

# Proving Properties of Rich Internet Applications

James Smith  
Imperial College  
London, United Kingdom  
jecs@imperial.ac.uk

## Abstract

We introduce application layer specifications, which allow us to reason about the state and transactions of rich Internet applications. We define variants of the state/event based logic UCTL\* along with two example applications to demonstrate this approach, and then look at a distributed, rich Internet application, proving properties about the information it stores and disseminates. Our approach enables us to justify proofs about abstract properties that are preserved in the face of concurrent, networked inputs by proofs about concrete properties in an Internet setting. We conclude that our approach makes it possible to reason about the programs and protocols that comprise the Internet’s application layer with reliability and generality.

## 1 Introduction

A glance at the existing literature shows that the Internet’s transport layer is well understood from a mathematical point of view, with the inductive analysis of the TLS protocol [15] being perhaps the best example. On the other hand, the application layer that sits above it encompasses a plethora of languages and frameworks, and understanding seems to be incomplete. At the highest level, the behaviour of Internet programs and protocols has been analysed [4, 13, 11], often for security reasons, but nevertheless there appears to be much work still to do.

These days, Internet applications are made up of both server side and client side parts. These so called “rich” Internet applications are now so commonplace as to be the norm. Any formalism that helps us to understand them mathematically seems a worthy goal, therefore. A complete formal treatment would include a formalisation of JavaScript [10] and much more, however, so here we attempt a less lofty goal, namely giving formal specifications and proofs for that part of their behaviour related to communication over HTTP. To give an example, consider a “shoot-em-up” style game implemented in JavaScript and run in a web page. The functionality related to the unfolding of the gameplay would encompass game logic and the rendering of each scene, functionality not directly related to the Internet. Suppose the game includes a high score table, however, with high scores stored on the server. Maintaining this table would encompass issues relating to concurrency and communication over HTTP. It is the formal specifications and proofs for this kind of behaviour that is the subject of this paper.

Our motivation in studying this kind of behaviour is a broad understanding of the programs and protocols that comprise the Internet’s application layer. With this in mind, this paper demonstrates that it is possible to formalise certain properties of rich Internet applications using a combination of mathematical models, logics and related techniques, which we call collectively application layer specifications. The approach is based on the conventional HTTP request-response mechanism whose participants are usually the web server on the one hand and the

web browser on the other. Crucially, however, since communication between the client side part of an Internet application with its server side part, usually carried out by way of the XMLHttpRequest object, uses the same HTTP request-response mechanism, application layer specifications encompass the behaviour of rich Internet applications, too.

Our contributions are therefore threefold. Firstly, we formalise HTTP transactions and Internet application state, and introduce the concept of transition rules which describe the relation between the two. This constitutes our model of the Internet’s application layer. Secondly, we define variants of the logic UCTL\* [19], whose semantics are defined over the variants of Kripke transition systems [14] that these transition rules generate. We use these logics to express properties of these transition systems. Thirdly, in the case study we present a formal treatment of a distributed, rich Internet application and prove properties about the information it stores and disseminates. Our approach is a general one in the sense that we do not rely on a particular framework or programming paradigm. The properties we prove are ones that could be exhibited by any HTTP-based Internet program or protocol.

## 2 HTTP transactions and Internet application state

To the uninitiated an Internet application may appear to persist. A web page loads and then remains in the browser in much the same way that a desktop application remains in front of the user. The real situation is somewhat different, however. The presence of the web page may give the impression of permanence but in reality the instance of the server part of the application that produced it is likely to be destroyed the moment the web page is produced. By contrast, the client side part of a rich Internet application does persist, it can communicate with the server side part alongside the browser and it has access to and can amend shared information stored by the browser. To complicate the picture still further, the freedom that the browser admits means that user interactions cannot be guaranteed to occur sequentially in relation to the unfolding of the application state. The user may open new tabs, refresh the page, type URLs directly into the address bar or, most notoriously, use the forward and back buttons. The effects of these actions are complex to discern to say the least.

Rather than attempting to model this cat’s cradle, we take a simplified approach. Our analysis begins with the working of the HTTP protocol and in particular the request-response pair, which we call an HTTP transaction. Although it is possible for the request-response pair to be broken, the connection may fail or the server may fault, for example, in our analysis it will be considered inviolable. Many transactions may happen, building up the complete picture, so we therefore take account of all possible transactions. Although if we make the assumption that no two transactions can occur at exactly the same time then they can be given a temporal ordering, such an ordering is of little practical use since it is extremely difficult if not impossible to establish any causal link between successive transactions. It is best therefore to think of the transactions to begin with as floating around in a kind of “soup”, with no connections between them at all.

At this point we make two observations. The first is that HTTP transactions can happen as a result of requests from both the browser and the client side part of the application, and therefore this soup may contain both types. The second is that no account is taken of the cause of each transaction. Each may happen as a result of the user simply clicking a link; pressing the back button and then confirming that they would like the browser to resend form information; manually typing an address into the address bar; the client side part of the application making

a call via the XMLHttpRequest object. By eschewing these causes we avoid the trap of failing to properly account for them. Furthermore, we do not have a model of browser architecture or attempt to discern the subtle interplay of user interactions with iframes, different tabs and the like. And if we lose something in taking this line what we gain is a complete model of what remains, namely the behaviour of the application as evinced by its transactions over HTTP.

This soup of transactions on its own is of little use. To begin with it takes no account of the state of the application, however this is defined. And some causal links must be established, some additional structure needs to be added. To do this we include the states of the application in our models, and associate them with transactions. More precisely, we say that if an application is in a certain state before an HTTP request, its state may change after the HTTP response. We call this process a transition, with that part of the transition visible over HTTP being precisely the transaction.

Our models therefore consist of the set of all possible application states together with sets of all possible transactions between any two states. We do not consider the causes of requests, and we assume that any request can happen at any time. This approach has precedents in the literature on the verification of network protocols, for example [16]. In such cases there is no user interface to be modelled, and verification consists of exhaustively “covering all the bases” with regard to all possible transactions and their effects on application state. And this approach also characterises some of the most stringent web security requirements in industry, including [2], where no assumptions are made about the types of requests that can be made in an “anything can happen” analysis. Our approach is a natural extension of these approaches.

### 3 A formal model and examples

In this section we define a formal model of HTTP transactions and Internet application state as well a transition rules, which describe the relation between the two. We then define two example applications based on this model. The first application uses cookies to enforce user flow through a site and demonstrates the persistence of application state in the browser. The second application maintains a server state in order to track visitors to a site.

The HTTP protocol is a request-response protocol that admits communication between a client, typically a web browser or a JavaScript program running in a browser, and a server, typically an instance of a script or program running on a web server. Both requests and responses consist of a number of headers together with a body. The request often contains a reference, called a uniform resource locator or URL for short, to a resource to be retrieved by the server. Scripts and programs on the server can also be invoked in this manner. The headers communicate cookies in both directions alongside other types information, such as the client’s IP address. Cookies consist of name-value pairs of textual information and allow stateful communication between clients and servers. They are initially passed by the server to the client in response to a request. They are then stored in the browser and passed back from the client to the server in subsequent requests. Usually the instance of a script or program that creates any response is destroyed the moment the response is produced and therefore cookies allow subsequent instances to recover information specific to a particular client. Typically an identifier unique to the client is stored in a cookie, which allows information specific to that client to be retrieved from a database on the server. This leads to the common misconception that the identifier itself persists there. Aside from this, state not related to any particular client may be held on the server. Also a JavaScript program running in the browser may hold state apart from the browser’s cookie store.

Cookies, which we model as atomic entities rather than name-value pairs, are taken from a set of cookies  $\mathcal{C}$ , with  $c, c'$  ranging over  $\mathcal{C} = \mathcal{P}(\mathcal{C})$ . Since the browser can effectively be instructed to both add and remove cookies from its store, we model them as being signed when passed to the browser, so we have a set of signed cookies  $\mathcal{C}^\pm = \{+x|x \in \mathcal{C}\} \cup \{-x|x \in \mathcal{C}\}$ , with  $c^\pm$  ranging over  $\mathcal{C}^\pm = \mathcal{P}(\mathcal{C}^\pm)$ . Global states are tuples containing browser's cookie store, which from now on we refer to as the browser's state, together with the client side and side server states:

$$(c, j, s) \in \mathcal{C} \times \mathcal{J} \times \mathcal{S}$$

Requests consist of a set of cookies  $c \in \mathcal{C}$  together with a URL  $u$  taken from a set of URLs  $\mathcal{U}$ . Responses consist of a set of signed cookies  $c^\pm \in \mathcal{C}^\pm$  together with a body  $b$  taken from a set of bodies  $B$ . We may refine requests and responses further as required, adding additional request headers, for example, or parameterising the response bodies. A transaction consists of a request and response in one tuple, which labels a transition between global states:

$$(c, j, s) \xrightarrow{(c, u, c^\pm, b)} (c', j', s')$$

Note that the browser's state  $c$  is echoed directly in the request, and therefore there is some redundancy in this formalism. We keep this, however, so that the transaction represents that part of the transition visible over HTTP in its entirety. To continue, when the browser receives signed cookies, it amends its state according to the following generic rule:

$$c' = c \cup \{x \in \mathcal{C} \mid +x \in c^\pm\} \setminus \{x \in \mathcal{C} \mid -x \in c^\pm\}$$

Transitions are governed by transition rules, which, given an initial state and request, determine both the response and final state:

$$(c, j, s) \xrightarrow[c^\pm, b]{c, u} (c', j', s')$$

Transition rules are typically obtained from static analysis. This process could be automated or might be an informal but nonetheless informed approximation, as in case study given later. In the case of the examples that follow, however, the transition rules are formally stated and effectively define the applications in question.

### 3.1 The 'agreement' example

When visiting a website, a user must first agree to the use of cookies, then agree to the terms and conditions, and only then are they given access to the welcome page. No assumptions can be made about the order in which the pages are requested, however. In order to enforce the correct flow, two nonce cookies are set when the user agrees to the relevant missives. An error page is shown if the user requests the pages out of order, say, or bookmarks a page in order to come back to it at a later stage when the cookies have been deleted. The user agrees both to the use of cookies and to the terms and conditions by way of forms presented on the first two pages. In both cases, only if the user ticks the checkbox and then submits the form are they allowed to continue. Upon successful submission of the forms nonce cookies are set, and it is the presence or otherwise of these cookies which determines which pages can be shown. No distinction is made between a page requested via the address bar and form submissions with the checkboxes left unticked, in which case the input is treated as invalid and the relevant page is shown again.

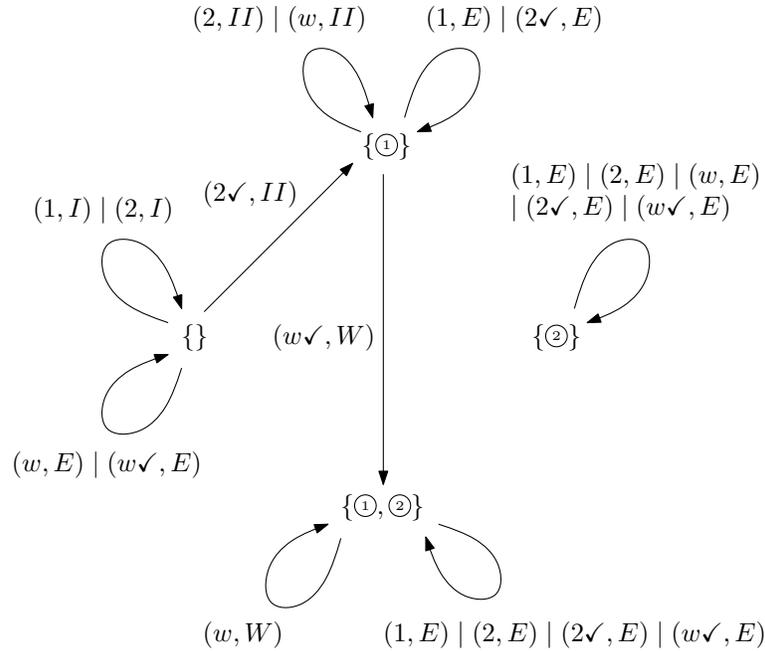


Figure 3.1: The ‘agreement’ application transition system

We model the two nonce cookies, plus the URLs that are used to request the three pages. To model valid form submissions, two additional URLs with ticks are added. We model the four page bodies returned, the first two containing the forms, then the welcome and error pages:

$$\mathcal{C} = \mathcal{P}(\{\textcircled{1}, \textcircled{2}\}) \quad \mathcal{U} = \{1, 2, w, 2\checkmark, w\checkmark\} \quad \mathcal{B} = \{I, II, W, E\}$$

The application’s behaviour is defined by the following set of transition rules:

$$\begin{aligned} \{\} &\xrightarrow[\{\} I]{\{\} 1|2} \{\} & \{\textcircled{1}\} &\xrightarrow[\{\} II]{\{\textcircled{1}\} 2|w} \{\textcircled{1}\} & \{\textcircled{1}, \textcircled{2}\} &\xrightarrow[\{\} W]{\{\textcircled{1}, \textcircled{2}\} w} \{\textcircled{1}, \textcircled{2}\} \\ & & \{\} &\xrightarrow[\{+\textcircled{1}\} II]{\{\} 2\checkmark} \{\textcircled{1}\} & \{\textcircled{1}\} &\xrightarrow[\{+\textcircled{2}\} W]{\{\textcircled{1}\} w\checkmark} \{\textcircled{1}, \textcircled{2}\} \\ \{\dots\} &\xrightarrow[\{\} E]{\{\textcircled{1}\}|\{\textcircled{2}\}|\{\textcircled{1}, \textcircled{2}\} 1} \{\dots\} & \{\dots\} &\xrightarrow[\{\} E]{\{\textcircled{2}\}|\{\textcircled{1}, \textcircled{2}\} 2} \{\dots\} & \{\dots\} &\xrightarrow[\{\} E]{\{\}|\{\textcircled{2}\} w} \{\dots\} \\ \{\dots\} &\xrightarrow[\{\} E]{\{\textcircled{1}\}|\{\textcircled{2}\}|\{\textcircled{1}, \textcircled{2}\} 2\checkmark} \{\dots\} & \{\dots\} &\xrightarrow[\{\} E]{\{\}|\{\textcircled{2}\}|\{\textcircled{1}, \textcircled{2}\} w\checkmark} \{\dots\} \end{aligned}$$

The first row represents the occasions when the user requests the correct page or submits a form in the correct order but without ticking the checkbox. The second row represents the occasions when the user correctly submits a form and a nonce cookie is set. The third and fourth rows represent the remaining occasions, on all of which error pages are shown. Note that both the client side and server side states are missing from these rules, since they are not present in this example. Note also that for brevity’s sake, the browser states are sometimes replaced by ellipses with the understanding that they can be inferred from the headers. Recall that the browser’s

state is echoed directly in the request. Finally, note that vertical bars allow several rules to be combined into one. Figure 3.1 shows the resulting transition system, which in this example is no more than a diagrammatic representation of the transition rules. The cookie headers are omitted, since they can be inferred from the initial and final states. Note that we also use the vertical bar rather than set notation, in keeping with the transition rules.

### 3.2 The ‘visitors’ example

A website maintains a counter on a landing page which is incremented every time a user requests that page for the first time. No cookies can be used and so in order to determine whether a request has originated from a new user, the request IP address is checked against the server state. If the request IP address cannot be found, the counter is incremented and the new IP address is added to the server state.

We model the lack of cookies, a single URL, a set of IP addresses and a countable set of page bodies. The server state is modelled as a subset of the set of the IP addresses paired with the states of the counter, which are precisely the natural numbers  $\mathbb{N}$ . In this case the application also has an initial state:

$$\mathcal{C} = \mathcal{P}(\{\}) \quad \mathcal{U} = \{U\} \quad \mathcal{A} = \{A_1, A_2, \dots, A_N\} \quad \mathcal{B} = \{B(n) \mid n \in \mathbb{N}\}$$

$$(a, n) \in \mathcal{S} \subseteq \mathcal{P}(\mathcal{A}) \times \mathbb{N} \quad (\{\}, 0) \in \mathcal{S}$$

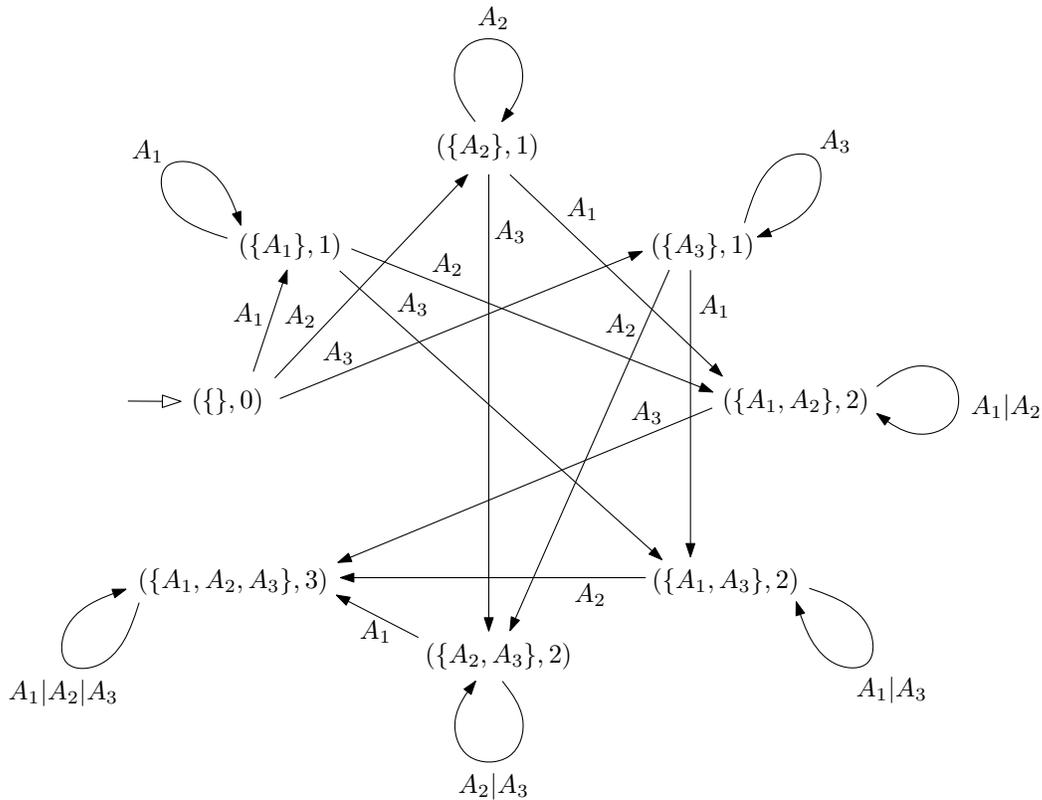


Figure 3.2: The ‘visitors’ application transition system

The application's behaviour is defined by the following two transition rules for  $1 \leq i \leq N$ :

$$(a, n) \xrightarrow[B(n)]{U, A_i} (a, n) \quad A_i \in a \qquad (a, n) \xrightarrow[B(n+1)]{U, A_i} (a \cup \{A_i\}, n+1) \quad A_i \notin a$$

Whereas the IP address is communicated alongside the URL in the request, on the other hand the counter is communicated as part of the page body. The returned page is essentially parameterised by the counter, therefore, reflecting the fact that the counter is embedded in the body of the page by some means. Note that the browser and client side states are missing from these rules, since neither is present in this example. Note also that since there is no browser state, no cookies are communicated in the request and response parts of the rules. Strictly speaking these two transition rules are not single rules but rule schemas defining two countable sets of rules for each  $(a, n) \in \mathcal{P}(\mathcal{A}) \times \mathbb{N}$ . Figure 3.2 shows the resulting transition system in the case when  $N = 3$ . The initial state is highlighted on the left with a white filled arrow. Since there is only request URL and one returned page, these are omitted. It is also assumed that the correct value of the counter is echoed in the returned page and therefore this information is not included. Note again that we use the vertical bar rather than set notation for the transition labels.

#### 4 A treatment of the examples using temporal logics

The two transition systems defined in the previous section are, with small differences, instances of Kripke transition systems, with information on both the states and transitions. In order to express properties of the applications in question we therefore use a variant of UCTL\* [19], a state/event based logic whose semantics are defined over these types of transition systems. In what follows we give the standard definition of Kripke transition systems [14], together with the syntax and semantics of UCTL\*, and then define variants of both for our purposes.

**Definition 4.1.** *A Kripke Transition System is a tuple  $(S, Act, \longrightarrow, AP, \mathcal{L})$  where:*

- $S$  is a set of states ranged over by  $s, s_0$  and  $s_1$ ,
- $Act$  is a set of actions, with  $2^{Act}$  ranged over by  $\alpha$ ,
- $\longrightarrow \subseteq S \times 2^{Act} \times S$  is the transition relation with  $(s_0, \alpha, s_1) \in \longrightarrow$ ,
- $AP$  is a set of atomic propositions ranged over by  $p$ ,
- $\mathcal{L}: S \times AP \longrightarrow \{true, false\}$  is an interpretation function that associates a value of true or false with each  $p \in AP$  for each  $s \in S$ ,
- For any two transitions,  $(s_0, \alpha_0, s_1), (s_0, \alpha_1, s_1) \in \longrightarrow \Rightarrow \alpha_0 = \alpha_1$ .

Note that since transitions carry sets of actions and not just one, there is no silent action, with the empty set  $\{\}$  being considered silent instead. Note also that we limit the number of transitions between any two states in any one direction to at most one.

Paths are sequences of transitions where the final state of one transition equals the initial state of the next, if there is one. They are ranged over by  $\sigma, \sigma'$  and  $\sigma''$ . Maximal paths are either infinite or their last state has no outgoing transitions. For the set of maximal paths starting at state  $s$  we write  $\mu path(s)$ . For the initial and final states of the first transition of a path  $\sigma$  we write  $\mathcal{I}\sigma$  and  $\mathcal{F}\sigma$ , respectively, and we write  $\mathcal{A}\sigma$  for the set of actions of the first transition of a path  $\sigma$ . A suffix  $\sigma'$  of a path  $\sigma$  is a path such that  $\sigma = \sigma''\sigma'$  for some possibly zero length path  $\sigma''$ . A proper suffix  $\sigma'$  of a path  $\sigma$  is a path such that  $\sigma = \sigma''\sigma'$  for some non-zero length path  $\sigma''$ . We write  $\sigma \leq \sigma'$  when  $\sigma'$  is a suffix of  $\sigma$  and  $\sigma < \sigma'$  when  $\sigma'$  is a proper suffix of  $\sigma$ .

**Definition 4.2.** *The syntax UCTL\* is:*

$$\phi ::= p \mid \neg\phi \mid \phi \wedge \phi' \mid \exists\pi \quad \pi ::= \phi \mid \neg\pi \mid \pi \wedge \pi' \mid X\pi \mid X_a\pi \mid \pi U\pi'$$

Here  $\phi$  is a state formula and  $\pi$  a path formula.

**Definition 4.3.** *The semantics of UCTL\* is:*

- $s \models p$  iff  $\mathcal{L}(s,p) = \text{true}$
- $s \models \neg\phi$  iff  $s \not\models \phi$
- $s \models \phi \wedge \phi'$  iff  $s \models \phi$  and  $s \models \phi'$
- $s \models \exists\pi$  iff  $\exists\sigma \in \mu\text{path}(s) : \sigma \models \pi$
- $\sigma \models \phi$  iff  ${}_S\sigma \models \phi$
- $\sigma \models \neg\pi$  iff  $\sigma \not\models \pi$
- $\sigma \models \pi \wedge \pi'$  iff  $\sigma \models \pi$  and  $\sigma' \models \pi'$
- $\sigma \models X\pi$  iff  $\exists(s, \alpha, s')\sigma'' : \sigma = (s, \alpha, s')\sigma'', \sigma'' \models \pi$
- $\sigma \models X_a\pi$  iff  $\exists(s, \alpha, s')\sigma'' : a \in \alpha, \sigma = (s, \alpha, s')\sigma'', \sigma'' \models \pi$
- $\sigma \models \pi U\pi'$  iff  $\exists\sigma' \geq \sigma : \sigma' \models \pi', \forall\sigma'' : \sigma \leq \sigma'' < \sigma' : \sigma'' \models \pi$

We now define a variant of UCTL\* for our purposes, replacing atomic propositions with propositions about typed variables.

**Definition 4.4.** *Let  $x$  be a typed variable. The type of  $x$ , written  $\tau_x$ , is the set of its permitted values. A set of typed variables is written  $X$ .*

**Definition 4.5.** *An interpretation of a set of typed variables  $X$  is a function  $I$  from  $X$  to  $\cup\{\tau_x \mid x \in X\}$  such that  $I(x) \in \tau_x$  for each  $x \in X$ . The set of all interpretations for a given set of typed variables  $X$  is written  $\mathbb{I}$ .*

**Definition 4.6.** *An (amended) Kripke Transition System a tuple  $(S, \text{Act}, \longrightarrow, X, \mathcal{L})$  where the set of atomic propositions  $AP$  and interpretation function  $\mathcal{L}$  replaced with the following:*

- $X$  is a set of typed variables,
- $\mathcal{L} : S \longrightarrow \mathbb{I}$  is an interpretation function that maps every state to an interpretation.

Action formulae remain the same, however state formulae must be amended in order to accommodate propositions of the form  $x = x'$  where  $x$  and  $x'$  are of the same type, or  $x = v$  where  $v \in \tau_x$ . In the spirit of [5], we also introduce a set  $K$  of typed global variables, extending the syntax to include propositions of the form  $x \leq k$  where  $k \in K$ ,  $x$  and  $k$  are of the same type, the type is a numerical type that supports inequalities and the value of  $k$  is in  $\tau_x$ .

**Definition 4.7.** *The (amended) syntax of UCTL\* state formulae  $\phi$  is:*

$$\phi ::= \text{true} \mid x = v \mid x = x' \mid x \leq k \mid \neg\phi \mid \phi \wedge \phi' \mid \exists\pi$$

**Definition 4.8.** *The (amended) semantics of UCTL\* state formulae  $\phi$  is:*

- $s \models x = v$  iff  $\mathcal{L}(s)(x) = v$
- $s \models x = x'$  iff  $\mathcal{L}(s)(x) = \mathcal{L}(s)(x')$
- $s \models x \leq k$  iff  $\mathcal{L}(s)(x) \leq k$

The relation  $s \models p$  is dropped with the other relations remaining.

If  $x$  has a numerical type, we also permit expressions of the form  $x = v + 1$  with the obvious semantics. Unlike [5], which introduces typed constants, we might close formulae containing typed global variables under existential or universal quantification. If they are not closed in such a way, we consider them constant. We also quantify over sets of constants.

Finally, we derive the “eventually”  $F$  and “always”  $G$  operators [8] in the usual way and add a  $T$  operator, the latter being a convenient shorthand for combining path formulae satisfied immediately and “nexttime” on a path:

$$F\pi' = \text{true}_{\text{true}}U_{\text{true}}\pi' \quad G\pi = \neg F\neg\pi \quad \pi T_\chi\pi' = \pi \wedge X_\chi\pi'$$

#### 4.1 The ‘agreement’ example

Here  $X = \{c\}$  and  $\tau_c = \{\{\}, \{\textcircled{1}\}, \{\textcircled{2}\}, \{\textcircled{1}, \textcircled{2}\}\}$ . For the sake of readability, formulae  $x = v$  are shortened to just  $v$ . For example,  $c = \{\textcircled{1}\}$  is written  $\{\textcircled{1}\}$ . The ideal user journey is the following path:

$$\{\} \xrightarrow{(1, I)} \{\} \xrightarrow{(2\checkmark, II)} \{\textcircled{1}\} \xrightarrow{(w\checkmark, W)} \{\textcircled{1}, \textcircled{2}\}$$

It is easy to produce a path with the same requests but the wrong responses, however:

$$\{\textcircled{1}, \textcircled{2}\} \xrightarrow{(1, E)} \{\textcircled{1}, \textcircled{2}\} \xrightarrow{(2\checkmark, E)} \{\textcircled{1}, \textcircled{2}\} \xrightarrow{(w\checkmark, E)} \{\textcircled{1}, \textcircled{2}\}$$

**Lemma 4.1.** *The following two formulae can be satisfied:*

$$\exists(\{\}T_{(1, I)}\{\}T_{(2\checkmark, II)}\{\textcircled{1}\}T_{(w\checkmark, W)}\{\textcircled{1}, \textcircled{2}\}) \quad \exists(X_{(1, E)}X_{(2\checkmark, E)}X_{(w\checkmark, E)}true)$$

*Proof.* These follow immediately from inspecting figure 3.1, which includes the above paths.  $\square$

#### 4.2 The ‘visitors’ example

Here  $X = \{a, n\}$ ,  $\tau_a = \mathcal{P}(\mathcal{A})$ ,  $\tau_n = \mathbb{N}$  and  $K = \{N_1, N_2, N\}$  where  $N_1, N_2, N \in \mathbb{N}$  and  $N_1, N_2 \leq N$ . We continue to use an ordered pair  $(a, n)$  for the state rather than a set.

**Lemma 4.2.** *The only state with  $n = 0$  and  $a = \{\}$  is the initial state.*

*Proof.* Only the second rule can be applied to the initial state, and there is no rule for subsequently decreasing  $n$  or reducing  $a$ .  $\square$

**Lemma 4.3.** *Apart from the initial state, all reachable states are of the form  $(a \cup \{A_i\}, n + 1)$ , where  $A_i \notin a$ , and are reached by way of a transition from state  $(a, n)$ .*

*Proof.* The first transition rule leaves states unchanged, therefore if a state is not the initial state it must be reached by the application of the second transition rule to some state, say  $(a, n)$ , and therefore must be of the form  $(a \cup \{A_i\}, n + 1)$ .  $\square$

**Lemma 4.4.** *For any reachable state  $(a, n)$ ,  $|a| = n$ .*

*Proof.* By induction on  $n$ . For  $n = 0$  we know from lemma 4.2 that the only state with  $n = 0$  is the initial state. Assume now that for all states  $(a, k)$  we have  $|a| = k$ . Consider some other, reachable state with  $n = k + 1$  which by lemma 4.3 we can write as  $(a \cup \{A_i\}, k + 1)$ , reachable from the state  $(a, k)$ . By the induction hypothesis  $|a| = k$  and therefore  $|a \cup \{A_i\}| = k + 1$ .  $\square$

**Lemma 4.5.** *For any reachable state  $(a, n)$ ,  $n \leq N$ . Formally,  $\forall G(n \leq N)$  is valid.*

*Proof.* We simply observe that  $n = |a| \leq |\mathcal{A}| = N$ .  $\square$

We omit the proof that there is a path that leads to a state  $(a, n)$  for any  $a \in \mathcal{P}(\mathcal{A})$  but note that this fact, in tandem with lemma 4.4, defines the set  $\mathcal{S}$  of all reachable states.

**Lemma 4.6.** *The value of the counter returned to the user is no greater than  $N$ .*

*Proof.* For the first rule, the response is  $B(n)$  with  $n \leq N$ . For the second rule, the response is  $B(n + 1)$  but  $n = |a|$  by lemma 4.4,  $a \subset \mathcal{A}$  and since  $|\mathcal{A}| = N$ ,  $n < N$  and therefore  $n + 1 \leq N$ .  $\square$

**Lemma 4.7.** *No user can increment the counter twice. Formally, the following is not satisfiable:*

$$\exists A_i, N_1, N_2 : \exists(((n = N_1)T_{A_i}(n = N_1 + 1)) \wedge F((n = N_2)T_{A_i}(n = N_2 + 1)))$$

*Proof.* Suppose that the formula is satisfied. The path in question must be of the following form:

$$(a_1, N_1) \xrightarrow{A_i} (a_1 \cup \{A_i\}, N_1 + 1) \dots (a_2, N_2) \xrightarrow{A_i} (a_2 \cup \{A_i\}, N_2 + 1)$$

We first note that  $a_1 \cup \{A_i\} \subseteq a_2$ , since all transitions are governed by two transition rules, one of which leaves the set  $a$  of any state  $(a, n)$  alone and the other of which adds an element. Then we note that in order for the second rule to be applied at the end of the path we must have  $A_i \notin a_2$ . However,  $a_2 \supseteq a_1 \cup \{A_i\}$  and  $A_i \in a_1 \cup \{A_i\}$  implies  $A_i \in a_2$ , a contradiction.  $\square$

## 5 The QuICDoc case study

QuICDoc [3] is a small, distributed, rich Internet application that uses the Concur algorithm [18] to merge concurrent changes to a shared document. A user simply types into a textarea and a client-side part of the application communicates the changes, what are called “diffs”, to the server side part, which disseminates them to all the other client side parts after amending them in such a way as to ensure that they can be merged consistently. In [18] it is proved that this is possible but no account is taken of the effects of implementing the algorithm in an Internet setting. We bridge this gap in this section.

We consider just two clients, defining the global state as a tuple  $(j_1, j_2, s) \in \mathcal{J}_1 \times \mathcal{J}_2 \times \mathcal{S}$  encompassing both client side states alongside the server side state. In order to relate the formal model to the implementation we define states as named tuples with values associated with the relevant heap variables, thus  $j_i(\text{gUid}, \text{gWorkingDoc}, \text{gTempDoc})$  and  $s(\text{gLastUid}, \text{gDocument}, \text{gDiffss})$  for  $i = 1, 2$ . The transition rules are broken down into client side and server side rules, and are obtained from the source code [3] by manual static analysis:

$$\begin{aligned} j_i(\epsilon, \epsilon, \epsilon) &\xrightarrow[\text{uid}, \text{doc}]{\text{GET\_DOC}} j_i(\text{uid}, \text{doc}, \epsilon) \\ j_i(\text{uid}, \text{doc}, \text{temp}) &\xrightarrow[\text{diffs}]{\text{GET\_DIFFS}, \text{uid}} j_i(\text{uid}, \text{applyDiffs}(\text{doc}, \text{diffs}), \text{temp}) \\ j_i(\text{uid}, \text{doc}, \text{temp}) &\xrightarrow{\text{PUT\_DIFFS}, \text{uid}, \text{makeDiffs}(\text{doc}, \text{temp})} j_i(\text{uid}, \text{temp}, \text{temp}) \\ s(\text{luid}, \text{doc}, \text{diffss}) &\xrightarrow[\text{luid} + 1, \text{doc}]{\text{GET\_DOC}} s(\text{luid} + 1, \text{doc}, \text{resetDiffs}(\text{diffss}, \text{uid})) \\ s(\text{luid}, \text{doc}, \text{diffss}) &\xrightarrow[\text{getDiffs}(\text{diffss}, \text{uid})]{\text{GET\_DIFFS}, \text{uid}} s(\text{luid}, \text{doc}, \text{resetDiffs}(\text{diffss}, \text{uid})) \\ s(\text{luid}, \text{doc}, \text{diffss}) &\xrightarrow{\text{PUT\_DIFFS}, \text{uid}, \text{diffs}} s(\text{luid}, \text{amendDoc}(\text{doc}, \text{uid}, \text{diffs}), \text{amendDiffss}(\text{diffss}, \text{uid}, \text{diffs})) \end{aligned}$$

Here  $\epsilon$  represents an undefined value and a side condition for all the rules is that  $\text{uid} \neq \epsilon$ . We note in passing that breaking the transition rules in this way provides a means of partially specifying the behaviour of each part of the application, and hence taken together can be considered as a

partial specification of the application as a whole. To continue, the application has an initial global state, which is again broken down into client side and server side parts, thus  $j_i(\epsilon, \epsilon, \epsilon)$  and  $s(0, [[], []], \text{“”})$  for  $i = 1, 2$ . Here [...] represents an array.

To make use of these rules we recombine them. To do this we note that the client side part produces a request and consumes a response whilst the server side part consumes a request and produces a response. We therefore replace the elements of the request part of the server side rules with matching elements of the client side rules and vice versa. We implement this by collecting the elements of the requests and responses into tuples, with elements being either variables, constants or terms. When collecting elements we add the relevant subscripts, for example the  $uid$  variable in client side rules becomes  $uid_i$ , and so on. We apply a partial matching function recursively in a similar vein to [9], with a substitution being returned if a match can be found.

**Definition 5.1.** *The partial matching function for pairing transition rules is defined by structural induction by means of the following partial functions:*

$$\begin{aligned} \text{match}((e_1, e_2 \dots), (e'_1, e'_2 \dots)) &= \{\text{match}(e_1, e'_1)\} \cup \text{match}((e_2 \dots), (e'_2 \dots)) \\ \text{match}(v, v') &= v/v & \text{match}(c, c) &= \emptyset \\ \text{match}(v, c') &= v/c' & \text{match}(( ), ( )) &= \{ \} \\ \text{match}(v, t') &= v/t' \end{aligned}$$

Here  $e_1, e'_1$ , etc are arbitrary elements,  $v, v'$  variables,  $c, c'$  constants and  $t'$  terms.

To match the requests, we make the server side tuple the first argument of the matching function since it is this tuple that contains the variables that must be replaced in the substitutions that follow. To match the responses, we make the client side tuple the first argument. By way of an example we show the matching of the requests and responses of the GET\_DIFFS rules:

$$\begin{aligned} &\text{match}((\text{GET\_DIFFS}, uid_s), (\text{GET\_DIFFS}, uid_i)) && \text{match}((diffs_i), (\text{getDiffs}(diffs_s, uid_s)_s)) \\ = &\{\emptyset\} \cup \text{match}((uid_s), (uid_i)) && = \{\text{match}(diffs_i, \text{getDiffs}(diffs_s, uid_s)_s)\} \cup \text{match}(( ), ( )) \\ = &\{\emptyset\} \cup \{\text{match}(uid_s, uid_i)\} && = \{diffs_i/\text{getDiffs}(diffs_s, uid_s)_s\} \\ = &\{\emptyset\} \cup \{uid_s/uid_i\} \cup \text{match}(( ), ( )) \\ = &\{\emptyset, uid_s/uid_i\} \end{aligned}$$

Here the rules are paired as expected by matching the constants. If the requests and responses of two rules match, we apply the resulting substitutions to obtain a global rule. Here we show just the part of the GET\_DIFFS rule corresponding to the transaction, omitting the global states:

$$\frac{\{diffs_i/\text{getDiffs}(diffs_s, uid_s)_s\} \cdot \{\emptyset, uid_s/uid_i\} \cdot (\text{GET\_DIFFS}, uid_s)}{\{\emptyset, uid_s/uid_i\} \cdot \{diffs_i/\text{getDiffs}(diffs_s, uid_s)_s\} \cdot (diffs_i)} \xrightarrow{\text{GET\_DIFFS}, uid_i} \text{getDiffs}(diffs_s, uid_i)_s$$

Note that the terms may contain elements that were not matched but are nevertheless subject to substitutions at this stage. The process is not a recursive one, however, since each set of substitutions only needs to be applied once.

### 5.1 A temporal logic with parameterised action formulae

Consider the proof of lemma 6.1 in [18]. To begin with we note that the assumption that the client cannot put  $diffs$  on the server before getting the document is intrinsic to the client side rules, since  $uid \neq \epsilon$  and the third rule cannot therefore be employed before the first. Aside from this assumption, the proof relies on two facts. The first is that the document is communicated from server to client intact, which we express as a temporal formula.

**Lemma 5.1.**  $X_{((\text{GET\_DOC}), (uid, doc))} (doc_i = doc_s)$  is satisfiable.

Here the action formula becomes a pair of tuples similar to those used when combining transition rules. We borrow the subscript notation from the global rules, too. In fact in order to satisfy this formula we must match it against these global transition rules, since we do not have the transition system to hand. The global rule in question is the following:

$$(j_i(\epsilon, \epsilon, \epsilon), s(luid, doc, diffss)) \xrightarrow[(luid+1)_s, doc_s]{\text{GET\_DOC}} (j_i((luid+1)_s, doc_s, \epsilon), s(luid+1, doc, diffss))$$

We again collect the elements of the requests and responses in the transition rule into tuples, applying a matching function recursively to both in order to match the formula against the rule.

**Definition 5.2.** The partial matching function for matching formulae against transition rules is defined by structural induction by means of the following partial functions:

$$\begin{aligned} \text{match}((e_1, e_2 \dots), (e'_1, e'_2 \dots)) &= \{\text{match}(e_1, e'_1)\} \cup \text{match}((e_2 \dots), (e'_2 \dots)) \\ \text{match}(v, v') &= \emptyset & \text{match}(\underline{v}, c') &= v/c' \\ \text{match}(v, c') &= \emptyset & \text{match}(\underline{v}, t') &= v/t' \\ \text{match}(v, t') &= \emptyset & \text{match}(\underline{v}, \underline{v}') &= \{v\} \\ \text{match}(c, c') &= \emptyset & \text{match}(\underline{v}, \underline{v}') &= \{v\} \end{aligned}$$

Here  $e_1, e'_1$ , etc are arbitrary elements,  $v, v'$  variables,  $c, c'$  constants,  $t'$  terms and  $\underline{v}$  binding variables.

The aforementioned temporal formula and transition rule do indeed match, returning the substitutions  $\{\emptyset\}$  and  $\{\emptyset, \emptyset\}$ , and since  $j_i((luid+1)_s, doc_s, \epsilon) \models (doc_i = doc_s)$ , the formula is satisfied.

The second fact is that when diffs are put on the server, they are suitably amended and then applied to the server's document. Transition rules allow the values of variables in the final state to be calculated from those in the initial state. In order to capture this mechanism in path formulae we must therefore save the value of variables in the initial state for use in the final path formula somehow. We do this by way of a universally quantified global variables, the values of which we fix in the initial proposition in the temporal formula.

**Lemma 5.2.**  $\forall d : (doc_s = d) T_{((\text{PUT\_DIFFS}, uid, diff_s), ())} (doc_s = \text{amendDoc}(d, uid, diff_s))$  is satisfiable.

In order to satisfy this formula we match it against the following global rule, which shows the relevant parts of the server side state only:

$$s(doc) \xrightarrow{\text{PUT\_DIFFS}, uid_i, \text{makeDiffs}(doc_i, temp_i)_i} s(\text{amendDoc}(doc, uid_i, \text{makeDiffs}(doc_i, temp_i)_i))$$

The match gives the substitution  $\{uid/uid_i, diff_s/\text{makeDiffs}(doc_i, temp_i)_i\}$ , which results in the final proposition:

$$(doc_s = \text{amendDoc}(d, uid_i, \text{makeDiffs}(doc_i, temp_i)_i))$$

A final state  $s(\text{amendDoc}(d, uid_i, \text{makeDiffs}(doc_i, temp_i)_i))$  results from replacing the server side variable  $doc$  with the global variable  $d$ , and the formula is satisfied.

In order to make use of substitutions we have effectively replaced the  $X_\chi \phi$  operator with an  $X_\zeta \phi$  operator, where  $\chi$  is a parameterised action formula.

**Definition 5.3.** The semantics of the  $X_\zeta \phi$  operator is:

- $\sigma \models X_\zeta \phi$  iff  $\sigma = \sigma' \sigma''$  with  $\sigma'_\top \models \zeta \triangleright \rho$  and  $\rho.(\sigma'') \models \phi$

We now formally define the syntax and semantics of parameterised action formulae.

**Definition 5.4.** *The syntax of parameterised action formulae is:*

$$\begin{array}{ll}
\zeta ::= \underline{a} \mid \chi & \chi ::= \tau \mid a \mid \neg\chi \mid \chi \wedge \chi' \\
\underline{a} ::= (\text{request}, \text{response}) & a ::= (\text{request}, \text{response}) \\
\text{request}, \text{response} ::= (\underline{e} \dots) & \text{request}, \text{response} ::= (e \dots) \\
\underline{e} ::= \underline{v} \mid v \mid c \mid t & e ::= v \mid c \\
t ::= \text{name}(e \dots) & \\
e ::= v \mid c &
\end{array}$$

**Definition 5.5.** *The semantics of parameterised action formulae is:*

- $\alpha \models \underline{a} \triangleright \rho$  iff  $\exists a \in \alpha : \text{match}(\underline{a}, a) = \rho$
- $\alpha \models \chi \triangleright \emptyset$  iff  $\alpha \models \chi$

*The satisfaction relations for  $\alpha \models \chi$  otherwise remain the same.*

We conclude this section with two further temporal formulae, the first expresses the fact that when diffs are put on the server, all pending diffs are suitably amended, the second expresses the fact that when pending diffs are gotten by a client, they are applied to the client's document. The proofs are omitted to save space, but follow along the lines of the ones already given.

**Lemma 5.3.** *The following formulae are satisfiable:*

$$\begin{aligned}
& \forall d : (\text{diffs}_s = d) T_{((\text{PUT\_DIFFS}, \text{uid}, \text{diffs}), ())} (\text{diffs}_s = \text{amendDiffs}(d, \text{uid}, \text{diffs})_s) \\
& X_{((\text{GET\_DIFFS}, \text{uid}_i), (\underline{\text{diffs}}))} (\text{applyDiffs}(\text{doc}_i, \text{diffs})_i = \text{doc}_s)
\end{aligned}$$

To satisfy the second of these formula we note that the previous results justify the equivalences  $s * \delta * \delta'' = \text{doc}_i$ ,  $\delta'' . \Delta' = \text{getDiffs}(\text{uid}_i)_s$  and  $s' * \Delta' . \delta'' = \text{doc}_s$ . Thus the path formula becomes  $s' * \Delta' . \delta'' = s * \delta * \delta'' * \delta'' . \Delta'$ , which holds by lemma 6.2 in [18].

## 6 Related work

The approach taken in [4] formalises significant parts of the HTTP protocol as well as key web concepts such as the browser and the server. A concept of non-linear time is adopted in order to deal with the problem of the browser's back button, with the authors not being "...concerned with the actual temporal order between unrelated actions", an approach strongly espoused here.

In [12], significant emphasis is placed on user interactions. The opinion that "...any verification tool for the web that does not account for user operations is not only incomplete, but potentially misleading" could be construed as being somewhat the opposite to our own. In [7], emphasis is placed on the structure of the web site in terms of its interrelated links, with transitions being seen as taking place between web pages. In effect a temporal order is imposed on actions, with users choosing only from a set of inputs prescribed by each web page, in marked contrast to the approach taken in [4] and our own. The variants of CTL and CTL\* developed therein, employing free and bounds variables as well universal closure, bear some resemblance to the variants defined here, however.

By eschewing a model of user interactions, the approach here bears more resemblance to approaches used in the analysis of web services rather than programs [16]. Although our approach is not a model checking one, for example, our model results in the same number of actions from each state, an exhaustive treatment that typifies model checking. Our approach also owes much

to that of [9]. The variants of the logic UCTL\* defined here borrow much from the logic SoCL, which is defined over similar transition systems to the ones defined here, there called doubly labelled transition systems [6]. In particular the concepts of binding variables and the matching of action formulae are taken directly and adopted here.

## 7 Conclusions and future work

The model of HTTP transactions and Internet application state we have defined has allowed us to prove properties of a rich, distributed Internet application. Because this model is founded only on the HTTP protocol and Internet application states, and not on any particular framework or programming paradigm, it could be used to prove properties of any Internet program or protocol, provided we can obtain the transition rules.

We coined the term application layer specification in the introduction but keep its precise meaning deliberately vague. It could refer to the transition systems we generate, which are useful visual aids and provide the semantics for the logics we employ. It could refer to the temporal formulae, which capture information about application states and transactions. Or it could refer to the transition rules which, broken down into rules for each part of the application, provide a useful starting point for design and specification. All of these concepts provide insight into the behaviour of Internet applications.

The case study chosen was a relatively simple rich Internet application, with both server side and client side parts written in JavaScript. Our long term goal is more ambitious, however. In future work we plan to use application layer specifications to formalise properties of larger HTTP-based programs and protocols with, specifically, the Open ID Connect protocol in mind [1]. In future we plan to develop a more realistic model of HTTP transactions, too. In particular we need to break the rule that they are inviolable and label transitions with either requests or responses but not necessarily both at once.

The choice of the logic UCTL\* we motivated by the need for formulae with nested next time operators but, rather than defining the  $\phi T_\chi \phi'$  operator as shorthand for  $\phi \wedge X_\chi \phi'$ , we could define its semantics directly and otherwise do away with the need for a fully branching logic.

Finally, we make some comments on the use of terms in the transition rules in the case study, and the fact that temporal formulae are satisfied against these rules rather than the transition system that they generate. Both approaches are somewhat informal, however a formal treatment of the former can be found in [17], which includes transition systems generated by rules, which are gratifyingly called transition system specifications. Moreover, semantics that encompasses both the denotational meaning of terms and their operational meaning over transition systems are developed therein, indeed these two semantics are shown to be equivalent by means of the concept of mixed terms. For example, the specification  $\mathcal{R}_P$  of a system is formally defined as:

$$\mathcal{R}_P = \mathcal{R} \cup \{p \xrightarrow{a} q \mid \langle a, q \rangle \in \psi(p)\}$$

Here  $\mathcal{R}$  are the rules and  $\psi: P \rightarrow \mathcal{P}_f(A \times P)$  is a transition system defined as a coalgebra, hence both the rules and all possible transitions are considered as axioms. We conclude therefore by looking forward to a full formal treatment based on this approach.

### 7.1 Acknowledgements

The author thanks the anonymous reviewers for their many helpful comments and corrections.

## References

- [1] *OpenID Connect*. Available at <http://openid.net/connect/>.
- [2] *PCI SSC Data Security Standards Overview*. Available at [https://www.pcisecuritystandards.org/security\\_standards/](https://www.pcisecuritystandards.org/security_standards/).
- [3] *QuICDoc*. Available at <http://www.doc.ic.ac.uk/~jes204/quicdoc.zip>.
- [4] Devdatta Akhawe, Adam Barth, Peifung E. Lam, John C. Mitchell & Dawn Song (2010): *Towards a Formal Foundation of Web Security*. In: *CSF*, pp. 290–304. Available at <http://doi.ieeecomputersociety.org/10.1109/CSF.2010.27>.
- [5] Jürgen Bohn, Werner Damm, Orna Grumberg, Hardi Hungar & Karen Laster (1998): *First-Order-CTL Model Checking*. In: *FSTTCS*, pp. 283–294. Available at <http://dx.doi.org/10.1007/b71635>.
- [6] Rocco De Nicola & Frits W. Vaandrager (1990): *Three Logics for Branching Bisimulation (Extended Abstract)*. In: *LICS*, pp. 118–129.
- [7] Alin Deutsch, Liying Sui & Victor Vianu (2007): *Specification and verification of data-driven Web applications*. *J. Comput. Syst. Sci.* 73(3), pp. 442–474. Available at <http://dx.doi.org/10.1016/j.jcss.2006.10.006>.
- [8] E. Allen Emerson & Joseph Y. Halpern (1983): *"Sometimes" and "not never" revisited: on branching versus linear time (preliminary report)*. In: *Proceedings of the 10th ACM SIGACT-SIGPLAN symposium on Principles of programming languages, POPL '83*, ACM, New York, NY, USA, pp. 127–140, doi:10.1145/567067.567081. Available at <http://doi.acm.org/10.1145/567067.567081>.
- [9] Alessandro Fantechi, Stefania Gnesi, Alessandro Lapadula, Franco Mazzanti, Rosario Pugliese & Francesco Tiezzi (2012): *A logical verification methodology for service-oriented computing*. *ACM Trans. Softw. Eng. Methodol.* 21(3), p. 16. Available at <http://doi.acm.org/10.1145/2211616.2211619>.
- [10] Philippa Gardner, Sergio Maffei & Gareth David Smith (2012): *Towards a program logic for JavaScript*. In: *POPL*, pp. 31–44. Available at <http://doi.acm.org/10.1145/2103656.2103663>.
- [11] Zef Hemel, Danny M. Groenewegen, Lennart C. L. Kats & Eelco Visser (2011): *Static consistency checking of web applications with WebDSL*. *J. Symb. Comput.* 46(2), pp. 150–182. Available at <http://dx.doi.org/10.1016/j.jsc.2010.08.006>.
- [12] Daniel R. Licata & Shriram Krishnamurthi (2004): *Verifying Interactive Web Programs*. In: *ASE*, pp. 164–173. Available at <http://doi.ieeecomputersociety.org/10.1109/ASE.2004.10054>.
- [13] Michael C. Martin, V. Benjamin Livshits & Monica S. Lam (2005): *Finding application errors and security flaws using PQL: a program query language*. In: *OOPSLA*, pp. 365–383. Available at <http://doi.acm.org/10.1145/1094811.1094840>.
- [14] Markus Müller-Olm, David A. Schmidt & Bernhard Steffen (1999): *Model-Checking: A Tutorial Introduction*. In: *SAS*, pp. 330–354. Available at [http://dx.doi.org/10.1007/3-540-48294-6\\_22](http://dx.doi.org/10.1007/3-540-48294-6_22).
- [15] Lawrence C. Paulson (1999): *Inductive Analysis of the Internet Protocol TLS*. *ACM Trans. Inf. Syst. Secur.* 2(3), pp. 332–351. Available at <http://doi.acm.org/10.1145/322510.322530>.
- [16] Anders P. Ravn, Jiri Srba & Saleem Vighio (2010): *A Formal Analysis of the Web Services Atomic Transaction Protocol with UPPAAL*. In: *ISoLA (1)*, pp. 579–593. Available at [http://dx.doi.org/10.1007/978-3-642-16558-0\\_47](http://dx.doi.org/10.1007/978-3-642-16558-0_47).
- [17] Jan Rutten & Daniele Turi (1994): *Initial Algebra and Final Coalgebra Semantics for Concurrency*.
- [18] James Smith (2013): *Concur - An Algorithm for Merging Concurrent Changes without Conflicts*. Available at <http://arxiv.org/abs/1303.7462>.
- [19] James Smith (2013): *State-event based versus purely Action or State based Logics*. Available at <http://arxiv.org/abs/1303.7459>.