

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

# *Modular and Flexible Causality Control on the Web*

Paul Leger — Éric Tanter — Rémi Douence

**N° 7742**

June 2011

Distributed Systems and Services




ISSN 0249-6399 ISRN INRIA/RR--7742--FR+ENG

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

centre de recherche
RENNES - BRETAGNE ATLANTIQUE

# Modular and Flexible Causality Control on the Web

Paul Leger , Éric Tanter * , Rémi Douence †

Theme : Distributed Systems and Services
Networks, Systems and Services, Distributed Computing
Équipes-Projets Ascola

Rapport de recherche n° 7742 — June 2011 — 26 pages

**Abstract:** Ajax has allowed JavaScript programmers to create interactive, collaborative, and user-centered Web applications, known as Web 2.0 Applications. These Web applications behave as distributed systems because processors are user machines that are used to send and receive messages between one another. Unsurprisingly, these applications have to address the same causality issues present in distributed systems like the need *a)* to control the causality between messages sent and responses received and *b)* to react to distributed causal relations. JavaScript programmers overcome these issues using rudimentary and alternative techniques that largely ignore the distributed computing theory. In addition, these techniques are not very flexible and need to intrusively modify these Web applications. In this paper, we study how causality issues affect these applications and present WeCa, a practical library that allows for modular and flexible control over these causality issues in Web applications. In contrast to current proposals, WeCa is based on (stateful) aspects, message ordering strategies, and vector clocks. We illustrate WeCa in action with several practical examples from the realm of Web applications. For instance, we analyze the flow of information in Web applications like Twitter using WeCa.

**Key-words:** WeCa; AspectScript; OTM; JavaScript; Web 2.0 applications; distributed computing; aspect-oriented programming.

* PLEIAD Laboratory, Computer Science Department (DCC), University of Chile, Chile, etanter@dcc.uchile.cl
† ASCOLA group, INRIA, Nantes, France, first.last@inria.fr

Centre de recherche INRIA Rennes – Bretagne Atlantique
IRISA, Campus universitaire de Beaulieu, 35042 Rennes Cedex
Téléphone : +33 2 99 84 71 00 — Télécopie : +33 2 99 84 71 71

# Modular and Flexible Causality Control on the Web

**Résumé :** Ajax a permis aux programmeurs JavaScript de créer des applications web interactives, collaboratives et centrées sur les utilisateurs, connues sous le nom d'application Web 2.0. Ces applications se comportent comme des systèmes distribués où les machines des utilisateurs qui s'échangent des messages jouent le rôle des processeurs. Sans surprise, ces applications doivent gérer les même problèmes de causalité que ceux présents dans les systèmes distribués comme a) la nécessité de contrôler la causalité entre les messages envoyés et les réponses reçues et b) de réagir aux relations causales distribuées. Les programmeurs JavaScript gérent ces problèmes à l'aide de techniques rudimentaires qui ignorent la théorie des systèmes distribués. De plus, ces techniques ne sont pas très flexibles et demandent de modifier les applications Web. Dans ce rapport, nous étudions comment les problèmes de causalité affectent ces applications et nous présentons WeCa, une bibliothèque qui permet un contrôle modulaire et flexible et ces problèmes dans les applications Web. Par rapport aux propositions actuelles, WeCa est basé sur les aspects à états, les stratégies d'ordonnancement de messages et les vecteurs d'horloges logiques. Nous illustrons WeCa avec plusieurs exemples concrets du domaine des applications Web. Par exemple, nous analysons le flot d'information dans un application Web comme Twitter à l'aide de WeCa.

**Mots-clés :** WeCa; AspectScript; OTM; JavaScript; applications Web 2.0; calcul distribué; programmation par aspects.

# 1 Introduction

There is a strong trend towards the use of Web 2.0 Applications (WebApps for now) [23], interactive, collaborative, and user-centered Web applications, like Twitter and Facebook. For the development of these applications, the JavaScript language and Ajax technologies [7] are widely used. This is so because JavaScript, a dynamic prototype-based language with higher-order functions, is supported by most modern browsers and Ajax technologies allow WebApps to send and receive messages from a server or other applications[1] asynchronously. This latter feature converts WebApps into distributed systems because processors are user machines that are used to send and receive messages between one another. As a consequence, these applications have to address the causality issues of distributed systems like the need *a)* to control the causality between messages sent and responses received [21] and *b)* to react to distributed causal relations [13].

The need to control the causality between messages sent and responses received arises when a WebApp retrieves various server responses in an arbitrary order. Instead, the need to react to distributed causal relations arises, for example, when it is necessary to analyze the flow of information at runtime that occurs among WebApps [3]. Surprisingly, there is not much technical support that addresses the previous issues and even less modularly and flexibly. The current proposals [5, 25, 29] allow JavaScript programmers to overcome these issues using rudimentary techniques like explicit function wrappers and post-mortem techniques like dynamic graphs. In addition, these techniques largely ignore the distributed computing theory and/or their inflexible uses end up scattered throughout many places in the code of a WebApp entangled with other concerns. In this paper, we present WeCa, a practical library that allows for modular and flexible control over causality on the Web.

In contrast to current proposals, WeCa is based on Aspect-Oriented Programing (AOP) [11] and distributed computing concepts [6]. In particular, WeCa uses stateful aspects [4], Lamport's vector clocks [13, 19] and message ordering strategies [21]. Our proposal allows for modular and flexible definition of *a)* message ordering strategies that control the causality between Ajax requests and server responses and *b)* monitors that react to distributed causal relations. Currently, the WeCa implementation uses AspectScript [26], an aspect-oriented extension of JavaScript, to define aspects that enforce message ordering strategies in WebApps. In addition, the WeCa implementation uses OTM [14, 15], an Open Trace-based Mechanism like tracematches [1] for JavaScript, to define stateful aspects that react to distributed causal relations in WebApps.

The rest of this paper is organized as follows. Section 2 presents and illustrates some causality issues on the Web through different WebApps. Section 3 introduces key concepts on which WeCa is based: (stateful) aspects, message ordering strategies, and vector clocks. Section 4 introduces WeCa, which combines the aforementioned concepts to address modularly and flexibly causality issues on the Web. Section 5 presents how our proposal addresses the causality issues described in Section 2. Section 6 discusses related work and Section 7 concludes.

[1] Using a server as intermediary.

# 2 AJax & Web 2.0 Applications

Ajax [7], a shorthand for *Asynchronous JavaScript and XML*, is a group of interrelated Web technologies used on the client-side to create interactive Web applications. Using Ajax, Web applications can send and retrieve data from a server or other applications asynchronously without reloading the current Web page. The following piece of code written in JavaScript shows a simple request to a server using Ajax:

```
var request = new XMLHttpRequest();
request.open("GET","server.com");
request.onreadystatechange = function() {
  if (this.readyState == 4)
    updateWebPage(this.responseText);
}
request.send(parameters);
```

The request object represents the *Ajax request* to the server. The open method configures the communication with the server, and the send method sets the parameters of the Ajax request and sends it. The server responds with an arbitrary delay. When the server responds, the onreadystatechange method is executed[2]. The responseText instance variable of request contains the *server response*, which can be used, for example, to update the Web page.

Using Ajax technologies, JavaScript programmers create interactive, collaborative, and user-centered Web Applications, known as Web 2.0 Applications [23]. For example, Housing Maps, an application for finding a house for sale, is created from server responses that come from different sources. Another example is Twitter, an application for social network and microblogging, which allows users to send and receive messages. When a WebApp uses Ajax technologies, some needs arise, such as the need *a)* to control the causality between messages and responses received [21] and *b)* to react to distributed causal relations [13].

## 2.1 Controlling Message Causality

A WebApp can send several Ajax requests to different servers. However, the application can retrieve and process the server responses in an arbitrary order; meaning that this application may behave nondeterministically due to the lack of control of the causality between Ajax requests and server responses. For example, Figure 1 shows that if a WebApp sends two Ajax requests to a server, two scenarios are possible: the server returns server response 1 followed by server response 2 or vice versa. Depending on the expected behavior of the WebApp, different strategies to control the message causality can be used. We now present three WebApps that require three different strategies:

***A FIFO strategy for a mashup application.*** A mashup application is created from the combination of information retrieved from different servers, *e.g.* Housing Maps. Programmers have to overcome the issue of creating an incorrect Web page due to an arbitrary (and unexpected) order of server responses. For instance, consider a Web page that is created with two Ajax

[2]This method is actually executed every time that the value of readyState changes. We only want to take an action when the server response is available for the application, *i.e.* when readystate takes the value 4.

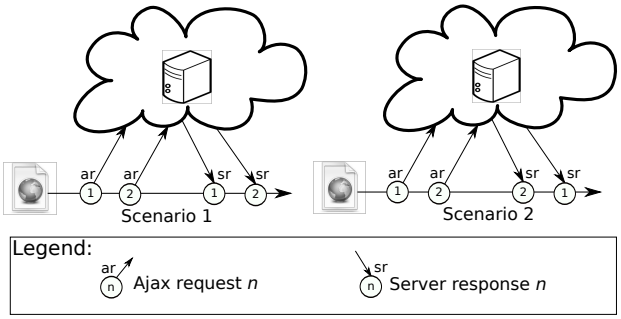



Figure 1: Two possible scenarios when a Web application sends two Ajax requests.

requests. If the second server response is processed before the first server response, the Web page is created incorrectly. A *FIFO* strategy, which processes the server responses in the same order as the Ajax requests are sent, ensures that the Web page is always created correctly.

***A Discard Late strategy for the visualization of news threads.*** WebApps typically update data from servers using Ajax (*e.g.* news, thread posts, and streaming medias). Programmers have to overcome the issue of processing obsolete data due to late server responses. For example, consider a Web page of news threads that shows the last updates about several news. These updates are retrieved from server responses. If some server responses are retrieved and displayed in an incorrect order, the reader could misunderstand these last updates of a news thread. A *Discard Late* strategy that discards the late (and obsolete) server responses always shows really the last updates of a news thread. Therefore, the last updates of a news thread are exposed in correct order for the reader[3].

***A Discard Early strategy for the visualization of thread posts.*** Another issue in the visualization arises due to the processing of early server responses. For example, consider a forum Web page that is frequently updated with threads from a server using Ajax, *i.e.* this Web page only increases with more threads that are located on the top. Every thread is shown with its title and its first posts, and the reader can click on a button "read more" to see the whole discussion of a given thread. Every post that is below the thread title is loaded dynamically using an independent Ajax request[4]. If server responses, which contain the posts, are retrieved and displayed in an incorrect order, the beginning of the discussion of the thread could be misunderstood. A solution could be to wait for all late posts of every thread. However, this solution could seriously delay the updating of the forum Web page. A better solution is to discard the early posts in order to correctly show the first posts of the thread. The discarded posts can be fetched later when the reader clicks the given "read more" button. A *Discard Early* strategy, which discards early server responses,

[3]At the WeCa website [16], we use the Discard Late to show a visually more extractive example, which shows a streaming video in a correct flow.

[4]A thread commonly contains a (large) variable number of posts.

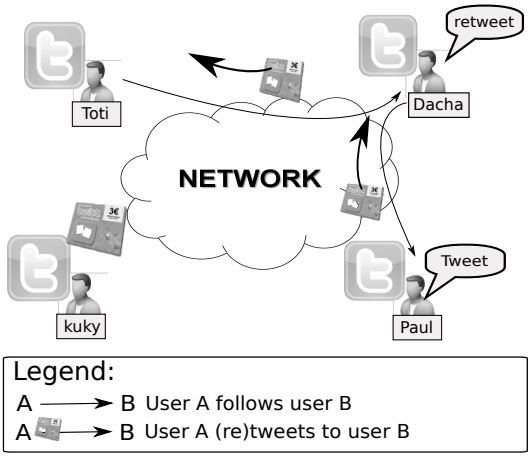



Figure 2: Retweeting a tweet.

can ensure the beginning of the discussion of every thread is read in a correct order.

## 2.2 Reacting to Distributed Causal Relations

Social network applications like Twitter and Facebook are another kind of WebApp. Nowadays, these applications are widely used, making the analysis of their flow of information an active research topic [3, 8, 30, 28]. This flow of information is analyzed through the messages sent and received between users of these applications. Such an analysis is complex due to the need to observe and react to distributed causal relations that occur among user interactions. As an example of the analysis of the flow of information in WebApps, consider the calculation of the popularity of user *tweets*[5] in Twitter:

***Tweet popularity.*** This feature in Twitter [17] allows a user to know the popularity of every tweet published by him or her, which is measured by the number of *retweets*[6] of direct and indirect followers. For example, Figure 2 shows four Tweeter users: Toti, Dacha, Kuky, and Paul. Toti follows Dacha and Dacha follows Paul; Kuky follows nobody and nobody follows Kuky. The figure shows that Paul publishes a tweet and Dacha receives this tweet and retweets it. The figure also shows that Kuky publishes a tweet and nobody receives it. Based on the popularity measurement, the popularity of Paul's tweet is 1 and that of Kuky is 0. Although Kuky and Paul would have published the same tweet[7], the popularity of Kuky's tweet is 0 because his tweet did not cause any retweet. An analysis based on the distributed causal relations observed between

tweets and retweets can determine how many users retweet a given tweet. For example, Paul's tweet caused Dacha's retweet.

## 2.3 State of the Practice

State of the practice provides JavaScript practitioners have at their disposal a number of JavaScript libraries and postmortem tools to solve the kind of issues we described previously.

Some lightweight JavaScript libraries [5, 25] are used to control the causality between Ajax requests and server responses. However, these libraries do not modularly control the message causality because programmers need to explicitly wrap every Ajax request in order to use these libraries. In addition, these libraries only provide a limited set of strategies that are not customizable. For example, the following piece of code that uses the AjaxManager library [5] enforces a server response to follow the FIFO strategy. The FIFO strategy is enabled with a boolean value (queue). In addition, the Ajax request associated to the server response that must follow the FIFO strategy ("myFIFO") has to be rewritten manually.

```
//deploying the FIFO strategy
managerAjax.create("myFIFO",{
  queue: true, //this property means that FIFO is enabled
});

//An Ajax request that follows the previous FIFO strategy
manageAjax.add("myFIFO",{
  success: function(serverResponse) {
    updateWebPage(serverResponse);
  },
  url: "server.com"
});
```

Postmortem tools based on dynamic graphs [29] are used to observe distributed causal relations. As these tools are postmortem, they cannot be used to react to distributed causal relations at runtime.

## 2.4 WeCa Overview

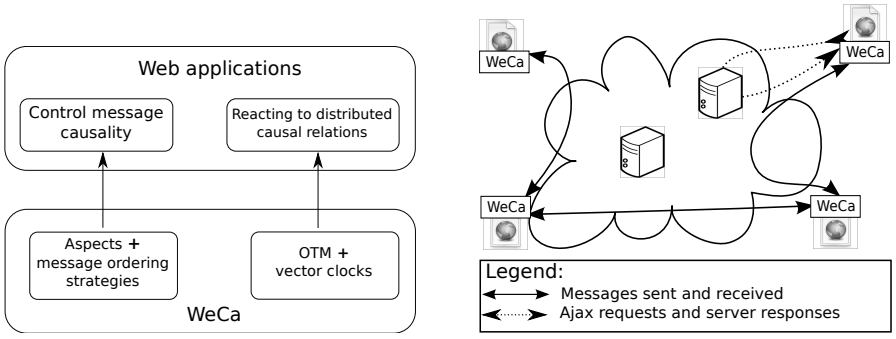
WeCa is a practical library that allows for modular and flexible control over causality as required in the examples presented previously (see Figure 3(a)). By *modular* we mean that the control over causality is addressed in a separate module at the code level, and by *flexible* we mean that the definition of the control over causality is customizable using the power of the base language. The runtime of WeCa observes every Ajax request with its server response to enforce a certain strategy to deal with server responses (*e.g.* Discard Early). In addition, this runtime observes every message sent and received between these applications to detect distributed causal relations at runtime in interactive communication of WebApps (*e.g.* the popularity of a tweet).

Figure 3(b) shows how WeCa works in a nutshell. Every WebApp has a WeCa runtime instance. Each instance can use an *aspect* [11] to match every

[5]Messages posted via Twitter containing 140 characters or fewer. This word is indistinctly used as a noun and a verb.

[6]Tweets reposted by another user. This word is indistinctly used as a noun and a verb.

[7]Two tweets are equal if both contain the same string or have the same url associated. In this paper, we say two tweets are equal if they contain the same string.




(a) The solution proposed by WeCa. (b) WeCa overview.

Figure 3: WeCa.

server response, and using *message ordering strategies* [21], the aspect enforces that (some) server responses follow a certain message ordering strategy. In addition, each instance can use *stateful aspects* [4] to match and react to distributed computations. These stateful aspects utilize *vector clocks* [13, 19] to match and react to distributed computations that satisfy certain causal relations.

# 3 Aspect-Oriented Programming & Distributed Computing

WeCa combines concepts from AOP and distributed computing [6]. In this section, we summarize these necessary concepts in order to clarify our proposal.

## 3.1 Aspect-Oriented Programming

Aspect-oriented programming makes it possible to modularize *crosscutting concerns*, like the ability to react to distributed causal relations. Specifically, in the pointcut-advice model for aspect-oriented programming [18, 27], crosscutting behavior is defined by means of *pointcuts* and *advice*. Execution points at which advice may be executed are called *(dynamic) join points*. A pointcut matches a set of join points, and a piece of advice is the action to be taken *before*, *around*, or *after* the matched join point. Around advice can invoke the computation of the join point matched, known as the *proceed* invocation. An *aspect* is a module that encompasses pointcuts and advice.

### 3.1.1 AspectScript in a Nutshell

AspectScript [26] is an aspect-oriented extension of JavaScript, which follows the pointcut-advice model. We introduce AspectScript using an example: consider a Web page that is frequently updated through Ajax and the following aspect updates the Web page only when it retrieves a new server response from a server:

```
var aspSelectiveUpdate = {
  pointcut: function(jp,env) {
    return jp.isExec() && jp.target instanceof XMLHttpRequest &&
```

```
          jp.fun == jp.target.onreadystatechange?
        env.bind("serverResponse",jp.target.responseText): false;
  },
  advice: function(jp,env) {
    if (this.lastServerResponse != env.serverResponse) {
      jp.proceed();
      this.lastServerResponse = env.serverResponse;
    }
   },
   kind: AROUND,
   lastServerResponse: null //initial value
};

AspectScript.deploy(aspSelectiveUpdate); //deployment
```

In AspectScript, an aspect is a plain JavaScript object that defines at least three properties. The pointcut and advice properties are plain JavaScript functions parametrized by the jp and env objects: jp represents the current join point and env is an environment that is used to pass information from the pointcut to the corresponding advice. As a consequence, a join point is a first-class value, where context exposure and original computation (*i.e.* proceed) are properties of the join point. The pointcut function matches the jp join point; the advice function takes the action before, around, or after (according to the kind property) of the join point matched. In our example, the aspSelectiveUpdate aspect defines a pointcut that returns an environment that contains the server response when this pointcut matches the execution of the onreadystatechange method; the around advice of the aspect only executes the proceed method if the server response differs from the last one, *i.e.* if the server response is new.

***Pointcut Model.*** Unlike many aspect-oriented extensions like AspectJ [9], where a pointcut is defined by a pattern expressed in a domain-specific language, AspectScript uses the base language, a JavaScript function, to express a pointcut. Following standard practice [18], a pointcut can return *a)* an environment if it matches the current join point or *b)* false if it does not[8].

### 3.1.2 OTM in a Nutshell

The pointcut of an aspect refers to the current join point. Therefore, a pointcut cannot refer to a trace of join points. For example, consider the implementation of an aspect that prevents a malicious application from inserting (random) credentials several times until logging successfully[9]:

```
var lg = function(jp,env) {
  return jp.isExec() && jp.fun == login? env:false;
};

var aspLogin = {
  pointcut: lg,
  advice: function(jp,env) {
    if (++this.counter > 3)
      throw "Only three attempts are permitted";
  },
  kind: BEFORE,
  counter: 0
}

AspectScript.deploy(aspLogin); //deployment
```

[8] A contrast between the standard practice and the pointcut model of AspectScript is the environment passed as parameter to AspectScript pointcuts. This environment is used to pass information between pointcuts, *e.g.* cflow.

[9] Using Ajax technologies, malicious applications can insert several credentials without reloading the current Web page.

To match the execution trace of three logins, the aspLogin aspect maintains a counter to keep track of the number of login attempts. Every time the login function is called, we increase the counter by one; and if this counter reaches four, the exception is triggered. Note how the state of the matching process (*i.e.* the counter) is explicit in this aspect. In this example, the explicit matching process leads to only minor complications, however, the burden of such state maintenance is often much greater. For example, the solution of a counter is not sufficient if we need to match a sequence of different pointcuts.

Trace-based mechanisms like tracematches [1] support the definition of *stateful aspects* [4] that match and react to an execution trace of an application. OTM [14, 15], a seamless AspectScript extension, is an Open Trace-based Mechanism for JavaScript. For example, consider the previous example using OTM:

```
var sAspLogin = {
  sequence: seqn(lg,4),
  advice: function(jp,env) {
    throw "Only three attempts are permitted";
  },
  kind: BEFORE
};
OTM.deploy(sAspLogin); //deployment
```

Similar to AspectScript, a stateful aspect in OTM is a plain JavaScript object that defines at least three properties. The sequence and advice properties are also plain JavaScript functions parametrized by a join point and an environment. The sequence function matches a trace of join points of the application execution; the advice and kind properties are used for the same purpose as in AspectScript. In our example, the sAspLogin stateful aspect defines a sequence that matches four executions of the login function and a piece of advice that triggers an exception before the fourth execution of login.

***Sequence Model.*** The function that represents a sequence can return *a)* an environment if it matches a trace of join points, *b)* a pair composed of a function and an environment if the sequence advances in its matching, or *c)* false otherwise. The pair returned by the sequence establishes what the next step will be within the matching process. Based on the definition of a sequence, an AspectScript pointcut is also a sequence because a pointcut returns an environment or false, which is a subset of what a sequence returns.

```
var seqn = function(subSeq,n) {
  return function(jp,env) {
    var mainSeq = subSeq;
    for (var i = 0; i < n; ++i)
      mainSeq = seq(subSeq,mainSeq);
    return mainSeq;
} };
```

The piece of code above shows the definition of the seqn sequence designator (*i.e.* a function that returns a sequence). This sequence designator is just a convenient abstraction on the top of the seq binary sequence designator. For example, the evaluation of seqn(lg,4) becomes seq(lg,seq(lg,seq(lg,lg)). We illustrate the sequence model of OTM with the seq sequence designator. This sequence designator takes two (sub)sequences as parameters and returns a sequence that matches a (sub)sequence followed by another. As shown in the piece of code below, the sequence returns a pair composed of the right function and the result environment if left matches. The sequence returns false if left does not match. Note that the sequence never returns an environment (*i.e.* the sequence matches)

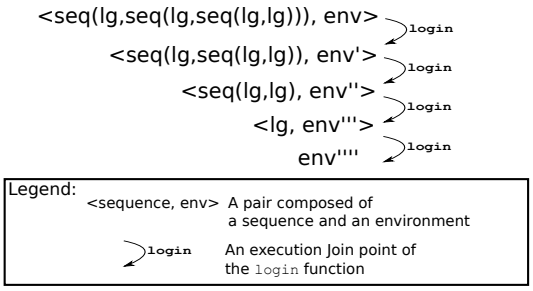



Figure 4: Matching an execution trace.

because the responsibility to match is delegated to the right sequence, which is evaluated with the result environment.

```
var seq = function(left,right) {
  return function(jp,env) {
    var result = left(jp,env);
    if (isEnv(result))
      return [right,result];
    return false;
} };
```

Figure 4 shows how the pair composed of a sequence and an environment of the sAspLogin aspect varies throughout the matching of a trace of four executions of the login function. This pair varies every time the sequence matches a login execution. In the beginning, sAspLogin begins with the pair composed of the sequence expressed by the programmer and an empty environment. The first time login is executed, the pair varies to a new pair composed of a sequence that only matches the three login executions and an environment possibly modified. The third time login is executed, the sequence of the pair is only the lg sequence. Finally, the fourth time login is executed, there is only an environment, meaning that the whole sequence matches; therefore, the stateful aspect executes its advice.

***Openness.*** OTM follows open implementation principles [10]. Thereby, OTM allows developers to customize the crucial semantics of stateful aspects, like the spawning of *matchers*. A matcher is an internal component of a stateful aspect that carries out the task of matching a sequence. A stateful aspect can have several matchers to match several sequences at time (*e.g.* multiple users trying to login at time). The semantics of spawning decides when a stateful aspect adds a new matcher (more on this in Section 5.2). Apart from the spawning semantics, other semantics of stateful aspects can be customized [14], however, these are omitted for reason space in this paper.

***Contribution for WeCa.*** This section showed the flexibility to express (stateful) aspects in AspectScript and OTM through the modularization of crosscutting concerns in WebApps. However, WeCa still needs some distributed computing concepts to address causality issues appropriately. The next section explains these concepts.

## 3.2 Distributed Computing

In distributed systems, processes communicate with each other using messages that are sent over the network. The sending and receiving of these messages as well as beginnings and ends of executions of functions are considered events of a distributed computation. Whereas we can observe a total order among events of a single process, there is no global clock (or perfectly synchronized local clocks) that allows us to observe a total order between events of different processes. Nevertheless, it is possible to observe a *partial order* between events if we use the *happened before* model proposed by Lamport [13].

***Aspects for distributed computations.*** In AOP, join points are execution points which correspond to events. Traditionally (stateful) aspects react to join points of a single process. To react to *distributed* traces of join points, they have to match join points of different processes; and to react to distributed causal relations, (stateful) aspects have to consider the causal relations among these join points. In this section, we explain distributed computing concepts using AOP terminology in order to understand how aspects can react to distributed causal relations. For example, we explain the happened before model using join points instead of events:

$$jp_1 \rightarrow jp_2$$

The happened before model defines a relation between two join points noted with an arrow. The arrow indicates a join point $jp_1$ causes another join point $jp_2$, meaning there is a causal relation from $jp_1$ to $jp_2$. For example, Figure 5 shows in both scenarios that the join point $jp_{s1}$ causes the join points $jp_{s2}$, $jp_{r1}$, and $jp_{r2}$ because the join point $jp_{s1}$ *happened* before.

The happened before model makes it possible to address causality issues such as the need *a)* to control the causality between messages sent and responses received and *b)* to react to distributed casual relations.

### 3.2.1 Controlling message causality

Distributed systems are difficult to test because of their nondeterministic nature, that is, these systems may exhibit multiple behaviors for the same input. This nondeterminism is caused by an arbitrary ordering of messages received by processes in different distributed computations. Figure 5 shows that a process $P_1$ sends two messages, represented by $jp_{s1}$ and $jp_{s2}$ join points, to a process $P_2$. $P_2$ can receive these messages, represented by $jp_{r1}$ and $jp_{r2}$ join points, in two different orders arbitrarily.

Fortunately, message ordering strategies [21] can be used to control the message causality. We now detail some of these strategies:

***FIFO.*** Any two messages from a process $P_i$ to $P_j$ are received in the same order as they are sent. In a formal manner, let $jp_{s1}$ and $jp_{s2}$ be any two join points sent and $jp_{s1} \rightarrow jp_{s2}$, then the $jp_{r2}$ join point cannot be observed before the $jp_{r1}$ join point by any process.

$$jp_{s1} \rightarrow jp_{s2} \Rightarrow \neg(jp_{r2} \rightarrow jp_{r1})$$

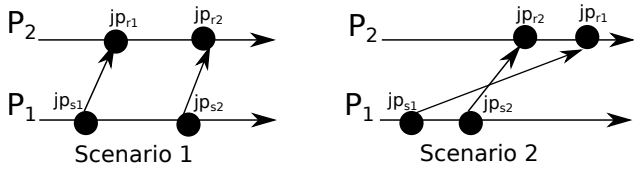



Figure 5: Two possible distributed computations for the same input.

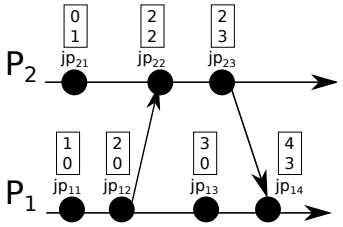



Figure 6: Join points of a distributed computation tagged with vector clocks.

***Discard Late.*** Any two messages from a process $P_i$ to $P_j$ are received in the same order as they are sent or only the most recent message is received. In a formal manner, this strategy is just an extension of FIFO: if $jp_{r2}$ is observed first, $jp_{r1}$ cannot be observed by any process.

$$jp_{s1} \rightarrow jp_{s2} \Rightarrow \neg(jp_{r2} \rightarrow jp_{r1}) \vee jp_{r2}$$

***Discard Early.*** Any two messages from a process $P_i$ to $P_j$ are received in the same order as they are sent or only the oldest message is received. In a formal manner, this strategy is just an extension of FIFO: if $jp_{r1}$ is observed first, $jp_{r2}$ cannot be observed by any process.

$$jp_{s1} \rightarrow jp_{s2} \Rightarrow \neg(jp_{r2} \rightarrow jp_{r1}) \vee jp_{r1}$$

### 3.2.2 Reacting to distributed causal relations

To react to distributed causal relations, stateful aspects have to match distributed traces of join points and to consider causal relations between these join points. Sadly, the absence of a total order among these join points does not allow stateful aspects to consider causal relations.

Fortunately, the algorithm of Lamport's vector clocks [13, 19] can allow stateful aspects to observe a partial order. We now explain a straightforward adaptation of this algorithm to allow stateful aspects to consider causal relations between join points of distributed traces.

To support vector clocks, every join point $jp$ is tagged with an array of counters of size $n$, where $n$ is the number of processes. This array represents a vector clock $V$, which is accessed by $V(jp)$, and is filled according to the following algorithm:

1. Initially, $V_i[k] = 0$ for $k = 1, ..., n$.
2. On each join point that does not represent the sending or receiving of a message, process $P_i$ increases $V_i$ as follows: $V_i[i] = V_i[i] + 1$.

3. On each join point $jp$ that represents the sending of a message $m$, process $P_i$ updates $V_i$ as in *2)* and attaches the new vector clock to $jp$.

4. On each join point $jp$ that represents the receiving of a message $m$, process $P_i$ increases $V_i$ as in *2)*. Then, $P_i$ updates its current $V_i$ as follows: $V_i = sup\{V_i, V(jp)\}$.

If all join points are tagged according to this algorithm, it is possible to define rules that verify if two join points satisfy distributed causal relations such as *causal* and *concurrent*:

***Causal rule.*** This rule allows us to verify if a join point $jp_1$ *caused* another join point $jp_2$. For example, Figure 6 shows that $jp_{12}$ caused $jp_{22}$ and $jp_{23}$. In a formal manner, the causal rule is defined by the following relation:

$$jp_1 \rightarrow jp_2 \Longleftrightarrow V(jp_1) < V(jp_2) \tag{1}$$

Where:

$$V(jp_1) < V(jp_2) \Longleftrightarrow \forall k[V(jp_1)[k] \leq V(jp_2)[k]] \wedge \exists k'[V(jp_1)[k'] < V(jp_2)[k']]$$

***Concurrent rule.*** This rule allows us to verify that the $jp_1$ and $jp_2$ join points are *concurrent*: $jp_1$ does not cause $jp_2$ and vice versa. For example, Figure 6 shows that $jp_{22}$ and $jp_{13}$ are concurrent. In a formal manner, the concurrent rule is defined by the following relation:

$$jp_1 \parallel jp_2 \Longleftrightarrow \neg(jp_1 \rightarrow jp_2) \wedge \neg(jp_2 \rightarrow jp_1) \tag{2}$$

***Contribution for WeCa.*** This section showed the necessary concepts to work on causality. The implementation of causality in distributed systems is crosscutting concern because it is necessary to intercept, modify, and (possibly) postpone the evaluation of every join point. For example, the FIFO strategy postpones the evaluation of early join points. For this reason, the use of (stateful) aspects of AspectScript and OTM allows WeCa to modularize these causality concerns.

# 4 WeCa

This section describes WeCa, a practical library that allows for modular and flexible control over causality on the Web. This library combines stateful aspects, message ordering strategies, and vector clocks to allow JavaScript programmers to modularly and flexibly define *a)* strategies that control the causality between Ajax requests and server responses and *b)* monitors that react to distributed causal relations.

Figure 3(b) shows that every WebApp has a WeCa runtime instance. Every instance observes join points generated on the application and other connected applications. If a programmer defines a message ordering strategy, the WeCa runtime enforces that server responses follow this strategy. In addition, if the programmer expresses a distributed causal relation pattern, the WeCa runtime frequently observes the distributed trace of join points to react whenever this trace is matched by such a pattern. We now explain how WeCa enforces message ordering strategies and reacts to distributed causal relations.

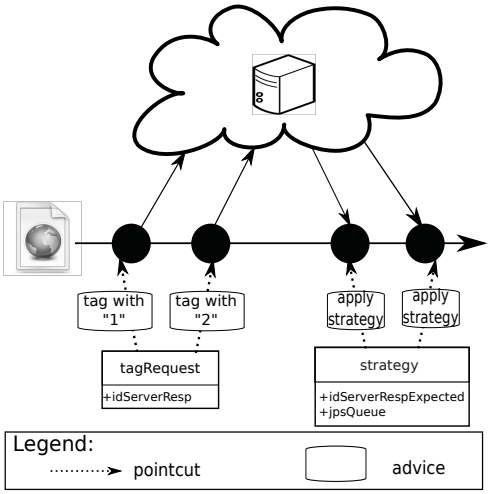



Figure 7: Two aspects to apply a message ordering strategy.

## 4.1 Controlling Message Causality

As mentioned in Section 2.1, a WebApp can send several Ajax requests to servers. However, this application can retrieve and process the server responses in an arbitrary order, therefore, this application behaves nondeterministically. Unlike that presented in Section 3.2.1, the problem does not arise between the messages sent by different processes, the problem arises when a WebApp retrieves server responses (Figure 1). Fortunately, message ordering strategies can also be used to control the causality between Ajax requests and server responses, and AspectScript aspects allow programmers to flexibly and modularly define these strategies. Using aspects and message ordering strategies, WeCa can be used, for example, to enforce server responses retrieved from a server following the Discard Early strategy (Section 2.1), which is provided as a library function in WeCa:

```
WeCa.deployStrategy(discardEarly,function(jp) {
  var request = jp.target;
  return request.url == "http://server.com";
});
```

In this piece of code, the deployStrategy method deploys the Discard Early strategy, which enforces server responses that come from http://server.com to follow this strategy. The two parameters of the deployStrategy method are plain JavaScript functions: the first represents the strategy and the second represents its scope. As message ordering strategies are plain JavaScript functions, the programmer can use the full power of the base language, in particular higher-order functions, to define a strategy. The second function of deployStrategy is parametrized by the join point that represents the execution of the onreadystatechange method, which processes the current server response. This function returns true if the server response must follow the strategy.

We now explain how WeCa implements and allows programmers to deploy flexible these strategies.

***Implementation details.*** Figure 7 shows how WeCa only needs two aspects to support message ordering strategies. The first aspect, tagRequest, matches calls to the send method, which represents an Ajax request, of a request object. The piece of advice of this aspect tags the request object with a fresh and incremental identifier idServerResp that is generated to identify the associated server response. The second aspect, strategy, matches the executions of the onreadystatechange method, which processes the server response of a request object. The piece of advice of this aspect is actually the message ordering strategy. This advice is defined by the programmer, who can use the idServerRespExpected and jpsQueue aspect instance variables to implement the strategy. The idServerRespExpected binding identifies the next server response expected, and jpsQueue is a queue of join points that contains methods to execute join point proceeds. As an example, the implementation of the Discard Early strategy is:

```
var discardEarly = function(jp,env) {
  var request = jp.target;
  if (this.idServerRespExpected >= request.idServerResp) {
    jp.proceed();
    this.idServerRespExpected = request.idServerResp + 1;
} };
```

The jp join point represents the execution of the onreadystatechange method that processes the current server response. The proceed method is only executed if the server response is equal or older than the server response expected, *i.e.* if the server response is not early. Apart from the Discard Early strategy, WeCa provides other message ordering strategies as library functions.

## 4.2 Reacting to Distributed Casual Relations

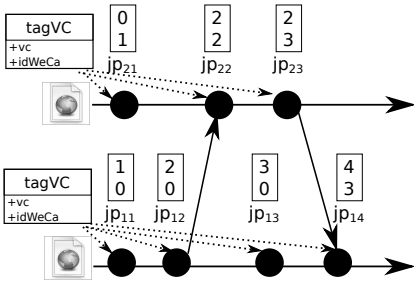
As mentioned in Section 2.2 the analysis of the flow of information of social network applications is an active research topic. This flow of information is analyzed through the messages sent and received between users of these applications, meaning that it is necessary to observe and react to distributed causal relations given by these messages.

Fortunately, Lamport's vector clocks can also be used to observe distributed causal relations in WebApps, and OTM stateful aspects allow programmers to react to these distributed causal relations in a flexible and modular manner. Using stateful aspects and vector clocks, WeCa can be used, for example, to deploy a stateful aspect that shows a message every time a follower retweets some tweet:

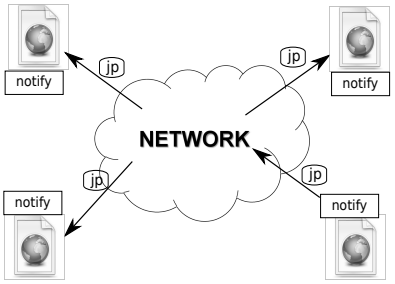
```
var callTweet = function(jp,env) {
  if (jp.isCall() && jp.fun == tweet) {
    var tweet = jp.args[0]; //1st argument to tweet
    return env.bind("tweet",tweet);
  }
  return false;
};

var notifyRetweet = function(jp,env) {
  if (jp.isCustom("notification-call") && jp.fun == retweet) {
    var retweet = jp.args[0]; //1st argument to retweet
    return retweet == env.tweet? env: false;
  }
  return false;
};
//necessary (sub)sequences to create the stateful aspect

var sAspNotifyRetweet = {
  sequence: causalSeq(callTweet,notifyRetweet),
```




(a) The tagVC aspect that tags every join points with a vector clock.




(b) The notify aspect that notifies every join point to other applications.

Figure 8: Aspects to support the matching of distributed causal relations.

```
  advice: function(jp,env){
    alert("One of your follower retweeted:"+env.tweet);
  },
  kind: AFTER
};

WeCa.OTM.deploy(sAspNotifyRetweet); //deployment
```

The sAspNotifyRetweet stateful aspect matches every time a tweet causes a retweet. The callTweet sequence returns an environment that contains the tweet, passed as parameter to the tweet function, if this sequence matches the call to tweet. The notifyRetweet sequence matches a *notification call join point*, a call join point that occurs in another connected application, of the call to the retweet function if both the tweet and retweet are equal. Finally, we use the causalSeq sequence designator to return a sequence that matches the two previous (sub)sequences if both satisfy the causal rule (Section 3.2.2).

We now explain how WeCa implements notification join points and sequences that match distributed causal relations.

***Implementation details.*** WeCa only needs two aspects to react to distributed causal relations. Figure 8(a) shows the first aspect, tagVC, which tags every join point with the corresponding vector clock. Figure 8(b) shows the second aspect, notify, which notifies other WeCa runtime instances for every join point generated on the computation of a connected WebApp. When a WeCa runtime instance receives the information about a remote join point, this instance generates a notification join point of the corresponding kind[10], which can be matched by any stateful aspect. The context information of a notification join point includes the join point generated plus the WeCa instance identifier.

As we show in the previous piece of code, to react to distributed causal relations, WeCa uses OTM to define stateful aspects that react when they match distributed causal relations. In OTM, to define sequences that match distributed causal relations using the causal or concurrent rule (Section 3.2.2), it is necessary to have sequence designators like causalSeq to create sequences

[10]Like Ptolemy [24], AspectScript allows developers to explicitly trigger customized join points.

that consider causal relations. To achieve this, we only have to extend the seq sequence designator (Section 3.1.2):

```
var seq = function(left,right) {
  return function(jp,env) {
    var result = left(jp,env);
    if (isEnv(result))
      return [function(jpNext,env) {
        if ( vectorClockCondition(jp,jpNext) )
          return right(jpNext,env);
        return false;},result];
    return false;
} };
```

In Section 3.1.1, the seq sequence designator returns a sequence that if left matches this sequence directly returns right. Instead, this new seq returns a sequence that if left matches this sequence returns a function that first verifies if the jp and jpNext join points satisfy some causal relation before evaluating right. The relation causal is verified through the vectorClockCondition function. Based on the piece of above, we can create sequence designators like causalSeq, concSeq, or seq using higher-order functions:

```
var makeSeq = function(vectorClockCondition) {
  return function seq(left,right) {
    //the piece of code above
} } };

var causalSeq = makeSeq(caused);
var concSeq = makeSeq(areConcurrent);
var seq = makeSeq(function(jp,jpNext) {return true;});
```

The caused function verifies whether jp causes jpNext or not , and the areConcurrent function verifies whether both join points are concurrent or not. The functions caused and areConcurrent are developed according to the relations 1 and 2 described in Section 3.2.2. Finally, if we need a seq sequence designator that does not consider any causal relation, the makeSeq function is called with a function that always returns true. These sequence designators are provided as library functions in WeCa.

### 4.3 Summary

WeCa uses only four aspects in order to work. These four aspects extend AspectScript and OTM to address causality issues on the Web. Two aspects are used to support the definition of flexible message ordering strategies, and two aspects are used to react to distributed causal relations. Using the power of the base language, developers can define message ordering strategies and sequences that match distributed causal relations.

## 5 Revisiting Web 2.0 Applications

In Section 2, we have shown examples of how causality issues affect WebApps. This section presents how to address these issues using WeCa. First, we use WeCa to define appropriate message ordering strategies that address the different needs of WebApps described in Section 2.1. Second, we use WeCa to define a stateful aspect that determines the popularity of every tweet of a user (Section 2.2).

## 5.1 Controlling Message Causality

A WebApp can send several Ajax requests to different servers. However, this application can retrieve and process server responses in an arbitrary order. In Section 2.1, we presented three WebApps that needed three different message ordering strategies. This section presents and describes these strategies, which are provided as library functions in WeCa. As mentioned in Section 4.1, these strategies are really advice of an aspect. In addition, we extend one of these strategies to show the usefulness of the flexibility provided by WeCa.

***A FIFO strategy for a mashup application.*** Programmers have to overcome the issue of creating an incorrect Web page due to an arbitrary order of server responses. The FIFO message ordering strategy, which processes the server responses in the same order as Ajax requests are sent, can ensure the Web page is always created correctly. This strategy postpones the use of an early server response until it is the expected one. We now present the function that describes the FIFO strategy:

```
var FIFO = function(jp,env){
  var request = jp.target;
  if (this.idServerRespExpected != request.idServerResp){
    this.jpsQueue.push(jp);
  }
  else{
    jp.proceed();
    this.idServerRespExpected = request.idServerResp + 1;

    //executing jp proceeds that can be executed now
    this.idServerRespExpected =
      this.jpsQueue.execRecEarlyJPs(this.idServerRespExpected);
  }
}

WeCa.deployStrategy(FIFO);
```

As mentioned in Section 4.1, the jp join point represents the execution of the onreadystatechange function. If this server response is not the expected one, jp is added to a queue and its proceed execution is postponed. Instead, if this server response is the expected one, the join point proceed is executed and the queue tries executing proceeds of join points stored due to some of them can be the expected one now. This is so because if a join point proceed is executed inside of the queue, it can permit the execution of another join point proceed. As the scope is not specified in this deployment, the scope is global for this strategy, meaning that all server responses follow this strategy.

***A Discard Late strategy for the visualization of news threads.*** In the visualization of news threads, programmers have to overcome the issue of using obsolete data due to late server responses. A Discard Late strategy, which discards late server responses, always show the last updates of a news thread. We now present the function that describes the Discard Late strategy:

```
var discardLate = function(jp,env) {
  var request = jp.target;
  if (this.idServerRespExpected <= request.idServerResp) {
    jp.proceed();
    this.idServerRespExpected = request.idServerResp + 1;
  }
};

WeCa.deployStrategy(discardLate,function(jp) {
```

```
  var request = jp.target;
  return request.url == "http://news.com";
});
```

If the server response is not late, the proceed of the jp join point is executed. If the server response is late, the join point proceed is not executed, meaning that the execution of the onreadystatechange method that processes this server response is discarded. Given the scope of the deployment, server responses that only come from http://news.com must follow this strategy.

***A Discard Early strategy for the visualization of thread posts.*** In the visualization, programmers also have to overcome the issue of using early data due to early server responses. Remembering the forum Web page which updates posts of last threads using Ajax requests. If some server responses are retrieved and displayed in an incorrect order, the beginning of the discussion of a thread may be misunderstood. A *Discard Early* strategy, which discards early server responses, can ensure the reader better understands the beginning of the discussion of every thread. In Section 4.1, we presented the Discard Early strategy; we now present an extension of this strategy, which executes a callback function when it discards a server response.

```
var deShowReadMore = discardEarlyCallback(function callback(jp,env) {
   //show the "read more" button
});

WeCa.deployStrategy(deShowReadMore, function (jp) {
  var request = jp.target;
  return request.url == "http://forum.com/threadlist";
});
```

We use the callback function to show a button "read more" when a server response is discarded. The piece of code below shows that the evaluation of the discardEarlyCallback function returns a function that uses the Discard Early strategy and executes a callback function, when a server response is discarded.

```
var discardEarlyCallback = function(callback) {
  return function(jp,env) {
   //invoke the discard early strategy

    var request = jp.target;

    //is it discarded?
    if (this.idServerRespExpected < request.idServerResp)
      callback(jp,env);
} };
```

## 5.2 Reacting to Distributed Causal Relations

The analysis of the flow of information of WebApps is analyzed through the messages sent and received among users of these applications. In this section, we use a stateful aspect to calculate the popularity of tweets in Twitter (Section 2.2):

***Tweet popularity.*** This feature allows a user to know the popularity of its tweets, which is measured by the number of retweets. As Figure 2 shows, if a user publishes a tweet that is retweeted by a follower, the popularity of this tweet is increased by one. An analysis based on the distributed causal relations between tweets and retweets can determine how many users retweeted a tweet.

The following stateful aspect counts the retweets of a tweet until the user clicks a button "popularity":

```
var callTweet = //... as Section 4.2

var retweetCount = function (jp,env){
  var result = notifyRetweet(jp,env); //notifyRetweet as Section 4.2
  if (isEnv(result))
    return env.counter == undefined?
      env.bind("counter",0): env.bind("counter",env.counter + 1);
  return false;
};

var callPopularity = function(jp,env) {
  return jp.isCall() && jp.fun == clickPopularityButton;
}
//necessary (sub)sequences to create the stateful aspect

var sAspTweetPopularity = {
  sequence: causalSeq(callTweet,repeatUntil(retweetCount,callPopularity)),
  advice: function(jp,env){
    addPopularity(env.tweet,env.counter);
  },
  kind: AFTER,
};

WeCa.OTM.deploy(sAspTweetPopularity);
```

The sAspTweetPopularity stateful aspect is only an extension of the stateful aspect shown in Section 4.2. This new stateful aspect counts the number of retweets of a tweet and adds this counter to the analyzed tweet. The sequence begins the matching when the callTweet (sub)sequence matches and continues with matchings of the retweetCount (sub)sequence until the user clicks the "popularity" button. Every time retweetCount matches, the counter of retweet is increased by one. The repeatUntil sequence designator, which is provided as a library function in OTM, returns a sequence that matches the first (sub)sequence several times until the second (sub)sequence matches.

Although this stateful aspect works, it can only count the popularity of one tweet. A solution to count the popularity of every tweet could be to have a stateful aspect for every tweet. However, this solution is not very efficient because Twitter users commonly publish the same tweet several times[11]. A better solution would be that the same stateful aspect uses a different matcher[12] of the sequence for every different tweet. As mentioned in Section 3.1.2, OTM allows developers to customize *when* a new matcher of the sequence is spawned. To customize the spawning of matchers in OTM, the spawn property is added to a stateful aspect definition. This property is a function that returns true if a new matcher is spawned and false otherwise. This function is parametrized by the candidate new matcher and current matchers of stateful aspects.

```
var sAspTweetPopularity = {
  //sequence, advice, and kind properties remain the same
  spawn: function(candidateMatcher,matchers){
    var candidateEnv = candidateMatcher.getEnv();
    var currentEnvs = getEnvs(matchers);

    return !currentEnvs.some(function(currentEnv){
      return currentEnv.tweet == candidateEnv.tweet;
    });
};
```

[11] http://holykaw.alltop.com/the-art-of-the-repeat-tweet.

[12] As explained in Section 3.1.2, a stateful aspect can have several matchers to match several sequences at the same time.

In our example, a new matcher is spawned when its associated environment binds a new tweet. The spawn function first gathers the environment of the candidate matcher and those of the current matchers. Then, spawn compares every tweet of current matchers to the tweet of the candidate matcher, and if the latter tweet is different, the candidate matcher is spawned.

# 6 Related Work

Although the relation between AOP and distributed computing to address the causality issues on the Web is a distinguishing contribution of this paper, there are already proposals that address these issues using different techniques. We now review JavaScript libraries to control the causality between Ajax requests and server responses and postmortem tools to observe distributed causal relations. In addition, we also review a proposal to address causality issues in distributed systems using AOP.

***Libraries to control the message causality.*** A number of lightweight JavaScript libraries have been developed by programmers that rely on function wrappers. In the simplest case, Rico [25], a domain-specific library implicitly uses the FIFO message ordering strategy to sort server responses that are shown in an HTML table. A more elaborate case is AjaxManager [5][13], a general-purpose library that offers a fixed set of message ordering strategies. As mentioned in Section 2.3, every Ajax request of a WebApp must manually be written to follow a strategy in AjaxManager. WeCa uses the AOP interception to transparently enforce strategies, *i.e.* Ajax requests are not manually written. In addition, WeCa provides as library functions the strategies available in AjaxManager, and our proposal also allows programmers to customize these strategies to add new variants (*e.g.* the Discard Early strategy with callback shown in Section 5.2).

***Frameworks to observe distributed causal relations.*** On the one hand, a large number of postmortem tools based on dynamic graphs are available to observe distributed causal relations in interactive WebApps [29][14]. On the other hand, Kossinets *et al.* [12] use a modified version of vector clocks, which uses timestamp instead of counters, to analyze the minimum time required for information to spread from one user to another. However, these proposals are postmortem. Therefore, they cannot react to distributed causal relations at runtime like WeCa does.

***Using AOP in distributed systems.*** The extension of AWED [22], a system that explicitly supports the monitoring of distributed computations in Java, takes into consideration distributed causal relations in tasks of debugging and testing of middleware [2]. Although AWED is not developed to address causality issues on the Web, it is a related proposal for WeCa because AWED uses vector clocks and stateful aspects to react to distributed causal relations. Unlike the flexibility provided by WeCa, AWED stateful aspects are expressed using a

[13] AjaxManager is an active project: https://github.com/aFarkas/Ajaxmanager/graphs/traffic.

[14] This reference refers to a list of these tools.

(limited) domain-specific language[15]; therefore, the stateful aspects cannot be defined using the full power of the base language. In addition, AWED only uses the FIFO strategy in an implicit way and does not present a modular (or semantic) separation between enforcing message ordering strategies and reacting to distributed causal relations.

# 7 Conclusions

We have presented WeCa, a practical library that addresses flexibly and modularly some causality issues on the Web through stateful aspects, ordering message strategies, and vector clocks. Concretely, our proposal allows JavaScript developers *a)* to control the causality between Ajax requests and server responses and *b)* to react to distributed causal relations in WebApps. To the best of our knowledge, WeCa is the first (practical) attempt to address these causality issues on the Web in a modular and flexible manner. The use of previous proposals are crosscutting, inflexible, and/or largely ignore the distributed computing theory.

We showed that, contrary to current proposals, WeCa uses general-purpose approaches instead of ad hoc solutions. The key element for this is the flexibility of AspectScript and OTM to define aspects and stateful aspects respectively. This flexibility makes it possible the extension of AspectScript and OTM using plain aspects that enable the supporting of message ordering strategies and vector clocks. To validate our proposal, we used WeCa with several practical examples from the realm of WebApps. For instance, we analyzed the flow of information in WebApps like Twitter.

While many WebApps are highly interactive in nature and may not be so performance-sensitive, it is important to consider the WeCa performance. We plan to improve performance through the use of partial evaluation techniques for higher-order languages like JavaScript [20].

***Availability.*** WeCa, along with the examples presented in this paper, is available online at [16]. Our proposal currently supports the Mozilla Firefox browser without needing any extension.

# References

[1] Chris Allan, Pavel Avgustinov, Aske Simon Christensen, Laurie Hendren, Sascha Kuzins, Ondrej Lhoták, Oege de Moor, Damien Sereni, Ganesh Sittampalam, and Julian Tibble. Adding trace matching with free variables to AspectJ. In *Proceedings of the 20th ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA 2005)*, pages 345–364, San Diego, California, USA, October 2005. ACM Press. ACM SIGPLAN Notices, 40(11).

[2] Luis Daniel Benavides Navarro, Rémi Douence, and Mario Südholt. Debugging and testing middleware with aspect-based control-flow and causal

[15] A regular expression language.

patterns. In *Proceedings of the 9th ACM/IFIP/USENIX International Middleware Conference*, volume 5346 of *Lecture Notes in Computer Science*, pages 183–202, Leuven, Belgium, December 2008. Springer-Verlag.

[3] Fabrício Benevenuto, Tiago Rodrigues, Meeyoung Cha, and Virgílio Almeida. Characterizing user behavior in online social networks. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, IMC '09, pages 49–62, New York, NY, USA, 2009. ACM.

[4] Rémi Douence, Pascal Fradet, and Mario Südholt. Trace-based aspects. In Robert E. Filman, Tzilla Elrad, Siobhán Clarke, and Mehmet Akşit, editors, *Aspect-Oriented Software Development*, pages 201–217. Addison-Wesley, Boston, 2005.

[5] Alexander Farkas. A JavaScript library to block, abort, queue, and synchronize Ajax requests, 2011.

[6] Vijay K. Garg, Ph.D. *Elements of distributed computing*. John Wiley & Sons, Inc., New York, NY, USA, 2002.

[7] Jesse James Garrett. Ajax: A new approach to Web applications, 2005.

[8] Jing Jiang, Christo Wilson, Xiao Wang, Peng Huang, Wenpeng Sha, Yafei Dai, and Ben Y. Zhao. Understanding latent interactions in online social networks. In *Proceedings of the 10th annual conference on Internet measurement*, IMC '10, pages 369–382, New York, NY, USA, 2010. ACM.

[9] Gregor Kiczales, Erik Hilsdale, Jim Hugunin, Mik Kersten, Jeffrey Palm, and William Griswold. An overview of AspectJ. In Jorgen L. Knudsen, editor, *Proceedings of the 15th European Conference on Object-Oriented Programming (ECOOP 2001)*, number 2072 in Lecture Notes in Computer Science, pages 327–353, Budapest, Hungary, June 2001. Springer-Verlag.

[10] Gregor Kiczales, John Lamping, Cristina V. Lopes, Chris Maeda, Anurag Mendhekar, and Gail Murphy. Open implementation design guidelines. In *Proceedings of the 19th International Conference on Software Engineering (ICSE 97)*, pages 481–490, Boston, Massachusetts, USA, 1997. ACM Press.

[11] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina V. Lopes, Jean-Marc Loingtier, and John Irwin. Aspect-oriented programming. In Mehmet Akşit and Satoshi Matsuoka, editors, *Proceedings of the 11th European Conference on Object-Oriented Programming (ECOOP 97)*, volume 1241 of *Lecture Notes in Computer Science*, pages 220–242, Jyväskylä, Finland, June 1997. Springer-Verlag.

[12] Gueorgi Kossinets, Jon Kleinberg, and Duncan Watts. The structure of information pathways in a social communication network. In *Proceeding of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, KDD '08, pages 435–443, New York, NY, USA, 2008. ACM.

[13] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, 1978.

[14] Paul Leger and Éric Tanter. An open trace-based mechanism. In Jonathan Aldrich and Ricardo Massa, editors, *Proceedings of the 14th Brazilian Symposium on Programming Languages (SBLP 2010)*, Salvador - Bahia, Brazil, September 2010.

[15] Paul Leger and Éric Tanter. Towards an open trace-baced mechanism. In Gary T. Leavens, Shmuel Katz, and Mira Mezini, editors, *Proceedings of the Ninth Workshop on Foundations of Aspect-Oriented Languages (FOAL 2010)*, pages 25–30, Rennes and Saint Malo, France, March 2010.

[16] Paul Leger, Éric Tanter, and Rémi Douence. WeCa: A practical library to address causality issues on the web, 2011.

[17] Mashable. Website for news in social and digital media, technology and web culture, 2010.

[18] Hidehiko Masuhara, Gregor Kiczales, and Christopher Dutchyn. A compilation and optimization model for aspect-oriented programs. In G. Hedin, editor, *Proceedings of Compiler Construction (CC2003)*, volume 2622 of *Lecture Notes in Computer Science*, pages 46–60. Springer-Verlag, 2003.

[19] Friedemann Mattern. Virtual time and global states of distributed systems. In Cosnard M. et al., editor, *Proceding Workshop on Parallel and Distributed Algorithms*, pages 215–226, North-Holland / Elsevier, 1989.

[20] Matthew Might, Yannis Smaragdakis, and David Van Horn. Resolving and exploiting the k-CFA paradox: illuminating functional vs. object-oriented program analysis. In *Proceedings of the 2010 ACM SIGPLAN conference on Programming language design and implementation*, PLDI '10, pages 305–315, New York, NY, USA, 2010. ACM.

[21] V.V. Murty and V.K. Garg. Characterization of message ordering specifications and protocols. In *17th IEEE International Conference on Distributed Computing Systems (ICDCS '97)*, pages 492–499, Los Alamitos, CA, USA, may 1997. IEEE Computer Society.

[22] Luis Daniel Benavides Navarro, Mario Südholt, Wim Vanderperren, Bruno De Fraine, and Davy Suvée. Explicitly distributed aop using awed. In *Proceedings of the 5th ACM International Conference on Aspect-Oriented Software Development (AOSD 2006)*, pages 51–62, Bonn, Germany, March 2006. ACM Press.

[23] Tim O'Reilly. What is Web 2.0: Design Patterns and Business Models for the Next Generation of Software. *Communications & Strategies*, 2005.

[24] Hridesh Rajan and Gary T. Leavens. Ptolemy: A language with quantified, typed events. In Jan Vitek, editor, *Proceedings of the 22nd European Conference on Object-oriented Programming (ECOOP 2008)*, number 5142 in Lecture Notes in Computer Science, pages 155–179, Paphos, Cyprus, july 2008. Springer-Verlag.

[25] Rico. A JavaScript library for rich internet applications, 2011.

[26] Rodolfo Toledo, Paul Leger, and Éric Tanter. AspectScript: Expressive aspects for the Web. In *Proceedings of the 9th ACM International Conference on Aspect-Oriented Software Development (AOSD 2010)*, pages 13–24, Rennes and Saint Malo, France, March 2010. ACM Press.

[27] Mitchell Wand, Gregor Kiczales, and Christopher Dutchyn. A semantics for advice and dynamic join points in aspect-oriented programming. *ACM Transactions on Programming Languages and Systems*, 26(5):890–910, September 2004.

[28] Stanley Wasserman and Katherine Faust. *Social network analysis: methods and applications*. Structural analysis in the social sciences, 8. Cambridge University Press, 1 edition, nov 1994.

[29] Wikipedia. Social network analysis software, 2011.

[30] Christo Wilson, Bryce Boe, Alessandra Sala, Krishna P.N. Puttaswamy, and Ben Y. Zhao. User interactions in social networks and their implications. In *Proceedings of the 4th ACM European conference on Computer systems*, EuroSys '09, pages 205–218, New York, NY, USA, 2009. ACM.



Centre de recherche INRIA Rennes – Bretagne Atlantique
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399