corresponds to a continuation, and the user chooses the continuation he wants by clicking on the page. This corresponds to a *Continuation Passing programming Style* (CPS), and has the advantage that it no longer needs control operators (no saving of the stack is required). Strangely, this style of programming, usually considered unnatural, is closer to what we are used to doing in traditional Web programming.

² In functional programming, a function is represented by its source code (or a pointer to it), together with an *environment* containing the values used by the function.

The use of dynamic services is a huge step in the understanding of Web interaction, and an huge gain of expressiveness. Up until now, very few tools have used these ideas. The only commercial framework we are aware of that implements this is Seaside [9]. Continuation-based Web programming is also used for example in PLT Scheme [17], Hop [29] or Links [6], and obviously Eliom.

Implementation of dynamic services The implementation of dynamic services (in CPS) is usually done by registering closures in a table on the server. It associates to an automatically generated key a function and its environment, which contains all the data needed by the service. Note that the cost (in terms of memory or disk space consumption) is about the same as with usual Web programming: no copy of the stack, one instance of the data, plus one pointer to the code of the function. An alternative approach used for example in Links [6] is to serialize the closures and send them to the client. Either you serialize the environment and a pointer to the function, or even the environment and the full code of the function. This has the advantage that no space is required on the server to save the closures. But serializing functions is not easy and this solution may require sending large amounts of data to the client, with potentially several copies of some state information. And obviously security issues must be considered, as the server is executing code sent by the client.

3.3 Finding the right service

The very few experimental frameworks which are proposing some kind of dynamic services impose usually too much rigidity in the way they handle URLs. This section is devoted to showing how it is possible to define a notion of service identification that keeps all the possibilities described in section 2.

The important thing to take care of is: how to do the association between a request and a service? For example if the service is associated to an URL, where, in this URL, is the service to be called encoded?

To make this as powerful as possible, we propose to delegate to the server the task of decoding and verifying parameters, which is traditionally done by the service itself. This has the obvious advantage of reducing a lot the work of the service programmer. Another benefit is that the choice of the service to be called can depend on parameters.

Let us first speak about distant (server side) bookmarkable services, i.e. services called by sending a GET request to a server. We will speak later about client side services, and hidden services.

Hierarchical services One obvious way to associate a service to an URL is by looking at the path (or one part of it). We will call these services *hierarchical services*. These kinds of services are usually the main entry points of a Web site. They may take parameters, in the *query* part of the URL, or in the path.

One way to distinguish between several hierarchical services registered on the same path is to look at parameters. For example the first registered service whose expected parameters exactly match the URL will answer.

Coservices Most of the time one probably wants dynamic services to share their path with a hierarchical service, at least those which last for only a short time (result of a search for example). Also you may want two services to share the same hierarchical position on the Web site.

We will call *coservices* services that are not directly associated to a path, but to a special parameter. From a semantic point of view, the difference is that hierarchical services are the entry points of the site. They must last forever, whereas coservices may have a timeout, and you probably want to use the associated main service as fallback when the coservice has expired.

We will distinguish between *named coservices* and *anonymous coservices*, the difference being the value of the special parameter. Named coservices have a fixed parameter value (the name of the coservice), whereas this value is generated automatically for anonymous coservice.

Like all other services, coservices may take parameters, that will be added to the URL. There must be a way to distinguish between parameters for this coservice and parameters of the original service. This can be done by adding automatically a prefix to coservice parameters.

Attached and non-attached coservices Actually, we will also distinguish between coservices *attached* to a path and *non-attached* coservices. The key for finding an attached coservice is the path completed by a special parameter, whereas non-attached coservices are associated to a parameter, whatever the path in the URL! This feature is not so common and we will see in section 3.4 how powerful it is.

Distant hidden services A distant service is said to be *hidden* when it depends on POST data sent by the browser. If you come back later, for example after having made a bookmark, it will not answer again, but another service, not hidden, will take charge of the request. We will speak about *bookmarkable* services, for services that are not hidden.

Hidden services may induce an URL change. Actually, we can make exactly the same distinction as for bookmarkable services: there are hierarchical hidden services (attached to a path), hidden attached coservices (attached to a path, and a special POST parameter), and hidden non-attached coservices (called by a special POST parameter).

It is important to allow the creation of hidden hierarchical services or coservices only if there is a bookmarkable (co)service registered at the same path. This service will act as a fallback when the user comes back to the URL without POST parameters. This is done by specifying the fallback instead of the path when creating a hidden service. Note that it is a good idea to do the same for bookmarkable coservices.

Registering a hidden service on top of a bookmarkable service with parameters allows to have both GET and POST parameters for the same service. But bear in mind that their roles are very different.

Client side services Client side service calls have the same status as hidden service calls. In a framework that allows to program both the server and client sides using the same language, we would like to see local function calls as non-attached (hidden) coservices. Hierarchical hidden services and (hidden) attached coservices correspond to local functions that would change the URL, without making any request to the server.

Non-localized parameters One additional notion that is interesting in concrete cases is to enable parameters that are not related to any service at all. The server does not take into account their presence to choose the service, and services do not have to declare them, but can access them if they want. This avoids declaring the same optional parameters for each service when you want the whole Web site to be parametrized by the same optional parameters.

3.4 Taxonomy of services

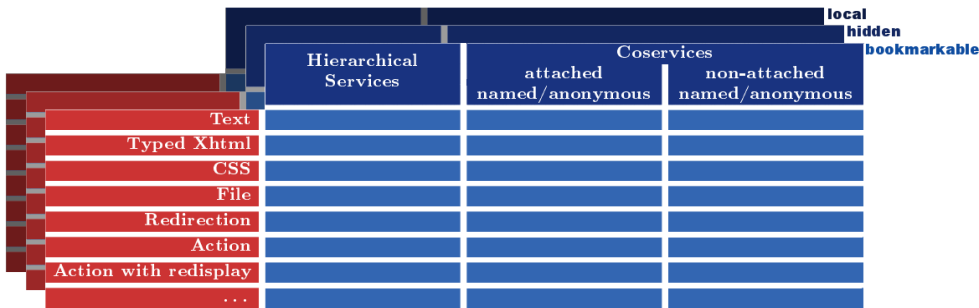


Fig. 1. Full taxonomy of services.

Summary of service kinds Figure 1 summarizes the full taxonomy of services we propose. Distant services can be registered in a public or session table. This set is obviously *complete* with respect to technical possibilities (as traditional services are part of the table). And it is powerful enough for describing in very few lines of code lots of features we want for Web sites, and does not induce any limitations with respect to the needs of Web developers. Note that current Web programming frameworks usually implement a small subset of these possibilities. For example “page-based” Web programming (like PHP or CGI scripts) does not allow for non-attached coservices at all. And even among “non-page-based” tools, very few allow for dynamic (anonymous) coservice creation. To our knowledge, none (but Eliom) is implementing actions on non-attached services as primary notions (even if all the notions can obviously be simulated).

Example cases We have already seen some examples of dynamic service creation: if a user creates a blog in a subdirectory of his or her personal site, one

possibility is to add dynamically a hierarchical service to the right path (and it must be recreated every time you relaunch the server). If you want to display the result of a search, for example plane ticket booking, you will create dynamically a new anonymous coservice (hidden or not), probably with a timeout. Without dynamic services, you would need to save manually the search keyword or the result list in a table.

Coservices are not always dynamic. Suppose you want a link towards the main page of your site, that will close the session. You will use a named hidden attached coservice (named, so that the coservice key is always the same).

We will now give an example where non-attached hidden coservices allow to reduce significantly the number of lines of code. Consider a site with several pages. Each page has a connected version and a non-connected version, and you want a connection box on each non-connected page. But you don't want the connection box to change the URL. You just want to log in and stay on the same URL, in connected version. Without non-attached services (and thus with almost all Web programming tools), you need to create a version with POST parameters of each of your hierarchical services to take into account the fact that each URL may be called with user credentials as POST parameters.

Using our set of services, you just need to define only one non-attached (hidden) coservice for the connection. At first sight, that service only performs an action (as defined in section 3.1): saving user information in a session table. But you probably want to return a new page (connected version of the same page). This can be done easily by returning a redirection to the same URL. Another solution if you don't want to pay the cost of a redirection, is to define a new kind of output: "*action with redisplay*" that will perform the action, then make an internal (server side) request as if the browser had done the redirection. The solution with redirection has one advantage: the browser won't try to resend POST data if the user reloads the page.

Now say for example that you want to implement a wiki, where each editable box may occur on several pages. Clicking on the *edit* button goes to a page with an edit form, and submitting the form must turn back to the original page. One dirty solution would be to send the original URL as hidden parameter in the edit link. But there is now a simpler solution: just do not change the path! The edit form is just a page registered on a non-attached service . . .

Our reflexion on services, also allows to express clearly a solution to the real estate site described in section 1.2. Use for example one bookmarkable hierarchical service for displaying one piece of advertisement, with additional (hidden or not) parameters to recall the information about the search.

Expressiveness The understanding of these notions and their division into very elementary ones induces a significant gain in expressiveness. This is particularly true for actions with redisplay. They are very particular service return values and seem to be closely related to non-attached coservices at first sight. But the separation of these two concepts introduces a new symmetry to the table, with new cells corresponding to very useful possibilities (see the example of the wiki

above). It is noteworthy that all cells introduced in this table have shown to be useful in concrete cases.

Priority rules One may wonder what succeeds when one gets apparently conflicting information from the request, for example both a hidden coservice name and a URL coservice name. In most cases, the solution is obvious. First, hidden services always have priority over URL services.

Secondly, as non-attached coservices parameters are added to the URL, there may be both attached and non-attached coservice parameters in the URL. In that case, the non-attached service must be applied (because its parameters have necessarily been added later).

The last ambiguous case is when you want to use a hidden non-attached coservice when you already have non-attached parameters in the URL: do you want to keep them or not? Either behaviour is possible and the service must declare which is the right one.

4 Typing Web interaction

4.1 Typing service parameters

Another advantage to our way of building Web sites is that it becomes possible to perform more static checking on the site and thus avoid lots of mistakes. We've already seen that some very strong guarantees (no broken links) are ensured just by the abstraction of the notion of services and links. But static types can help us to make things even safer.

When defining a service, it is easy to add type information to each parameter. This has three advantages:

- The server will be able to dynamically convert the data it receives into the type expected by the service (and also check that the data corresponds to what is expected, which saves a lot of time while writing the service).
- It is possible to check statically the types of parameters given to the links you create towards your services.
- It is also possible to do some static verification of forms (see next section).

It should also be possible to use type inference to avoid declaring manually the types of parameters.

On first sight, the type system for services parameters seems rather poor. But declaring basic types like string, int and float is not enough. We need more complex types. Just think about variable length forms, for example a page displaying a list of people with just a checkbox for each of them. The implementation of this is a tedious work with traditional Web programming tools, because you need to give different names to each parameter (for example by using numbers) to be able to find back on server side the person associated to the checkbox. Obviously you don't want to worry about such details when programming a Web

application. Your service just wants to get an association table from a name to a boolean value, however they are encoded in parameters.

Another example is when you want to send an (unordered) set of values to a service, or optional values.

Unfortunately, most of this is sometimes difficult to implement, due to the way parameters are encoded in the URL, and the way browsers send them. One example of this is the way browsers handle unchecked boxes: they send no parameter at all, instead of sending a *false* value. Thus there is no way to distinguish between an unchecked box and no parameter at all.

Sets of base types data may be implemented using the same parameter name for each member of the set. But an implementation of sets of more complex data requires more thinking.

In conclusion, the types of parameters for Web pages is highly constrained by current technology.

4.2 Forms

To be correct relative to the service it leads to, a form must respect these conditions:

- the names of the fields of the form must correspond to the names expected by the service,
- the types of the fields must correspond to the types expected by the service,
- all required fields must be present, and the right number of times.

The first item may be solved by taking the parameter names from the data representing the service in memory. The static verification of the adequacy of types may be done by putting type information in the type of names (instead of using just the *string* type, we use an abstract type with phantom type information [19] on the type of the parameter). The third condition is more difficult to encode in types.

One idea interesting for some particular cases of services is to create dynamically the service that will answer to the form from the form itself (as for example in Seaside).

5 Abstracting sessions

Sessions are a way to maintain a state during several interactions with one user. The abstraction of sessions is something more common in usual Web programming tools than the abstraction of services. It is now standard to have a way to save some data for one user and recover it each time the user comes back.

5.1 Session data and session services

Opening a session on server side may be done transparently when the Web site decides to register some data for one user. A session identifier is then generated

automatically and a cookie is set, and sent back by the browser in the header of each request for the site.

This is a common feature in current Web programming tools. But we propose to add the ability to dynamically register coservices or services in a session table, that is, for one user. One obvious reason for this is to allow the use of certain anonymous coservices only for the user who requested them. You probably do not want to give access to these kind of coservices to everybody, especially because they may contain private data in the closure.

We also propose to allow to re-register some existing services in a session table. When the server is receiving a request, it first looks in the session table to see if there is a private version of the service, and if not, looks in the public table. Using this feature, it is possible to save all the session data in the closure of services. From a theoretical point of view, this makes session data tables useless. Basically when a user logs in you just need to register in the session service table a new version of each service, specialized for this user.

For the sake of flexibility, it is a good idea to allow both session data and session services.

Thus, there is now a new possibility for creating each of our service kinds: registering them in the session table. This corresponds to a fourth dimension in the table of figure 1.

5.2 Session duration

When implementing such kinds of session, one must be very careful about the duration of sessions and timeout for services. Note that if you want your sessions to survive a restarting of the server, you need a way to serialize data (for session data) and closures (for session services).

5.3 Session names and session groups

To make the session system more powerful, we may want to add two more notions, namely *session names* and *session groups*.

Naming sessions allows to use several sessions in the same site. Think for example about one site that allows some features for non connected users, which requires some private coservices to be created. If the user logs in, this opens a new session, but must not close the previous one. To avoid that, we use another session by specifying another session name.

Technically speaking, session naming can be implemented by recording the name of the session in the cookie name.

The idea of session groups³ is complementary to that of the session name. While session naming allows for a single session to have multiple buckets of data associated with it, session grouping allows multiple sessions to be referenced together. For most uses, the session group is the user name. It allows to implement features like “close all sessions” for one user (even those opened on other

³ Session groups have been suggested by Dario Teixeira

browsers), or to limit the number of sessions one user may open at the same time (for security reasons).

6 Implementation

To implement the features presented above, we had two possible solutions: either we created our own language, which would have given us the full freedom in implementation. Or we needed to choose a language expressive enough to encode most of the features presented here.

We chose the second solution, for two reasons:

- We think that Web programming is not a task for a domain specific language. We need the full power of a general purpose language, because we do not want only to speak about Web interaction and typing of pages, but also for example about database interaction. We also want code structuring and separation, and we need programming environments and libraries. The cost of creating our own new language would have been much too high.
- We knew one language, namely Objective Caml (OCaml) that has almost all features we want to implement the concepts of this paper, most notably a powerful typing system able to encode most of the properties we wanted to check statically. Moreover, this language now has a strong basis of users, in academia and industry, and a large set of libraries.

We made this choice with the goal in mind to write not only a research prototype but a full framework usable for real applications. We benefited greatly from the free software development model, which enabled us to have a powerful framework very quickly, thanks to the growing community of users.

Our implementation takes the form of a module for the *Ocsigen* [23] Web server, called *Eliom* [2]. More implementation details may be found in [1] (but for an old version of Eliom that was using a more basic model of services).

6.1 Static type checking of pages

As the return value of services is completely independent of services, it is important to make the creation of new kinds of output types easy. This is realized through the use of OCaml's module language, which allows parametrized modules (called functors). To create a new output module, you apply a functor that will create the registration functions for your output type.

Eliom does not impose one way to write the output, but proposes several predefined modules. One allows text output, as in usual Web programming, if you don't want any type-checking of pages.

Another one is using an extension of OCaml, called OCamlduce [10], that adds XML types. This is really powerful, as the typing is very strict, and corresponds exactly to what you need to take into account all features of DTDs. It also has the advantage of allowing to easily create new output modules for other XML types, just from a DTD. Furthermore, it allows to parse and transform

easily incoming XML data. The drawback is that is not part of the standard OCaml compiler.

As an alternative, we are using a second typing method based on OCaml's *polymorphic variants* (see [20]) used as *phantom types* [19]. See [1] for more information about that technique. Both typing techniques have shown to be very helpful, and relieves the programmer from thinking about validation of her pages.

6.2 Defining services

Creation and registration of services Creating a new service means two things: first filling a data structure with all the information about the service (that will be used to create forms and links towards this service), and registering in a table the “handling” function implementing the service.

In Eliom, these two operations have been disconnected. You first create the service, then register a function on it. This may be considered dangerous, as some services may be created and not registered, which can lead to broken links. We were forced to do so because services are usually highly mutually recursive (every page may contain links towards any other one), and it is not easy in OCaml to have mutually recursive data in several different modules.

The solution we consider is to use a syntax extension to the language to put back together creation and registration. But the concrete realization of that idea is not straightforward.

Types of parameters One important thing to check statically when registering a function on a service is that the type of this function corresponds to the type expected by the service. This makes the type of the registration function dependent on the value of one of its parameters (the service). This requires very sophisticated type systems. As explained in [1], it can be done either using *functional unparsing* [8] or using *generalized algebraic data types* [24, 32].

Ideally, one would expect to declare page parameters only once when defining both the service and its handling function. As we separate definition and registration, it is not possible to do so. But even without this separation it is probably not possible to avoid this duplication without a syntax extension for the language, for two reasons:

- We need the names of parameters to live both in the world of variable names (static) and in the world of data (dynamic), to be used as parameter names for pages.
- We need dynamic types (that is keeping typing information during the execution) to enable the server to convert received page parameters into the type expected by the handling function.

But in any case, we must keep in mind that the types of service parameters are not OCaml types, and it is difficult to use sophisticated OCaml types for services as mentioned in section 4.1.

6.3 Links and forms

We are using functions to create links and forms. Obviously the adequacy of the service to its parameters is checked statically while creating a link.

Non-localized parameters have not been implemented for now, but it is not difficult to do: we just need a way to build a new service data structure by combining together a service and non-localized parameters.

We mentioned in section 4.2 that performing statically all the verification on the form would require a very sophisticated type system, that would be difficult to mix with the already complex typing model we use to check the validity of pages. We decided to restrict the static verifications to the names and types of fields. To do that, we use an abstract type for names, as explained in section 4.2, and we get parameters names from the service. This is done by giving as parameter to our form building function, a *function* that will build the content of the form from parameters names.

We implemented some sophisticated types for service parameters, like lists or sets. In the case of lists, the names of parameters are generated automatically by an iterator.

7 Conclusions

7.1 Related work

Most of the problems this papers addresses are well known and modern Web programming frameworks are already trying to propose solutions, but not in a so coherent and comprehensive way. They often provide some abstraction for some concepts, but most of them preserve some historical habits related to technical constraints. It is impossible to make a full review of such tools, as there are numerous. But one can try to make a classification of existing Web frameworks with respect to the way they do service identification.

The old method is what we called “page-based Web programming”, where one path corresponds to one file. Modern tools are all more flexible and make service identification and selection independent of the physical organization of components in the Web server (for example JSP assigns an URL to a service from a configuration file). But very few belong to the third group, that allows dynamic services. Among them: Seaside [9], Links [6] and Hop [29], Wash/CGI [30]. Their service models are more basic. Some of them are using an abstraction of forms [30, 7] that is fully compatible with our model.

There have been few attempts to formalize Web interaction. The most closely related work is by Paul T. Graunke, Robert Bruce Findler, Shriram Krishnamurthi and Matthias Felleisen [12, 16]. Their work is more formal but does not take into account all the practical cases we speak about. In particular their service model is much simpler and does not fully take into account the significance of URLs.

We think our approach is compatible with more data driven approaches [28], component based interfaces [33], or even code generation techniques. One interesting work would be to see how they can be mixed together.

7.2 Evolution of technologies and standards

This reflection about Web programming techniques has shown that Web technologies suffer from some limitations that slow down the evolution towards really dynamic applications. Here are a few examples:

- As mentioned above, the format of page parameters and the way browsers send them from form data does not allow for sophisticated parameters types. This may be solved by technologies like XForms [31].
- (X)HTML forms cannot mix GET and POST methods. It is possible to send URLs parameters in the `action` attribute of a form that is using the POST method, but it is not possible to take them from the form itself. This would open many new possibilities.
- A link from HTTP towards the same site in HTTPS is always absolute. This breaks the discipline we have to use only relative links (for example to behave correctly behind a reverse proxy).
- There is no means to send POST data through a link, and it is difficult to disguise a form into a link. Links and forms should probably be unified into one notion, that would allow to make a (POST or GET) request from a click on any part of the page.
- CSS is not powerful enough to allow a full separation between content, layout and computation.
- Having the ability to put several `id` attributes for one tag would be very useful for automatically generated dynamic pages.
- And probably one of the main barriers to the evolution of the Web today is the impossibility to run fast code on the browser (without plug-ins), even with recent implementations of Javascript. When thinking about a Web application as a complex application distributed between a server and a client, we would often like to perform computationally intensive parts of the execution on the client, which is not feasible for now.

7.3 Concluding words and future works

This paper presents a new programming style for Web interaction which simplifies a lot the programming work and reduces the possibilities of semantical errors and bad practices. The principles we advocate are summarized here:

1. Services as first class values
2. Decoding and verification of parameters done by the server
3. Dynamic creation of services
4. Full taxonomy of services for precise service identification
5. Same language on server and client sides
6. Symmetry between local and distant services

Beyond just presenting a new Web programming model, this paper defines a new vocabulary for describing the behaviour of Web sites, on a semantic basis. It is a first step towards a formalization of Web interaction. We started from an

analysis of existing Web sites and we extracted from this observation the underlying concepts, trying to move away as much as possible from non-essential technical details. This allowed a better understanding of the important notions but above all to bring to light some new concepts that were hidden by technical details or historical habits. The main feature that allowed this is the introduction of dynamic services, and also forgetting the traditional page-based Web programming. There exist very few frameworks with these features, and none is going as far as we do, especially in the management of URLs.

Besides the gain in expressiveness, we put the focus on reliability. This is made necessary by the growing complexity of Web applications. The concepts we propose allow for very strong static guarantees, like the absence of broken links. But more static checks can be done, for example the verification of adequacy of links and forms to the service they lead to. These static guarantees have not been developed here because of space limitation. They are summarized by the following additional principles:

7. Static type checking of generated data
8. Static type checking of links and forms

This paper does not present an abstract piece of work: all the concepts we present have been inspired by our experience in programming concrete Web sites, and have been implemented. Please refer to Eliom's manual [2] and source code for information about the implementation. Some implementation details may also be found in [1] (describing an old version of Eliom that was using a more basic model of services). Several content management systems are currently being written with Eliom (for example Ocsimor [23], Nurpawiki [22], Litiom [21] or Lambdium [18]). These concrete experiences showed that the programming style we propose is very convenient for Web programmers and reduces a lot the work to be done on Web interaction.

We are now working intensively on client-side programming [4] in order to allow to write both sides of the application using the same language, and with the same strong static guarantees. As we have seen, our notions of services also apply to client side functions. Obviously we want to be able to use the same typing system for services but also for HTML. It is not easy to guarantee that a page will remain valid if it can evolve over time. We are also addressing the problem of communication between the server and the client. On a more theoretical point of view, we are working on a formal description of our services.

8 Acknowledgments

Many acknowledgements are due to Jean-Vincent Loddó, Jérôme Vouillon and all the people who took part in Ocsigen and Eliom development. I also want to thank Russ Harmer, Boris Yakobowski and Yann Régis-Gianas for their helpful remarks about this paper.

References

1. V. Balat. Ocsigen: Typing web interaction with objective caml. In *ML'06: Proceedings of the 2006 ACM SIGPLAN workshop on ML*, 2006.
2. V. Balat. Eliom programmer's guide. Technical report, Technical report, Laboratoire PPS, CNRS, université Paris-Diderot, 2007.
3. V. Benzaken, G. Castagna, and A. Frisch. CDuce: An XML-centric general-purpose language. In *Proceedings of the International Conference on Functional Programming (ICFP)*, pages 51–63, 2003.
4. B. Canou, V. Balat, and E. Chailloux. O'browser: objective caml on browsers. In *ML '08: Proceedings of the 2008 ACM SIGPLAN workshop on ML*, pages 69–78, New York, NY, USA, 2008. ACM.
5. Cherry py. <http://www.cherrypy.org/>.
6. E. Cooper, S. Lindley, P. Wadler, and J. Yallop. Links: Web programming without tiers. In *In 5th International Symposium on Formal Methods for Components and Objects (FMCO)*, page 10. Springer-Verlag, 2006.
7. E. Cooper, S. Lindley, P. Wadler, and J. Yallop. The essence of form abstraction. In *Sixth Asian Symposium on Programming Languages and Systems*, 2008.
8. O. Danvy. Functional unparsing. *Journal of Functional Programming*, 8(6):621–625, 1998.
9. S. Ducasse, A. Lienhard, and L. Renggli. Seaside – a multiple control flow web application framework. In *Proceedings of ESUG Research Track 2004*, 2004.
10. A. Frisch. Ocaml + xduce. In *Proceedings of the international conference on Functional programming (ICFP)*, pages 192–200. ACM, 2006.
11. P. Graham. Beating the averages <http://www.paulgraham.com/avg.html>.
12. P. T. Graunke, R. B. Findler, S. Krishnamurthi, and M. Felleisen. Modeling web interactions. In *European Symposium on Programming (ESOP)*, April 2003.
13. H. Hosoya and B. C. Pierce. XDuce: A statically typed XML processing language. *ACM Transactions on Internet Technology*, 3(2):117–148, May 2003.
14. J. Hughes. Generalising monads to arrows. *Science of Computer Programming*, 37(1–3):67–111, 2000.
15. Javaserer pages technology. <http://java.sun.com/products/jsp/index.jsp>.
16. S. Krishnamurthi, R. B. Findler, P. Graunke, and M. Felleisen. Modeling web interactions and errors. In *In Interactive Computation: The New Paradigm*. Springer Verlag, 2006.
17. S. Krishnamurthi, P. W. Hopkins, J. McCarthy, P. T. Graunke, G. Pettyjohn, and M. Felleisen. Implementation and use of the plt scheme web server. In *Higher-Order and Symbolic Computation*, 2007.
18. Lambdium. <http://www.dse.nl/~dario/projects/lambdium-light/>.
19. D. Leijen and E. Meijer. Domain specific embedded compilers. In *Domain-Specific Languages*, pages 109–122, 1999.
20. X. Leroy, D. Doligez, J. Garrigue, D. Rémy, and J. Vouillon. The objective caml system release 3.10 documentation and user's manual. Technical report, Inria, may 2007.
21. Litiom. <http://litiom.forge.ocamlcore.org/>.
22. Nurpawiki. <http://code.google.com/p/nurpawiki/>.
23. The Ocsigen project. <http://www.ocsigen.org>.

24. F. Pottier and Y. Régis-Gianas. Stratified type inference for generalized algebraic data types. In *Proceedings of the 33rd ACM Symposium on Principles of Programming Languages (POPL'06)*, pages 232–244, Charleston, South Carolina, Jan. 2006.
25. C. Queinnec. The influence of browsers on evaluators or, continuations to program web servers. In *International conference on Functional programming (ICFP)*, pages 23–33, Montréal (Canada), Sept. 2000.
26. C. Queinnec. Inverting back the inversion of control or, continuations versus page-centric programming. *ACM SIGPLAN Notices*, 38(2):57–64, Feb. 2003.
27. C. Queinnec. Continuations and web servers. *Higher-Order and Symbolic Computation*, 17(4):277–295, Dec. 2004.
28. Ruby on rails. <http://www.rubyonrails.com/>.
29. M. Serrano, E. Gallesio, and F. Loitsch. Hop, a language for programming the web 2.0. In *Dynamic Languages Symposium*, Oct. 2006.
30. P. Thiemann. Wash/cgi: Server-side web scripting with sessions and typed, compositional forms. In *Practical Aspects of Declarative Languages (PADL'02)*, 2002.
31. W3C. Xforms 1.0 – W3C recommendation
<http://www.w3.org/tr/xforms/>.
32. H. Xi, C. Chen, and G. Chen. Guarded recursive datatype constructors. In *Proceedings of the 30th ACM SIGPLAN Symposium on Principles of Programming Languages*, pages 224–235, New Orleans, January 2003.
33. J. Yu, B. Benatallah, R. Saint-Paul, F. Casati, F. Daniel, and M. Matera. A framework for rapid integration of presentation components. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, New York, NY, USA, 2007. ACM.